

From ???@??? Fri Jun 05 08:19:37 1998

Received: from rhein.villa-bosch.de (whaleshark.villa-bosch.de [194.25.153.3])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id CAA02946

for <phayes@coginst.uwf.edu>; Fri, 5 Jun 1998 02:07:02 -0500 (CDT)

Received: by whaleshark.villa-bosch.de with Internet Mail Service (5.0.1457.3)
id <1KVAS1JB>; Thu, 5 Feb 1998 09:09:56 +0100

Message-ID:

<21C49639A20DD111842C0060B0684B2A2957@whaleshark.villa-bosch.de>

From: Beate Keller <Beate.Keller@kts.villa-bosch.de>

To: "'Pat Hayes'" <phayes@coginst.uwf.edu>

Subject: AW: travel information- URGENT

Date: Thu, 5 Feb 1998 09:09:53 +0100

X-Priority: 3

MIME-Version: 1.0

X-Mailer: Internet Mail Service (5.0.1457.3)

Content-Transfer-Encoding: 8bit

X-MIME-Autoconverted: from quoted-printable to 8bit by nuts.coginst.uwf.edu id
CAA02946

Content-Type: text/plain

Content-Length: 1774

Status:

Dear Pat,

sorry for not answering earlier we had some serious problems with our
e-mail.

Easiest way to come to Heidelberg from Frankfurt airport is to take the
Lufthansabus which is driving almost every hour departing directly in
front of the main entrance at the airport. The bus brings you to the
Heidelberg Renaissance Hotel and from there you take a taxi to the
Holiday Inn. At the Holiday Inn you will receive a schedule for the
Ontology Meeting and also information about transportation service to
the meeting in the morning and in the evening.

Bus from Frankfurt takes about 1-1,5 hours. Taxi from Renaissance to
Holiday Inn about 10 min.

Costs for the bus about 40 Marks, Taxi about 10 marks.

Don't hesitate to ask me if you have anymore questions.

Best regards, Beate

Klaus Tschira Stiftung

Schloss-Wolfsbrunnenweg 33

69118 Heidelberg

Tel. 06221/533-101

Fax.06221/533-199

Email: Beate.Keller@kts.villa-bosch.de

-----Ursprüngliche Nachricht-----

Von: Pat Hayes [SMTP:phayes@coginst.uwf.edu]
Gesendet am: Donnerstag, 5. Februar 1998 08:46
An: Beate Keller
Betreff: travel information- URGENT

Frankfurt Greetings. Please, can you tell me how I should get from
airport
to the Ontology meeting? I need to know approximately
how much it will
cost, and how long the trip will take, so that I can
decide exactly when
I
need to arrive in Frankfurt in order to be there on
Wednesday morning.

I would be very grateful for a quick reply

Many thanks

Pat Hayes

8903 home IHMC, University of West Florida (850)434
2091 office 11000 University Parkway (850)474
fax Pensacola, FL 32514 (850)474 3023

phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Mon Feb 16 10:51:35 1998
Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id AAA19365
for <phayes@nuts.coginst.uwf.edu>; Mon, 16 Feb 1998 00:07:38 -0600 (CST)
Received: from scratchy (scratchy [207.207.8.118])
by catbert.cyc.com (8.8.8/8.8.8) with SMTP id AAA01609;
Mon, 16 Feb 1998 00:04:24 -0600 (CST)
Message-Id: <3.0.32.19980216000523.00a79b88@catbert.cyc.com>
X-Sender: fritz@catbert.cyc.com
X-Mailer: Windows Eudora Pro Version 3.0 (32)
Date: Mon, 16 Feb 1998 00:05:28 -0600
To: phayes@nuts.coginst.uwf.edu

From: Fritz Lehmann <fritz@cyc.com>
Subject: FYI - I came across this reference
Cc: fritz@cyc.com, dmac@research.att.com
Mime-Version: 1.0
Content-Type: text/plain; charset="us-ascii"
Content-Length: 919
Status:

Dear Pat,

I stumbled on the reference below. I surmise that it either refutes, or is refuted by, or illustrates some important limit of, your proposed "recursive" method of defining the predicate of connectedness of finite structures in First-Order logic (in an email message of yours a few years ago).

I haven't seen or read the article.

```
-----  
@article{GV85,  
  author={H. Gaifman and M.~Y. Vardi},  
  title={A simple proof that connectivity is not first-order},  
  journal= {Bulletin of the European Association for  
    Theoretical Computer Science},  
  volume=26,  
  month=jun,  
  year=1985,  
  pages={43--45}  
}
```

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

```
=====  
From ???@??? Mon Apr 20 14:31:32 1998  
Received: from rhein.villa-bosch.dex (rhein.villa-bosch.de [194.25.153.3])  
  by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA05656  
  for <phayes@picayune.coginst.uwf.edu>; Mon, 20 Apr 1998 12:33:15 -0500
```

(CDT)

```
Received: from linux3.villa-bosch.de by rhein.villa-bosch.dex with SMTP (Microsoft Ex-  
change Internet Mail Service Version 5.0.1457.7)  
  id 20XCYBDJ; Mon, 20 Apr 1998 19:35:30 +0200
```

```
Received: by linux3.villa-bosch.de with Microsoft Mail
```

id <01BD6C93.D1CE1C20@linux3.villa-bosch.de>; Mon, 20 Apr 1998 19:37:57
+0100

Message-ID: <01BD6C93.D1CE1C20@linux3.villa-bosch.de>

From: "Prof. Andreas Reuter" <reuter@villa-bosch.de>

To: "AFarquhar" <Adam_Farquhar@ksl.stanfrod.edu>,

"EFeigenbaum"

<feigenbaum@KSL.Stanford.EDU>,

"JPustejovski"

<jamesp@cs.brandeis.edu>,

"JTsuji" <tsujii@is.s.u-tokyo.ac.jp>,

"MvdBerg" <vdberg@us.ibm.com>,

"PHayes" <phayes@nuts.coginst.uwf.edu>

To: "PSimons" <p.m.simons@leeds.ac.uk>

Cc: Beate Keller

</o=Klaus.Tschira.Foundation/ou=VILLABOSCH/cn=Recipients/cn=keller@villa-bosch.d
e>,

"guarino" <guarino@ladseb.pd.cnr.it>, "hovv" <hovv@isi.edu>,

"miller" <geo@clarity.princeton.edu>,

"peters" <peters@csl.stanford.edu>,

"polanyi" <polanyi@pal.xerox.com>

Cc: "self" <Andreas.Reuter@eml.villa-bosch.de>,

"sowa"

<sowa@west.poly.edu>, "spillers" <skydog@pacbell.net>,

"tschira"

<Klaus.Tschira@ktf.villa-bosch.de>,

"Vossen"

<Piek.Vossen@let.uva.nl>

Subject: Re: Ontology Workshop in Heidelberg

Date: Mon, 20 Apr 1998 19:37:56 +0100

MIME-Version: 1.0

Content-Transfer-Encoding: 8bit

X-MIME-Autoconverted: from quoted-printable to 8bit by nuts.coginst.uwf.edu id

MAA05656

Content-Type: text/plain; charset="us-ascii"

Content-Length: 2197

Status:

Dear colleagues,

let me briefly introduce myself: My name is Andreas Reuter, Scientific Director of the European Media Lab (EML) at Heidelberg. Together with our Managing Director, Klaus Tschira, I have the pleasure of hosting a workshop on ontological issues in our lab, which is located in "Villa Bosch", a beautiful old villa right behind the Heidelberg Castle. In the meantime you should have been approached by one of the colleagues mentioned in the CC-list, and from what we know, you agreed to participate.

Let me restrict this first mail to organisational issues; the technical matters will be filled in by those among the organizers who work in the field, which neither Klaus Tschira nor I do.

The workshop starts on June 10, 1998, and will go on for one week (weekend included), ending on June 16. We have booked an number of hotel rooms for that period, but since most of you will not arrive at the day before opening and leave on the last day of the workshop, we would like to know your exact travel schedules as soon as possible. If you need help in making arrangements, please let us know.

It will be a small workshop, focussing on work in small groups and plenary discussions. There will be between 20 and 25 participants.

If you need special equipment or want to suggest that certain books be kept available, please let us know in advance.

Travel expenses related to the workshop that are not picked up by your employer will be covered by Klaus-Tschira-Stiftung, the foundation supporting the EML.

For all matters related to the workshop, please use my email address. Of course, as we get closer to the workshop and things get more specific, other people from our lab will contact you as well, but you can simply use that one address. During the next couple of days we will install a web site for the workshop that will reflect the current state of preparation, both from an organizational and a technical perspective.

Please understand that because of the small size of the workshop, the whole event is strictly by invitation only.

I am glad that you agreed to accept the invitation, and I am looking forward to a stimulating and rewarding workshop.

Best wishes

Andreas Reuter

From ???@??? Mon Apr 27 13:22:50 1998

Received: from rhein.villa-bosch.de (whaleshark.villa-bosch.de [194.25.153.3])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA24001
for <phayes@picayune.coginst.uwf.edu>; Mon, 27 Apr 1998 12:27:11 -0500

(CDT)

Received: from linux3.villa-bosch.de by rhein.villa-bosch.de with SMTP (Microsoft Exchange Internet Mail Service Version 5.0.1457.7)
id JMB9JTB0; Mon, 27 Apr 1998 19:29:22 +0200

Received: by linux3.villa-bosch.de with Microsoft Mail
id <01BD7213.1FCC7500@linux3.villa-bosch.de>; Mon, 27 Apr 1998 19:31:49
+0100

Message-ID: <01BD7213.1FCC7500@linux3.villa-bosch.de>

From: "Prof. Andreas Reuter" <reuter@villa-bosch.de>

To: "AReuter" <Andreas.Reuter@eml.villa-bosch.de>,
"BSpillers"
<skydog@pacbell.net>, "EHovy" <hovy@isi.edu>,
"GMiller"
<geo@clarity.princeton.edu>,
"JSowa" <sowa@west.poly.edu>,
"KTschira" <Klaus.Tschira@ktf.villa-bosch.de>

To: "LPolanyi" <polanyi@pal.xerox.com>,

"NGuarino"
<guarino@ladseb.pd.cnr.it>,
"PVossen" <Piek.Vossen@let.uva.nl>,
"SPeters" <peters@csl.stanford.edu>,
"AFarquhar" <Adam_Farquhar@ksl.stanford.edu>,
"CFellbaum" <fellbaum@clarity.princeton.edu>
To: "DSkuce" <doug@csi.uottawa.ca>,
"EFeigenbaum"
<feigenbaum@ksl.stanford.edu>,
"FLehmann" <fritz@cyc.com>, "JMcCarthy" <jmc@cs.stanford.edu>,
"JPustejovski" <jamesp@cs.brandeis.edu>,
"JTsuji" <tsujii@is.s.u-tokyo.ac.jp>
To: "MvdBerg" <vdberg@us.ibm.com>,
"PHayes"
<phayes@nuts.coginst.uwf.edu>,
"PSimons"
<p.m.simons@leeds.ac.uk>,
"WWahlster" <wahlster@dfki.uni-sb.de>
Cc: B%rbel Mack

</o=Klaus.Tschira.Foundation/ou=VILLABOSCH/cn=Recipients/cn=mack@villa-bosch.de>,
e>

Beate Keller

</o=Klaus.Tschira.Foundation/ou=VILLABOSCH/cn=Recipients/cn=keller@villa-bosch.de>
e>

Subject: Re: Ontology Workshop

Date: Mon, 27 Apr 1998 19:31:48 +0100

MIME-Version: 1.0

Content-Transfer-Encoding: 8bit

X-MIME-Autoconverted: from quoted-printable to 8bit by nuts.coginst.uwf.edu id
MAA24001

Content-Type: text/plain; charset="us-ascii"

Content-Length: 366

Status:

Dear colleagues,

the web site is online now (URL: <http://www.ontology.villa-bosch.de>), with a minor exception to the explanation I gave before: When entering the page, you have provide the string "ontology" as the user name, independent of your real last name, and then "villa-bosch" as the password. Works here, hopefully works on your end, too.

Best wishes

Andreas

From ???@??? Thu Apr 30 00:11:11 1998

Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id VAA09642

for <phayes@nuts.coginst.uwf.edu>; Wed, 29 Apr 1998 21:51:37 -0500 (CDT)

Received: from scratchy (scratchy [207.207.8.118])
by catbert.cyc.com (8.8.8/8.8.8) with SMTP id VAA12616;
Wed, 29 Apr 1998 21:37:57 -0500 (CDT)
Message-Id: <3.0.32.19980429213914.00a2fb20@catbert.cyc.com>
X-Sender: fritz@catbert.cyc.com
X-Mailer: Windows Eudora Pro Version 3.0 (32)
Date: Wed, 29 Apr 1998 21:39:43 -0500
To: webmaster@villa-bosch.de
From: Fritz Lehmann <fritz@cyc.com>
Subject: The top 40 concepts
Cc: fritz@cyc.com, Andreas.Reuter@eml.villa-bosch.de, skydog@pacbell.net,
hovy@isi.edu, geo@clarity.princeton.edu, sowa@west.poly.edu,
Klaus.Tschira@ktf.villa-bosch.de, polanyi@pal.xerox.com,
guarino@ladseb.pd.cnr.it, Piek.Vossen@let.uva.nl,
peters@csl.stanford.edu, Adam_Farquhar@ksl.stanford.edu,
fellbaum@clarity.princeton.edu, doug@csi.uottawa.ca,
feigenbaum@ksl.stanford.edu, fritz@cyc.com, jmc@cs.stanford.edu,
jamesp@cs.brandeis.edu, tsujii@is.s.u-tokyo.ac.jp, vdberg@us.ibm.com,
phayes@nuts.coginst.uwf.edu, p.m.simons@leeds.ac.uk,
wahlster@dfki.uni-sb.de, doug@csi.uottawa.ca
Mime-Version: 1.0
Content-Type: text/plain; charset="us-ascii"
Content-Length: 601
Status:

Dear Sir or Madam,

The web page for the Heidelberg Ontology Workshop Agenda mentions Piek Vossen's "Top 40" concepts. I think it would be a good idea to distribute a list of the Top 40, with any existing definitions, comments and axioms, to all Workshop participants now. That way, we can give these 40 concepts some specific consideration in advance.

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

=====
=====

From ???@??? Fri May 01 11:07:43 1998

Received: from cclsun01.let.uva.nl (cclsun01.let.uva.nl [145.18.228.21])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id HAA11332
for <phayes@nuts.coginst.uwf.edu>; Thu, 30 Apr 1998 07:57:57 -0500 (CDT)

Received: from anpisani (uva36.remote.uva.nl) by cclsun01.let.uva.nl with SMTP id AA25265

(5.67a/IDA-1.5 for <phayes@nuts.coginst.uwf.edu>); Thu, 30 Apr 1998 14:33:44 +0200

Message-Id: <199804301233.AA25265@cclsun01.let.uva.nl>

From: "piek vossen" <piek.vossen@let.uva.nl>

To: "John F. Sowa" <sowa@west.poly.edu>, <fritz@cyc.com>, <webmaster@villa-bosch.de>

Cc: <Adam_Farquhar@ksl.stanford.edu>, <Andreas.Reuter@eml.villa-bosch.de>, <Klaus.Tschira@ktf.villa-bosch.de>, <doug@csi.uottawa.ca>, <feigenbaum@ksl.stanford.edu>, <fellbaum@clarity.princeton.edu>, <geo@clarity.princeton.edu>, <guarino@ladseb.pd.cnr.it>, <hovy@isi.edu>, <jamesp@cs.brandeis.edu>, <jmc@cs.stanford.edu>, <p.m.simons@leeds.ac.uk>, <peters@csl.stanford.edu>, <phayes@nuts.coginst.uwf.edu>, <polanyi@pal.xerox.com>, <skydog@pacbell.net>, <tsujii@is.s.u-tokyo.ac.jp>, <vdberg@us.ibm.com>, <wahlster@dfki.uni-sb.de>

Subject: Re: The top 40 concepts

Date: Thu, 30 Apr 1998 14:27:48 +0200

X-Msmail-Priority: Normal

X-Priority: 3

X-Mailer: Microsoft Internet Mail 4.70.1155

Mime-Version: 1.0

Content-Transfer-Encoding: 7bit

Content-Type: text/plain

Content-Length: 815

Status:

First of all, I would like to stress that the selection is more than 40 concepts. I will make a nice list and specification of it which can be put on the web. However, it is a national holiday here until Wednesday next week (May 6th). I cannot access the machine with the data until then. I will provide the data and a description by the end of next week.

Piek.

> Onderwerp: Re: The top 40 concepts

> Datum: donderdag 30 april 1998 12:12

>

> The Top 40 concepts were circulated by email a while ago, but it would be

> much more convenient if they (and all other related info) were available

> on a web site. That could be the site for Villa Bosch, or it could be on

> some other site, but all the sites that are relevant to this workshop

> should be accessible via links from the Villa Bosch site.

>

> John Sowa

From ???@??? Fri May 01 11:07:47 1998

Received: from rhein.villa-bosch.de (whaleshark.villa-bosch.de [194.25.153.3])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id FAA15554

for <phayes@picayune.coginst.uwf.edu>; Fri, 1 May 1998 05:54:00 -0500 (CDT)

Received: from linux3.villa-bosch.de by rhein.villa-bosch.de with SMTP (Microsoft Exchange Internet Mail Service Version 5.0.1457.7)

id J5MM94TW; Fri, 1 May 1998 12:56:29 +0200

Received: by linux3.villa-bosch.de with Microsoft Mail

id <01BD7500.E75EA030@linux3.villa-bosch.de>; Fri, 1 May 1998 12:58:57

+0100

Message-ID: <01BD7500.E75EA030@linux3.villa-bosch.de>

From: "Prof. Andreas Reuter" <reuter@villa-bosch.de>

To: "AReuter" <Andreas.Reuter@eml.villa-bosch.de>,

"BSpillers"

<skydog@pacbell.net>, "EHovy" <hovy@isi.edu>,

"GMiller"

<geo@clarity.princeton.edu>,

"JSowa" <sowa@west.poly.edu>,

"KTschira" <Klaus.Tschira@ktf.villa-bosch.de>

To: "LPolanyi" <polanyi@pal.xerox.com>,

"NGuarino"

<guarino@ladseb.pd.cnr.it>,

"PVossen" <Piek.Vossen@let.uva.nl>,

"SPeters" <peters@csl.stanford.edu>,

"AFarquhar" <Adam_Farquhar@ksl.stanford.edu>,

"CFellbaum" <fellbaum@clarity.princeton.edu>

To: "DSkuce" <doug@csi.uottawa.ca>,

"EFeigenbaum"

<feigenbaum@ksl.stanford.edu>,

"FLehmann" <fritz@cyc.com>, "JMcCarthy" <jmc@cs.stanford.edu>,

"JPustejovski" <jamesp@cs.brandeis.edu>,

"JTsuji" <tsujii@is.s.u-tokyo.ac.jp>

To: "MvdBerg" <vdberg@us.ibm.com>,

"PHayes"

<phayes@nuts.coginst.uwf.edu>,

"PSimons"

<p.m.simons@leeds.ac.uk>,

"WWahlster" <wahlster@dfki.uni-sb.de>

Cc: Susanne Winkelmann

</o=Klaus.Tschira.Foundation/ou=VILLABOSCH/cn=Recipients/cn=susanne@villa-bosch.de>

Subject: Re: Ontology Workshop - Top level concepts

Date: Fri, 1 May 1998 12:58:56 +0100

MIME-Version: 1.0

Content-Transfer-Encoding: 8bit

X-MIME-Autoconverted: from quoted-printable to 8bit by nuts.coginst.uwf.edu id
FAA15554

Content-Type: text/plain; charset="us-ascii"

Content-Length: 514

Status:

Dear colleagues,

responding to a suggestion that was made repeatedly: The list of top level concepts provided by Piek Vossen can be accessed through our workshop web site. In a recent mail, Piek has promised to prepare a "nice list"; as soon as we have it, we will put it there instead of the one we copied from his mail about a month ago. Note there are some links back and forth in the concept list; this is an attempt to document the discussions among Piek and others about those concepts.

Best wishes

Andreas

From ???@??? Mon May 04 10:10:45 1998

Received: from smtp.ontek.com (ontek.ontek.com [199.107.111.10])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id RAA18404

for <phayes@nuts.coginst.uwf.edu>; Fri, 1 May 1998 17:01:54 -0500 (CDT)

Received: from [199.107.111.211] by smtp.ontek.com with ESMTP (Eudora Internet Mail Server 2.0.1); Fri, 1 May 1998 14:21:44 -0700

X-Sender: phlpms@lucs-mac.leeds.ac.uk

Message-Id: <v0313031ab16fd54cb608@[199.107.111.211]>

Mime-Version: 1.0

Date: Fri, 1 May 1998 14:16:21 -0700

To: Andreas.Reuter@eml.villa-bosch.de, "BSpillers" <skydog@pacbell.net>,

"EHovy" <hovoy@isi.edu>, "GMiller" <geo@clarity.princeton.edu>,

"JSowa" <sowa@west.poly.edu>,

"KTschira" <Klaus.Tschira@ktf.villa-bosch.de>,

"LPolanyi" <polanyi@pal.xerox.com>,

"NGuarino" <guarino@ladseb.pd.cnr.it>,

"PVossen" <Piek.Vossen@let.uva.nl>,

"SPeters" <peters@csl.stanford.edu>,

"AFarquhar" <Adam_Farquhar@ksl.stanford.edu>,

"CFellbaum" <fellbaum@clarity.princeton.edu>,

"DSkuce" <doug@csi.uottawa.ca>,

"EFeigenbaum" <feigenbaum@ksl.stanford.edu>,

"FLehmann" <fritz@cyc.com>, "JMcCarthy" <jmc@cs.stanford.edu>,

"JPustejovski" <jamesp@cs.brandeis.edu>,

"JTsuji" <tsujii@is.s.u-tokyo.ac.jp>,

"MvdBerg" <vdberg@us.ibm.com>,

"PHayes" <phayes@nuts.coginst.uwf.edu>,

"PSimons" <p.m.simons@leeds.ac.uk>,

"WWahlster" <wahlster@dfki.uni-sb.de>

From: Peter Simons <p.m.simons@leeds.ac.uk>

Subject: Top 40 "concepts"

Content-Type: text/plain; charset="us-ascii"

Content-Length: 431

Status:

I sympathise with Pat Hayes's puzzlement about the Top 40 list. If definitions of words using other words are in question then what we have here is a partial dictionary, and people have been doing those with large teams for centuries. The OED is the best example. Other information such as synonymy is carried in dictionaries of synonyms and antonyms. Why do the work again sketchily that others have already done thoroughly?

From ???@??? Mon May 04 13:44:45 1998

Received: from rhein.villa-bosch.de (whaleshark.villa-bosch.de [194.25.153.3])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA25824
for <phayes@picayune.coginst.uwf.edu>; Mon, 4 May 1998 12:30:22 -0500

(CDT)

Received: from reuter.villa-bosch.de by rhein.villa-bosch.de with SMTP (Microsoft Exchange Internet Mail Service Version 5.0.1457.7)
id J5MM9VBW; Mon, 4 May 1998 19:32:42 +0200

Received: by reuter.villa-bosch.de with Microsoft Mail
id <01BD7792.D6D29F20@reuter.villa-bosch.de>; Mon, 4 May 1998 19:28:38
+0100

Message-ID: <01BD7792.D6D29F20@reuter.villa-bosch.de>

From: Andreas Reuter <andreas.reuter@eml.villa-bosch.de>

To: "AReuter" <Andreas.Reuter@eml.villa-bosch.de>,
"BSpillers"
<skydog@pacbell.net>, "EHovy" <hovy@isi.edu>,
"GMiller"

<geo@clarity.princeton.edu>,
"JSowa" <sowa@west.poly.edu>,
"KTschira" <Klaus.Tschira@ktf.villa-bosch.de>

To: "LPolanyi" <polanyi@pal.xerox.com>,
"NGuarino"
<guarino@ladseb.pd.cnr.it>,
"PVossen" <Piek.Vossen@let.uva.nl>,
"SPeters" <peters@csl.stanford.edu>,
"AFarquhar" <Adam_Farquhar@ksl.stanford.edu>,
"CFellbaum" <fellbaum@clarity.princeton.edu>

To: "CMenzel" <cmenzel@tamu.edu>, "DSkuce" <doug@csi.uottawa.ca>,
"EFeigenbaum" <feigenbaum@ksl.stanford.edu>,
"FLehmann" <fritz@cyc.com>, "JMcCarthy" <jmc@cs.stanford.edu>,

"JPustejovski" <jamesp@cs.brandeis.edu>
To: "JTsuji" <tsujii@is.s.u-tokyo.ac.jp>, "MvdBerg" <vdberg@us.ibm.com>,
"PHayes" <phayes@nuts.coginst.uwf.edu>,
"PSimons" <p.m.simons@leeds.ac.uk>,
"WWahlster" <wahlster@dfki.uni-sb.de>
Cc: B%rbel Mack

</o=Klaus.Tschira.Foundation/ou=VILLABOSCH/cn=Recipients/cn=mack@eml.villa-bosch.de>,
Beate Keller

</o=Klaus.Tschira.Foundation/ou=VILLABOSCH/cn=Recipients/cn=keller@eml.villa-bosch.de>,
Susanne Winkelmann

</o=Klaus.Tschira.Foundation/ou=VILLABOSCH/cn=Recipients/cn=susanne@eml.villa-bosch.de>

Subject: Re: Heidelberg Ontology Workshop - Mailing List

Date: Mon, 4 May 1998 19:28:37 +0100

MIME-Version: 1.0

Content-Transfer-Encoding: 8bit

X-MIME-Autoconverted: from quoted-printable to 8bit by nuts.coginst.uwf.edu id
MAA25824

Content-Type: text/plain; charset="us-ascii"

Content-Length: 361

Status:

Dear all,

this is to inform you about an new participant: Chris Menzel (cmenzel@tamu.edu).
The above TO-list describes the current set of participants (organizers and invited contributors). Pat Hayes asked for such a list, and we will keep it current with each additional confirmation we get. A copy of that list can be obtained from the web site.

Best

Andreas

From ???@??? Tue May 05 17:28:52 1998

Received: from mail-gw6.pacbell.net (mail-gw6.pacbell.net [206.13.28.41])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id QAA10629

for <phayes@coginst.uwf.edu>; Tue, 5 May 1998 16:40:05 -0500 (CDT)

Received: from pacbell.net (ppp-206-170-6-179.rdcy01.pacbell.net [206.170.6.179]) by
mail-gw6.pacbell.net (8.8.8/8.7.1+antispam) with ESMTP id OAA07007; Tue, 5 May
1998 14:34:56 -0700 (PDT)

Message-ID: <354F8609.F2CC1FB7@pacbell.net>

Date: Tue, 05 May 1998 14:35:05 -0700

From: Robert Spillers <skydog@pacbell.net>

X-Mailer: Mozilla 4.04 [en] (Win95; I)

MIME-Version: 1.0

To: Andreas Reuter <andreas.reuter@EML.villa-bosch.de>

CC: "W Wahlster" <wahlster@dfki.uni-sb.de>,

Beate Keller <"Beate,Keller"@kts.villa-bosch.de>,
Chris Menzel <cmenzel@tamu.edu>,
Christiane Fellbaum <fellbaum@clarity.Princeton.EDU>,
Doug Skuce <doug@site.uottawa.ca>,
Ed Feigenbaum <eaf@KSL.Stanford.edu>, Eduard Hovy <hovy@isi.edu>,
Fritz Lehmann <fritz@cyc.com>,
George Miller <geo@clarity.Princeton.EDU>,
James Pustsejovski <jamesp@cs.brandeis.edu>,
John McCarthy <jmc@cs.stanford.edu>, John Sowa <sowa@west.poly.edu>,
Junidi Tsujii <tsujii@is.s.u-tokyo.ac.jp>,
Klaus Tschira <klaus.tschira@ktf.villa-bosch.de>,
Livia Polanyi <polanyi@pal.xerox.com>,
Martin van den Berg <vdberg@let.uva.nl>,
Nicola Guarino <guarino@ladseb.pd.cnr.it>,
Pat Hayes <phayes@coginst.uwf.edu>,
Peter Simons <p.m.simons@leeds.ac.uk>,
Piek Vossen <piek.vossen@let.uva.nl>,
Stanley Peters <peters@csl.stanford.edu>

Subject: Invitation - Larry Reeker at NSF

Content-Type: multipart/alternative; boundary="-----
-883EB2917B64EA11C8EBF9AA"

Content-Length: 1037

Status:

<x-html><!x-stuff-for-pete base="" src="" id="0"><HTML>

Andreas,

I invited Larry Reeker who is Program Director of the Knowledge and
Cognitive Systems Program at the National Science Foundation. He
is heads the program at NSF that most closely corresponds to our workshop.
Larry will attend, but may not be able to attend all of the sessions.
His email address is <U>lreeker@nsf.gov </U>, telephone is 703-306-1926.

<P>Bob</HTML>

</x-html>

From ???@??? Mon May 04 10:10:48 1998

Received: from mail-gw.pacbell.net (mail-gw.pacbell.net [206.13.28.25])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id WAA19566

for <phayes@nuts.coginst.uwf.edu>; Fri, 1 May 1998 22:29:51 -0500 (CDT)

Received: from pacbell.net (ppp-206-170-7-25.rdcy01.pacbell.net [206.170.7.25]) by
mail-gw.pacbell.net (8.8.8/8.7.1+antispam) with ESMTP id UAA13818; Fri, 1 May 1998
20:25:57 -0700 (PDT)

Message-ID: <354A9256.94981C49@pacbell.net>

Date: Fri, 01 May 1998 20:26:14 -0700

From: Robert Spillers <skydog@pacbell.net>

X-Mailer: Mozilla 4.04 [en] (Win95; I)

MIME-Version: 1.0

To: Fritz Lehmann <fritz@cyc.com>, Pat Hayes <phayes@nuts.coginst.uwf.edu>

CC: Andreas Reuter <andreas.reuter@EML.villa-bosch.de>,

Ed Feigenbaum <eaf@KSL.Stanford.edu>, Eduard Hovy <hovy@isi.edu>,

George Miller <geo@clarity.princeton.edu>,

John Sowa <sowa@west.poly.edu>,

Klaus Tschira <klaus.tschira@ktf.villa-bosch.de>,

Livia Polanyi <polanyi@pal.xerox.com>,

Nicola Guarino <guarino@ladseb.pd.cnr.it>,

Piek Vossen <piek.vossen@let.uva.nl>,

Stanley Peters <peters@csl.stanford.edu>

Subject: [Fwd: WordNet treelike]

Content-Type: multipart/mixed; boundary="-----B25C8A818AFDDC0E96D2B111"

Content-Length: 3900

Status:

Fritz / Pat,

In my view a major task of the workshop is to find a way to combine extremely useful, if imperfect, linguistic structures with conceptual ontologies (which of course are of all without flaw ;-). This is the spirit of the work done by Ed Hovy and Fritz that resulted in the ANSI Reference Ontology (available on the Stanford KSL server). I hope that one of the items that result from this workshop is agreement on a clear methodology of how to create such an ontology.

Bob

Return-Path: fritz@cyc.com

Received: from mail-gw5.pacbell.net (mail-gw5.pacbell.net [206.13.28.23]) by mail-sf1.pacbell.net (8.8.8/8.7.1) with ESMTP id TAA04642 for <skydog@mail-sf1.pacbell.net>; Fri, 1 May 1998 19:27:54 -0700 (PDT)

Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5]) by mail-gw5.pacbell.net (8.8.8/8.7.1+antispam) with ESMTP id TAA10123 for <skydog@pacbell.net>; Fri, 1 May 1998 19:27:48 -0700 (PDT)

Received: from scratchy (scratchy [207.207.8.118]) by catbert.cyc.com (8.8.8/8.8.8) with SMTP id VAA20256; Fri, 1 May 1998 21:25:39 -0500 (CDT)

Message-Id: <3.0.32.19980501212716.00fc9da8@catbert.cyc.com>

X-Sender: fritz@catbert.cyc.com

X-Mailer: Windows Eudora Pro Version 3.0 (32)

Date: Fri, 01 May 1998 21:27:38 -0500

To: phayes@coginst.uwf.edu

From: Fritz Lehmann <fritz@cyc.com>

Subject: WordNet treelike

Cc: fritz@cyc.com, Andreas.Reuter@eml.villa-bosch.de, skydog@pacbell.net, hovy@isi.edu, geo@clarity.princeton.edu, sowa@west.poly.edu,

Klaus.Tschira@ktf.villa-bosch.de, polanyi@pal.xerox.com,
guarino@ladseb.pd.cnr.it, onto-std@KSL.Stanford.EDU,
Piek.Vossen@let.uva.nl, peters@csl.stanford.edu,
Adam_Farquhar@KSL.Stanford.EDU, fellbaum@clarity.princeton.edu,
doug@csi.uottawa.ca, feigenbaum@KSL.Stanford.EDU, jmc@cs.stanford.edu,
jamesp@cs.brandeis.edu, tsujii@is.s.u-tokyo.ac.jp, vdberg@us.ibm.com,
phayes@nuts.coginst.uwf.edu, p.m.simons@leeds.ac.uk,
wahlster@dfki.uni-sb.de, geo@clarity.princeton.edu

Mime-Version: 1.0

Content-Type: text/plain; charset="us-ascii"

Dear Pat,

In looking at WordNet, I've noticed that it's almost a single-inheritance tree, in which almost all synsets each have exactly one hypernym. There are a few exceptions, like the noun piano having both stringed instrument and percussion instrument as hypernyms, but this is rare. They had a half-hearted dedication to "freedom". The choice of such a tree-like structure was a mistake, in my view. I think it should have a non-tree-like poset structure (partially ordered set, DAG) with much more multiple inheritance from multiple hypernyms. Putting concepts into a tree forces you to make some silly decisions as to which of several salient superclasses should be designated as "the" hypernym.

The higher you go in WordNet, the more you get into controversial or dubious linkings; the lower levels are more obviously reliable and WordNet is quite useful in its lower reaches. And it's big.

Cyc has established over 6,000 links to WordNet (Cyc "constants" linked to WordNet "synsets") with many more to go, and I believe these cover all 3000+ of the Cyc-based, publicly released "reference ontology". (In case you don't know, those are on Cycorp's web page at <http://www.cyc.com> --- see the section on the Upper Cyc Ontology.)

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

=====
=====

From ???@??? Mon May 04 13:44:42 1998

Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA17179;
Mon, 4 May 1998 11:22:30 -0500 (CDT)
Date: Mon, 4 May 1998 11:22:30 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a06b17352946c63@[143.88.7.118]>
In-Reply-To: <354A9256.94981C49@pacbell.net>
Mime-Version: 1.0
To: Robert Spillers <skydog@pacbell.net>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: [Fwd: WordNet treelike]
Cc: Andreas.Reuter@eml.villa-bosch.de, "BSpillers" <skydog@pacbell.net>,
"EHovy" <hovy@isi.edu>, "GMiller" <geo@clarity.princeton.edu>,
"JSowa" <sowa@west.poly.edu>,
"KTschira" <Klaus.Tschira@ktf.villa-bosch.de>,
"LPolanyi" <polanyi@pal.xerox.com>,
"NGuarino" <guarino@ladseb.pd.cnr.it>,
"PVossen" <Piek.Vossen@let.uva.nl>,
"SPeters" <peters@csl.stanford.edu>,
"AFarquhar" <Adam_Farquhar@ksl.stanford.edu>,
"CFellbaum" <fellbaum@clarity.princeton.edu>,
"DSkuce" <doug@csi.uottawa.ca>,
"EFeigenbaum" <feigenbaum@ksl.stanford.edu>,
"FLehmann" <fritz@cyc.com>, "JMcCarthy" <jmc@cs.stanford.edu>,
"JPustejovski" <jamesp@cs.brandeis.edu>,
"JTsuji" <tsujii@is.s.u-tokyo.ac.jp>,
"MvdBerg" <vdberg@us.ibm.com>,
"PHayes" <phayes@nuts.coginst.uwf.edu>,
"PSimons" <p.m.simons@leeds.ac.uk>,
"WWahlster" <wahlster@dfki.uni-sb.de>
Content-Type: text/plain; charset="us-ascii"
Content-Length: 1429
Status:

>Fritz / Pat,
>In my view a major task of the workshop is to find a way to combine
>extremely useful, if imperfect, linguistic structures with conceptual
>ontologies (which of course are of all without flaw ;-). This is the
>sprit of the work done by Ed Hovy and Fritz that resulted in the ANSI
>Reference Ontology (available on the Stanford KSL server). I hope that
>one of the items that result from this workshop is agreement on a clear
>methodology of how to create such an ontology.
>

Bob, thanks for your message, but I fail to see what your point is. Fritz is probably fairly well acquainted with the Fritz/Hovy collaboration. The issue of tree-likeness (or, if you prefer, of whether or not to allow

multiple inheritance) is a key issue in this methodology, and it cuts across the linguistic/conceptual distinction. For example, is wood to be classified as a substance, or as a structural material, or as a plant product, or ...? The natural answer seems to be, all of them.

Is there any *linguistic* evidence for the word/concept heirarchy being tree-like, as opposed to having multiple routes of classification inheritance?

Pat Hayes

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Mon May 04 13:44:44 1998
Received: from LADSEB.LADSEB.PD.CNR.IT (ladseb.ladseb.pd.cnr.it [150.178.2.3])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id MAA24750
for <phayes@coginst.uwf.edu>; Mon, 4 May 1998 12:11:56 -0500 (CDT)
Received: from [150.178.2.93] by 150.178.2.93 with SMTP;
Mon, 4 May 1998 19:08:28 +0200
X-Sender: guarino@ladseb.ladseb.pd.cnr.it
Message-Id: <v0310280eb173a06738a4@[150.178.2.93]>
In-Reply-To: <v04003a06b17352946c63@[143.88.7.118]>
References: <354A9256.94981C49@pacbell.net>
Mime-Version: 1.0
Date: Mon, 4 May 1998 19:09:53 +0200
To: Pat Hayes <phayes@coginst.uwf.edu>, Robert Spillers <skydog@pacbell.net>
From: Nicola Guarino <guarino@ladseb.pd.cnr.it>
Subject: Re: [Fwd: WordNet treelike]
Cc: Andreas.Reuter@eml.villa-bosch.de, "BSpillers" <skydog@pacbell.net>,
"EHovy" <hovv@isi.edu>, "GMiller" <geo@clarity.princeton.edu>,
"JSowa" <sowa@west.poly.edu>,
"KTschira" <Klaus.Tschira@ktf.villa-bosch.de>,
"LPolanyi" <polanyi@pal.xerox.com>,
"PVossen" <Piek.Vossen@let.uva.nl>,
"SPeters" <peters@csl.stanford.edu>,
"AFarquhar" <Adam_Farquhar@ksl.stanford.edu>,
"CFellbaum" <fellbaum@clarity.princeton.edu>,
"DSkuce" <doug@csi.uottawa.ca>

"EFeigenbaum" <feigenbaum@ksl.stanford.edu>,
"FLehmann" <fritz@cyc.com>, "JMcCarthy" <jmc@cs.stanford.edu>,
"JPustejovski" <jamesp@cs.brandeis.edu>,
"JTsuji" <tsujii@is.s.u-tokyo.ac.jp>,
"MvdBerg" <vdberg@us.ibm.com>,
"PHayes" <phayes@nuts.coginst.uwf.edu>,
"PSimons" <p.m.simons@leeds.ac.uk>,
"WWahlster" <wahlster@dfki.uni-sb.de>,
cmenzel@tamu.edu (Chris Menzel)

Content-Transfer-Encoding: 8bit

X-MIME-Autoconverted: from quoted-printable to 8bit by nuts.coginst.uwf.edu id

MAA24750

Content-Type: text/plain

Content-Length: 3210

Status:

At 11:22 AM -0500 5/4/98, Pat Hayes wrote:

>>Fritz / Pat,

>>In my view a major task of the workshop is to find a way to combine
>>extremely useful, if imperfect, linguistic structures with conceptual
>>ontologies (which of course are of all without flaw ;-). This is the
>>sprit of the work done by Ed Hovy and Fritz that resulted in the ANSI
>>Reference Ontology (available on the Stanford KSL server). I hope that
>>one of the items that result from this workshop is agreement on a clear
>>methodology of how to create such an ontology.

>>

>

>Bob, thanks for your message, but I fail to see what your point is. Fritz
>is probably fairly well acquainted with the Fritz/Hovy collaboration. The
>issue of tree-likeness (or, if you prefer, of whether or not to allow
>multiple inheritance) is a key issue in this methodology, and it cuts
>across the linguistic/conceptual distinction. For example, is wood to be
>classified as a substance, or as a structural material, or as a plant
>product, or ...? The natural answer seems to be, all of them.

>

>Is there any *linguistic* evidence for the word/concept heirarchy being
>tree-like, as opposed to having multiple routes of classification
>inheritance?

>

I address a question *very* similar to this one in my latest paper (to be presented at the Granada conference), entitled "Some Ontological Principles for Designing Upper Level Lexical Resources". The difference is that I focus on *ontological* evidence for a tree-like concept hierarchy. Such an evidence is based on (meta-level) distinctions among unary predicates, i.e. on a formal ontology of (unary) universals. The most important distinction is between "types" and "roles": "substance" is a type, while "structural material"

is a role. By focusing on types (which I can prove form a tree according to the assumptions discussed in the paper) we can isolate a "basic backbone" useful for various purposes.

This paper is retrievable (in multiple file formats) from our recently restructured web site (see below). The printing problems found in the past should have now disappeared. I intend to base on this paper my own contribution to the Heidelberg workshop.

Being absorbed by the FOIS'98 organization, I do not have the time for entering in the previous debate right now. For those of you who are going to attend FOIS, however, I would like to remind you two things:

1) Register now if you want a good accommodation in Trento. Due to other parallel events, we are running short of hotel rooms downtown. Drop me a line after you have faxed your registration (form retrievable from the FOIS web page below).

2) In order to go from Trento to Heidelberg, there is a train leaving from Trento at 9:54 and arriving at Heidelberg at 17:51 (via Munich). I suggest to take that train.

Cheers,

-- Nicola

Nicola Guarino
National Research Council phone: +39 49 8295751
LADSEB-CNR fax: +39 49 8295763
Corso Stati Uniti, 4 email: guarino@ladseb.pd.cnr.it
I-35127 Padova
Italy

Home page: ** updated 27/4/98 **
<http://www.ladseb.pd.cnr.it/infor/ontology/ontology.html>

FOIS'98 home page:
<http://mnemosyne.itc.it:1024/fois98/>

From ???@??? Mon May 04 16:28:41 1998
Received: from mail-gw3.pacbell.net (mail-gw3.pacbell.net [206.13.28.55])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id QAA04674

for <phayes@coginst.uwf.edu>; Mon, 4 May 1998 16:21:28 -0500 (CDT)
Received: from pacbell.net (ppp-206-170-7-87.rdcy01.pacbell.net [206.170.7.87]) by mail-gw3.pacbell.net (8.8.8/8.7.1+antispam) with ESMTP id OAA14653; Mon, 4 May 1998 14:18:39 -0700 (PDT)

Message-ID: <354E30A1.A39747E@pacbell.net>

Date: Mon, 04 May 1998 14:18:25 -0700

From: Robert Spillers <skydog@pacbell.net>

X-Mailer: Mozilla 4.04 [en] (Win95; I)

MIME-Version: 1.0

To: Pat Hayes <phayes@coginst.uwf.edu>

CC: Andreas.Reuter@eml.villa-bosch.de, "EHovy" <hovy@isi.edu>,

"GMiller" <geo@clarity.princeton.edu>,

"JSowa" <sowa@west.poly.edu>,

"KTschira" <Klaus.Tschira@ktf.villa-bosch.de>,

"LPolanyi" <polanyi@pal.xerox.com>,

"NGuarino" <guarino@ladseb.pd.cnr.it>,

"PVossen" <Piek.Vossen@let.uva.nl>,

"SPeters" <peters@csl.stanford.edu>,

"AFarquhar" <Adam_Farquhar@ksl.stanford.edu>,

"CFellbaum" <fellbaum@clarity.princeton.edu>,

"DSkuce" <doug@csi.uottawa.ca>,

"EFeigenbaum" <feigenbaum@ksl.stanford.edu>,

"FLehmann" <fritz@cyc.com>, "JMcCarthy" <jmc@cs.stanford.edu>,

"JPustejovski" <jamesp@cs.brandeis.edu>,

"JTsuji" <tsujii@is.s.u-tokyo.ac.jp>,

"MvdBerg" <vdberg@us.ibm.com>,

"PHayes" <phayes@nuts.coginst.uwf.edu>,

"PSimons" <p.m.simons@leeds.ac.uk>,

"WWahlster" <wahlster@dfki.uni-sb.de>, cmenzel@tamu.edu

Subject: Re: [Fwd: WordNet treelike]

References: <v04003a06b17352946c63@[143.88.7.118]>

Content-Transfer-Encoding: 7bit

Content-Type: text/plain; charset=us-ascii

Content-Length: 3016

Status:

Pat,

Perhaps my point was obscure. I don't disagree that multiple inheritance is a significant problem.

What I meant to say is that in the past, Fritz, Ed and others have found ways to incorporate

large chunks of WordNet into conceptual ontologies. This did require surgery, but I believe

they thought the results were useful - at least as part of a first effort to build a reference ontology

(Ed/Fritz comments?).

One purpose of the workshop is to find agreement on methods of constructing WordNet /

EuroWordNet like structures so that they can be easily incorporated into conceptual ontologies

- or better - that they are constructed as an ontology. I don't think many people doubt their

utility. Since EuroWordNet is still under construction, I believe its authors would welcome

suggestions that give their work greater impact (Piek comments?).

Although WordNet is at a different stage, I think the same comment applies (George/Christiane?).

>From an ontological point of view, there are problems with both WordNet and EuroWordNet.

I believe both George and Piek would concur. I hope that (along with other things) this

workshop will devise a strategy and a methodology that allows George, Piek, Christiane and

their colleagues to solve these problems and add another (ontological) dimension to their work .

The current email discussion is useful - it brings out the issues.

Suggestions of how to avoid / solve these problems would also help.

Bob

Pat Hayes wrote:

>>Fritz / Pat,

>>In my view a major task of the workshop is to find a way to combine

>>extremely useful, if imperfect, linguistic structures with conceptual

>>ontologies (which of course are of all without flaw ;-). This is the

>>spirit of the work done by Ed Hovy and Fritz that resulted in the ANSI

>>Reference Ontology (available on the Stanford KSL server). I hope that

>>one of the items that result from this workshop is agreement on a clear

>>methodology of how to create such an ontology.

>>

>

> Bob, thanks for your message, but I fail to see what your point is. Fritz

> is probably fairly well acquainted with the Fritz/Hovy collaboration. The

> issue of tree-likeness (or, if you prefer, of whether or not to allow

> multiple inheritance) is a key issue in this methodology, and it cuts
> across the linguistic/conceptual distinction. For example, is wood to be
> classified as a substance, or as a structural material, or as a plant
> product, or ...? The natural answer seems to be, all of them.
>
> Is there any *linguistic* evidence for the word/concept heirarchy being
> tree-like, as opposed to having multiple routes of classification
> inheritance?
>
> Pat Hayes
>
> -----
> IHMC, University of West Florida (850)434 8903 home
> 11000 University Parkway (850)474 2091 office
> Pensacola, FL 32514 (850)474 3023 fax
> phayes@ai.uwf.edu
> http://www.coginst.uwf.edu/~phayes

From ???@??? Mon May 11 11:05:28 1998
Received: from cclsun01.let.uva.nl (cclsun01.let.uva.nl [145.18.228.21])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id GAA13596
for <phayes@coginst.uwf.edu>; Wed, 6 May 1998 06:16:34 -0500 (CDT)
Received: from cclpc102.let.uva.nl by cclsun01.let.uva.nl with SMTP id AA00092
(5.67a/IDA-1.5 for <phayes@coginst.uwf.edu>); Wed, 6 May 1998 13:12:10 +0200
Message-Id: <1.5.4.32.19980506120716.006a3ad0@mail.let.uva.nl>
X-Sender: piek@mail.let.uva.nl
X-Mailer: Windows Eudora Light Version 1.5.4 (32)
Mime-Version: 1.0
Date: Wed, 06 May 1998 13:07:16 +0100
To: Robert Spillers <skydog@pacbell.net>, Pat Hayes <phayes@coginst.uwf.edu>
From: Piek Vossen <Piek.Vossen@let.uva.nl>
Subject: Re: [Fwd: WordNet treelike]
Cc: Andreas.Reuter@eml.villa-bosch.de, "EHovy" <hovy@isi.edu>,
"GMiller" <geo@clarity.princeton.edu>,
"JSowa" <sowa@west.poly.edu>,
"KTschira" <Klaus.Tschira@ktf.villa-bosch.de>,
"LPolanyi" <polanyi@pal.xerox.com>,
"NGuarino" <guarino@ladseb.pd.cnr.it>,
"PVossen" <Piek.Vossen@let.uva.nl>,
"SPeters" <peters@csl.stanford.edu>,
"AFarquhar" <Adam_Farquhar@ksl.stanford.edu>,
"CFellbaum" <fellbaum@clarity.princeton.edu>,
"DSkuce" <doug@csi.uottawa.ca>,
"EFeigenbaum" <feigenbaum@ksl.stanford.edu>,

"FLehmann" <fritz@cyc.com>, "JMcCarthy" <jmc@cs.stanford.edu>,
"JPustejovski" <jamesp@cs.brandeis.edu>,
"JTsuji" <tsujii@is.s.u-tokyo.ac.jp>,
"MvdBerg" <vdberg@us.ibm.com>,
"PHayes" <phayes@nuts.coginst.uwf.edu>,
"PSimons" <p.m.simons@leeds.ac.uk>,
"WWahlster" <wahlster@dfki.uni-sb.de>, cmenzel@tamu.edu

X-Attachments: C:\Piek\EuroWordNet\Papers\Papers98\Granada\VossenBloksma
.rtf;

Content-Type: multipart/mixed; boundary="====_894452836=="

Content-Length: 330620

Status: RO

At least in EuroWordNet (and I guess also Wordnet) we agree that we need multiple inheritance. We are encoding it where appropriate. However, just having a link to any possible classification is not sufficient. We try to make a difference between the purpose of the classification. We distinguish between two main purposes:

1. to represent an inference scheme
2. substitution of words in text: general words (animal) that can be used to replace more specific words (horse)

Encoding of inference schemes is what we would like to do in an ontology. In the design of the EuroWordNet database the wordnets are linked to an index which gives access to such a shared ontology (currently the EuroWordNet top-ontology). Via the index, it is possible to recover the inferences for any word in all the languages from the ontology. This has the advantage that we can limit the inference schemes for the ontology to what is explicitly defined in e.g. the ANSI group or the Reference Ontology.

The wordnets are then encoding 'substitution patterns of words in a semantic network'. In the wordnets we encode the hyponymic relations between the lexicalized units of languages. Among these lexicalized classes we find many words which generalize over things and express some conceptualization but are not conventionally considered as 'classifications': threat, winner, favourite, failure, investment, breeder, draught animal, riding animal, pet. Strictly speaking these words can all be used to refer to 'horses' as well but they are much more 'circumstantial' than other classifications. We are trying to differentiate between conventional hyperonyms and circumstantial hyperonyms by labelling hyponymic links. Another difference is that some hyperonyms are disjoint (animal, human, plant) while others (especially the circumstantials) are non-exclusive, which means that they can cross-classify with other co-hyponyms. This difference is also encoded by labels in the hierarchy.

We have a paper in the Granada LREC conference in which we describe our position with respect to multiple hyperonyms. I want to use this paper as a starting point for Heidelberg as well. I will attach this paper to the mail. It is in Word RTF format. Perhaps Andreas can put it on the WEB site in addition to Nicola's paper.

Furthermore, I agree very much with Bob that we should not be too negative about the work done in Wordnet. It is because of Wordnet that we can have this workshop and these discussions. Because Wordnet is available as an example we can now ask questions such as: is this good or bad; why is something else better. Wordnet represents the starting point from which we can move onwards. Any new ontology first has to prove that it is better than wordnet and has the same coverage.

best wishes,

Piek.

At 02:18 PM 5/4/98 -0700, Robert Spillers wrote:

>Pat,

>Perhaps my point was obscure. I don't disagree that multiple inheritance is a >significant problem.

>What I meant to say is that in the past, Fritz, Ed and others have found ways >to incorporate

>large chunks of WordNet into conceptual ontologies. This did require

>surgery, but I believe

>they thought the results were useful - at least as part of a first effort to

>build a reference ontology

>(Ed/Fritz comments?).

>

>One purpose of the workshop is to find agreement on methods of constructing

>WordNet /

>EuroWordNet like structures so that they can be easily incorporated into

>conceptual ontologies

>- or better - that they are constructed as an ontology. I don't think many

>people doubt their

>utility. Since EuroWordNet is still under construction, I believe its

>authors would welcome

>suggestions that give their work greater impact (Piek comments?).

>

>Although WordNet is at a different stage, I think the same comment applies

>(George/Christiane?).

>

>From an ontological point of view, there are problems with both WordNet and

>EuroWordNet.

>I believe both George and Piek would concur. I hope that (along with other

>things) this
>workshop will devise a strategy and a methodology that allows George, Piek,
>Christiane and
>their colleagues to solve these problems and add another (ontological)
>dimension to their work .
>
>The current email discussion is useful - it brings out the issues.
>Suggestions of how to avoid /
>solve these problems would also
>help.
>
>Bob
>
>
>Pat Hayes wrote:
>
>> >Fritz / Pat,
>> >In my view a major task of the workshop is to find a way to combine
>> >extremely useful, if imperfect, linguistic structures with conceptual
>> >ontologies (which of course are of all without flaw ;-). This is the
>> >sprit of the work done by Ed Hovy and Fritz that resulted in the ANSI
>> >Reference Ontology (available on the Stanford KSL server). I hope that
>> >one of the items that result from this workshop is agreement on a clear
>> >methodology of how to create such an ontology.
>> >
>>
>> Bob, thanks for your message, but I fail to see what your point is. Fritz
>> is probably fairly well acquainted with the Fritz/Hovy collaboration. The
>> issue of tree-likeness (or, if you prefer, of whether or not to allow
>> multiple inheritance) is a key issue in this methodology, and it cuts
>> across the linguistic/conceptual distinction. For example, is wood to be
>> classified as a substance, or as a structural material, or as a plant
>> product, or ...? The natural answer seems to be, all of them.
>>
>> Is there any *linguistic* evidence for the word/concept heirarchy being
>> tree-like, as opposed to having multiple routes of classification
>> inheritance?
>>
>> Pat Hayes
>>
>> -----
>> IHMC, University of West Florida (850)434 8903 home
>> 11000 University Parkway (850)474 2091 office
>> Pensacola, FL 32514 (850)474 3023 fax
>> phayes@ai.uwf.edu
>> <http://www.coginst.uwf.edu/~phayes>

>
>
>
>

Content-Type: application/rtf; charset="us-ascii"

Attachment converted: lonestar:Untitled 1 (?????/----) (0000425F)

Piek Vossen
Universiteit van Amsterdam
Spuistraat 134
1012 VBAmsterdam
The Netherlands

tel. +31 20 525 4669
fax. +31 20 525 4429

From ???@??? Mon May 11 11:05:53 1998

Received: from mail-gw2.pacbell.net (mail-gw2.pacbell.net [206.13.28.53])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id FAA00281

for <phayes@coginst.uwf.edu>; Fri, 8 May 1998 05:39:49 -0500 (CDT)

Received: from pacbell.net (ppp-206-170-6-12.rdcy01.pacbell.net [206.170.6.12]) by mail-gw2.pacbell.net (8.8.8/8.7.1+antispam) with ESMTP id DAA14937; Fri, 8 May 1998 03:33:24 -0700 (PDT)

Message-ID: <3552DF80.931A25A@pacbell.net>

Date: Fri, 08 May 1998 03:33:37 -0700

From: Robert Spillers <skydog@pacbell.net>

X-Mailer: Mozilla 4.04 [en] (Win95; I)

MIME-Version: 1.0

To: Andreas Reuter <andreas.reuter@EML.villa-bosch.de>

CC: Chris Menzel <cmenzel@tamu.edu>,

Christiane Fellbaum <fellbaum@clarity.Princeton.EDU>,

Doug Skuce <doug@site.uottawa.ca>,

Ed Feigenbaum <eaf@KSL.Stanford.edu>, Eduard Hovy <hovy@isi.edu>,

Fritz Lehmann <fritz@cyc.com>,

George Miller <geo@clarity.Princeton.EDU>,

James Pustsejovski <jamesp@cs.brandeis.edu>,

John McCarthy <jmc@cs.stanford.edu>, John Sowa <sowa@west.poly.edu>,

Junidi Tsujii <tsujii@is.s.u-tokyo.ac.jp>,

Klaus Tschira <klaus.tschira@ktf.villa-bosch.de>,

Livia Polanyi <polanyi@pal.xerox.com>,

Martin va den Berg <vdberg@let.uva.nl>,

Nicola Guarino <guarino@ladseb.pd.cnr.it>,

Pat Hayes <phayes@coginst.uwf.edu>,

Peter Simons <p.m.simons@leeds.ac.uk>,

Piek Vossen <piek.vossen@let.uva.nl>,

Stanley Peters <peters@csl.stanford.edu>,
W Wahlster <wahlster@dfki.uni-sb.de>

Subject: Invitations

Content-Type: multipart/alternative; boundary="-----
-76C5148C88CDE695AAEB4FEC"

Content-Length: 1647

Status:

<x-html><!x-stuff-for-pete base="" src="" id="0"><HTML>

Andreas,

I invited:

<P>Graheme Hirst who will attend if he can manage to work it into his schedule. He thought he would be able to attend. Graheme's email address is <U>gh@cs.toronto.edu</U>. His phone number is 416-978-8747.

<P>Lee Auspitz who will attend. His email address is

<U>lee@textwise.com</U>.

Lee is a philosopher and a member of the Board of Directors of TextWise.

<P>Giovanni Varile who will attend some of the sessions. Dr. Varile is Vice Chairman of the Directorate General of the European Commission that most closely corresponds to our workshop. His email address is <U>gv@lux.dg13.cec.be</U> . His phone is 352-4301-32867.

 </HTML>

</x-html>

From ???@??? Mon May 11 11:05:59 1998

Received: from cclsun01.let.uva.nl (cclsun01.let.uva.nl [145.18.228.21])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id KAA10372

for <phayes@nuts.coginst.uwf.edu>; Mon, 11 May 1998 10:24:32 -0500 (CDT)

Received: from cclpc102.let.uva.nl by cclsun01.let.uva.nl with SMTP id AA20348

(5.67a/IDA-1.5 for <phayes@nuts.coginst.uwf.edu>); Mon, 11 May 1998 17:16:57 +0200

Message-Id: <1.5.4.32.19980511161316.006cac94@mail.let.uva.nl>

X-Sender: piek@mail.let.uva.nl

X-Mailer: Windows Eudora Light Version 1.5.4 (32)

Mime-Version: 1.0

Date: Mon, 11 May 1998 17:13:16 +0100

To: Josiah Lee Auspitz <lauspitz@world.std.com>, jmc@cs.Stanford.EDU,

Patrick Cassidy <micra@tiger.jvnc.net>, schwartz@NU.cs.fsu.edu,

"Martin H. v.d. Berg" <Martin.v.d.Berg@let.uva.nl>

From: Piek Vossen <Piek.Vossen@let.uva.nl>

Subject: Top 40

Cc: lee@textwise.com, sowa@west.poly.edu, Adam_Farquhar@KSL.Stanford.EDU,

Andreas.Reuter@eml.villa-bosch.de, Klaus.Tschira@ktf.villa-bosch.de,
Piek.Vossen@let.uva.nl, doug@csi.uottawa.ca,
feigenbaum@KSL.Stanford.EDU, fellbaum@clarity.princeton.edu,
fritz@cyc.com, geo@clarity.princeton.edu, guarino@ladseb.pd.cnr.it,
hovy@isi.edu, jamesp@cs.brandeis.edu, lee@sabre.org,
onto-std@KSL.Stanford.EDU, p.m.simons@leeds.ac.uk,
peters@csl.stanford.edu, phayes@coginst.uwf.edu,
phayes@nuts.coginst.uwf.edu, polanyi@pal.xerox.com, skydog@pacbell.net,
tsujii@is.s.u-tokyo.ac.jp, vdberg@us.ibm.com, wahlster@dfki.uni-sb.de

Content-Type: text/plain; charset="us-ascii"

Content-Length: 1013

Status:

There are many requests from people for the Top-40. Please, you have to be a bit more patient. The first lists contained 150 concepts and not 40. This was a first trial. I promised a new version last week but could not do it because of other work. I will be working on this list this week and send around to both lists the updated list. If you cannot wait for that I would like to refer to the EuroWordNet home-page: <http://www.let.uva.nl/~ewn> where you can find a list of 1024 concepts. This list was used as a starting point to derive the Top-150 (which will be further reduced). There you will also find a Top-Ontology which has been used to classify the 1024 concepts. There is a deliverable which describes the ontology, the selection of the 1024 concepts etc.. All this gives some background for the selection which I will provide by the end of this week.

best wishes,

Piek.
Piek Vossen
Universiteit van Amsterdam
Spuistraat 134
1012 VBAmsterdam
The Netherlands

tel. +31 20 525 4669
fax. +31 20 525 4429

From ???@??? Mon May 11 18:35:33 1998
Received: from cclsun01.let.uva.nl (cclsun01.let.uva.nl [145.18.228.21])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id PAA14553
for <phayes@coginst.uwf.edu>; Mon, 11 May 1998 15:45:51 -0500 (CDT)
Received: from anpisani (uva67.remote.uva.nl) by cclsun01.let.uva.nl with SMTP id
AA00686
(5.67a/IDA-1.5 for <phayes@coginst.uwf.edu>); Mon, 11 May 1998 22:43:20 +0200
From: "Piek Vossen" <piek.vossen@let.uva.nl>

To: "Pat Hayes" <phayes@coginst.uwf.edu>
Subject: Re: Top 40
Date: Mon, 11 May 1998 22:25:03 +0200
Message-Id: <01bd7d1a\$e0fadac0\$621d1291@anpisani>
Mime-Version: 1.0
Content-Transfer-Encoding: 7bit
X-Priority: 3
X-Msmail-Priority: Normal
X-Mailer: Microsoft Outlook Express 4.71.1712.3
X-Mimeole: Produced By Microsoft MimeOLE V4.71.1712.3
Content-Type: text/plain
Content-Length: 1736
Status:

Hi Pat,

Our selection is not based on 'primitiveness' but on lexicalization. It thus only means that we have relatively many words directly below this level in at least 3 European languages. I think that is different from what I think you are thinking of. Furthermore, the set it occurred in is 150 concepts big not 40.

Piek.

-----Oorspronkelijk bericht-----

Van: Pat Hayes <phayes@coginst.uwf.edu>
Aan: Piek Vossen <Piek.Vossen@let.uva.nl>
Datum: maandag 11 mei 1998 21:24
Onderwerp: Re: Top 40

>Piek, greetings

>

>Many thanks for the Euronet reference and all your hard work. Please dont
>take my messages as being critical. They expressed rather my frustration at
>being unable to comprehend what the point of the enterprise was supposed to
>be.

>

>On the top-40 and my point about 'building', I wasnt meaning to suggest
>that any particular top-level heirarchy was proper, only that whatever one
>you used, it seemed unlikely that anything as concrete (forgive the pun) as
>'building' was likely to be found in the top 40 concepts, except by chance.
>But this may be wrong, if you tell us that several different languages have
>it located very high in their implicit subset heirarchies. If so, then I
>guess my comment must be changed to it being very remarkable that

>'building' should be so distinguished from all the other concepts at about
>the same ontological level (of which there must be, at a guess, at least a
>thousand?)

>
>Pat Hayes

>
>-----
>IHMC, University of West Florida (850)434 8903 home
>11000 University Parkway (850)474 2091 office
>Pensacola, FL 32514 (850)474 3023 fax
>phayes@ai.uwf.edu
><http://www.coginst.uwf.edu/~phayes>

>
>
>

From ???@??? Tue May 12 15:35:25 1998
Received: from snapdragon.textwise.com (snapdragon.Textwise.com [199.100.96.2])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id PAA21577
for <phayes@coginst.uwf.edu>; Tue, 12 May 1998 15:26:00 -0500 (CDT)
Received: from textwise.com by snapdragon.textwise.com (SMI-8.6/SMI-SVR4)
id QAA00794; Tue, 12 May 1998 16:10:37 -0400
Received: from daisy.Textwise.com by textwise.com (SMI-8.6/SMI-SVR4)
id QAA06881; Tue, 12 May 1998 16:21:08 -0400
Received: from localhost by daisy.Textwise.com (SMI-8.6/SMI-SVR4)
id QAA13415; Tue, 12 May 1998 16:21:03 -0400
Date: Tue, 12 May 1998 16:21:03 -0400 (EDT)
From: Josiah Lee Auspitz <lee@textwise.com>
X-Sender: lee@daisy
Reply-To: Josiah Lee Auspitz <lee@textwise.com>
To: Pat Hayes <phayes@coginst.uwf.edu>
cc: "John F. Sowa" <sowa@west.poly.edu>
Subject: Error
In-Reply-To: <Pine.GSO.3.96.980512075823.10784A-100000@daisy>
Message-ID: <Pine.GSO.3.96.980512160310.13307B-100000@daisy>
MIME-Version: 1.0
Content-Type: TEXT/PLAIN; charset=US-ASCII
Content-Length: 574
Status:

I just sent out a message responding to your query about a sentence of mine in which the notion of relation (from McCarthy) is equated with rules. It errs in suggesting that you have read it out of context from a previous message of May 12. In fact, it comes from one on May 7. The two messages are on two different e-mail accounts, so having checked them both

I now see that the notion that you have taken it out of context is false.
It is I who have put it in the wrong context.

The apology for having been loose with McCarthy's notion of relation still applies.

From ???@??? Wed May 13 12:36:22 1998
Received: from cclsun01.let.uva.nl (cclsun01.let.uva.nl [145.18.228.21])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id MAA26603
for <phayes@nuts.coginst.uwf.edu>; Wed, 13 May 1998 12:09:09 -0500 (CDT)
Received: from cclsun01 (cclsun01.let.uva.nl) by cclsun01.let.uva.nl with SMTP id
AA08782
(5.67a/IDA-1.5 for <phayes@nuts.coginst.uwf.edu>); Wed, 13 May 1998 19:06:19
+0200
Date: Wed, 13 May 1998 19:06:18 +0200 (MET DST)
From: "Martin H. v.d. Berg" <Martin.v.d.Berg@let.uva.nl>
X-Sender: vdberg@cclsun01
To: vdberg@let.uva.nl
Cc: Susanne.Winkelmann@eml.villa-bosch.de, Beate.Keller@kts.villa-bosch.de,
Baerbel.Mack@eml.villa-bosch.de, wahlster@dfki.uni-sb.de,
p.m.simons@leeds.ac.uk, phayes@nuts.coginst.uwf.edu, lreeker@nsf.gov,
tsujii@is.s.u-tokyo.ac.jp, jamesp@cs.brandeis.edu, jmc@cs.stanford.edu,
fritz@cyc.com, feigenbaum@ksl.stanford.edu, doug@csi.uottawa.ca,
cmenzel@tamu.edu, fellbaum@clarity.princeton.edu,
Adam_Farquhar@ksl.stanford.edu, peters@csl.stanford.edu,
Piek.Vossen@let.uva.nl, guarino@ladseb.pd.cnr.it,
polanyi@pal.xerox.com, Klaus.Tschira@ktf.villa-bosch.de,
sowa@west.poly.edu, geo@clarity.princeton.edu, hovy@isi.edu,
skydog@pacbell.net, Andreas.Reuter@eml.villa-bosch.de
Subject: From Livia Polanyi: Ontology Workshop Organization (fwd)
Message-Id: <Pine.SOL.3.91.980513190219.8562A-100000@cclsun01>
Mime-Version: 1.0
Content-Type: TEXT/PLAIN; charset=US-ASCII
Content-Length: 2356
Status:

LS

Livia Polanyi asked me to forward this to the list

Cheers
Martin

Time until the Workshop is getting short, and the full agenda is still not set. So I have decided to re-send my response to John Sowa's May 6th message suggesting that people who had not been attending the ontology meetings during the past few years be invited to present papers during day 2 and 3 of the Workshop. This mail did not reach the general list when I first sent it due to some mailer problems, so I have asked Martin van den Berg to forward this message in hopes of eliciting more discussion of how workshop activities should be scheduled.

Livia

John F. Sowa wrote:

>
> At the organizational meeting in January, we scheduled some talks on the first day by some of
> the people who had been attending the ontology meetings during the past two years.
>
> The additional people who have been invited to the Heidelberg meeting have done a lot of work
> on related issues, and they have many ideas that need to be included in the mix.
>
> I think that it would be appropriate to hear from them on the second and/or third days of
> the meeting. We definitely do not want this to be just another week of prepared talks that
> are unrelated to one another, but we do need to give the people we invited a chance to say
> what they have been doing and how they believe it relates to what we have been doing.
>
> John Sowa
John,

As one of those people who has not been scheduled to give a talk, I would like to suggest that we *not*

>hear from them on the second and/or third days of
>the meeting.

The reason is simple. Time.

For 15 people to give a one half hour talk with, let's say, 15 minutes for questions -- an appropriate length of time for a small workshop -- would mean we are looking at 10 hours of talks or so after the first

day. 5 hours a day plus time to think about the talks that were given will consume days 2 and 3. By the time we settle down to begin to work in small groups, we will have used almost half our effective work time.

What I would suggest instead is that everyone post a short position paper of one to three pages on the conference website, or, alternatively, contribute a paper or piece of ontological work. That way we will all have the opportunity to become acquainted with each other's work before the meeting.

Livia

From ???@??? Thu May 21 11:31:14 1998

Received: from rhein.villa-bosch.de (whaleshark.villa-bosch.de [194.25.153.3])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id IAA05382
for <phayes@picayune.coginst.uwf.edu>; Thu, 14 May 1998 08:25:42 -0500

(CDT)

Received: by whaleshark.villa-bosch.de with Internet Mail Service (5.0.1457.3)
id <K55C2BBW>; Thu, 14 May 1998 15:28:03 +0200

Message-ID:

<21C49639A20DD111842C0060B0684B2A088CA9@whaleshark.villa-bosch.de>

From: Andreas Reuter <Andreas.Reuter@eml.villa-bosch.de>

To: AReuter <Andreas.Reuter@eml.villa-bosch.de>,

BSpillers

<skydog@pacbell.net>, EHovy <hovy@isi.edu>,

GMiller

<geo@clarity.princeton.edu>, JSowa <sowa@west.poly.edu>,

KTschira

<Klaus.Tschira@ktf.villa-bosch.de>,

LPolanyi <polanyi@pal.xerox.com>, NGuarino <guarino@ladseb.pd.cnr.it>,

PVossen <Piek.Vossen@let.uva.nl>, SPeters <peters@csl.stanford.edu>,

AFarquhar

<Adam_Farquhar@ksl.stanford.edu>,

CFellbaum

<fellbaum@clarity.princeton.edu>,

CMenzel <cmenzel@tamu.edu>, DSkuce

<doug@csi.uottawa.ca>,

FLehmann <fritz@cyc.com>, GHirst

<gh@cs.toronto.edu>,

GVarile <giovanni.varile@lux.dg13.cec.be>,

JMcCarthy <jmc@cs.stanford.edu>, JPustejovski <jamesp@cs.brandeis.edu>,

JTsujii <tsujii@is.s.u-tokyo.ac.jp>, LAuspitz <lee@textwise.com>,

LReeker

<lreeker@nsf.gov>, MvdBerg <vdberg@us.ibm.com>,

NLawler

<E6NL001@coe.coppin.umd.edu>,

PHayes <phayes@nuts.coginst.uwf.edu>, PSimons <p.m.simons@leeds.ac.uk>, WWahlster <wahlster@dfki.uni-sb.de>
Cc: Beate Keller <Beate.Keller@kts.villa-bosch.de>, Bärbel Mack <Baerbel.Mack@eml.villa-bosch.de>
Subject: Re: Heidelberg Ontology Workshop (HOW)
Date: Thu, 14 May 1998 15:28:01 +0200
X-Priority: 3
MIME-Version: 1.0
X-Mailer: Internet Mail Service (5.0.1457.3)
Content-Type: text/plain
Content-Length: 1821
Status:

Dear colleagues,

June 9, the starting date of the Heidelberg Ontology Workshop is getting closer, so we would like to make the final travel arrangements for the participants. In particular, we need the following information ASAP:

Arrival (date, time, train station or airport)

Departure (date, time, train station or airport)

Since we expect many of you to arrive at (or depart from) the same location at roughly the same time, we plan to arrange shuttle services, but those need to be scheduled well in advance, which is why we ask you to tell us about your plans
v e r y s o o n.

In response to several questions about the reimbursement policy, the Klaus Tschira Foundation (KTF) has asked us to make the following clarifications:

Hotel reservations have been made for all participants at the Heidelberg Holiday Inn; rooms will be paid for by KTF.

All the meals during the workshop will be provided courtesy of KTF.

Shuttle service (where applicable) will be organised soon after we receive your travel information and will be paid for by KTF.

For those of you who travel directly to the Heidelberg Ontology Workshop with the only purpose of attending that workshop KTF will pay for your travel cost, unless your employer picks up the bill.

For those of you who combine the trip to Heidelberg with other business (e.g. attend the Trento workshop), we assume that the other activities are paid by somebody else, hence KTF will only pay for extra travel cost incurred by the Heidelberg - part of the trip.

We are trying to assist you in the preparations for your attendance of the workshop in the best possible way; in order to do that, we need the information mentioned above and therefore hope to get your responses in the very near future.

Thank you.

Beate Keller, Baerbel Mack, Andreas Reuter

From ???@??? Thu May 21 11:31:36 1998
Received: from LADSEB.LADSEB.PD.CNR.IT (ladseb.ladseb.pd.cnr.it [150.178.2.3])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id MAA16249
for <phayes@picayune.coginst.uwf.edu>; Fri, 15 May 1998 12:59:37 -0500
(CDT)

Received: from [150.178.2.93] by 150.178.2.93 with SMTP;
Fri, 15 May 1998 19:53:32 +0200

X-Sender: guarino@ladseb.ladseb.pd.cnr.it

Message-Id: <v0310280fb1822ca7fafa@[150.178.2.93]>

Mime-Version: 1.0

Date: Fri, 15 May 1998 19:55:28 +0200

To: Andreas Reuter <Andreas.Reuter@eml.villa-bosch.de>,
BSpillers <skydog@pacbell.net>, EHovy <hovy@isi.edu>,
GMiller <geo@clarity.princeton.edu>, JSowa <sowa@west.poly.edu>,
KTschira <Klaus.Tschira@ktf.villa-bosch.de>,
LPolanyi <polanyi@pal.xerox.com>, PVossen <Piek.Vossen@let.uva.nl>,
SPeters <peters@csl.stanford.edu>,
AFarquhar <Adam_Farquhar@ksl.stanford.edu>,
CFellbaum <fellbaum@clarity.princeton.edu>, CMenzel <cmenzel@tamu.edu>,
DSkuce <doug@csi.uottawa.ca>,
EFeigenbaum <feigenbaum@ksl.stanford.edu>, FLehmann <fritz@cyc.com>,
GHirst <gh@cs.toronto.edu>, GVarile <gv@lux.dg13.cec.be>,
JMcCarthy <jmc@cs.stanford.edu>, JPustejovski <jamesp@cs.brandeis.edu>,
JTsuji <tsujii@is.s.u-tokyo.ac.jp>, LAuspitz <lee@textwise.com>,
LReeker <lreeker@nsf.gov>, MvdBerg <vdberg@us.ibm.com>,
PHayes <phayes@nuts.coginst.uwf.edu>, PSimons <p.m.simons@leeds.ac.uk>,
WWahlster <wahlster@dfki.uni-sb.de>

From: Nicola Guarino <guarino@ladseb.pd.cnr.it>

Subject: Trying to follow the discussion...

Content-Type: text/plain; charset="us-ascii"

Content-Length: 1459

Status:

Dear all,

Trying to manage a couple of weeks of backlog, I have spent some time making order in the very interesting discussion developed so far. I have identified a number of distinct threads, and I have tried to cut and paste the various statements in order to get something readable.

The various threads I have identified, grouped in clusters, are the following:

- Methodology and purpose of this workshop
- Role of an upper-level ontology

- Language vs. ontology
- Wordnet and Eurowordnet
- Multiple inheritance
- Structuring relations (relations used to "explain concepts")

- Ontology of (non-unary) relations

- Static concepts
- Things, events, situations
- Identity of events

- Use of logic

- Vagueness and precision

Each of these threads corresponds to a single ascii file, which I will mail to Andreas in order to put it on the web site. I hope this may help those who haven't been tracking the discussion. It also can help to start drafting the program of the workshop. I am planning to comment on some of these threads in the next days.

-- Nicola

Nicola Guarino

National Research Council phone: +39 49 8295751

LADSEB-CNR fax: +39 49 8295763

Corso Stati Uniti, 4 email: guarino@ladseb.pd.cnr.it

I-35127 Padova
Italy

Home page: ** updated 4/5/98 **
<http://www.ladseb.pd.cnr.it/infor/ontology/ontology.html>

FOIS'98 home page:
<http://mnemosyne.itc.it:1024/fois98/>

From ???@??? Thu May 21 11:31:52 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id UAA26882
for <phayes@nuts.coginst.uwf.edu>; Sun, 17 May 1998 20:57:37 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21]) by hobbes.poly.edu (8.7.3/8.7.3)
with SMTP id VAA00039; Sun, 17 May 1998 21:50:22 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
id VAA16844; Sun, 17 May 1998 21:49:14 -0400
Date: Sun, 17 May 1998 21:49:14 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199805180149.VAA16844@west>
To: Adam_Farquhar@ksl.stanford.edu, Andreas.Reuter@eml.villa-bosch.de,
E6NL001@coe.coppin.umd.edu, Klaus.Tschira@ktf.villa-bosch.de,
Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fellbaum@clarity.princeton.edu, fritz@cyc.com,
geo@clarity.princeton.edu, gh@cs.toronto.edu,
giovanni.varile@lux.dg13.cec.be, guarino@ladseb.pd.cnr.it,
hovy@isi.edu, jamesp@cs.brandeis.edu, jmc@cs.stanford.edu,
lee@textwise.com, Ireeker@nsf.gov, onto-std@ksl.stanford.edu,
p.m.simons@leeds.ac.uk, peters@csl.stanford.edu,
phayes@nuts.coginst.uwf.edu, polanyi@pal.xerox.com, skydog@pacbell.net,
tsujii@is.s.u-tokyo.ac.jp, vdberg@us.ibm.com, wahlster@dfki.uni-sb.de
Subject: NCITS T2 Standards Meeting
Cc: Baerbel.Mack@eml.villa-bosch.de, Beate.Keller@kts.villa-bosch.de,
Susanne.Winkelmann@eml.villa-bosch.de, sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 3136
Status:

For the past two years, the ontology workshops have been held as ad hoc meetings sponsored by the NCITS T2 Committee on Information Interchange and Interpretation. That committee is also in charge of the ANSI standards for KIF and conceptual graphs (which by the way should be sent out for letter ballot in June, as soon as some editing details are finished).

>From May 13 to 15, the T2 meeting was held in Pennsylvania (near the intersection of PA, NY, and NJ) in conjunction with the NCITS L8 Committee on Data Elements, which includes several people who have also attended the T2 ontology meetings). One topic discussed was a possible merger of T2 and L8 into a single committee. Both groups agreed that such a merger would be highly desirable. From the perspective of the ontology work, that merger has several advantages:

1. There is a large overlap of interests, which include ontology, languages like CGs and KIF, which can be used to represent ontologies, the mapping of ontologies to computational elements, and the tools and facilities for managing all of the above.
2. Recent reorganizations in ISO have transferred the international projects assigned to T2 and L8 into the same working group, whose official name is JTC1 SC32 WG2 on Metadata. Following is the path through the ISO hierarchy leading to that group:

Top: ISO/IEC Joint Technical Committee 1, Information Technology

Next: Standing Committee 32, Data Management and Interchange

Next: Working Group 2, Metadata
3. The larger, merged committee would have more clout and visibility in dealing with other organizations that are consumers or developers of ontologies and tools for dealing with them.
4. To avoid making people travel to multiple meetings on topics that might be of marginal interest to them, the merged committee would meet for only one plenary session per year. That would be the only one that anyone would need to attend in order to maintain voting rights. For various projects in that committee, such as ontology, additional working sessions could be scheduled at other times and places during the year.
5. If and when any proposed standards are developed for ontologies, the merged NCITS committee and the ISO Metadata working group would be the natural place to submit them. SC32, which is the parent committee of WG2, also includes working groups for database languages (i.e. SQL and whatever it evolves into) and data interchange. Those groups are potential allies and customers for ontologies.

As the next step in the merger, the organizers will contact NCITS (which by the way is pronounced EN-SIGHTS and stands for National

Committee on Information Technology Standards) to determine the administrative procedures. The most likely name for the new committee is Metadata, which is the name of the ISO working group. Another suggestion is Metadata and Interchange, which adds a word from the name of SC32 (but adding that word might draw some flack from other committees whose primary task is interchange).

John Sowa

From ???@??? Mon May 04 10:10:48 1998

Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id VAA19448
for <phayes@coginst.uwf.edu>; Fri, 1 May 1998 21:29:49 -0500 (CDT)

Received: from scratchy (scratchy [207.207.8.118])
by catbert.cyc.com (8.8.8/8.8.8) with SMTP id VAA20256;
Fri, 1 May 1998 21:25:39 -0500 (CDT)

Message-Id: <3.0.32.19980501212716.00fc9da8@catbert.cyc.com>

X-Sender: fritz@catbert.cyc.com

X-Mailer: Windows Eudora Pro Version 3.0 (32)

Date: Fri, 01 May 1998 21:27:38 -0500

To: phayes@coginst.uwf.edu

From: Fritz Lehmann <fritz@cyc.com>

Subject: WordNet treelike

Cc: fritz@cyc.com, Andreas.Reuter@eml.villa-bosch.de, skydog@pacbell.net,
hovy@isi.edu, geo@clarity.princeton.edu, sowa@west.poly.edu,
Klaus.Tschira@ktf.villa-bosch.de, polanyi@pal.xerox.com,
guarino@ladseb.pd.cnr.it, onto-std@KSL.Stanford.EDU,
Piek.Vossen@let.uva.nl, peters@csl.stanford.edu,
Adam_Farquhar@KSL.Stanford.EDU, fellbaum@clarity.princeton.edu,
doug@csi.uottawa.ca, feigenbaum@KSL.Stanford.EDU, jmc@cs.stanford.edu,
jamesp@cs.brandeis.edu, tsujii@is.s.u-tokyo.ac.jp, vdberg@us.ibm.com,
phayes@nuts.coginst.uwf.edu, p.m.simons@leeds.ac.uk,
wahlster@dfki.uni-sb.de, geo@clarity.princeton.edu

Mime-Version: 1.0

Content-Type: text/plain; charset="us-ascii"

Content-Length: 1492

Status:

Dear Pat,

In looking at WordNet, I've noticed that it's almost a single-inheritance tree, in which almost all synsets each have exactly one hypernym. There are a few exceptions, like the noun piano having both stringed instrument and percussion instrument as hypernyms, but this is rare. They had a half-hearted dedication to "freedom". The choice of such a tree-like structure was a mistake, in my view. I think it should have a non-tree-like poset structure (partially ordered set, DAG) with much more

multiple inheritance from multiple hypernyms. Putting concepts into a tree forces you to make some silly decisions as to which of several salient superclasses should be designated as "the" hypernym.

The higher you go in WordNet, the more you get into controversial or dubious linkings; the lower levels are more obviously reliable and WordNet is quite useful in its lower reaches. And it's big.

Cyc has established over 6,000 links to WordNet (Cyc "constants" linked to WordNet "synsets") with many more to go, and I believe these cover all 3000+ of the Cyc-based, publicly released "reference ontology". (In case you don't know, those are on Cycorp's web page at <http://www.cyc.com> --- see the section on the Upper Cyc Ontology.)

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

=====
=====

From ???@??? Tue May 05 14:58:33 1998

Received: from news.uni-kl.de (news.uni-kl.de [131.246.137.51])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id NAA09567

for <phayes@coginst.uwf.edu>; Tue, 5 May 1998 13:42:44 -0500 (CDT)

Received: from uklirb.informatik.uni-kl.de by news.news.uni-kl.de id ae05191;
5 May 98 20:39 MET DST

Date: Tue, 5 May 98 20:37:29 MET DST

From: Harold Boley <boley@informatik.uni-kl.de>

To: Fritz Lehmann <fritz@cyc.com>

cc: phayes@coginst.uwf.edu, fritz@cyc.com, Andreas.Reuter@eml.villa-bosch.de,
skydog@pacbell.net, hovy@isi.edu, geo@clarity.princeton.edu,
sowa@west.poly.edu, Klaus.Tschira@ktf.villa-bosch.de,
polanyi@pal.xerox.com, guarino@ladseb.pd.cnr.it,
onto-std@ksl.stanford.edu, Piek.Vossen@let.uva.nl,
peters@csl.stanford.edu, Adam_Farquhar@ksl.stanford.edu,
fellbaum@clarity.princeton.edu, doug@csi.uottawa.ca,
feigenbaum@ksl.stanford.edu, jmc@cs.stanford.edu,
jamesp@cs.brandeis.edu, tsujii@is.s.u-tokyo.ac.jp, vdberg@us.ibm.com

Subject: Re: WordNet treelike

Organization: University of Kaiserslautern, DFKI, Germany

Message-ID: <9805052037.aa09416@uklirb.informatik.uni-kl.de>

Mime-Version: 1.0

Content-Type: text/plain; charset=us-ascii

Content-Length: 3245

Status:

Dear Fritz Lehmann,

I agree that multiple hypernyms are a good choice when there is no naturally distinguished hypernym. In general, if there are N unary predicates p_1, \dots, p_N applicable to X , the ontology standard should permit us to express, without a priori bias, all of the following (where " \wedge " is ordinary conjunction and " $\overset{\sim}{\wedge}$ " is hypernym 'intersection'):

No hypernym: $p_1(X) \wedge p_2(X) \wedge \dots \wedge p_H(X) \wedge p_{H+1}(X) \wedge \dots \wedge p_N(X)$

One hypernym: $p_2(X:p_1) \wedge \dots \wedge p_H(X:p_1) \wedge p_{H+1}(X:p_1) \wedge \dots \wedge p_N(X:p_1)$

$1 < H < N$ hypernyms: $p_{H+1}(X:p_1 \overset{\sim}{\wedge} p_2 \overset{\sim}{\wedge} \dots \overset{\sim}{\wedge} p_H) \wedge \dots \wedge p_N(X:p_1 \overset{\sim}{\wedge} p_2 \overset{\sim}{\wedge} \dots \overset{\sim}{\wedge} p_H)$

Only hypernyms: $\text{true}(X:p_1 \overset{\sim}{\wedge} p_2 \overset{\sim}{\wedge} \dots \overset{\sim}{\wedge} p_H \overset{\sim}{\wedge} p_{H+1} \overset{\sim}{\wedge} \dots \overset{\sim}{\wedge} p_N)$ or $X:p_1 \overset{\sim}{\wedge} p_2 \overset{\sim}{\wedge} \dots \overset{\sim}{\wedge} p_H \overset{\sim}{\wedge} p_{H+1} \overset{\sim}{\wedge} \dots \overset{\sim}{\wedge} p_N$

Of course, we have to clarify the semantics of hypernym 'intersection':

It might correspond to set intersection on the predicate extensions or to a less obvious inheritance strategy. While we may consider the extensional semantics of the four above possibilities to be the same, we should be able to cope with the remaining difference of, say,

$\text{black-instrument}(X) \wedge \text{stringed-instrument}(X) \wedge \text{percussion-instrument}(X)$

and (the more natural)

$\text{black-instrument}(X:\text{stringed-instrument} \overset{\sim}{\wedge} \text{percussion-instrument})$

where we selected stringed-instrument and percussion-instrument to be 'sortal' predicates and black-instrument to be a 'non-sortal' predicate.

As mentioned by John Sowa, a hypernym hierarchy should also be permitted for non-unary predicates. And the above remarks should apply to it, too.

In "ONTOFILE: Ontological Modelling of Local and URL File Systems", to be presented at the "Eighth European-Japanese Conference on Information Modelling and Knowledge Bases", 26-29 May 1998 in Finland, I felt the need for this flexibility for a real-world application.

Abstract:

The conceptual-modelling language ONTOFILE is introduced to cope with the ontological complexity of file systems. On the basis of a functionally extended logic, files are described by exterior and interior ontologies for the respective structuring of their manifest and underlying features. The declarative representation of manifest file attributes and relations as well as underlying file entities and properties is discussed with an information-systems example. Exterior

ONTOFILE attributes and relations can be parameterized; single-valued and multiple-valued attributes are modelled by deterministic and non-deterministic functions, respectively. Interior entities and properties are modelled by subsumption hierarchies; property-to-entity applications return the files in which they hold. These descriptions can be employed at the same time, like in fact retrieval, as a knowledge base summarizing the content of files and, like in document retrieval, as an index for the names of files containing detailed information. Besides retrieval, ONTOFILE hierarchies permit two kinds of inference, inheritance and expansion. This modelling approach is applied consistently to local file systems and to URL-addressed WWW pages.

URL of full paper: <http://www.dfki.uni-kl.de/~boley/filekb.ps>

BTW, I'm a little unclear about which contributions make it to me via onto-std@ksl.stanford.edu (the thread's context suggests that I have missed the ones by Pat Hayes and John McCarthy).

Greetings, Harold Boley.

From ???@??? Tue May 05 17:28:52 1998

Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id PAA10148;

Tue, 5 May 1998 15:41:25 -0500 (CDT)

Date: Tue, 5 May 1998 15:41:25 -0500 (CDT)

X-Sender: phayes@mail.coginst.uwf.edu

Message-Id: <v04003a05b174d4fe7296@[143.88.7.118]>

In-Reply-To: <9805052037.aa09416@uklirb.informatik.uni-kl.de>

Mime-Version: 1.0

To: Harold Boley <boley@informatik.uni-kl.de>

From: Pat Hayes <phayes@coginst.uwf.edu>

Subject: Re: WordNet treelike

Cc: <Andreas.Reuter@eml.villa-bosch.de>, "BSpillers" <skydog@pacbell.net>,

"EHovy" <hovy@isi.edu>, "GMiller" <geo@clarity.princeton.edu>,

"JSowa" <sowa@west.poly.edu>,

"KTschira" <Klaus.Tschira@ktf.villa-bosch.de>,

"LPolanyi" <polanyi@pal.xerox.com>,

"NGuarino" <guarino@ladseb.pd.cnr.it>,

"SPeters" <peters@csl.stanford.edu>,

"AFarquhar" <Adam_Farquhar@ksl.stanford.edu>,

"CFellbaum" <fellbaum@clarity.princeton.edu>,

"DSkuce" <doug@csi.uottawa.ca>,

"EFeigenbaum" <feigenbaum@ksl.stanford.edu>,

"FLehmann" <fritz@cyc.com>, "JMcCarthy" <jmc@cs.stanford.edu>,

"JPustejovski" <jamesp@cs.brandeis.edu>,

"JTsuji" <tsujii@is.s.u-tokyo.ac.jp>,

"MvdBerg" <vdberg@us.ibm.com>,
"PHayes" <phayes@nuts.coginst.uwf.edu>,
"PSimons" <p.m.simons@leeds.ac.uk>,
"WWahlster" <wahlster@dfki.uni-sb.de>

Content-Type: text/plain; charset="us-ascii"

Content-Length: 2809

Status:

Greetings Harold

I agree that sort heirarchies are needed for nonunary predicates, however things do get complicated rather rapidly. Take as an example the sorts <human>, <man>, <lady> , with the first as the hypernym, and the binary relation Married. In most states in the USA, the sort of Married is (man X lady) U (lady X man), but excludes (man X man) and (lady X lady). In most states of the USA one can be jailed for getting this wrong, so it is important that the relational sort heirarchy keeps things straight.

We could allow a 'universal' Married relation with sort (human X human) ie (man U lady) X (lady U man) to be the hypernym, but the relationship between regular marriages and same-sex marriages is rather harder to state. They are exclusive, for example, and neither of them can be determined by specifying the sorts of their arguments; that is, there are no sorts S, T such that the sort of these relations is S X T. Also, there are the sorts (human X man) and (human X lady) and their reflections, which exist in the abstract sort structure but are rarely discussed, presumably because their sort structure is not preserved under the axiom of symmetry which is known to be satisfied by the Married relation. (Contrast Parent, where we have a collections of words for making exactly these distinctions: father, mother, son, daughter.) This suggests that whether or not a combination of sorts is to be allowed in the heirarchy must depend in part on what properties the relation is supposed to have.

These issues have been discussed at length by others, in several contexts. Sorted logics with disjoint-union sorts (like Married) were discussed at length by Tony Cohn in his PhD thesis, and the theory of computation has long had to deal with multiply-sorted languages in which functions and relation sorting is 'overdetermined' in this way.

Pat Hayes

PS. Why do you say that

black-instrument(X:stringed-instrument^percussion-instrument)

is *more natural* than

black-instrument(X) \wedge stringed-instrument(X) \wedge percussion-instrument(X) ?

Does it have something to do with the stringedness and percussiveness of the instrument being somehow more intrinsic to its functional role *as an instrument* than its color? If so, is this something that could be described in a theory of instruments (that they make sounds, say, and that the sound of a string is different in musical quality than the sound of a plate) sufficiently well that the naturalness of the former could be figured out by examining the theory?

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Mon May 11 11:05:41 1998
Received: from news.uni-kl.de (mmdf@news.uni-kl.de [131.246.137.51])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id PAA16357
for <phayes@coginst.uwf.edu>; Wed, 6 May 1998 15:44:01 -0500 (CDT)
Received: from uklirb.informatik.uni-kl.de by news.news.uni-kl.de id aa10648;
6 May 98 22:40 MET DST
Date: Wed, 6 May 98 22:37:02 MET DST
From: Harold Boley <boley@informatik.uni-kl.de>
To: Pat Hayes <phayes@coginst.uwf.edu>
cc: Harold Boley <boley@informatik.uni-kl.de>,
Andreas.Reuter@eml.villa-bosch.de, "BSpillers" <skydog@pacbell.net>,
"EHovy" <hovya@isi.edu>, "GMiller" <geo@clarity.princeton.edu>,
"JSowa" <sowa@west.poly.edu>,
"KTschira" <Klaus.Tschira@ktf.villa-bosch.de>,
"LPolanyi" <polanyi@pal.xerox.com>,
"NGuarino" <guarino@ladseb.pd.cnr.it>,
"SPeters" <peters@csl.stanford.edu>,
"AFarquhar" <Adam_Farquhar@ksl.stanford.edu>,
"CFellbaum" <fellbaum@clarity.princeton.edu>,
"DSkuce" <doug@csi.uottawa.ca>,
"WWahlster" <wahlster@dfki.uni-sb.de>, onto-std@ksl.stanford.edu
Subject: Re: WordNet treelike
Organization: University of Kaiserslautern, DFKI, Germany
Message-ID: <9805062237.aa01305@uklirb.informatik.uni-kl.de>
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii

Content-Length: 2873

Status: O

Hi Pat,

I forwarded your yesterday's email to onto-std@ksl.stanford.edu, since this was also included in my reply to Fritz, whose recipients I just overtook (well, we should install an intelligent recipient-management system ...).

Here some quick remarks concerning your points on nonunary inheritance.

1. We can distinguish between hierarchies over the argument positions of nonunary predicates (as in your 'universal' Married relation), hierarchies over nonunary predicates with fixed argument sorts (see example below), and the combination of both. Varying your example, let us consider this little 'couple' hierarchy:

```
couple(human, human)
  concubinage(human, human)
  universally-married(human, human)
    regular-married(human, human)
      honeymoon-married(human, human)
        after-honeymoon-married(human, human)
          same-sex-married(human, human)
```

Obviously, you can inherit useful information for, say, the honeymoon-married relation from its super-relations, regular-married and above.

2. Perhaps we can treat some of the problems you mention by introducing a canonical order over symmetric relation arguments, as by permitting only regular-married(man, lady) or, better, regular-married(husband:man, wife:lady). We can now instantiate universally-married(uhusband:human, uwife:human) to either regular-married(uhusband:man, uwife:lady) or to the first alternative same-sex-married(uhusband:man, uwife:man) or to the second alternative same-sex-married(uhusband:lady, uwife:lady), where uhusband/uwife denote 'universal' marriage roles. Introducing roles instead of fixed argument positions corresponds to proceeding from Herbrand terms to Ait-Kaci's psi-terms. His work on inheritance between psi-term-described concepts seems to be quite relevant for our current efforts in ontology.

3. Concerning your original problem formulation, one might ask whether the relational sort hierarchy (the 'taxonomy') or some additional integrity constraints (the 'axioms') should be responsible to keep things straight. If you permit the full power of FOPC (or more) for the axioms, you can even control the use of very irregular inheritance schemes. I think the KL-ONE experience has shown that one should carefully balance the taxonomic and

axiomatic part of ontologies.

Cheers, Harold.

PS: I very much like your interpretation of the "black piano" example:
A theory of instruments should separate sound-relevant and -irrelevant
properties. The example was inspired by Nicola Guarino's "red apple"
example (GCG94), now with *two* sortal predicates.

(GCG94)

%0 Book Section

%A Guarino, N.

%A Carrara, M.

%A Giaretta, P.

%D 1994

%T An Ontology of Meta-Level Categories

%B Principles of Knowledge Representation and Reasoning: Proceedings of the
Fourth International Conference (KR94)

%E J., Doyle

%E Sandewall, E.

%E Torasso, P.

%I Morgan Kaufmann

%C San Mateo, CA

%P 270-280

From ???@??? Thu May 21 11:31:58 1998

Received: from news.uni-kl.de (news.uni-kl.de [131.246.137.51])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id OAA29212

for <phayes@coginst.uwf.edu>; Mon, 18 May 1998 14:52:27 -0500 (CDT)

Received: from uklirb.informatik.uni-kl.de by news.news.uni-kl.de id ag09366;

18 May 98 21:49 MET DST

Date: Mon, 18 May 98 21:43:36 MET DST

From: Harold Boley <boley@informatik.uni-kl.de>

To: Pat Hayes <phayes@coginst.uwf.edu>

cc: Harold Boley <boley@informatik.uni-kl.de>

Subject: Re: WordNet treelike

Organization: University of Kaiserslautern, DFKI, Germany

Message-ID: <9805182143.aa10410@uklirb.informatik.uni-kl.de>

Mime-Version: 1.0

Content-Type: text/plain; charset=us-ascii

Content-Length: 7792

Status:

Hi Pat,

first, in the ontological spirit of the discussion group's context,
I distinguish a (central) taxonomy from (additional) axioms, hoping

to obtain a not only epistemologically, but also heuristically adequate representation in your sense of McCarthy/Hayes69 (year? is this online?).

CONCERNING A CANONICAL ORDER OVER SYMMETRIC RELATION ARGUMENTS YOU WROTE:

... suppose Married has its arguments ordered, and I define general-married(x,y) to mean (Married(x,y) or Married(y,x)) . You have to somehow prevent this happening.

I think there are several issues here:

* If we define some ordering over the sorts such that man < lady will be enforced for all humans, we need to store only facts of the form Married(man,lady), because a call like Married(tina,fred) would be 'reduced' to Married(fred,tina) before/during unification with this fact, using self-normalizing arguments/equational unification. Then, general-married(tina,fred) would be Married(tina,fred) or Married(fred,tina), but both disjuncts would again be equivalent to just Married(fred,tina). Similarly for general-married(fred,tina). While equational unification is well-known, self-normalizing arguments can be viewed as if we called Married({tina,fred}), where the active ordered-set constructor "{...}" canonically orders its elements according to "<", obtaining {fred,tina}; this works nicely for ground calls, but would, at least, require constraint-programming-like delays for non-ground calls.

* Your distinction of a capitalized "Married" predicate from a lower-case "general-married" predicate suggests that we may be crossing a boundary between sortal and non-sortal binary predicates here: while "Married" was assumed to be part of the taxonomy, "general-married" may be part of the axioms operating over it. While in your original example, the union (man X lady) U (lady X man) was inherently taxonomic, the disjunction in the general-married definition appears to belong to the axioms. So you are right, for non-unary predicates it is much more difficult to keep the boundary between the taxonomy and the axioms straight: We don't any more have the intuitive necessary condition for a unary predicate being 'taxonomic', namely that it can be used to constrain individual arguments of axioms; an N-ary predicate could, however, be used (even more generally) to constrain argument combinations, as in the tiny 4-ary taxonomy

buy(buyer,seller,object,price)
regular-buy(buyer,seller,object,price)
favorable-buy(buyer,seller,object,price)
unfavorable-buy(buyer,seller,object,price)

where good-price(object,price) could be used as a/the defining binary

constraint for favorable-buy(buyer,seller,object,price).

* For the `role' version we would have to allow a (somewhat strange) mixture of role-using and normal predicates, if we want to define, as you suggest, general-married(x,y) as Married(husband:x,wife:y) OR Married(husband:y,wife:x) because a role-using general-married would again fix x as the husband, as in general-married(husband:x,wife:y), thus preventing the second disjunct.

WRT ROLES YOU THEN WROTE: ... labelling the argument places by role names certainly seems useful but I dont think it solves the basic complications. It also introduces a new kind of classification heirarchy eg consider the role names 'husband' (as opposed to the binary relation husband-of), 'wife', and 'spouse' which generalises the first two.

Well, since role names allow us to forget about argument order and to introduce additional roles incrementally down the inheritance hierarchy, they make N-ary relations more concept-like. Principles of unary predicates may such be easier transferred to non-unary ones. The work of Ait-Kaci focussed, however, on psi-term unification, and you can define something like same-aged-couple(husband:man(age:n),wife:lady(age:n)), where n is a free logical variable (AK93). You are right with the doubling of hierarchies, but perhaps we could construct a simple (one-to-one?) mapping between, e.g.

spouse		spouse-of
husband	and	husband-of
wife		wife-of

ABOUT THE TAXONOMY-AXIOM SEPARATION YOU WROTE: ... I dont see the need to make this distinction; or perhaps better, it seems to be a distinction between axiomatic styles and perhaps mechanisms for using them, rather than between different *parts* of an ontology. All the same, I agree that we need to find general-pupose inheritance mechanisms rather than just hack the logic to suit each different case.

Ok, I originally asked for flexibility wrt taxonomic and/or axiomatic formalizations in the upcoming ontology standard. But even more important would now appear to me to try to formalize a given (benchmark?) example in a purley axiomatic style, a highly taxonomic style, and any mixtures thereof. We could then begin to compare issues of heuristic adequatedness. BTW, if you separate taxonomic and axiomatic definitions, you can always reconstruct a purely axiomatic version, but not vice versa (compared to a good mixture, the purely axiomatic version may thus be regarded as an epistemologically adequate but heuristically impoverished version). >From your earlier contributions (to the interlingua list?) I seem to

remember that you would now say that we should maintain and evolve the purely axiomatic version, and construct special-purpose mixtures as needed heuristically. I like the pureness, but where, then, do we maintain and evolve the sortal information lost in the pure axioms? (If you have a taxonomic-axiomatic mixture, you attach the sortal information directly 'in place'.)

FINALLY, YOU QUESTIONED THE GENERALITY OF 'RELEVANT/ESSENTIAL' PROPERTIES:

P is a property of an object of a type T, where the definition (?) of T involves some 'function' F whose definition 'crucially' involves a concept(?) C. Then P is 'essential' to T just when the definition of P also 'involves' C. (Glossary: P=black/loud, T=musical-instrument, F= making-sound, C=sound)

But that seems pretty weak. I have no idea what any of the quoted words mean! Can you do better?

In general, I also have problems in selecting the heuristically most adequate superconcepts. Maybe principles of maximally economic concept grouping apply here, so as to have semantically related concepts clustered in topological or even metric neighbourhoods. These may provide a meaning to your quoted words 'crucially', 'essential', and 'involves'. E.g., the attribute color would be topologically/metrically more distant to your 'function' F, than would be the attribute sound, as in the very much *simplified* taxonomy:

```
instrument(function:{coloring,making-sound},
           color:{black,white},
           sound:{loud,calm},...)
musical-instrument(function:making-sound,sound:{loud,calm},...)
paint-brush(function:coloring,color:{black,white},...)
```

While it seems clear here that the function making-sound is closer to the attribute sound than to the attribute color, in general we may well have a kind of "frame problem": additional attributes may turn out to become relevant (as when someone would discover that black paint changes wooden surfaces so as to produce better sound). But the discovery of new relevant attributes also constitutes part of the progress of science, and as such must necessarily lead to some restructuring of our taxonomies.

Cheers, Harold.

(AK93) Hassan Ait-Kaci, "An Introduction to LIFE - Programming with Logic, Inheritance, Functions, and Equations," Proceedings of the 10th International Logic Programming Symposium, Vancouver, BC, Canada,

October 1993, pp.1-17.

For further info also see: <http://www.isg.sfu.ca/~hak/prof.html>

From ???@??? Thu May 21 11:32:31 1998

Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id CAA08349

for <phayes@coginst.uwf.edu>; Thu, 21 May 1998 02:07:43 -0500 (CDT)

Received: from west (west.poly.edu [128.238.20.21]) by hobbes.poly.edu (8.7.3/8.7.3) with SMTP id DAA20525; Thu, 21 May 1998 03:00:13 -0400 (EDT)

Received: by west (SMI-8.6/SMI-SVR4)

id CAA08110; Thu, 21 May 1998 02:58:41 -0400

Date: Thu, 21 May 1998 02:58:41 -0400

From: sowa@west.poly.edu (John F. Sowa)

Message-Id: <199805210658.CAA08110@west>

To: onto-std@ksl.stanford.edu, phayes@coginst.uwf.edu, piek.vossen@let.uva.nl

Subject: Knowledge Bus and RDF

Cc: Adam_Farquhar@ksl.stanford.edu, Andreas.Reuter@eml.villa-bosch.de,

Klaus.Tschira@ktf.villa-bosch.de, chezewiz@erols.com,

doug@csi.uottawa.ca, feigenbaum@ksl.stanford.edu,

fellbaum@clarity.princeton.edu, fritz@cyc.com,

geo@clarity.princeton.edu, guarino@ladseb.pd.cnr.it, hovy@isi.edu,

jamesp@cs.brandeis.edu, jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk,

peters@csl.stanford.edu, phayes@nuts.coginst.uwf.edu,

polanyi@pal.xerox.com, skydog@pacbell.net, sowa@west.poly.edu,

tsujii@is.s.u-tokyo.ac.jp, vdberg@us.ibm.com, wahlster@dfki.uni-sb.de

Mime-Version: 1.0

Content-Type: text/plain; charset=us-ascii

Content-Length: 6915

Status:

At the NCITS T2 meeting last week, there were presentations on two important examples of the use of ontologies in computer applications. The first was on the Knowledge Bus by Bill Andersen from DoD, and the second was on the Resource Description Framework (RDF) by Frank Olken from the Lawrence Berkeley Laboratory. Either directly or indirectly, both topics are related to Cyc, but they are just as relevant to any ontology that may be developed by or be incorporated in the ontology work we are considering.

Knowledge Bus:

Bill Andersen's talk on the Knowledge Bus was a preview of a paper that will be presented at the 5th KRDB Workshop in Seattle on 31 May 1998. The title is "Knowledge Bus: Generating Application-Focused Databases from Large Ontologies" by B. J. Peterson, W. A. Andersen, and J. Engel. A PostScript version of it can be downloaded from

<http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-10/>

The Knowledge Bus is a system that generates database definitions and programming interfaces (APIs) from the Cyc ontology. It doesn't map the entire Cyc knowledge base into a database, but only that subset that is accessible from a specific context or `_microtheory_`. Instead of the 500,000 or so axioms of Cyc, it extracts about 5,000 that are relevant to some application domain.

In this case, the Knowledge Bus was used to "develop databases for the Department of Defense, which are now in operational use in complex decision-support applications." The APIs are the Java class definitions and interfaces, which are generated automatically from the Cyc ontology. The programming details in the Java methods are filled in by a human programmer, but they use straightforward programming techniques that might someday be automated.

Cyc is used only in developing and testing the ontology and the associated axioms. Cyc is not involved in the operational system, which uses Java programs and a deductive database query engine, XSB. The XSB system is a Prolog-like engine with well-founded semantics that was developed at SUNY Stony Brook.

One interesting point is that Cyc uses full first-order logic with default reasoning, but XSB only supports the Horn-clause subset of FOL for deduction. As it turns out, about 98% of the FOL axioms in Cyc are already in Horn-clause form, from which they can be automatically translated to XSB rules. The other 2% of the axioms are not thrown away; instead, they are used as integrity checks on the database.

Computationally, that approach is significant: Horn-clause deductions, as in Prolog, are highly efficient, but full FOL theorem proving may take an exponential amount of time. The non-Horn 2% of the axioms are not used for deduction, but for integrity checks, which can also be done efficiently: the truth or falsity of any FOL statement can be evaluated in terms of a given database in polynomial time by the equivalent of an ordinary SQL query. Although full FOL may be inefficient for arbitrary deductions, it can still be used efficiently for other kinds of applications. That is a point I have been emphasizing for years: efficiency depends primarily on what you do with the logic and only secondarily on the structure of the logical formulas.

Some people have advocated a restricted version of logic for specifying ontologies. However, that seems to be short-sighted because we cannot know in advance what users will want to do with the ontologies. The experience with Knowledge Bus shows that automated tools can extract

an efficiently computable subset from an ontology stated in full FOL. The ontology developers should provide as much knowledge as they can in whatever notation is appropriate for the domain experts. Then the application developers can select whatever subset they need and translate it to any form their tools require.

RDF:

Frank Olken's talk was about the Resource Description Framework (RDF), which evolved from the Meta Context Framework (MCF), which was developed at Apple by R. V. Guha, the former associate director of Cyc, who is now at Netscape. One of the other people involved at Apple was Larry Tesler, who was the coauthor of the first paper that Roger Schank published on his conceptual dependency theory (at IJCAI in 1969). Given that heritage, it is not surprising that RDF happens to be a semantic network that could be translated directly to a subset of conceptual graphs. RDF has now been adopted by the W3 consortium as the primary language for specifying resources on the Internet.

Following is a brief description of an RDF database by Guha et al.:

- > 1. a set of labels, also referred to as property types
- > 2. a set of nodes
- > 3. a set of arcs where each arc is a triple consisting of two nodes
 - > (the source and target) and a label. Arcs are also referred to as
 - > properties. Often, we will refer to an arc with a certain source
 - > as a `_property of that source_`. Similarly we will refer to the
 - > target of the arc as the `_value of the property_`.

Following is an example from the RDF specification:

- > An RDF expression is represented pictorially in text with nodes
- > in '[...]' and arcs in '`--...-->`' as follows:
 - > `[resource R]---PropertyType P-->[value V]`.
 - > This is read "V is the value of the property type P for resource R";
 - > or left-to-right, "R has property type P with value V." Consider
 - > as a simple example the statement:
 - > Ora Lassila is the author of the resource <http://www.w3.org/People/Lassila>
 - > This statement can be represented as follows:

> [<http://www.w3.org/People/Lasilla>]---Author-->"Ora Lassila"

> where the notation '[URI]' denotes the node representing the resource
> identified by URI and quotation marks (") denote an atomic value.

All of this happens to look like a version of the linear notation for conceptual graphs. In fact, RDF is essentially the "simple graph" subset of CGs, which was defined in *Conceptual Structures* as CGs with no negations, nested contexts, or quantifiers other than the default existential.

What makes RDF important is not its theoretical sophistication, but the fact that it has been adopted by the W3 consortium, which is supported by all the big players, including IBM, Netscape, Microsoft, etc. Technical reports that describe RDF and related topics can be viewed or downloaded from the W3 web site: <http://www.w3.org/TR/>

For a two-page introduction to RDF at the "executive summary" level,

<http://www.w3.org/TR/NOTE-rdf-simple-intro>

For the current working draft of the RDF definition and syntax,

<http://www.w3.org/TR/WD-rdf-syntax/>

For an older paper on MCF by Guha et al., which is now obsolete for the details, but is more interesting for the underlying rationale,

<http://www.w3.org/TR/NOTE-MCF-XML/>

For the latest working draft on RDF schemas by Guha of Netscape and Andrew Layman of Microsoft,

<http://www.w3.org/TR/WD-rdf-schema/>

From ???@??? Wed May 27 10:56:38 1998

Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id KAA00941
for <phayes@nuts.coginst.uwf.edu>; Wed, 27 May 1998 10:46:16 -0500 (CDT)

Received: from scratchy (scratchy [207.207.8.118])
by catbert.cyc.com (8.8.8/8.8.8) with SMTP id KAA24855;
Wed, 27 May 1998 10:41:36 -0500 (CDT)

Message-Id: <3.0.32.19980527104249.00f41298@catbert.cyc.com>

X-Sender: fritz@catbert.cyc.com

X-Mailer: Windows Eudora Pro Version 3.0 (32)
Date: Wed, 27 May 1998 10:43:00 -0500
To: phayes@nuts.coginst.uwf.edu
From: Fritz Lehmann <fritz@cyc.com>
Subject: continuity and glass time
Cc: fritz@cyc.com
Mime-Version: 1.0
Content-Type: text/plain; charset="us-ascii"
Content-Length: 5759
Status:

Dear Pat,

I noticed the following message on the Conceptual Graphs mailing list. I'm forwarding it to you only because the discussion of continuity towards the end reminds me of your notion of "glass time".

===== begin quoted matter =====
Date: Wed, 27 May 1998 01:07:37 -0400
From: sowa@west.poly.edu (John F. Sowa)
To: cg@pluto.cs.uah.edu
Subject: Re: CG: Peirce's comments on EGs (fwd)
Sender: owner-cg@pluto.cs.uah.edu
Reply-To: cg@pluto.cs.uah.edu
Content-Length: 4937

Norman,

Those excerpts from Peirce's manuscript raise all sorts of questions, and it indicates just one more reason why we should hope that more of them become more widely accessible soon. A few comments on your comments:

>I have a query and a comment on your interesting excerpts from
>unpublished Peirce.

Actually, those excerpts were by Christian Kloesel, who has been working on and editing many of those manuscripts for many years. There is much more that I (and most other people) still haven't seen.

>> "Now besides necessary Reasoning (since the reasoning of the doctrine of
>> chances is merely mathematical reasoning about Probability and all
>> mathematical reasoning is Necessary,) there is only Abduction, or
>> conjectural reasoning, and Induction, or experimental reasoning; and both
>> of these depend upon Necessary Reasoning. At least Induction does. It

>> follows that all reasoning may be represented by Graphs. . . .

>Is there an unwarranted leap from a necessary condition to a sufficient
>condition here?

It seems so. But there is undoubtedly more that we haven't seen, so it's hard to tell what other points he might have mentioned.

>> "Now Existential Graphs furnish us the best diagram of Thought that has
>> ever yet been invented. And do not forget that I have only developed one
>> department of it. There are countless Objects of consciousness that words
>> cannot express; such as the feelings a symphony inspires or that which is
>> in the soul of the furiously angry man in presence of his enemy. But all
>> these can perfectly be expressed in Graphs. Let us call all that ever
>> could be present to the mind in any way or any sense, when taken
>> collectively, the Phaneron. Then every thought is a Constituent of the
>> Phaneron, and much besides that would not ordinarily be called a thought.

>Comment: the work of Rodney Brooks and his team on embodied intelligence
>asserts that the 'feelings a symphony inspires' is not something that
>is expressible in any syntactic formalism, but is buried deep in the
>neural architecture. I do not know if he is right, but it gives me
>pause.

Yes, I am especially suspicious of the word "perfectly". It is the kind of term that mathematicians and philosophers use when they are trying to convince someone else of something that they aren't sure of themselves. (See footnote below.)

On the other hand, Peirce had written extensively about continuity, and he felt that Cantor's claim that the cardinality of the real numbers was equal to the cardinality of points in a line was dubious. In particular, he raised the old question about what happens to the mid-point when you break a line segment in two. According to Euclidean geometry, it should be possible to bisect a line in two exactly congruent halves. But by Cantor's hypothesis, the midpoint can only be in one of the two segments: one of them would be a closed line segment, and the other would be half open.

This paradox also troubled Kurt Goedel. Peirce believed that a truly continuous line must have a much greater cardinality than Cantor claimed.

That indicates that Peirce might have some deeper concept of graphs, which might include continuous ones, which might be able to capture more than what could be representable in any discrete notation. But those thoughts of his were also among his later musings, most of which are still unpublished, unedited, and mostly unavailable.

John

*Note: A typical Norbert Wiener anecdote illustrates the mathematical use of the terms "perfectly," "obvious," or "perfectly obvious." At MIT, a favorite term for something especially obscure was "perfectly obvious to the most casual observer."

One day, Weiner was giving a lecture in mathematics, which consisted of his writing equation after equation on the blackboard while the audience was copying everything furiously in the hope of getting something that they might be able to decipher after many hours of sweat and tears. Then after writing one equation on the board, he said "From this, it is obvious that..." and he wrote down another equation.

But then he stopped and stared at the board for 10 minutes without saying a word. Then he went to the side of the board, which he hid with his rather large bulk, and went clickety-clack with the chalk while writing many more equations that no one else could see. Finally, he erased it all, turned around to the audience, and proclaimed triumphantly, "I was right! It is obvious!"

=====
=====

To post a message, send mail to cg@cs.uah.edu.
To unsubscribe, send mail to majordomo@cs.uah.edu with the command 'unsubscribe cg' in the message body.

=====
=====

For help or administrative assistance, mail to owner-cg@cs.uah.edu

=====
=====

=====
=====

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

=====
=====

=====
=====

From ???@??? Tue May 12 10:49:22 1998
Received: from quark.isi.edu (quark.isi.edu [128.9.208.208])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id VAA16278
for <phayes@coginst.uwf.edu>; Mon, 11 May 1998 21:43:17 -0500 (CDT)
Received: from [128.9.208.73] (skydome.isi.edu [128.9.208.73])
by quark.isi.edu (8.8.7/8.8.6) with ESMTP id TAA27313;
Mon, 11 May 1998 19:40:18 -0700 (PDT)

Date: Mon, 11 May 1998 19:40:18 -0700 (PDT)
X-Sender: hovy@quark.isi.edu
Message-Id: <v0302090ab17cf566f46f@[128.9.208.73]>
In-Reply-To: <5030100020362828000002L082*@MHS>
Mime-Version: 1.0
Content-Transfer-Encoding: 8bit
To: Martin Van Den Berg <vdberg@us.ibm.com>
From: Eduard Hovy <hovy@ISI.EDU>
Subject: Goal of the workshop
Cc: <phayes@coginst.uwf.edu>, <Andreas.Reuter@eml.villa-bosch.de>, <skydog@pacbell.net>, <hovy@ISI.EDU>, <geo@clarity.princeton.edu>, <sowa@west.poly.edu>, <Klaus.Tschira@ktf.villa-bosch.de>, <polanyi@pal.xerox.com>, <guarino@ladseb.pd.cnr.it>, <Piek.Vossen@let.uva.nl>, <peters@csl.stanford.edu>, <Adam_Farquhar@ksl.stanford.edu>, <fellbaum@clarity.princeton.edu>, <cmenzel@tamu.edu>, <doug@csi.uottawa.ca>, <feigenbaum@ksl.stanford.edu>, <fritz@cyc.com>, <jmc@cs.stanford.edu>, <jamesp@cs.brandeis.edu>, <tsujii@is.s.u-tokyo.ac.jp>, <lreeker@nsf.gov>, Martin Van Den Berg <vdberg@us.ibm.com>, <p.m.simons@leeds.ac.uk>, <wahlster@dfki.uni-sb.de>, <Baerbel.Mack@eml.villa-bosch.de>, <Beate.Keller@kts.villa-bosch.de>, <Susanne.Winkelmann@eml.villa-bosch.de>
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2422
Status:

Hello all,

At 12:56 AM -0000 5/12/98, Martin Van Den Berg wrote:
>I am getting a bit confused about what the goal of the workshop is supposed
>to be. Surely it cannot be that we plan to solve the whole of Philosophy
>and Semantics in this week *and* have an excursion in the middle of it.

Absolutely right! :-)

We do not want to spoil philosophers' fun for the next few centuries.

The major goals of the workshop were spelled out at our meeting at CSLI in January:

GOALS

Create the topmost region of an ontology,
building on all experience present,
using this exercise to:

1. identify problematic issues;

2. propose solutions and methods;
3. test their workability.

The 'top 40' EuroWordNet concepts are simply a convenient (or perhaps not) starting point for discussion.

In particular, the major issue is to flesh out the following questions:

1. How are candidate concepts generated?
 - Lattice-based feature combination;
 - Formal ontological distinctions;
 - Language(s)-based lexicon;
 - Inference-based concept generalizations on domain models.
2. How are they validated?
 - Linguistic sentence forms;
 - Formal semantics;
 - KR/AI inference: deductive consistency;
 - Philosophical agreement;
 - Psychology tests on people.
3. Source data:
 - objects, phrases, qualities, etc.;
 - lexicons;
 - previous work in Philosophy (Formal Ontology) and Linguistics (Lexical Semantics);
 - intuition.

By running the top ontology creation exercise, we will inevitably come up with a set of particular methods and sources for the three above questions. Our goal, as I see it, is to record these particular methods and sources as exactly as we can, so as to produce a reliable guide for future creators of the Reference Ontology.

It is of less importance *exactly* which concepts we pick, and whether we get them *exactly* right (since we know we can't, let's not waste time). What we *can* do is develop these guidelines (and clarify our own internal methods of ontologization).

E

Eduard Hovy
email: hovy@isi.edu USC Information Sciences Institute
tel: 310-822-1511 ext 731 4676 Admiralty Way
fax: 310-823-6714 Marina del Rey, CA 90292-6695

project homepage: <http://www.isi.edu/natural-language/nlp-at-isi.html>

From ???@??? Thu May 28 09:58:00 1998

Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id RAA04887;
Wed, 27 May 1998 17:37:00 -0500 (CDT)

Date: Wed, 27 May 1998 17:37:00 -0500 (CDT)

X-Sender: phayes@mail.coginst.uwf.edu (Unverified)

Message-Id: <v04003a01b17e7524f819@[143.88.7.118]>

In-Reply-To: <v0302090ab17cf566f46f@[128.9.208.73]>

References: <5030100020362828000002L082*@MHS>

Mime-Version: 1.0

To: Eduard Hovy <hovy@ISI.EDU>

From: Pat Hayes <phayes@coginst.uwf.edu>

Subject: Re: Goal of the workshop

Cc: AReuter <Andreas.Reuter@eml.villa-bosch.de>,
BSpillers <skydog@pacbell.net>, EHovy <hovy@ISI.EDU>,
GMiller <geo@clarity.princeton.edu>, JSowa <sowa@west.poly.edu>,
KTschira <Klaus.Tschira@ktf.villa-bosch.de>,
LPolanyi <polanyi@pal.xerox.com>, NGuarino <guarino@ladseb.pd.cnr.it>,
PVossen <Piek.Vossen@let.uva.nl>, SPeters <peters@csl.stanford.edu>,
AFarquhar <Adam_Farquhar@ksl.stanford.edu>,
CFellbaum <fellbaum@clarity.princeton.edu>, CMenzel <cmenzel@tamu.edu>,
DSkuce <doug@csi.uottawa.ca>,
EFeigenbaum <feigenbaum@ksl.stanford.edu>, FLehmann <fritz@cyc.com>,
JMcCarthy <jmc@cs.stanford.edu>, JPustejovski <jamesp@cs.brandeis.edu>,
JTsuji <tsujii@is.s.u-tokyo.ac.jp>, LReeker <lreeker@nsf.gov>,
MvdBerg <vdberg@us.ibm.com>, PHayes <phayes@nuts.coginst.uwf.edu>,
PSimons <p.m.simons@leeds.ac.uk>, WWahlster <wahlster@dfki.uni-sb.de>,
Bärbel Mack <Baerbel.Mack@eml.villa-bosch.de>,
Beate Keller <Beate.Keller@kts.villa-bosch.de>,
Susanne Winkelmann <Susanne.Winkelmann@eml.villa-bosch.de>

Content-Type: text/plain; charset="us-ascii"

Content-Length: 1441

Status:

...

>

>In particular, the major issue is to flesh out the following questions:

>

>1. How are candidate concepts generated?

> - Lattice-based feature combination;

> - Formal ontological distinctions;

> - Language(s)-based lexicon;

> - Inference-based concept generalizations on domain models.

We should ask some psychologists to get involved on this topic. There are whole methodologies devoted to accessing 'plausible' concepts. One technique is called, I believe, difference analysis (?). One presents people with three objects A, B and C and asks them to think of a dimension on which B is most like A and most unlike C, ie where the position of B on a scale from A to C is most extremely skewed. The results can be then be subjected to dimensional analysis to discover a smallest number of dimensions which can span the space. This technique has been used in knowledge engineering, social psychology and also by linguists attempting to locate the conceptual space underlying spatial preposition meanings.

I'm sure that there are many other empirical techniques out there that make us all look like amateurs. Maybe we need to get professional help.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Jun 03 09:45:55 1998
Received: from cclsun01.let.uva.nl (cclsun01.let.uva.nl [145.18.228.21])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id CAA18806
for <phayes@nuts.coginst.uwf.edu>; Wed, 3 Jun 1998 02:48:12 -0500 (CDT)
Received: from cclpc102.let.uva.nl by cclsun01.let.uva.nl with SMTP id AA05703
(5.67a/IDA-1.5 for <phayes@nuts.coginst.uwf.edu>); Wed, 3 Jun 1998 09:38:15 +0200
Message-Id: <1.5.4.32.19980603084236.006976d8@mail.let.uva.nl>
X-Sender: piek@mail.let.uva.nl
X-Mailer: Windows Eudora Light Version 1.5.4 (32)
Mime-Version: 1.0
Date: Wed, 03 Jun 1998 09:42:36 +0100
To: Andreas Reuter <Andreas.Reuter@eml.villa-bosch.de>,
AReuter <Andreas.Reuter@eml.villa-bosch.de>,
BSpillers <skydog@pacbell.net>, EHovy <hovy@isi.edu>,
GMiller <geo@clarity.princeton.edu>, JSowa <sowa@west.poly.edu>,
KTschira <Klaus.Tschira@ktf.villa-bosch.de>,
LPolanyi <polanyi@pal.xerox.com>, NGuarino <guarino@ladseb.pd.cnr.it>,
PVossen <Piek.Vossen@let.uva.nl>, SPeters <peters@csl.stanford.edu>,
AFarquhar <Adam_Farquhar@ksl.stanford.edu>,
CFellbaum <fellbaum@clarity.princeton.edu>, CMenzel <cmenzel@tamu.edu>,

DSkuce <doug@csi.uottawa.ca>,
EFeigenbaum <feigenbaum@ksl.stanford.edu>, FLehmann <fritz@cyc.com>,
JMcCarthy <jmc@cs.stanford.edu>, JPustejovski <jamesp@cs.brandeis.edu>,
JTsujii <tsujii@is.s.u-tokyo.ac.jp>, LReeker <lreeker@nsf.gov>,
MvdBerg <vdberg@us.ibm.com>, NLawler <E6NL001@coe.coppin.umd.edu>,
PHayes <phayes@nuts.coginst.uwf.edu>, PSimons <p.m.simons@leeds.ac.uk>,
WWahlster <wahlster@dfki.uni-sb.de>

From: Piek Vossen <Piek.Vossen@let.uva.nl>
Subject: Re: Heidelberg Ontology Workshop
Cc: Bärbel Mack <Baerbel.Mack@eml.villa-bosch.de>,
Beate

Keller <Beate.Keller@kts.villa-bosch.de>,
Susanne Winkelmann <Susanne.Winkelmann@eml.villa-bosch.de>

Content-Type: text/plain; charset="us-ascii"

Content-Length: 461

Status:

Due to very serious personal issues I have to return a day earlier from the Heidelberg workshop. I cannot stay for the final day, Tuesday June 16 and have to leave at Monday June 15th in the evening.

Please accept my apologies for this and I hope I can still sufficiently contribute in the remaining 6 days.

best wishes,

Piek.
Piek Vossen
Universiteit van Amsterdam
Spuistraat 134
1012 VBAmsterdam
The Netherlands

tel. +31 20 525 4669

fax. +31 20 525 4429

From ???@??? Tue May 26 14:01:22 1998

Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA24913;
Tue, 26 May 1998 11:23:10 -0500 (CDT)

Date: Tue, 26 May 1998 11:23:10 -0500 (CDT)

X-Sender: phayes@mail.coginst.uwf.edu

Message-Id: <v04003a10b189d1100aae@[143.88.7.118]>

Mime-Version: 1.0

To: AReuter <Andreas.Reuter@eml.villa-bosch.de>,
BSpillers <skydog@pacbell.net>, EHovy <hovy@isi.edu>,

GMiller <geo@clarity.princeton.edu>, JSowa <sowa@west.poly.edu>, KTschira <Klaus.Tschira@ktf.villa-bosch.de>, LPolanyi <polanyi@pal.xerox.com>, NGuarino <guarino@ladseb.pd.cnr.it>, PVossen <Piek.Vossen@let.uva.nl>, SPeters <peters@csl.stanford.edu>, AFarquhar <Adam_Farquhar@ksl.stanford.edu>, CFellbaum <fellbaum@clarity.princeton.edu>, CMenzel <cmenzel@tamu.edu>, DSkuce <doug@csi.uottawa.ca>, EFeigenbaum <feigenbaum@ksl.stanford.edu>, FLehmann <fritz@cyc.com>, JMcCarthy <jmc@cs.stanford.edu>, JPustejovski <jamesp@cs.brandeis.edu>, JTsuji <tsuji@is.s.u-tokyo.ac.jp>, LReeker <lreeker@nsf.gov>, MvdBerg <vdberg@us.ibm.com>, PHayes <phayes@nuts.coginst.uwf.edu>, PSimons <p.m.simons@leeds.ac.uk>, WWahlster <wahlster@dfki.uni-sb.de>, Bärbel Mack <Baerbel.Mack@eml.villa-bosch.de>, Beate Keller <Beate.Keller@kts.villa-bosch.de>, Susanne Winkelmann <Susanne.Winkelmann@eml.villa-bosch.de>

From: Pat Hayes <phayes@coginst.uwf.edu>

Subject: upperlevels and orthogonality

Content-Type: text/plain; charset="us-ascii"

Content-Length: 2579

Status:

(First, a minor request: I seem to be getting two of many of these messages. Could you please take a second to check your Cc list and delete one of me? Thanks.)

Im glad that John Sowa and I agree on the need for catholicism in the upper ontological levels. Let me suggest that we ask the CYC group (or anyone else with data) for feedback on a related issue that may be central to the task of the workshop.

I suspect that the highest levels of an ontology are in fact the least important, since they deal with such high-level abstractions that not a lot can be usefully said about them, and what can be said can be relatively harmlessly translated into things said about an alternative upper level inspired by a different metaphysical opinion. Middle-level concepts like 'table' and 'building' are useful precisely because there is quite a lot to be usefully said (or believed) about such things. If we were to carve up nature's joints differently at this level, things would get very hard to handle: or at any rate, it would require a very different strategy for organizing the knowledge. On the other hand, whether or not we take a table to be primarily a physical-thing, or a concrete-object or a temporally-extended-entity or an object-with-a-social-function, etc. , seems to be important only if we have enough to say about these various classifications; and the higher we go, the less there is to usefully say. (In wordnets, the definitions become circular.)

Now, do the CYCers have any statistics on how often a query couched in 'middle-level' concepts needs information couched in 'higher-level' concepts in order to be answered, and how far up the isa heirarchy from the level at which the query is made is the information usually found? To put the question a different way, how far up (from the 'middle') does inferencing typically or usually need to look? Or to put it yet another way, how far down the heirarchy do properties typically get inherited?

This kind of question seems to be important in knowing how to design the upper level(s). It would surely be better to abstract them from the middle levels than to impose them a priori, but how much insight does the middle give us into what the top must be like? If the answer is, not much, then we may waste time in the clouds.

Pat Hayes

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed May 27 10:37:39 1998

Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id BAA03026

for <phayes@coginst.uwf.edu>; Wed, 27 May 1998 01:27:35 -0500 (CDT)

Received: from west (west.poly.edu [128.238.20.21]) by hobbes.poly.edu (8.7.3/8.7.3) with SMTP id CAA17390; Wed, 27 May 1998 02:19:59 -0400 (EDT)

Received: by west (SMI-8.6/SMI-SVR4)

id CAA11676; Wed, 27 May 1998 02:18:30 -0400

Date: Wed, 27 May 1998 02:18:30 -0400

From: sowa@west.poly.edu (John F. Sowa)

Message-Id: <199805270618.CAA11676@west>

To: Adam_Farquhar@ksl.stanford.edu, Andreas.Reuter@eml.villa-bosch.de,
Baerbel.Mack@eml.villa-bosch.de, Beate.Keller@kts.villa-bosch.de,
Klaus.Tschira@ktf.villa-bosch.de, Piek.Vossen@let.uva.nl,
Susanne.Winkelmann@eml.villa-bosch.de, cmenzel@tamu.edu,
doug@csi.uottawa.ca, feigenbaum@ksl.stanford.edu,
fellbaum@clarity.princeton.edu, fritz@cyc.com,
geo@clarity.princeton.edu, guarino@ladseb.pd.cnr.it, hovy@isi.edu,
jamesp@cs.brandeis.edu, jmc@cs.stanford.edu, lreeker@nsf.gov,
p.m.simons@leeds.ac.uk, peters@csl.stanford.edu,
phayes@coginst.uwf.edu, polanyi@pal.xerox.com, skydog@pacbell.net,

tsujii@is.s.u-tokyo.ac.jp, vdberg@us.ibm.com, wahlster@dfki.uni-sb.de
Subject: Re: upperlevels and orthogonality
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 3732
Status:

Pat,

I agree that the upper levels of many, if not most ontologies are not important to the way the ontologies are used. But that is not evidence for the claim that a *good* upper level would not be widely used.

An example of what I would consider a good distinction that should appear somewhere in almost every ontology is the one I call Physical vs. Abstract. The labels you pick for those two categories are useful as mnemonic aids, but to avoid getting hung up on the labels, I would be just as happy with less readable labels like P vs. A.

A distinction like that is so fundamental that you cannot give a definition in terms of anything more primitive. Instead, you just have to take your labels, say "P" and "A", as undefined primitives. But that certainly does not mean that they are devoid of meaning. On the contrary, their meaning consists in the axioms associated with them:

1. For P, the primary axiom is "Having a position in space-time", which you can, if you like, translate to whatever symbolism you prefer. That axiom is not a definition, since position, space, and time are less fundamental terms that have not yet been defined. But it is a placeholder for a family of very important inferences for everything that is classified under P. Another axiom would be "Having mass or energy". That is similar to the first one, since it also depends on terms that are much less fundamental than P. In fact, the notion of P has occurred to people long before the notions of mass and energy were clarified in the terms of modern physics.
2. For A, there is the negative axiom "Does not have a position in space-time." That is not as satisfactory, since negative statements aren't as useful in clarifying meaning or pinning it down. But there are other axioms that are more characteristic of A: "Can be encoded in some entity that is P." But this also introduces words like "encode", which are far less fundamental than A or P. A related axiom is that "Any A encoded in an entity at one location can be transmitted to an encoding of A in another entity at another location without changing the total mass or energy at either location." This of course, introduces many other terms that are less fundamental than A or P, but it is an

example of a property that is characteristic of everything that is abstract.

Distinctions like P vs. A are nothing new. Heraclitus used the terms _physis_ (nature) vs. _logos_ (translated variously as word, speech, reason, account, etc., etc.). Modern computer types make the distinction between atoms and bits, and they are groping towards very much the same intuition.

This distinction shows up again and again in everyday inferences. If someone asks you "Send me a copy of Tolstoy's _War and Peace_", you can send an A version by email, but you have to use snail mail for a P version.

Doug Skuce, Nicola Guarino, and I have been talking to one another at many ontology gatherings, and we are firmly convinced that there exist a number of such fundamental distinctions, whose cross product forms a top-level ontology that is inherited throughout every level. Furthermore, such distinctions lead to inferences that are fundamental to the most mundane kinds of actions (such as sending someone a book). However, I can't claim that Doug, Nicola, and I all agree on the same list of fundamental distinctions -- that is another issue to be discussed in Heidelberg.

John Sowa

PS Re duplicate copies: To send this note, I typed "R" in response to your note, and I noticed that there were two copies of your email address in the header, but with minor variations. I deleted one of them.

From ???@??? Wed May 27 17:21:45 1998

Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id PAA04181;
Wed, 27 May 1998 15:49:29 -0500 (CDT)

Date: Wed, 27 May 1998 15:49:29 -0500 (CDT)

X-Sender: phayes@mail.coginst.uwf.edu

Message-Id: <v04003a07b191a533efb6@[143.88.7.118]>

In-Reply-To: <199805270618.CAA11676@west>

Mime-Version: 1.0

To: sowa@west.poly.edu (John F. Sowa)

From: Pat Hayes <phayes@coginst.uwf.edu>

Subject: Re: upperlevels and orthogonality

Cc: Andreas.Reuter@eml.villa-bosch.de, skydog@pacbell.net, hovy@isi.edu,
geo@clarity.princeton.edu, sowa@west.poly.edu,
Klaus.Tschira@ktf.villa-bosch.de, polanyi@pal.xerox.com,
guarino@ladseb.pd.cnr.it, Piek.Vossen@let.uva.nl,

peters@csl.stanford.edu, Adam_Farquhar@ksl.stanford.edu,
fellbaum@clarity.princeton.edu, cmenzel@tamu.edu, doug@csi.uottawa.ca,
feigenbaum@ksl.stanford.edu, fritz@cyc.com, jmc@cs.stanford.edu,
jamesp@cs.brandeis.edu, tsujii@is.s.u-tokyo.ac.jp, vdberg@us.ibm.com,
phayes@nuts.coginst.uwf.edu, p.m.simons@leeds.ac.uk,
wahlster@dfki.uni-sb.de

Content-Type: text/plain; charset="us-ascii"

Content-Length: 10882

Status:

>I agree that the upper levels of many, if not most ontologies are not
>important to the way the ontologies are used. But that is not evidence
>for the claim that a *good* upper level would not be widely used.

Who could argue with such a claim? As you know, I am far less enthusiastic than you about the wisdom of trying to *standardise* any kind of upper level, no matter how good it might be. But maybe we should put that issue aside for now.

>An example of what I would consider a good distinction that should appear
>somewhere in almost every ontology is the one I call Physical vs. Abstract.
>The labels you pick for those two categories are useful as mnemonic aids,
>but to avoid getting hung up on the labels, I would be just as happy with
>less readable labels like P vs. A.

OK, I'm happy with this: I also think that some kind of such distinction is central, and I wholeheartedly endorse your overall plan of trying to find a number of such distinctions and then form an upper level by multiplication.

However....Im not sure that your A/P distinction is quite the same as mine. The distinction I would urge is between things that are temporally located (and maybe spatially, though thats less important) and things that arent. Let me use the terminology T/NT . Examples of T things include physical objects, events, headaches and ocean waves. T things, characteristically, have a temporal location and duration; it makes sense to talk of when they begin and when they end. It makes sense to use tensed language to refer to them. Examples of NT things include numbers, velocities, geometric figures like a square, and mathematical objects like functions. Interestingly, times are NT; it doesnt make sense to ask when 3pm began, for example, and thinking of a time as temporally located quickly leads to paradox. (In general, the coordinate of something in a coordinate system cannot itself be positioned in that same coordinate system: locations dont have a location, times dont have a time, etc.)

This seems like your P/A distinction, except that yours seems to have a lot more extra baggage. For example, I wouldnt want to say that every NT object

can be represented by a T one; some NT things may have no representation (noncomputable functions, say, and most of the real numbers) and one may want to allow one NT thing to represent another (the number three representing the third object on a list, say) ; but in any case, this temporal distinction doesn't seem to have anything particularly to do with representation: that seems to be an orthogonal matter. Similarly for energy and mass. Something may be (spatio)temporally located but not *physical* in this 'massive' sense: examples include events such as conversations and the having of an idea (eg Kekule's famous dream of the snake which led him to think of the benzene ring.)

There are some debateable cases. Consider for example 'the universe'; is this T or NT? My own preference for NT (since it's hard to see what it means to say that the universe has a position in spacetime) but of course this doesn't imply that the universe isn't physical.

Another distinction I would suggest is the Real/Fictional one, by the way. This is quite orthogonal to the T/NT contrast. There are fictional temporal objects (Santa Claus, Moby Dick) and fictional nontemporal things (the largest integer.) The reason we need to allow the last category is to allow proofs by reductio. Consider for example the ancient argument for there being no largest integer: Suppose there were a largest, call it L; then L+1 is larger than L; contradiction; ergo, there is no such L. But notice that the argument requires us to *consider* the possibility of L, if only temporarily, and then reason about it. If there is no category for 'fictional abstract' objects, then this reasoning becomes incoherent.

Even here there are some debateable cases. Some things seem half-real. Consider for example Good King Wenceslas (hero of a popular Xmas carol). There really was a king in ancient Saxony whose charitable nature gave rise to the legend which eventually was celebrated in this song, but the link between the actual king and the character in the song is tenuous, to say the least. Or consider the Julius Caesar in Shakespeare's play: do we want to say that there are two Caesars, one real and one fictional? Or is it better to say that there is one, but some of the things said about him by Shakespeare might not be accurate?

>A distinction like that is so fundamental that you cannot give a definition
>in terms of anything more primitive. Instead, you just have to take your
>labels, say "P" and "A", as undefined primitives. But that certainly does
>not mean that they are devoid of meaning. On the contrary, their meaning
>consists in the axioms associated with them:

Im glad you have come to agree with me on this point :-)

> 1. For P, the primary axiom is "Having a position in space-time", which

- > you can, if you like, translate to whatever symbolism you prefer. That
- > axiom is not a definition, since position, space, and time are less
- > fundamental terms that have not yet been defined. But it is a place
- > holder for a family of very important inferences for everything that
- > is classified under P. Another axiom would be "Having mass or energy".
- > That is similar to the first one, since it also depends on terms that
- > are much less fundamental than P. In fact, the notion of P has occurred
- > to people long before the notions of mass and energy were clarified
- > in the terms of modern physics.

There's a danger here in talking about modern physics. Modern physics is wildly unintuitive: so unintuitive that it is hard to even describe it in English, which makes ontological assumptions which are violated by it. Space and time are inseparable in relativity theory, for example, and the vacuum is full of energy according to QED. There are things with negative mass and (arguably) backward time dimensions; and so on. I think if we appeal to modern physics and try to reconcile it with anything like intuitive thought or language use we will run into insuperable difficulties. For example, you claim that P can be characterised by 'having a position in space-time' (which is also my T category), and that P is an ancient notion. But "space-time" was invented by Minkowski, who was Einstein's teacher: the very concept would have been alien to Newton, let alone the ancients. So whatever Heraclitus was talking about, this wasn't it.

- > 2. For A, there is the negative axiom "Does not have a position in space-
- > time." That is not as satisfactory, since negative statements aren't
- > as useful in clarifying meaning or pinning it down. But there are
- > other axioms that are more characteristic of A: "Can be encoded in
- > some entity that is P." But this also introduces words like "encode",
- > which are far less fundamental than A or P. A related axiom is that
- > "Any A encoded in an entity at one location can be transmitted to an
- > encoding of A in another entity at another location without changing
- > the total mass or energy at either location." This of course, introduces
- > many other terms that are less fundamental than A or P, but it is an
- > example of a property that is characteristic of everything that is
- > abstract.

Is that really true? Bear in mind that according to thermodynamics, any transmission of information must use energy. And it isn't true, of course, that information can really be transmitted between locations arbitrarily.

- > Distinctions like P vs. A are nothing new. Heraclitus used the terms
- > _physis_ (nature) vs. _logos_ (translated variously as word, speech,
- > reason, account, etc., etc.). Modern computer types make the distinction
- > between atoms and bits, and they are groping towards very much the same

>intuition.

>

>This distinction shows up again and again in everyday inferences.

>If someone asks you "Send me a copy of Tolstoy's _War and Peace_",

>you can send an A version by email, but you have to use snail mail

>for a P version.

Well now that's a VERY debateable claim! Here you seem to be confusing (?) the physical/abstract distinction with the token/type distinction. I'd suggest that token/type belongs much lower in the hierarchy: it applies only to texts (things with a syntax) which are a very particular kind of thing. Also, by the way, a thought: although it is tempting to say that a type (unlike a token) is abstract, it isn't clear whether or not types can be temporally located. It seems to make sense, for example, to speak of when a certain alphabet first came into use, or to speak of 'dead languages', whose typology no longer exists. If statements like this make sense, then types would seem to belong on the T rather than the NT side. There is always the alternative of saying that the type exists in a timeless Platonic space, along with the mathematical stuff, but this has the slightly odd consequence that this timeless domain then has to contain all the languages that will exist in the future, or even, arguably, all the languages that *could* exist in a possible future. Also, on this view, one needs to make a three-way distinction between the type (NT), the token (T) and something like the region of cultural usage of tokens of that type (also T, but rather broader in extent, eg consider the timeperiod during which ancient Sumerian was in use.)

But in any case, surely the thing that is emailed is a physical token? It may be encoded electronically, but it is just as physical as a piece of paper: it has a spatiotemporal location, for example (it moves from one place at one time to another at another time with a measurable velocity) and even uses energy. The fact that it can be copied more easily doesn't seem enough to make it A instead of P. On that criterion, for example, all printed documents ought to have been transferred from P to A when Xerox copiers became widely available.

>Doug Skuce, Nicola Guarino, and I have been talking to one another

>at many ontology gatherings, and we are firmly convinced that there

>exist a number of such fundamental distinctions, whose cross product

>forms a top-level ontology that is inherited throughout every level.

I'm glad you three agree with the ancient philosophers on this.

...

>However, I can't claim that Doug, Nicola, and I all agree on the same

>list of fundamental distinctions...

Just like the ancient philosophers, in fact.

>.. -- that is another issue to be discussed
>in Heidelberg.

Back to two thousand years in a week?

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Thu May 28 09:58:02 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
 by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id XAA07791
 for <phayes@coginst.uwf.edu>; Wed, 27 May 1998 23:35:17 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21]) by hobbes.poly.edu (8.7.3/8.7.3)
with SMTP id AAA03762; Thu, 28 May 1998 00:27:49 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
 id WAA16569; Wed, 27 May 1998 22:09:16 -0400
Date: Wed, 27 May 1998 22:09:16 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199805280209.WAA16569@west>
To: phayes@coginst.uwf.edu
Subject: Re: upperlevels and orthogonality
Cc: Adam_Farquhar@ksl.stanford.edu, Andreas.Reuter@eml.villa-bosch.de,
 Klaus.Tschira@ktf.villa-bosch.de, Piek.Vossen@let.uva.nl,
 cmenzel@tamu.edu, doug@csi.uottawa.ca, feigenbaum@ksl.stanford.edu,
 fellbaum@clarity.princeton.edu, fritz@cyc.com,
 geo@clarity.princeton.edu, guarino@ladseb.pd.cnr.it, hovy@isi.edu,
 jamesp@cs.brandeis.edu, jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk,
 peters@csl.stanford.edu, phayes@nuts.coginst.uwf.edu,
 polanyi@pal.xerox.com, skydog@pacbell.net, tsujii@is.s.u-tokyo.ac.jp,
 vdberg@us.ibm.com, wahlster@dfki.uni-sb.de
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 6607
Status:

Pat,

Again, I think you are primarily disagreeing with what you think I might be saying rather than with what I have been saying.

>Who could argue with such a claim? As you know, I am far less enthusiastic
>than you about the wisdom of trying to *standardise* any kind of upper
>level, no matter how good it might be. But maybe we should put that issue
>aside for now.

I have never thought that it would be possible or even desirable to have a single "standard" ontology that everyone should use or conform to. There are many people who have been involved in this ontology business who have made many different claims and suggestions. For the record, following is a brief summary of what I think could be done quite usefully within the scope of this ontology effort. It might not be finished in one week in Heidelberg, but I think that we can at least get some understanding of how we might proceed:

1. Take stock of some of the resources that are available, such as Cyc, WordNet, EDR, and various other projects that are represented by one or more of the people who will be in Heidelberg. Analyze what, if anything, they have that is usable for some sort of collaboration, and suggest ways in which those resources could be used, if the owners are willing to make them available.
2. Make a list of the basic distinctions, such as my P/A or your T/NT distinction, and the various ones that Nicola, Doug, and others have been proposing. Begin to analyze them and determine how they might be related, what axioms are implied by saying that some entity x belongs to category C, etc.
3. Determine what other distinctions are implicit (or explicit) in the top levels of the ontologies we are considering in point #1, and add them to the list in point #2.
4. Consider any other distinctions that anyone might suggest from the literature of philosophy, linguistics, lexicography, etc., over the past 2,500 years. A large number of them will very likely be present in one form or another in the lists developed for #1, #2, and #3, but we should allow new ones to be added at any time.
5. These lists developed in points #1-4 above should be open ended, so that anyone can suggest new distinctions that seem to be overlooked. These lists and the criticisms of them should all be made available on the WWW for open examination, analysis, and discussion by anyone

who might be interested. The only reason why the Heidelberg meeting has a limited attendance list is that the size of the meeting room cannot accommodate more than 25, and a larger number would make discussions difficult. But the meeting in June is only intended to be one of many that may be held in one forum or another, and one purpose of the Heidelberg meeting is to work out a way for us to continue the collaboration and keep it open to all interested parties.

6. Given a list of distinctions that have been suitably critiqued and analyzed. it should be possible for anyone at any time to take whatever selection seems appropriate, push a button on an appropriate tool (many prototypes of which have been implemented) and generate a new top level. See, for example, that paper on the Knowledge Bus that I recommended in an earlier note. That tool used Cyc as the resource for generating a suitable domain-specific ontology. But similar tools could be applied to the outputs of the efforts suggested in points #1-5.

>Another distinction I would suggest is the Real/Fictional one, by the way.
>This is quite orthogonal to the T/NT contrast. There are fictional temporal
>objects (Santa Claus, Moby Dick) and fictional nontemporal things (the
>largest integer.)....

Modality is certainly important, and real/fictional is one kind of modal distinction that must be considered somewhere. That gets into the kinds of issues we have been discussing at the workshops on context. The question of how context is related to the ontology is important and should be considered. For most purposes, existence in the real world or some possible or fictional one is orthogonal to the definition of the categories. We can talk about unicorns or space stations on Mars independent of their actual existence.

>... But "space-time" was invented by Minkowski, who was
>Einstein's teacher: the very concept would have been alien to Newton, let
>alone the ancients. So whatever Heraclitus was talking about, this wasn't
>it.

We will need to consider many kinds of distinctions that may depend on other distinctions. The P/A distinction that I was proposing is general enough that it could be understood quite well by Heraclitus, Aristotle, or Moses. It does, however, have many implications, which could be stated in axioms, if anyone would care to add them to some particular domain. The one about space-time might be dropped from some applications and be inserted for others. And I'm sure that there are other possible axioms that future physicists will discover. But none of their discoveries will make it irrelevant.

>Is that really true? Bear in mind that according to thermodynamics, any
>transmission of information must use energy. And it isn't true, of course,
>that information can really be transmitted between locations arbitrarily.

The transmission of information changes the entropy, and in an inefficient system (which most are) it would undoubtedly cause some minute energy change. This is an example of an axiom that would have to be qualified; the qualification, however, is typical of the kind that appear in all physical reasoning, which is only accurate up to a granularity determined by your measuring instruments.

>Well now that's a VERY debateable claim! Here you seem to be confusing (?)
>the physical/abstract distinction with the token/type distinction. I'd
>suggest that token/type belongs much lower in the hierarchy: it applies
>only to texts (things with a syntax) which are a very particular kind of
>thing....

The words "type" and "token" were introduced by Peirce to discuss two terms of one of his triads (mark, token, type). There are a lot of related issues in his semiotics, which is something else that is very significant to ontology. We have discussed parts of his semiotics many times, and I'm sure that we will discuss more of it in the future.

>Back to two thousand years in a week?

Why not? The amount of material in Cyc, WordNet, and EDR would probably fill more CD-ROMs than all of the texts of the major philosophers of the past 2,000 years.

John

From ???@??? Thu May 28 16:52:08 1998

Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id QAA02077;
Thu, 28 May 1998 16:30:23 -0500 (CDT)

Date: Thu, 28 May 1998 16:30:23 -0500 (CDT)

X-Sender: phayes@mail.coginst.uwf.edu

Message-Id: <v04003a04b192e7db9af2@[143.88.7.118]>

In-Reply-To: <199805280209.WAA16569@west>

Mime-Version: 1.0

To: sowa@west.poly.edu (John F. Sowa)

From: Pat Hayes <phayes@coginst.uwf.edu>

Subject: Re: upperlevels and orthogonality

Cc: AReuter <Andreas.Reuter@eml.villa-bosch.de>,
BSpillers <skydog@pacbell.net>, EHovy <hovy@isi.edu>,
GMiller <geo@clarity.princeton.edu>, JSowa <sowa@west.poly.edu>,
KTschira <Klaus.Tschira@ktf.villa-bosch.de>

LPolanyi <polanyi@pal.xerox.com>, NGuarino <guarino@ladseb.pd.cnr.it>, PVossen <Piek.Vossen@let.uva.nl>, SPeters <peters@csl.stanford.edu>, AFarquhar <Adam_Farquhar@ksl.stanford.edu>, CFellbaum <fellbaum@clarity.princeton.edu>, CMenzel <cmenzel@tamu.edu>, DSkuce <doug@csi.uottawa.ca>, EFeigenbaum <feigenbaum@ksl.stanford.edu>, FLehmann <fritz@cyc.com>, JMcCarthy <jmc@cs.stanford.edu>, JPustejovski <jamesp@cs.brandeis.edu>, JTsuji <tsujii@is.s.u-tokyo.ac.jp>, LReeker <lreeker@nsf.gov>, MvdBerg <vdborg@us.ibm.com>, PHayes <phayes@nuts.coginst.uwf.edu>, PSimons <p.m.simons@leeds.ac.uk>, WWahlster <wahlster@dfki.uni-sb.de>, Bärbel Mack <Baerbel.Mack@eml.villa-bosch.de>, Beate Keller <Beate.Keller@kts.villa-bosch.de>, Susanne Winkelmann <Susanne.Winkelmann@eml.villa-bosch.de>

Content-Type: text/plain; charset="us-ascii"

Content-Length: 12100

Status:

Hi John

>Again, I think you are primarily disagreeing with what you think I might
>be saying rather than with what I have been saying.

Then we agree more than we thought we did. Good!

.....

>

> 6. Given a list of distinctions that have been suitably critiqued and
> analyzed. it should be possible for anyone at any time to take whatever
> selection seems appropriate, push a button on an appropriate tool (many
> prototypes of which have been implemented) and generate a new top level.
> See, for example, that paper on the Knowledge Bus that I recommended
> in an earlier note. That tool used Cyc as the resource for generating
> a suitable domain-specific ontology. But similar tools could be
> applied to the outputs of the efforts suggested in points #1-5.

This raises an issue I wanted to raise independently, so here goes: are we to assume that these distinctions, whatever we ultimately decide they are, must be exhaustive? In the usual picture of defining an upper-level ontology, I presume this must be the case: if we make your A/P distinction, then everything below that point in the heirarchy must be either A or P , or maybe both; but it can't be *neither*. This is one reason I am leery of putting too much axiomatic baggage onto these high-level distinctions, as these are liable to exclude some things from either branch. (see below for some examples.) However, with your 'shopping-list' approach outlines above, it is perfectly sensible for someone to just say, Im not going to make this particular distinction; and this of course doesnt then exclude things that cannot be classified by it. So this shopping-list picture is more catholic

in an important respect.

I think we need to get this issue clear, in any case, as otherwise we are liable to find ourselves accidentally excluding all kinds of useful concepts. After our recent A/P or T/NT discussion I thought of the following, for example:

== holes (temporal and corporeal but of course having no mass)
(A lovely special case: a 'shed' in weaving. The shed is the space between the stretched-apart warp threads on a loom through which the shuttle is thrown by the weaver; but the same pices of space is considered a different 'shed' depending on the setting of the heddles. Thus in a simple weave the heddles are set, revealing the 'first shed' through which the shuttle passes, then they reverse, revealing the 'second shed' through which the shuttle returns, and so on. In more complex weaving there may be many sheds. These things are holes which exist intermittently: they come into and go out of existence as the heddles move.)

== surfaces

== textures

(Is a texture - say, 'denim' or 'rough' or 'metallic' - physical? There seems again to be something like the type/token distinction here. Any particular surface has a texture, and that surface, exhibiting its texture, is temporal and concrete; but the texture itself seems not to be spatiotemporally locatable. Maybe textures, like colors, are *properties* of surfaces; if so, then we have to allow these properties to be temporally indexed, since a surface may change its texture or color and yet be the same surface.)

(In general, for 'physical properties' like this, there seem to be two divergent ontological strategies. On one view, all properties are abstract, and one makes a distinction between the roughness-property of a rough surface (A/NT) and the particular piece of surface which is rough (P/T). However, a different strategy would be to say that the roughness was itself physical, temporally located, and so forth, and consider the rough surface as a kind of mereological sum of a 'plain' surface (with no particular texture) and the localised texture which it manifests. On this view, for example, so smoothen a rough surface would be to *remove* its roughness, rather like stripping away a layer of paint. (It occurs to me that one can buy 'textured paint'.) Its going to be important to get things like this clear, as the textures are definitely P/T in the second sense though probably A/NT in the first. Maybe a suitably catholic ontology should allow both of them? It wouldnt be hard to give some general axioms to link all such cases in a uniform way.)

== radiance

(the quality of light being given off by a self-illuminated object; similar distinction applies.)

== strength/fragility/rigidity/etc.

(Fragility is especially interesting, as Aaron Sloman pointed out to me: a thing is fragile if it is *potentially* liable to breakage. But a fragile object may endure for ever, in fact. The concept seems to have an essentially counterfactual component: it is fragile if, *were it to be hit*, it would break. Quite a lot of ordinary physical concepts have this character, including such notions as 'dangerous'. In general, concepts like this seem important because they indicate areas of our surroundings that require us to take more than ordinary care.)

== 'comings together' of various kinds: smooth landings, impacts, caresses, slaps, adhesions, etc. (Here's a nice snippet I saw on a label recently: "Warning: failure of adhesion may occur.")

== interpersonal attitudes: politeness, disdain, contempt, respect, etc. . These seem relational, but I think the same kind of distinction can be made between a 'property' approach and a 'mereological' approach where we refer to these attitudes as entities with properties themselves. ("His impeccable politeness was legendary, but the Abbot thought it masked an ungracious superciliousness.")

== spatial and temporal fields: gravity, winds, currents, alpha rhythms, the sound of an orchestra, and so forth. (What kind of thing is El Nino, for example?)

>> Another distinction I would suggest is the Real/Fictional one, by the way.
>> This is quite orthogonal to the T/NT contrast. There are fictional temporal
>> objects (Santa Claus, Moby Dick) and fictional nontemporal things (the
>> largest integer.)....

>

> Modality is certainly important, and real/fictional is one kind of modal
> distinction that must be considered somewhere. That gets into the kinds
> of issues we have been discussing at the workshops on context. The question
> of how context is related to the ontology is important and should be
> considered.

One way to talk about unicorns is to make assertions inside a context or the scope of a modality (it is possible that...), but an alternative way is to speak quite openly (at the top contextual level, or without modal qualification) about fictional objects, ie things classed by type as fictional. This is what we seem to do when we say that unicorns don't exist (for in the fictional context they *do* exist, so this would be false there; and if we can only speak of real objects, we can't even refer to

unicorns in this context.)

If our job is to design an *ontology* then this would seem the natural way to do it: we can't decree what kind of language the users of our ontology must or must not use, only the signature of the language; so the most conservative course is to incorporate those distinctions into the signature (or at any rate, in the spirit you outline, to allow them to be incorporated if the user wishes.)

>>.... But "space-time" was invented by Minkowski, who was
>>Einstein's teacher: the very concept would have been alien to Newton, let
>>alone the ancients. So whatever Heraclitus was talking about, this wasn't
>>it.

>

>We will need to consider many kinds of distinctions that may depend on
>other distinctions. The P/A distinction that I was proposing is general
>enough that it could be understood quite well by Heraclitus, Aristotle,
>or Moses.

Well, I disagree. First, I don't think any of these distinctions are particularly 'fundamental': they are just distinctions, that's all. Second, I suspect that the truth of the matter is that there are in fact a large number of slightly different distinctions, all roughly corresponding to the division between the clusters {physical, corporeal, temporal, concrete,...} and {abstract, symbolic, Platonic, linguistic,...}, and these different distinctions are ambiguously denoted by the P/A division, rather in the way that the English word 'bank' has about a dozen or more different but related meanings. But vagueness is not profundity.

(The inherent sloppiness of these old terms is attested to by the fact that thousands of people have died ugly deaths in the wars fought over differing interpretations of concepts such as 'incorporeal'. I bet you wouldn't have been able to get Moses and Heraclitus to agree on *anything* much: they would certainly have had very different notions of 'the Word'.)

It does, however, have many implications, which could be
>stated in axioms, if anyone would care to add them to some particular
>domain. The one about space-time might be dropped from some applications
>and be inserted for others. And I'm sure that there are other possible
>axioms that future physicists will discover. But none of their discoveries
>will make it irrelevant.

>>Is that really true? Bear in mind that according to thermodynamics, any
>>transmission of information must use energy. And it isn't true, of course,
>>that information can really be transmitted between locations arbitrarily.

>

>The transmission of information changes the entropy, and in an inefficient
>system (which most are) it would undoubtedly cause some minute energy
>change. This is an example of an axiom that would have to be qualified;
>the qualification, however, is typical of the kind that appear in all
>physical reasoning, which is only accurate up to a granularity determined
>by your measuring instruments.

But the point about entropy is fundamental in (real) physics. You see,
that's my point. Are we appealing to actual physics or to naive physics? You
seem to be using real physics some of the time and naive physics at other
times. There is no concept of entropy in naive physics, and I suspect it
fails to distinguish impact from energy, or heat from temperature. That's
fine with me, of course: but if we are going to write axioms then we need
to be clear whether we are thinking like Moses or like Einstein.

>>Well now that's a VERY debateable claim! Here you seem to be confusing (?)
>>the physical/abstract distinction with the token/type distinction. I'd
>>suggest that token/type belongs much lower in the hierarchy: it applies
>>only to texts (things with a syntax) which are a very particular kind of
>>thing....

>

>The words "type" and "token" were introduced by Peirce to discuss two
>terms of one of his triads (mark, token, type). There are a lot of
>related issues in his semiotics, which is something else that is very
>significant to ontology. We have discussed parts of his semiotics many
>times, and I'm sure that we will discuss more of it in the future.

John, you have a remarkable ability to change the subject! Whoever
introduced it, the type/token distinction is now widely used throughout
linguistics and philosophy and is perfectly clear. Most authors give credit
to Frege, I believe: but in any case, that's just a matter of historical
interest. I don't give a damn about Peircian history, and I think that to
get involved in it at Heidelberg is just a recipe for wasting valuable
time. Burch's recent book claimed to give a rigorous justification of
Peircian triadicity, but as I showed in my review, the 'triadicity'
property vanishes completely when one allows existential quantification. It
doesn't survive even a tiny change in the algebra (eg. introduce a 'merge'
operator of identity relation arcs.) It stems, in fact, from a misleading
and rather simplistic analogy between relational connection and chemical
valency which apparently Peirce found compelling. (It's a pity he didn't
know more about hydrogen bonding or benzene.)

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Fri May 29 10:13:54 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
 by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id AAA06755
 for <phayes@nuts.coginst.uwf.edu>; Fri, 29 May 1998 00:40:14 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
 by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id BAA05954;
 Fri, 29 May 1998 01:32:06 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
 id BAA23246; Fri, 29 May 1998 01:30:27 -0400
Date: Fri, 29 May 1998 01:30:27 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199805290530.BAA23246@west>
To: phayes@coginst.uwf.edu
Subject: Re: upperlevels and orthogonality
Cc: Adam_Farquhar@ksl.stanford.edu, Andreas.Reuter@eml.villa-bosch.de,
 Baerbel.Mack@eml.villa-bosch.de, Beate.Keller@kts.villa-bosch.de,
 Klaus.Tschira@ktf.villa-bosch.de, Piek.Vossen@let.uva.nl,
 Susanne.Winkelmann@eml.villa-bosch.de, cmenzel@tamu.edu,
 doug@csi.uottawa.ca, feigenbaum@ksl.stanford.edu,
 fellbaum@clarity.princeton.edu, fritz@cyc.com,
 geo@clarity.princeton.edu, guarino@ladseb.pd.cnr.it, hovy@isi.edu,
 jamesp@cs.brandeis.edu, jmc@cs.stanford.edu, lreeker@nsf.gov,
 p.m.simons@leeds.ac.uk, peters@csl.stanford.edu,
 phayes@nuts.coginst.uwf.edu, polanyi@pal.xerox.com, skydog@pacbell.net,
 sowa@west.poly.edu, tsujii@is.s.u-tokyo.ac.jp, vdberg@us.ibm.com,
 wahlster@dfki.uni-sb.de

Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 9912
Status:

Pat,

Some further comments on your comments:

>This raises an issue I wanted to raise independently, so here goes: are we
>to assume that these distinctions, whatever we ultimately decide they are,
>must be exhaustive? In the usual picture of defining an upper-level

>ontology, I presume this must be the case: if we make your A/P distinction,
>then everything below that point in the hierarchy must be either A or P ,

I just mentioned the A/P distinction as an example of something very high in the ontology that has implications at every level beneath it. But there are many distinctions that require others as a prerequisite. They would only subdivide the lattice at some lower levels.

>== holes (temporal and corporeal but of course having no mass)
>(A lovely special case: a 'shed' in weaving. The shed is the space between
>the stretched-apart warp threads on a loom through which the shuttle is....

Yes, there are many such examples at high levels and even more at lower levels. There are enormous numbers of distinctions that presuppose that the entity in question belongs to some particular category, such as human, plant, living, mineral, etc.

>If our job is to design an *ontology* then this would seem the natural way
>to do it: we can't decree what kind of language the users of our ontology
>must or must not use, only the signature of the language; so the most
>conservative course is to incorporate those distinctions into the signature
>(or at any rate, in the spirit you outline, to allow them to be
>incorporated if the user wishes.)

I would say that our job is much narrower: we are not responsible for designing the best possible ontology, but rather a set of guidelines for designing ontologies. In that process, we should prove an existence theorem that shows that such ontologies exist, and the simplest way to do that is to give an example. But I would not want to claim that any example we construct next month would be the final one or even a particularly good or stable one.

>>.... But "space-time" was invented by Minkowski, who was
>>Einstein's teacher: the very concept would have been alien to Newton, let
>>alone the ancients. So whatever Heraclitus was talking about, this wasn't
>>it.

This gets into a lot of delicate issues about which we may very well disagree. One distinction that I am not especially happy with is the one between common-sense and scientific reasoning. I believe that there is a continuity, and that one century's scientific breakthroughs become the next century's common sense. I'd like to quote a comment by Whitehead about Newton's notion of absolute space and time, which is very far from being either commonsensical or scientifically true:

For the purposes of science, it was an extraordinarily clarifying statement, that is to say, for all the purposes of science within the next two hundred

years, and for most of its purposes since that period. But as a fundamental statement, it lies completely open to skeptical attack; and also, as Newton himself admits, diverges from common sense -- "the vulgar conceive those quantities under no other notions but from the relation they bear to sensible objects."

That passage is from Whitehead's *Process and Reality*, which he described as "an attempt to return to the conceptions of the vulgar." That does not mean to the unanalyzed conceptions of commonsense, but to a highly refined and deeply analyzed set of concepts that end up rather close to what the common "man in the street" might agree to. And surprisingly, that vulgar notion is closer to Einstein's relative than to Newton's absolute space-time.

>(The inherent sloppiness of these old terms is attested to by the fact that >thousands of people have died ugly deaths in the wars fought over differing >interpretations of concepts such as 'incorporeal'. I bet you wouldnt have >been able to get Moses and Heraclitus to agree on *anything* much: they >would certainly have had very different notions of 'the Word'.)

I don't believe that there is anything inherently wrong with vagueness. Recall that note I circulated a while ago with quotations from Peirce about vagueness. In particular, his point that a vague truth is often much more useful than a precise falsehood.

Re Moses and Heraclitus: The Jewish philosopher Philo of Alexandria, who lived across the divide between BC and AD, wrote quite a lot in his efforts to reconcile the Torah with Greek philosophy. In particular, his writings probably were a strong influence on St. John the Evangelist, whose gospel "In the beginning was the logos, and the logos was with God, and the logos was God..." has very strong echoes of Heraclitus, including many of the same word choices. I would cite those as examples of vague truths that we will probably never be able to pin down precisely, but I believe that any AI system that can claim to reach the level of human intelligence will have to be able to deal with. And such vagueness does not occur only in religion or only in writings that are thousands of years old. I'm sure that you will be able to find many examples of a similar level of vagueness in the email that has been circulated to this list. And I don't believe that it is necessarily bad -- on the contrary, it is inevitable in the early stage of any kind of scientific inquiry.

>But the point about entropy is fundamental in (real) physics. You see, >thats my point. Are we appealing to actual physics or to naive physics? You >seem to be using real physics some of the time and naive physics at other >times. There is no concept of entropy in naive physics, and I suspect it >fails to distinguish impact from energy, or heat from temperature. Thats >fine with me, of course: but if we are going to write axioms then we need

>to be clear whether we are thinking like Moses or like Einstein.

As I said before, I don't believe there is any difference between the kind of thinking of Moses and Einstein. In his philosophical musings, Einstein said many things that are at the same level of precision as the publications of Moses. And I'm sure that in his plans for leading the Israelites out of Egypt, Moses had some rather clever and precise strategems that would not be out of place in modern times. And in another century or two, I'm sure that much of "modern physics" will look rather naive, even to high school students of the day.

>John, you have a remarkable ability to change the subject! Whoever >introduced it, the type/token distinction is now widely used throughout >linguistics and philosophy and is perfectly clear. Most authors give credit >to Frege, I believe: but in any case, that's just a matter of historical >interest. I don't give a damn about Peircian history, and I think that to >get involved in it at Heidelberg is just a recipe for wasting valuable time.

No one who knows anything about the subject credits Frege. The terms were explicitly introduced into the literature by Peirce, and they were borrowed from Peirce's writing into modern linguistic literature by Roman Jakobson, who is considered by everyone except Noam Chomsky to be the greatest linguist of the twentieth century. I am citing Peirce not as a historical figure, but as someone whose ideas are among the avant garde of modern logic and linguistics. Most people who cite the type/token terms do not realize that they are the second and third terms of a Peircean triad. And that point is fundamental because such triads come up again, and again, and again at every level of the ontology.

As just one more example, consider the term "granularity", which is fundamental to any discussion of physical quantities. While I was writing about granularity in my KR book, I happened to notice that people use the term in many different ways. As I was classifying them, I realized that those uses formed a triad: actual granularity in the nature of the subject matter, such as atoms or photons; epistemic granularity in our ability to measure the subject; or intentional granularity in our decisions to ignore detail that is irrelevant to the application. This three way distinction helps to clarify much of the discussion about granularity, but other people didn't notice it because they didn't happen to be thinking in terms of Peircean triads.

> Burch's recent book claimed to give a rigorous justification of >Peircian triadicity, but as I showed in my review, the 'triadicity' >property vanishes completely when one allows existential quantification....

I am not about to defend Burch or his presentation. But when I give you

some of these historical references about Peirce, I am doing so primarily to emphasize that CSP would never have made such a trivial blunder as you are attributing to him. Remember that Peirce invented the algebraic notation for predicate calculus in its modern form in 1883-1885 without any knowledge of Frege's 1879 Begriffsschrift, which was never adopted by anyone, even his own very few students (among whom was Rudolf Carnap, who switched to the algebraic notation as soon as he learned of it).

To give you a few more Peircean "trivia", which are very far from trivial, he introduced the terms "existential quantifier" (with the symbol Sigma for repeated "logical sum") and "universal quantifier" (with the symbol Pi for "logical product"). He also introduced the terms "first-intentional logic" and "second-intentional logic", which Ernst Schroeder translated into German as "erste Ordnung" and "zweite Ordnung" and which Russell translated back into English as "first order" and "second order". In that same paper, he also used "second-intentional" logic to define the equality $x=y$ as

Forall P, $P(x) \Leftrightarrow P(y)$.

And here I'm not using CSP's symbols because I don't have Greek letters. As Casey Stengel used to say "Ya could look it up" -- American Journal of Mathematics, 1885.

John

From ???@??? Tue Jun 02 17:46:47 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id QAA16667;
Tue, 2 Jun 1998 16:45:50 -0500 (CDT)
Date: Tue, 2 Jun 1998 16:45:50 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a09b1944134b9bd@[143.88.7.118]>
In-Reply-To: <199805290530.BAA23246@west>
Mime-Version: 1.0
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: upperlevels and orthogonality
Cc: Andreas.Reuter@eml.villa-bosch.de, skydog@pacbell.net, hovy@isi.edu,
geo@clarity.princeton.edu, sowa@west.poly.edu,
Klaus.Tschira@ktf.villa-bosch.de, polanyi@pal.xerox.com,
guarino@ladseb.pd.cnr.it, Piek.Vossen@let.uva.nl,
peters@csl.stanford.edu, Adam_Farquhar@ksl.stanford.edu,
fellbaum@clarity.princeton.edu, cmenzel@tamu.edu, doug@csi.uottawa.ca,
feigenbaum@ksl.stanford.edu, fritz@cyc.com, jmc@cs.stanford.edu,
jamesp@cs.brandeis.edu, tsujii@is.s.u-tokyo.ac.jp, vdberg@us.ibm.com,

phayes@nuts.coginst.uwf.edu, p.m.simons@leeds.ac.uk,
wahlster@dfki.uni-sb.de

Content-Type: text/plain; charset="us-ascii"

Content-Length: 5411

Status:

>Pat,

>

>Some further comments on your comments:

>

>>This raises an issue I wanted to raise independently, so here goes: are we
>>to assume that these distinctions, whatever we ultimately decide they are,
>>must be exhaustive? In the usual picture of defining an upper-level
>>ontology, I presume this must be the case: if we make your A/P distinction,
>>then everything below that point in the heirarchy must be either A or P ,

>

>I just mentioned the A/P distinction as an example of something very high
>in the ontology that has implications at every level beneath it. But there
>are many distinctions that require others as a prerequisite. They would
>only subdivide the lattice at some lower levels.

>

>>= holes (temporal and corporeal but of course having no mass)
>>(A lovely special case: a 'shed' in weaving. The shed is the space between
>>the stretched-apart warp threads on a loom through which the shuttle is....

>

>Yes, there are many such examples at high levels and even more at lower
>levels. There are enormous numbers of distinctions that presuppose that
>the entity in question belongs to some particular category, such as human,
>plant, living, mineral, etc.

You seem to have missed my point. Of course lower categorisations may assume earlier ones and only be meaningful in their local part of the heirarchy. That is aside from my point, which (to repeat) was that once a split is put into the heirarchy, at whatever level, then everything below that must be classified into one or the other category, or maybe both. So we must take care not to define our categorical splits so as to *exclude* anything. Your A/P, seems to me, excludes holes and surfaces: they are situated in space and time, etc., so are not in A, but they have no mass so cannot be included in your P category as you describe it. They are neither A nor P: they don't have any place to be categorised; and yet they seem much too 'particular' to be higher in the heirarchy.

>>If our job is to design an *ontology* then this would seem the natural way
>>to do it:

>

>I would say that our job is much narrower: we are not responsible for

>designing the best possible ontology,

Read my words again. I said *AN* ontology, not the best possible.

....

>

>This gets into a lot of delicate issues about which we may very well
>disagree. One distinction that I am not especially happy with is the
>one between common-sense and scientific reasoning. I believe that there
>is a continuity, and that one century's scientific breakthroughs become
>the next century's common sense.

Yes, we do disagree (profoundly) on this one. Lots of physical common sense (what I've called naive physics) seems to be built into us: children as young as a few months old seem to have it. Moreover, even MIT physics graduates often have pre-Newtonian naive physics, so last centuries real physics isn't today's naive physics. It's true that there is a kind of gentle seepage of ideas from science into everyday life, but (unlike the flow of water from rainfall into underground aquifers) this process doesn't usually improve the quality.

>

>I don't believe that there is anything inherently wrong with vagueness.
>Recall that note I circulated a while ago with quotations from Peirce
>about vagueness. In particular, his point that a vague truth is often
>much more useful than a precise falsehood.

>

.....

> I'm sure that you will be able to find many examples of a similar
>level of vagueness in the email that has been circulated to this list.
>And I don't believe that it is necessarily bad -- on the contrary, it is
>inevitable in the early stage of any kind of scientific inquiry.

We need to be more precise about vagueness. Of course many of the concepts we are trying to include in our hierarchies will be vague in the sense that they are not fully defined or axiomatized. But we will rapidly discover, as everyone who tries to write axioms has already discovered, that even slight steps along the direction of trying to make general deductive assertions involving these ideas forces us to make distinctions that we did not initially consider. Your A/P and my T/NT seem at first to be similar, but rapidly diverge on examples like holes and conversations. Even such mundane claims as a carpet being in an office become controversial when one is forced to take a deductive stance on their truth (Some people think of a fitted carpet as *part of* the office. If this seems silly to you, ask whether the paint on the wall is *in* the office or *part of* it.)

I know that humans get along perfectly fine without getting involved in such debates; but then humans aren't being asked to come up with axiomatic criteria for membership in sets, criteria that are supposed to be applicable for entire universes of concepts. But that is what we plan to be doing.

....

>.... I don't believe there is any difference between the
>kind of thinking of Moses and Einstein. ...

I guess I have nothing to add to that one, which I will leave alone for others to ponder.

The rest of our discussion is about CS Peirce, so I will take it into a different message.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Jun 03 09:45:56 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
 by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id GAA19184
 for <phayes@coginst.uwf.edu>; Wed, 3 Jun 1998 06:53:17 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
 by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id HAA20527;
 Wed, 3 Jun 1998 07:45:13 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
 id HAA20371; Wed, 3 Jun 1998 07:43:31 -0400
Date: Wed, 3 Jun 1998 07:43:31 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199806031143.HAA20371@west>
To: fritz@cyc.com, phayes@coginst.uwf.edu
Subject: Re: upperlevels and orthogonality
Cc: Adam_Farquhar@ksl.stanford.edu, Andreas.Reuter@eml.villa-bosch.de,
 Klaus.Tschira@ktf.villa-bosch.de, Piek.Vossen@let.uva.nl,
 cmenzel@tamu.edu, doug@csi.uottawa.ca, feigenbaum@ksl.stanford.edu,
 fellbaum@clarity.princeton.edu, geo@clarity.princeton.edu,

guarino@ladseb.pd.cnr.it, hovy@isi.edu, jamesp@cs.brandeis.edu,
jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk, peters@csl.stanford.edu,
phayes@nuts.coginst.uwf.edu, polanyi@pal.xerox.com, skydog@pacbell.net,
sowa@west.poly.edu, tsujii@is.s.u-tokyo.ac.jp, vdberg@us.ibm.com,
wahlster@dfki.uni-sb.de

Mime-Version: 1.0

Content-Type: text/plain; charset=us-ascii

Content-Length: 4671

Status:

Pat,

Re holes and surfaces:

>,,,. They are neither
>A nor P: they don't have any place to be categorised; and yet they seem
>much too 'particular' to be higher in the heirarchy.

When I mentioned the English statements that could be axiomatized, I was just listing them quickly as illustrations without going into the details. In Chapter 2 of my book, I was a bit more careful:

If x is physical, then x has a location in space-time. If x is anything other than empty space, x must also have a positive mass or energy.

If x is abstract, then x has no mass, energy, or location.

>... Moreover, even MIT physics
>graduates often have pre-Newtonian anive physics, so last centuries real
>physics isnt todays naive physics. Its true that there is a kind of gentle
>seepage of ideas from science into everyday life, but (unlike the flow of
>water from rainfall into underground aquifiers) this process doesnt usually
>improve the quality.

Yes, naive physics is rather close to Aristotelian physics, which is the part that has not stood the test of time as a basis for building rockets, but it is still close to the way we talk about events.

>We need to be more precise about vagueness. Of course many of the concepts
>we are trying to include in our heirarchies will be vague in the sense that
>they are not fully defined or axiomatised. But we will rapidly discover, as
>everyone who tries to write axioms has already discovered, that even slight
>steps along the direction of trying to make general deducitve assertions
>involving these ideas forces us to make distinctions that we did not
>initially consider. Your A/P and my T/NT seem at first to be similar, but
>rapidly diverge on examples like holes and conversations. Even such mundane

>claims as a carpet being in an office become controversial when one is
>forced to take a deductive stance on their truth (Some people think of a
>fitted carpet as *part of* the office. If this seems silly to you, ask
>whether the paint on the wall is *in* the office or *part of* it.)

I completely agree with this passage. And I certainly did not mean to imply that A/P and T/NT were interchangeable. On the contrary, the issues are so important that I passed over them because any more attention would have led to another several volumes of email.

I also agree with you about the profound differences in the views of physics presented by Aristotle, Newton, Einstein, and Bohr. Each of them had a view that is very different from naive physics and from each other's. But there is a continuity in the methods of thinking. What Einstein thought about physics when he was doing his Gedanken experiments was very different from what he did when he was walking down the street, dodging cars. Yet his method of doing his Gedanken experiments was very much in the same tradition as Zeno and his paradoxes about motion (which Plato & Aristotle among others analyzed in depth with quite a bit of ingenuity).

I would say that what makes 20th-century science seem so strange in comparison to "common sense" is the strangeness of the subject matter, not the methods of thinking or reasoning. I would recommend 19th-century science as a better testbed for studying scientific methodology and reasoning.

In particular, I would recommend the development of organic chemistry and the germ theory of disease as major breakthroughs at the time, which we can reflect on from a perspective that doesn't require (for us, at least) any strange habits of thinking. Take a look, for example, at Lister's and Pasteur's reasoning processes, where Lister was using analogies between fermentation processes (which were known to be caused by living organisms) with the processes that caused putrefaction in eggs and meat (whose causes were unknown and as unobservable to them as atoms are to us).

Another text to look at is Richard Feynman's Lectures on Physics, which are intended to introduce freshmen to very sophisticated topics. With his great expository skills, Feynman is able to make strange ideas seem "commonsensical." Feynman is an example of an extraordinarily clear 20th-century expositor, who is able to make his readers think like 21st-century students.

But even the 21st-century professors will talk pretty much like Aristotelians when they carry on a conversation about everyday topics. Somehow, that language that has evolved over millennia of human experience with "naive physics" is so flexible that it can still be used to discuss the most profoundly radical views of modern science. The fact that people can deal

with it is an existence proof that perhaps there is a solution to our ontology problems.

John

From ???@??? Thu Jun 25 17:44:03 1998
Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id PAA19158
for <phayes@coginst.uwf.edu>; Thu, 25 Jun 1998 15:18:26 -0500 (CDT)
Received: from scratchy (scratchy [207.207.8.118])
by catbert.cyc.com (8.8.8/8.8.8) with SMTP id PAA18511
for <phayes@coginst.uwf.edu>; Thu, 25 Jun 1998 15:13:45 -0500 (CDT)
Message-Id: <3.0.32.19980625151455.00a4c1d8@catbert.cyc.com>
X-Sender: fritz@catbert.cyc.com
X-Mailer: Windows Eudora Pro Version 3.0 (32)
Date: Thu, 25 Jun 1998 15:14:55 -0500
To: Pat Hayes <phayes@coginst.uwf.edu>
From: Fritz Lehmann <fritz@cyc.com>
Subject: Re: whoops
Mime-Version: 1.0
Content-Type: text/plain; charset="us-ascii"
Content-Length: 941
Status:

At 02:08 PM 6/25/98 -0600, you wrote:
>Fritz, you forgot to include Nicola (guarino@ladseb.pd.cnr.it) in the CC
>list. As I did myself, until now.
>Nostra Culpa.
>Pat

Dear Pat,

I realized my culpa immediately and forwarded Guarino a copy of my note to Sowa. Incidentally, in that note I (again) took the liberty of forwarding someone else's words addressed to me without consulting him first. I think this is generally OK if the topic is non-personal and of general interest, but some people insist that it is a terrible sin. In this case John had alluded to y'all having accepted or agreed with his proposal (whatever it was); I would like to be cautious in accepting such claims of support.

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

=====
=====
From ???@??? Thu Jun 25 17:44:02 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id OAA18983;
Thu, 25 Jun 1998 14:57:43 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a05b1b83847b9c6@[143.88.7.118]>
In-Reply-To: <3.0.32.19980625015655.008e9418@catbert.cyc.com>
Mime-Version: 1.0
Date: Thu, 25 Jun 1998 14:59:55 -0600
To: Fritz Lehmann <fritz@cyc.com>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Abstract occurrents
Cc: fritz@cyc.com, sowa@west.poly.edu, p.m.simons@leeds.ac.uk,
phayes@nuts.coginst.uwf.edu, cmenzel@tamu.edu,
guarino@ladseb.pd.cnr.it, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 10763
Status:

[FL]>I'm copying this to the cited "top-down" people from Heidelberg.

Me too.

>At 02:48 PM 6/23/98 -0400, you wrote:

[JS]>>Fritz,

>>I'm sorry I missed you on the last day of the Heidelberg meeting.
>>That morning, I discussed the question of abstract continuants with
>>Pat, Peter, and a couple of others. My solution is that we define
>>the distinction of continuant vs. occurrent in a purely abstract way
>>that depends only on structure (time-like or not). Then all the
>>definitions and axioms can be inherited by both physical and abstract
>>entities. Peter, Pat, et al. agreed that was a reasonable approach.

As I recall, John's idea for abstract continuant/occurrents was that a sign is classified according to the type of the thing it signifies, so that an abstraction is called 'continuant' if it *means* something which is a continuant, etc.. This seems to me to be a reasonable idea, and of course it can be used to apply to any case where one might want to talk about an abstract thing of some type where that type is normally not used to classify abstracta. I have only two concerns with this idea. First, it is potentially confusing (a minor point :-). Second, more seriously, we might foresee a need to allow a signifier which itself - ie the sign itself - has physical properties, and then this convention might become awkward. John and I talked about this a little, and I tried to imagine something which

was 'abstract' in the nonphysical sense yet could be said to have a temporal structure: perhaps a dance, say, (as opposed to a particular physical performance of the dance), or _the behavior of an algorithm_. I can imagine that there might be a need for such things in an ontology of emerging interactive-media or virtual realities, for example. If so, then we might want to allow a distinction between the occurrent/continuant structure of the sign itself, and that of the thing it signifies, and these might not always be the same; this would pose real problems for John's proposal, I think.

>I don't understand the above. The problem I remember is: abstract things
>never change whereas the continuant/occurrent distinction depends mainly on
>change. Your Heidelberg solution was to interpret all "abstract
>continuants" as forever-unchanging abstract specifications of continuing
>concrete things, and to interpret "abstract occurrents" as
>forever-unchanging abstract specifications of "changing" concrete occurrent
>"things" like events or scripts. In this case, neither "abstract
>continuants" nor "abstract occurrents" can really inherit the axioms of
>continuants or occurrents. The opposite approach, treating "abstract
>continuants" and "abstract occurrents" as genuine continuants and
>occurrents, respectively, has a similar drawback: these then cannot inherit
>the "unchanging" and "isomorphism=identity" and other axioms from "abstract
>entity".

>

>The quoted paragraph above doesn't seem to me to offer any way out of this
>dilemma without considerable further elaboration of its key phrase "in a
>purely abstract way (time-like or not)". Structure itself is not enough to
>distinguish continuant from occurrent; the same abstract structure can
>describe a time series of events (within a compound event) and the series
>of encyclopaedia volumes on a bookshelf.

>

>I'm inclined to go the Peirce/semiosis route and openly acknowledge that
>so-called "abstract continuants" and "abstract occurrents" make no sense
>directly, and that they are abstract specifications that are not real
>continuants or occurrents. They are _signs_ of what they specify, and
>their abstract structural relations are often preserved in their concrete
>embodiments. This unfortunately wrecks the nice crystal but so be it (the
>continuant/occurrent distinction depends on a prior choice of concrete over
>abstract). I wish it were not so. Plato's pure forms are simply not to be
>found among Aristotle's pebbles and burps.

Well, let us keep a certain modesty in the face of the tidal wave of potential entities that some future ontologist may want to classify. Computers have already wreaked havoc with old distinctions between texts and objects, signs and tokens, etc., which seemed quite clear and sharp after millenia of human thought and were obviously beyond the imagination

of Plato and Aristotle, and maybe even Peirce. I don't find the idea of things that are as like abstracta as one wishes, except that they have temporal extent, for example, at all implausible. For example, consider an algorithm (say, quicksort.). There are textual embodiments of programs which implement an algorithm, but the algorithm itself must be abstract. You can't even write it down. Now, consider the notion of a *run* or a *behavior* of the algorithm, eg _the way that quicksort behaves_. This isn't any particular process of running a particular embodiment of the algorithm: it's much more abstract than that; yet, it seems to be essentially temporal in nature: it has beginnings and endings, for example, and can be compared in its temporal structure with runnings of other algorithms (like bubble-sort). Why not allow it to be both abstract and temporally structured? And if we do, then (rather as temporal databases must distinguish valid time from transaction time, and planners must keep track of the time of planning and the time planned about) we might want to allow a clear distinction between the temporal structure of the sign and the temporal structure of the signified. It's not hard to think of examples in every one of those four boxes.

>

>I think that a concrete thing can represent an abstract thing and that a >("grasped") abstract thing can represent a concrete thing (to an >interpretant, in both cases, of course).

Already I disagree (and so does our agreed consensual theory of mentality, in which the thing grasped is not an abstraction but a mental representation!)

>>For the definition, Sir Arthur Eddington coined the phrase "the arrow >>of time", which is defined by increasing entropy. This definition can >>be applied to either physical entities or abstract descriptions of >>entities to distinguish a notion of time or time-like-ness.

>

>I do not believe that entropy is the key to time-like-ness or change or >continuant-vs-occurrent (it may help define the _direction_ of time but the >direction is not the problem here). I do not believe that many people know >exactly what they mean by "entropy" or "order" anyway, other than >practitioners of macroscopic statistical mechanics. Also, your "abstract >descriptions" are just that: descriptions. They are signs.

I'm afraid I agree with John here. Fritz is right that we needn't and shouldn't get into talk of entropy in order to *define* time-order, but Eddington certainly had a good point. In fact, entropy, being defined in terms of information, could be said to be an abstract concept.

Look, if someone wants to *insist* that abstracta are timeless then nobody

can prove them wrong. Any case offered can always be divided into an eternal Platonic abstract aspect and a time-embedded aspect (which is, on this view, therefore not abstract by definition) plus some kind of relation between them whereby the second somehow captures or signifies or makes concrete the former. So OK Fritz, your strict Platonism is impregnable. All the same, I think we ought to consider that if someone wants to try an alternative way of thinking, they aren't somehow blocked from doing so by our conservatism.

>As I said in Heidelberg, I think continuant and occurrent relate to
>4-dimensional spacetime worms as follows: the continuant is characterized
>mainly by what stays the same over time, and a continuant can undergo
>change; an occurrent is characterized mainly by a difference over time; it
>is a change and cannot undergo change (although acceleration superficially
>appears like a possible counterexample, as Spillers pointed out). I
>realize that "characterized mainly" is mushy talk, but I still think the
>foregoing is correct.

Part of the above is simply incoherent, by confusing A and B series. If we are talking about spacetime 'worms', then *nothing* 'undergoes change'. In this way of talking, the best we can do is to talk about different ways to carve up the worms into parts, whether these are 'spacelike' or 'timelike'. See my message from yesterday on this, and there's another (better!) one coming tomorrow which integrates the temporal/spatial part distinction with identity criteria being translated from B to A language, in Nicola's spirit (though not quite coming to the same conclusions.)

>Cyc used to call the top continuant/occurrent categories:
>SomethingOccurring and SomethingExisting. The former is now called Event.
>The Heidelberg report left several big things undisposed-of --- especially
>individual time intervals. Cyc has TemporalThing _above_
>SomethingOccurring, Event, and TimeInterval. Cyc seems right to me, on
>this point. You "factorers" need to address this issue immediately.

I think CYC has this wrong. A timeinterval is not itself temporally located in the same way that events like wars and weddings are. If it were, we could ask, what interval does the interval occupy?, which is meaningless. Similarly for asking how fast does "now" change, and when does 3pm happen; all these are errors of mistaking the coordinate for the thing in the coordinate space. (One way out, which CYC may use, is to say that timeintervals are unique among temporally situated things in that they are their own temporal location. this seems to work, but it's ugly and counterintuitive.) In general, a dimension is not itself located in the space it is a dimension of.

One advantage of the CYC classification is that it provides a quick way of

talking about how one coordinate frame may be 'located' in another in terms of mappings between them (eg conversion of radial to cartesian coordinates). However I think this gets two different ideas muddled: (a) a mathematical object which is a mapping between coordinate frameworks, and (b) the use of a coordinate frame which is located relative to some object or place in the world, eg _the corner of the room_ or _the third door on the left_. It is tempting to subsume one of these to the other, but I'd need convincing that it is worth the potential trouble it causes. (The persuasion has to consist of more than just being told that CYC does it that way and that Fritz has come to like it. :-)

Pat Hayes

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Thu Jun 25 17:44:01 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id OAA18986;
Thu, 25 Jun 1998 14:57:57 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a06b1b857d72546@[143.88.7.118]>
Mime-Version: 1.0
Date: Thu, 25 Jun 1998 15:00:11 -0600
To: Fritz Lehmann <fritz@cyc.com>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: ; Whether sets have mass
Cc: fritz@cyc.com, sowa@west.poly.edu, p.m.simons@leeds.ac.uk,
phayes@nuts.coginst.uwf.edu, cmenzel@tamu.edu,
guarino@ladseb.pd.cnr.it, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 4058
Status:

[JS]>>WHETHER SETS HAVE MASS

>>

>>The discussion of various kinds of collections provides

>>a way of answering Lenat's question of whether sets have mass:

>>

>> - A mereological union or sum of multiple entities is the same kind

>> of entity as each of its parts. If the parts are physical, then
>> their union would be a physical entity whose mass would be the total
>> of all its parts. If the parts are all abstract, then their union
>> would also be abstract. But in all the usual versions of mereology,
>> the union of a physical entity like the cat Yojo with an abstract
>> entity like the number 7 is meaningless or undefined.
>>
>> - In set theory, however, a set is an entity that is distinct from the
>> union of its elements,
[FL]>
>... but not distinct from the union of their singleton sets ...

There's a pun on "union" here, of course. Maybe it would be better to use a different word for the mereological case. I'll try to use 'sum' consistently.

>
>> and the axioms impose no constraints on the
>> nature of the elements. It is possible to form the set {Yojo,7}, whose
>> elements are a cat and a number; since the types of the elements are
>> incompatible, the set could not be of the same type as its elements.
>> The most reasonable interpretation is that sets are abstractions,
>> independent of the nature of their elements, which may be physical,
>> abstract, or mixed. Therefore, Lenat would be correct in claiming
>> that sets have no mass.

>>
>>This distinction helps to solve puzzles that trouble philosophers
>>and logicians, but people who had not studied set theory would never
>>think of it. If the phrase a herd of sheep occurs in English
>>speech, it could be interpreted as a mereological union, which would
>>be a physical entity having mass. But if the phrase a set of parts
>>occurs in ordinary English, it should be considered ambiguous until
>>the nature of the speaker is determined.

>
>I agree. Cyc tries to be careful about sets versus Cyc "Groups" (compound
>relational structures of objects), but English uses "set" and "group"
>ambiguously for both. A "Group" in Cyc has parts but not set-elements; it
>is an individual not a set.

Its also not a mereological sum, i think. One mereopart of a flock of sheep consists entirely of muscle tissue, another entirely of wool, but neither of these is part of a flock (which has to be a group of *sheep*.) We need (at least) three notions of collection here: a set, a group (in the CYC sense?), and a mereosum. Both the latter have mass, and it is the same mass (at a particular instant, anyway) but they arent the same kind of thing. In particular, the same flock might lose or gain a sheep or two from time to

time. So for example a pile of bricks has mass which is the mass of all the brickstuff in the pile at any moment, but this might change even though a brick's mass is (let us suppose) constant. (Theres a remarkable number of special nouns in english for denoting these groups: flock, shoal, exhaltation, herd, gaggle, etc., eaqch of which identifies the kind of atom involved: sheep, fish, larks, cattle, geese. Anyone got any idea why?) (BTW, why are CYC groups called 'structured' ? I'd say its almost a hallmark of flocks, herds etc. that they *lack* structure. What distinguishes them from mereosums is their special atomicity, not any special relations among their members.)

> Your "puzzles that trouble philosophers and
>logicians" also trouble any practical automated inference system: the
>puzzles of getting the right inferences from a Knowledge Base. They are
>thus not mere "puzzles," in my parlance.

Hey, lets distinguish "puzzles" from "mere puzzles". Some puzzles are highly non-mere, for example the liar paradox.

Pat Hayes

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Fri Jul 03 12:19:55 1998
Received: from LADSEB.LADSEB.PD.CNR.IT (ladseb.ladseb.pd.cnr.it [150.178.2.3])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id IAA02301
for <phayes@coginst.uwf.edu>; Thu, 2 Jul 1998 08:46:29 -0500 (CDT)
Received: from [150.178.2.93] by LADSEB.LADSEB.PD.CNR.IT with ESMTP;
Thu, 2 Jul 1998 15:43:04 +0200
X-Sender: guarino@ladseb.ladseb.pd.cnr.it
Message-Id: <v03102804b1c1390b1965@[150.178.2.93]>
In-Reply-To: <v04003a00b1b9936375f7@[143.88.7.118]>
References: <199806260726.DAA00926@west>
Mime-Version: 1.0
Date: Thu, 2 Jul 1998 15:37:14 +0200
To: Pat Hayes <phayes@coginst.uwf.edu>, sowa@west.poly.edu (John F. Sowa)
From: Nicola Guarino <guarino@ladseb.pd.cnr.it>
Subject: Re: Whether sets have mass/Nicola might disagree
Cc: cmenzel@tamu.edu, doug@csi.uottawa.ca, fritz@cyc.com,

p.m.simons@leeds.ac.uk
Content-Transfer-Encoding: 8bit
X-MIME-Autoconverted: from quoted-printable to 8bit by nuts.coginst.uwf.edu id
IAA02301
Content-Type: text/plain
Content-Length: 1201
Status:

Hi guys,

back to the discussion after some relax... Organizational issues first.

At 1:05 PM -0600 6/26/98, Pat Hayes wrote:

>...

>>

>[JS]>I think it's fine to address it to those people who seem to be actively
>>contributing to the immediate discussion, but it should also be directed
>>to some bucket that will periodically be dumped on a web site such as
>>onto-std.

>

>Good idea, though I think it would be proper for now to keep the active
>distribution to the local (sub)group from the VillaBosch meeting, as we
>still have a lingering responsibility to come to a consensus on a document
>of some kind. (Or was this abandoned after we left??)

I am asking myself the same question. What's the status of this work? What
about the press announcement supposed to be released after the workshop?

-- Nicola

Nicola Guarino
National Research Council phone: +39 49 8295751
LADSEB-CNR fax: +39 49 8295763
Corso Stati Uniti, 4 email: guarino@ladseb.pd.cnr.it
I-35127 Padova
Italy

Home page: ** updated 4/5/98 **
<http://www.ladseb.pd.cnr.it/infor/ontology/ontology.html>

FOIS'98 home page:
<http://mnemosyne.itc.it:1024/fois98/>

From ???@??? Mon Jul 06 10:41:32 1998
Received: from philebus.tamu.edu (root@PHILEBUS.TAMU.EDU [165.91.161.22])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id OAA12440
for <phayes@coginst.uwf.edu>; Sun, 5 Jul 1998 14:25:53 -0500 (CDT)
Received: from philebus.tamu.edu (localhost [127.0.0.1])
by philebus.tamu.edu (8.8.5/8.8.5) with ESMTP id OAA05373;
Sun, 5 Jul 1998 14:18:34 -0500
Message-Id: <199807051918.OAA05373@philebus.tamu.edu>
X-Mailer: exmh version 2.0gamma 1/27/96
To: Pat Hayes <phayes@coginst.uwf.edu>
cc: sowa@west.poly.edu (John F. Sowa), doug@csi.uottawa.ca, fritz@cyc.com,
p.m.simons@leeds.ac.uk, phayes@nuts.coginst.uwf.edu,
guarino@ladseb.pd.cnr.it, cmenzel@philebus.tamu.edu
Subject: The meta-meta-question
In-reply-to: Your message of "Fri, 03 Jul 1998 12:20:53 MDT."
<v04003a01b1bbfd041fc7@[143.88.7.107]>
Mime-Version: 1.0
Date: Sun, 05 Jul 1998 14:18:34 -0500
From: Chris Menzel <cmenzel@philebus.tamu.edu>
Content-Type: text/plain; charset=us-ascii
Content-Length: 620
Status:

Pat wrote:

> OK, I have a meta-meta-question: who is in charge here?

and John wrote:

> No one is in charge. That is one part of the problem.

It's not quite finalized yet, but it looks like I will be going out to Stanford for the rest of the summer to help Bob Spillers out with a couple of things, notably *some* of the co-ordination of whatever publication comes out of the Heidelberg workshop. If so, beginning around 13 July I will be available to prod, edit, write, proof, critique, collect, collate, or whatever else to help speed things along. Ideas about how best to use me would be appreciated.

-chris

From ???@??? Mon Jul 06 10:41:39 1998
Received: from LADSEB.LADSEB.PD.CNR.IT (ladseb.ladseb.pd.cnr.it [150.178.2.3])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id GAA14713
for <phayes@coginst.uwf.edu>; Mon, 6 Jul 1998 06:11:39 -0500 (CDT)
Received: from [150.178.2.93] by LADSEB.LADSEB.PD.CNR.IT with ESMTP;
Mon, 6 Jul 1998 13:08:07 +0200
<v04003a01b1bbfd041fc7@[143.88.7.107]>
X-Sender: guarino@ladseb.ladseb.pd.cnr.it
Message-Id: <v03102802b1c659038445@[150.178.2.93]>
In-Reply-To: <199807051918.OAA05373@philebus.tamu.edu>
References: "Your message of Fri, 03 Jul 1998 12:20:53 MDT."
<v04003a01b1bbfd041fc7@[143.88.7.107]>
Mime-Version: 1.0
Date: Mon, 6 Jul 1998 13:08:35 +0200
To: Chris Menzel <cmenzel@philebus.tamu.edu>,
Pat Hayes <phayes@coginst.uwf.edu>
From: Nicola Guarino <guarino@ladseb.pd.cnr.it>
Subject: Re: The meta-meta-question
Cc: sowa@west.poly.edu (John F. Sowa), doug@csi.uottawa.ca, fritz@cyc.com,
p.m.simons@leeds.ac.uk
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2263
Status:

At 2:18 PM -0500 7/5/98, Chris Menzel wrote:

>Pat wrote:

>

>> OK, I have a meta-meta-question: who is in charge here?

>

>and John wrote:

>

>> No one is in charge. That is one part of the problem.

>

>It's not quite finalized yet, but it looks like I will be going out to
>Stanford for the rest of the summer to help Bob Spillers out with a couple
>of things, notably *some* of the co-ordination of whatever publication
>comes out of the Heidelberg workshop. If so, beginning around 13 July I
>will be available to prod, edit, write, proof, critique, collect, collate,
>or whatever else to help speed things along. Ideas about how best to use
>me would be appreciated.

>

>-chris

This is a VERY good new! On the last day of the Heidelberg week I got a very negative feeling regarding the outcome of the workshop, due to the way Bob managed to keep a strong and exclusive personal control of the whole initiative. Let's hope that at least on the theoretical framework we shall be able to express the result of a common ongoing work [a lot of work has

still to be done to present in a decent way what we have discussed about].

My major concern however regards the "political" initiatives bound to this project (especially, the foreseen meeting of the major funding agencies in September). Here is the area where Bob wants absolute personal independence, but, as I tried to argue on the last day, this does not appear to me as the most appropriate way to manage the kind of scientific initiative we have in mind: some of you had unfortunately already left, but it may be useful to know that when I proposed to appoint a small international committee to coordinate the lobbying initiatives Bob was definitely against, and nobody else dared to object...

Some comments regarding the technical discussion so far with the next message [early, I hope].

All the best,

-- Nicola

Nicola Guarino
National Research Council phone: +39 49 8295751
LADSEB-CNR fax: +39 49 8295763
Corso Stati Uniti, 4 email: guarino@ladseb.pd.cnr.it
I-35127 Padova
Italy

Home page: ** updated 4/5/98 **
<http://www.ladseb.pd.cnr.it/infor/ontology/ontology.html>

FOIS'98 home page:
<http://mnemosyne.itc.it:1024/fois98/>

From ???@??? Tue Jul 07 10:39:13 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id WAA21291
for <phayes@coginst.uwf.edu>; Mon, 6 Jul 1998 22:08:59 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id XAA22819;
Mon, 6 Jul 1998 23:05:41 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
id XAA29656; Mon, 6 Jul 1998 23:03:01 -0400

Date: Mon, 6 Jul 1998 23:03:01 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199807070303.XAA29656@west>
To: cmenzel@philebus.tamu.edu, guarino@ladseb.pd.cnr.it,
phayes@coginst.uwf.edu
Subject: Re: The meta-meta-question
Cc: doug@csi.uottawa.ca, fritz@cyc.com, p.m.simons@leeds.ac.uk
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 582
Status:

Nicola's concerns are very close to mine. I think that

Chris is an excellent choice for the editor, and I would be the first to second his nomination. But what bothers me is that none of us was asked to vote or was even informed that someone was even being considered.

As I said at the Heidelberg meeting, I would like to see strict parliamentary procedure being followed. Usually, I don't care for too much formality, but my experiences in the ANSI and ISO gatherings have convinced me of the importance of having procedures that ensure fairness to all points of view.

John

From ???@??? Tue Jul 07 10:39:14 1998
Received: from philebus.tamu.edu (root@PHILEBUS.TAMU.EDU [165.91.161.22])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id BAA28325
for <phayes@coginst.uwf.edu>; Tue, 7 Jul 1998 01:00:14 -0500 (CDT)
Received: from philebus.tamu.edu (localhost [127.0.0.1])
by philebus.tamu.edu (8.8.5/8.8.5) with ESMTP id AAA13046;
Tue, 7 Jul 1998 00:52:48 -0500
Message-Id: <199807070552.AAA13046@philebus.tamu.edu>
X-Mailer: exmh version 2.0gamma 1/27/96
To: sowa@west.poly.edu (John F. Sowa)
cc: guarino@ladseb.pd.cnr.it, phayes@coginst.uwf.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, cmenzel@philebus.tamu.edu
Subject: Re: The meta-meta-question
In-reply-to: Your message of "Mon, 06 Jul 1998 23:03:01 EDT."
<199807070303.XAA29656@west>
Mime-Version: 1.0
Date: Tue, 07 Jul 1998 00:52:48 -0500
From: Chris Menzel <cmenzel@philebus.tamu.edu>
Content-Type: text/plain; charset=us-ascii

Content-Length: 1175

Status:

> Nicola's concerns are very close to mine. I think that
> Chris is an excellent choice for the editor, and I would be
> the first to second his nomination.

Gracious words! I hope I can be of help -- but see next paragraph.

> But what bothers me is that
> none of us was asked to vote or was even informed that someone
> was even being considered.

Careful not to read more into the situation than is warranted, John. Bob asked me out to Stanford to help with a number of things, including things other than the intended publication, and I have promised him only about 7 weeks of my time (though I will likely to take on duties that will keep me involved in this and other matters come fall). In particular, the idea of serving as editor for whatever document issues from the "Heidelberg Summit" never came up. If you'll look at my first message, I was careful to use the word "co-ordinate" to describe my duties vis-a-vis the workshop. (Ok, I did use the word "edit", but I meant that only in the strict sense of proof reading, critiquing, etc.) The issue of the general editorship is wide open, and frankly I am not at all inclined to take on that task.

-chris

From ???@??? Tue Jul 07 10:39:20 1998

Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id CAA02511
for <phayes@coginst.uwf.edu>; Tue, 7 Jul 1998 02:42:10 -0500 (CDT)

Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id DAA23250;
Tue, 7 Jul 1998 03:38:56 -0400 (EDT)

Received: by west (SMI-8.6/SMI-SVR4)
id DAA00711; Tue, 7 Jul 1998 03:36:26 -0400

Date: Tue, 7 Jul 1998 03:36:26 -0400

From: sowa@west.poly.edu (John F. Sowa)

Message-Id: <199807070736.DAA00711@west>

To: cmenzel@philebus.tamu.edu

Subject: Re: The meta-meta-question

Cc: doug@csi.uottawa.ca, fritz@cyc.com, guarino@ladseb.pd.cnr.it,
p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu, sowa@west.poly.edu

Mime-Version: 1.0

Content-Type: text/plain; charset=us-ascii
Content-Length: 1936
Status:

Chris,

>Careful not to read more into the situation than is warranted, John. Bob
>asked me out to Stanford to help with a number of things, including things
>other than the intended publication....

That's fine, but the duties that are paying your salary should be distinguished from what is being done for the ontology project, which has been a volunteer effort governed by consensus. The concerns that Nicola raised are important, and we should not allow any precedent for decisions to be made outside the scope of our so-far nonexistent charter.

As a result of some recent discussions, Pat raised an important question: "Who's in charge here?" When something has to be accomplished, someone has to take responsibility for either doing it or getting other people to do it. We need some mechanism for assigning tasks to people and seeing that they are carried out.

> I was careful to use the word "co-ordinate" to describe my duties
> vis-a-vis the workshop....

The word "co-ordinate" bothers me much more than the word "edit". It is so vague that it can mask all the issues that Nicola and I were worried about. I would like to see clearly defined terms with associated responsibilities:

1. A chairman who pounds the gavel and calls people to order.
2. A secretary who keeps the minutes.
3. An editor who co-ordinates the tasks that lead to a publication.

But the authority to appoint people to various positions, including that of "co-ordinator", should reside in the voting power of the committee as a whole.

I agree with Robert Frost, who said "Good fences make good neighbors."
And I agree with the other Robert who said "Good rules make good committees."

I would just like to see all decisions made either by the committee voting as a whole or by someone who has been appointed by the committee

to carry out some clearly designated tasks, which must be reported back to the committee when accomplished.

John

From ???@??? Tue Jul 07 10:39:28 1998
Received: from philebus.tamu.edu (root@PHILEBUS.TAMU.EDU [165.91.161.22])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA22271
for <phayes@coginst.uwf.edu>; Tue, 7 Jul 1998 11:48:01 -0500 (CDT)
Received: from philebus.tamu.edu (localhost [127.0.0.1])
by philebus.tamu.edu (8.8.5/8.8.5) with ESMTP id LAA15413;
Tue, 7 Jul 1998 11:37:10 -0500
Message-Id: <199807071637.LAA15413@philebus.tamu.edu>
X-Mailer: exmh version 2.0gamma 1/27/96
To: sowa@west.poly.edu (John F. Sowa)
cc: doug@csi.uottawa.ca, fritz@cyc.com, guarino@ladseb.pd.cnr.it,
p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
cmenzel@philebus.tamu.edu
Subject: Re: The meta-meta-question
In-reply-to: Your message of "Tue, 07 Jul 1998 03:36:26 EDT."
<199807070736.DAA00711@west>
Mime-Version: 1.0
Date: Tue, 07 Jul 1998 11:37:10 -0500
From: Chris Menzel <cmenzel@philebus.tamu.edu>
Content-Type: text/plain; charset=us-ascii
Content-Length: 2465
Status:

>>Careful not to read more into the situation than is warranted, John. Bob
>>asked me out to Stanford to help with a number of things, including things
>>other than the intended publication....

>

> That's fine, but the duties that are paying your salary should be
> distinguished from what is being done for the ontology project,
> which has been a volunteer effort governed by consensus.

I'm not sure of your point here, but as a matter of fact any work I do on the ontology project will still be on a volunteer basis. I will not be getting paid for any of it.

> The concerns

> that Nicola raised are important, and we should not allow any precedent
> for decisions to be made outside the scope of our so-far nonexistent
> charter.

I hope you have no reason to think that I disagree. If you do, you are

laboring under a misconception.

> As a result of some recent discussions, Pat raised an important
> question: "Who's in charge here?" When something has to be accomplished,
> someone has to take responsibility for either doing it or getting other
> people to do it. We need some mechanism for assigning tasks to people
> and seeing that they are carried out.

Again, I completely agree, and, indeed, one of my functions might be to help put such mechanisms in place. Again if you look at my first message, I was simply putting myself at the disposal of the group.

>> I was careful to use the word "co-ordinate" to describe my duties
>> vis-a-vis the workshop....
>
> The word "co-ordinate" bothers me much more than the word "edit".
> It is so vague that it can mask all the issues that Nicola and I
> were worried about.
> ...
> But the authority to appoint people to various positions, including
> that of "co-ordinator", should reside in the voting power of the
> committee as a whole.

Sheesh. Maybe this was a bad idea. You don't like "co-ordinator". Fine. Give me a label then. Facilitator? Deferential and compliant servant? I am making myself available to the group, for crying out loud. I asked *you* and the others how best I might be of service. Does that *really* sound like a position to which someone needs to be *appointed*?

Look, I have no agenda, and I would not support anything that rode roughshod over the will of the others in the project. If you don't want my help I will be happy to spend my time working on any of several projects that would, frankly, be of greater benefit to me personally.

-chris

From ???@??? Tue Jul 07 10:39:30 1998

Received: from philebus.tamu.edu (root@PHILEBUS.TAMU.EDU [165.91.161.22])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA22352
for <phayes@coginst.uwf.edu>; Tue, 7 Jul 1998 11:55:27 -0500 (CDT)

Received: from philebus.tamu.edu (localhost [127.0.0.1])
by philebus.tamu.edu (8.8.5/8.8.5) with ESMTP id LAA15484
for <phayes@coginst.uwf.edu>; Tue, 7 Jul 1998 11:48:14 -0500

Message-Id: <199807071648.LAA15484@philebus.tamu.edu>

X-Mailer: exmh version 2.0gamma 1/27/96

To: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: The meta-meta-question
In-reply-to: Your message of "Mon, 06 Jul 1998 12:08:55 -0800."
<v04003a05b1c6d7cc3564@[143.88.7.108]>
Mime-Version: 1.0
Date: Tue, 07 Jul 1998 11:48:14 -0500
From: Chris Menzel <cmenzel@philebus.tamu.edu>
Content-Type: text/plain; charset=us-ascii
Content-Length: 1589
Status:

> OK, fearless Leader, you asked for it.

No leadership desired, if you please!! ;-)

> 1. How about some guidance as to what KIND of document we are supposed to
> be producing? (A book, a paper, a part of a longer book, a draft ISO
> Standard of Human Thought, a user's manual for some axioms, an Ontologist's
> Field Guide, ...?)

I see from a message that Bob forwarded to me that you have not stopped thinking about this question. I like your ideas a lot. Unfortunately, in order to get out to CA next week I have to devote every spare minute of this week to tying up loose ends and fulfilling other obligations, so I have not been able to put any time into thinking about these issues at length. But, like I said, I like your ideas a lot. I will be in touch with you about how we might refine and implement them as soon as I get out to Palo Alto.

> 2. Im not at all convinced that we all agree on everything yet, so should
> we be trying to continue our efforts by email? The trouble is that some of
> us seem to be more e-connected than others.

I think email is still productive at this point. We just have to recognize when we've reached the point of diminishing returns. One of the things I liked about your suggestions to Bob was that it seemed to provide a way to get people writing without there being complete consensus.

More later. I will probably be using you as my primary sounding board! (Supposing John doesn't still have his panties in a wad about the fact that no one *appointed* me to serve in whatever manner I end up serving... :-)

Best,

-chris

From ???@??? Wed Jul 08 10:36:03 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id SAA11767
for <phayes@coginst.uwf.edu>; Tue, 7 Jul 1998 18:59:36 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id TAA25535;
Tue, 7 Jul 1998 19:56:22 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
id TAA05131; Tue, 7 Jul 1998 19:53:52 -0400
Date: Tue, 7 Jul 1998 19:53:52 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199807072353.TAA05131@west>
To: cmenzel@philebus.tamu.edu
Subject: Re: The meta-meta-question
Cc: doug@csi.uottawa.ca, fritz@cyc.com, guarino@ladseb.pd.cnr.it,
p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 701
Status:

Chris,

Please don't quit. My motto is always take advantage of people who are willing to volunteer.

I hope that you will help us put together a revised version of the notes that Fritz started to write in Heidelberg together with further additions and modifications that the rest of us contribute. That would be something we would all be grateful for. There have been some other email discussions recently that should also go into the melting pot.

I was only echoing Nicola's concerns. If we keep everything open on email, we shouldn't have any problems. But I would also like to see the results put somewhere like the onto-std web site so that the whole world would be able to see them.

John

From ???@??? Wed Jul 08 10:36:04 1998
Received: from philebus.tamu.edu (root@PHILEBUS.TAMU.EDU [165.91.161.22])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id UAA15649
for <phayes@coginst.uwf.edu>; Tue, 7 Jul 1998 20:30:00 -0500 (CDT)
Received: from philebus.tamu.edu (localhost [127.0.0.1])

by philebus.tamu.edu (8.8.5/8.8.5) with ESMTP id UAA17446;
Tue, 7 Jul 1998 20:22:23 -0500
Message-Id: <199807080122.UAA17446@philebus.tamu.edu>
X-Mailer: exmh version 2.0gamma 1/27/96
To: sowa@west.poly.edu (John F. Sowa)
cc: doug@csi.uottawa.ca, fritz@cyc.com, guarino@ladseb.pd.cnr.it,
p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu
Subject: Re: The meta-meta-question
In-reply-to: Your message of "Tue, 07 Jul 1998 19:53:52 EDT."
<199807072353.TAA05131@west>
Mime-Version: 1.0
Date: Tue, 07 Jul 1998 20:22:23 -0500
From: Chris Menzel <cmenzel@philebus.tamu.edu>
Content-Type: text/plain; charset=us-ascii
Content-Length: 1144
Status:

> Please don't quit.

Oh I won't. You caught me at a bad moment and I over-reacted. Just please bear in mind that (i) my only goal in helping out this summer is to help the project to succeed; (ii) I firmly believe that the success of the project will require moving according to the general will of the group; and (iii) I have no interest in occupying any kind of authoritative, editorial position. That is not My Way.

> I hope that you will help us put together a revised version of the
> notes that Fritz started to write in Heidelberg together with further
> additions and modifications that the rest of us contribute. That would
> be something we would all be grateful for. There have been some other
> email discussions recently that should also go into the melting pot.

I have all the stuff, I think, but will check that.

> I was only echoing Nicola's concerns. If we keep everything open on
> email, we shouldn't have any problems. But I would also like to see
> the results put somewhere like the onto-std web site so that the
> whole world would be able to see them.

I think we all agree on this one.

Regards,

-chris

From ???@??? Wed Jul 08 14:58:05 1998
Received: from philebus.tamu.edu (root@PHILEBUS.TAMU.EDU [165.91.161.22])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id NAA26146
for <phayes@coginst.uwf.edu>; Wed, 8 Jul 1998 13:27:56 -0500 (CDT)
Received: from philebus.tamu.edu (localhost [127.0.0.1])
by philebus.tamu.edu (8.8.5/8.8.5) with ESMTP id NAA20918
for <phayes@coginst.uwf.edu>; Wed, 8 Jul 1998 13:20:38 -0500
Message-Id: <199807081820.NAA20918@philebus.tamu.edu>
To: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: The meta-meta-question
In-reply-to: Your message of "Wed, 08 Jul 1998 10:33:53 -0800."
<v04003a0cb1c8322fd388@[143.88.7.102]>
MIME-Version: 1.0
Content-ID: <20915.899922038.1@philebus.tamu.edu>
Date: Wed, 08 Jul 1998 13:20:38 -0500
From: Chris Menzel <cmenzel@philebus.tamu.edu>
Content-Type: text/plain; charset="us-ascii"
Content-Length: 405
Status:

> Hey, Im in Palo Alto for the rest of this month as well. Give me a
> ring ...

Will definitely do so!

I think your diagnosis of the situation with John is correct, and as for
the play, it was dead on the money! I was in stitches, esp when I read
Peter's line about the rippling moonlight -- perfect! Your ear is every
bit the equal of your eye when it comes to caricature! I am in awe! :-)

-chris

From ???@??? Wed Jul 08 14:57:34 1998
Received: from [143.88.7.102] (lanrover3.coginst.uwf.edu [143.88.7.103])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA23646;
Wed, 8 Jul 1998 12:38:58 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a0ab1c81ceed4b0@[143.88.7.102]>
In-Reply-To: <199807070736.DAA00711@west>
Mime-Version: 1.0
Date: Wed, 8 Jul 1998 10:33:47 -0800
To: cmenzel@tamu.edu, doug@csi.uottawa.ca, fritz@cyc.com,
p.m.simons@leeds.ac.uk, phayes@nuts.coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: The meta-meta-question

Cc: skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 3837
Status:

[JS:]

>I would just like to see all decisions made either by the committee
>voting as a whole or by someone who has been appointed by the committee
>to carry out some clearly designated tasks, which must be reported
>back to the committee when accomplished.

Chris Menzel (Chair): The Committee shall come to order. Gentlemen, if you please! ORDER!!! Thank you. The Secretary shall read the minutes of the last meeting.

Fritz Lehmann (Secretary): We haven't got an ontology for metric time yet, so there aren't any minutes.

Chair: OK, mark that down as something we must get around to eventually. Any other business?

Pat Hayes(member): I propose that the committee declare itself not to be a committee after all and that the members just get on with the work.

John Sowa(member): I object.

Chair: I have an objection...

Sowa(member): I OBJECT to that objection! I objected first!

Chair: Go ahead.

Sowa(member) I object that if the Committee just vanishes, there won't be anything to give the subcommittee any authority to act.

Hayes(member): We don't have any authority, and we are the subcommittee.

Chair: True. Objection overruled.

Sowa(member): I object to that overruling!

Chair: You can't object to an overruling, it's uncomputable in committee logic. Everybody knows that, Tarski proved it, Oh I don't know, the late fifties I think it was. Of course it depends on what semantics one uses...

Secretary: Never mind, get on.

Chair: Ah yes, of course. But here's another objection. If the Committee ceases to exist, there won't be anything for the members to report back to.

Hayes(member): When we have anything to report, we can re-form ourselves into a committee and report back to ourselves. That would be a good excuse for an international party, in any case.

Chair: Neat idea. I'll talk to Bob about it. Objection withdrawn, no more objections allowed. Now, is there a seconder?

Nicola Guarino(member): But wait! It is important that I speak now! If the committee becomes not a committee then according to the doctrine of universal individualized detachment, which I have described in my papers, the new committee which is formed when this committee forms itself again into a committee will not be the same committee as this committee, so we, that is, this committee, considered as a collection, will not then be able to report back to this committee, considered as a whole, because this committee will then not be that committee, even though it has the same mereological parts.

Peter Simons(member): Obviously, a committee is an aeoliaric phrasangial - in the vulgar tongue, an intermittent continuant. Another example would be the reflection of moonlight on the ripples of a deep, silent ocean.

Chair: How do you do that?

Simons(member): Do what?

Chair: Float in the air with your legs crossed.

Secretary: In CYC, anything can be intermittent or not. Or even both and neither, some of the time. I thought this was crazy, but it seems to work. I think.

Doug Skuce(member): OK, I've got all that, but the tree is now 13 layers deep and Windows-98 just crashed. Could everyone stop talking for about 30 minutes?

John Sowa(member): Can we please keep this meeting organized properly? According to Whitehead's Rules of Order, nobody should be able to say more than three things at once unless they have a two-ary relationship to the chairman.

Hayes(member): I think Whitehead is nasty and I'm not going to play ANY MORE and I'm taking my balls! (Member stamps foot and leaves meeting.)

Chair: I have this wierd feeling....

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Jul 08 14:57:36 1998
Received: from [143.88.7.102] (lanrover3.coginst.uwf.edu [143.88.7.103])
 by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA23648;
 Wed, 8 Jul 1998 12:39:07 -0500 (CDT)
<199807070736.DAA00711@west>
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a0bb1c830fa8abd@[143.88.7.102]>
In-Reply-To: <199807071637.LAA15413@philebus.tamu.edu>
References: Your message of "Tue, 07 Jul 1998 03:36:26 EDT."
<199807070736.DAA00711@west>
Mime-Version: 1.0
Date: Wed, 8 Jul 1998 10:33:50 -0800
To: Chris Menzel <cmenzel@philebus.tamu.edu>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: The meta-meta-question
Cc: cmenzel@tamu.edu, doug@csi.uottawa.ca, fritz@cyc.com,
 p.m.simons@leeds.ac.uk, phayes@nuts.coginst.uwf.edu,
 sowa@west.poly.edu, guarino@ladseb.pd.cnr.it
Content-Type: text/plain; charset="us-ascii"
Content-Length: 744
Status:

Chris, I hereby appoint you to be Master of the Universe, at least in a certain context to be determined later by plebiscite. You may wonder on what authority I make this appointment, and that's a good question but the answer cannot be understood by mere mortals, including myself, so there's no point in asking it. Nevertheless, please consider yourself to be Master of the Universe, King of the Hill, Lord of All the Ontologies and general Gopher-in-Chief, OK?

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Fri Jul 10 12:40:30 1998

Received: from LADSEB.LADSEB.PD.CNR.IT (ladseb.ladseb.pd.cnr.it [150.178.2.3])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id LAA22462

for <phayes@coginst.uwf.edu>; Fri, 10 Jul 1998 11:52:59 -0500 (CDT)

Received: from [150.178.2.93] by LADSEB.LADSEB.PD.CNR.IT with ESMTP;
Fri, 10 Jul 1998 18:49:26 +0200

Tue, 07 Jul 1998 03:36:26 EDT." <199807070736.DAA00711@west>

X-Sender: guarino@ladseb.ladseb.pd.cnr.it

Message-Id: <v0310280bb1cbf3ddb447@[150.178.2.93]>

In-Reply-To: <v04003a0bb1c830fa8abd@[143.88.7.102]>

References: <199807071637.LAA15413@philebus.tamu.edu> "Your message of
Tue, 07 Jul 1998 03:36:26 EDT." <199807070736.DAA00711@west>

Mime-Version: 1.0

Date: Fri, 10 Jul 1998 18:50:09 +0200

To: Pat Hayes <phayes@coginst.uwf.edu>, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, sowa@west.poly.edu

From: Nicola Guarino <guarino@ladseb.pd.cnr.it>

Subject: Re: The meta-meta-question

Content-Transfer-Encoding: 8bit

X-MIME-Autoconverted: from quoted-printable to 8bit by nuts.coginst.uwf.edu id
LAA22462

Content-Type: text/plain

Content-Length: 1342

Status:

>[JS:]

>>I would just like to see all decisions made either by the committee
>>voting as a whole or by someone who has been appointed by the committee
>>to carry out some clearly designated tasks, which must be reported
>>back to the committee when accomplished.

>

>-----

>

>Chris Menzel (Chair): The Committee shall come to order. Gentlemen, if you

>please! ORDER!!! Thank you. The Secretary shall read the minutes of the
>last meeting.
[....]

Pat,

your little picture of a typical discussion among us is the nicest
outcome of the Heidelberg workshop...

I am leaving for some holidays, I will go back to the discussion by
the end of the month [Mutton IS part of the flock of sheep, but is NOT a
member of the flock; a single sheep is also a part of the flock, which, by
having some special characteristics (namely, by being a whole) is ALSO a
member of the flock].

Cheers [I will try not to think at mereology during my holidays],

-- Nicola

Nicola Guarino
National Research Council phone: +39 49 8295751
LADSEB-CNR fax: +39 49 8295763
Corso Stati Uniti, 4 email: guarino@ladseb.pd.cnr.it
I-35127 Padova
Italy

Home page: ** updated 4/5/98 **
<http://www.ladseb.pd.cnr.it/infor/ontology/ontology.html>

FOIS'98 home page:
<http://mnemosyne.itc.it:1024/fois98/>

From ???@??? Mon Jul 20 21:56:50 1998
Received: from Hypatia.Stanford.EDU (Hypatia.Stanford.EDU [171.64.22.122])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id SAA17879
for <phayes@coginst.uwf.edu>; Mon, 20 Jul 1998 18:38:44 -0500 (CDT)
Received: from Turing.Stanford.EDU (Turing.Stanford.EDU [171.64.22.14])
by Hypatia.Stanford.EDU (8.9.0.Beta5/8.9.0.Beta5) with ESMTP id QAA28790
for <phayes@coginst.uwf.edu>; Mon, 20 Jul 1998 16:35:22 -0700 (PDT)

Received: from localhost (cmenzel@localhost)
by Turing.Stanford.EDU (8.8.6/8.8.5) with ESMTP id QAA06765
for <phayes@coginst.uwf.edu>; Mon, 20 Jul 1998 16:35:21 -0700 (PDT)
Message-Id: <199807202335.QAA06765@Turing.Stanford.EDU>
X-Authentication-Warning: Turing.Stanford.EDU: cmenzel owned process doing -bs
X-Mailer: exmh version 2.0.2 2/24/98
To: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Comments on notes so far
In-reply-to: Your message of "Fri, 17 Jul 1998 20:04:21 PST."
<v04003a04b1d57d318d0c@[143.88.7.101]>
Mime-Version: 1.0
Date: Mon, 20 Jul 1998 16:35:21 -0700
From: Christopher Menzel <cmenzel@csl.stanford.edu>
Content-Type: text/plain; charset=us-ascii
Content-Length: 245
Status:

Great Pat, thanks. Am currently working up a message to send to the "upper level" group on the organization of The Document (authorship, structure, etc) and then will print out and study your comments and distribute appropriately.

-chris

From ???@??? Thu Aug 06 12:16:08 1998
Received: from mserv1b.u-net.net (mserv1b.u-net.net [195.102.240.137])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id QAA29998
for <phayes@coginst.uwf.edu>; Wed, 5 Aug 1998 16:08:41 -0500 (CDT)
Received: from ([195.102.196.126]) [195.102.196.126]
by mserv1b.u-net.net with esmtp (Exim 1.82 #2)
id 0z4Ahy-0006r0-00; Wed, 5 Aug 1998 22:02:54 +0100
X-Sender: pmsimons.peewit@mail.u-net.com (Unverified)
Message-Id: <v03110700b1ee66743640@[195.102.196.104]>
Mime-Version: 1.0
Date: Wed, 5 Aug 1998 21:08:32 +0100
To: cmenzel@philebus.tamu.edu, sowa@west.poly.edu, axf@ksl.stanford.edu,
phayes@coginst.uwf.edu, doug@csi.uottawa.ca, fritz@cyc.com,
guarino@ladseb.pd.cnr.it, e6nl001@coe.coppin.umd.edu,
chezewiz@erols.com, jmc@cs.stanford.edu, cmenzel@tamu.edu
From: P Simons <p.m.simons@leeds.ac.uk>
Subject: Ontology Document Comments
Content-Type: text/plain; charset="us-ascii"
Content-Length: 1711
Status:

Gentlemen

Some comments on the recent discussion.

First: Thanks to Chris for following up his clearing-house and mediating position.

It is surely for those who have an interest in a standard to say what form they want a document to take. Sensibly it should record what consensus we have achieved plus indications of the alternatives where dissent exists. The popular saw that a camel is a horse designed by a committee will no doubt haunt us.

Whether DEPENDS ON is primitive or definable depends on other choices. Kit Fine thinks it is primitive and he may be right. The problem is then the axioms for it.

If one does not believe in necessary existents (as I for one do not) then "a depends on b iff necessarily if a exists then b exists and ..." as in PARTS is not so dusty. But this shows that even how one defines something is not theoretically neutral.

To my knowledge the first person to use 'trope' to mean what realists about universals might call an instance of a universal was Donald Carey Williams in his "The Elements of Being". ca. 1951. I believe he took the term from Santayana, but used it in this sense precisely because it was not in other use in philosophy. The term is now very popular in Australia and is more or less established in this usage.

When databasers use a noun in the singular with a capital letter they usually mean a kind. But not all kinds are properties or relations: 'Animal' for instance names a (higher) kind of thing, not a property, whereas e.g. 'square' names a property.

Peter

Professor Peter Simons
School of Philosophy
University of Leeds
Leeds LS2 9JT
England
Tel +44 113 233 3298
Fax +44 113 233 3265
<http://www.leeds.ac.uk/philosophy/html/simons.htm>

From ???@??? Fri Aug 07 08:33:42 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id BAA00966
for <phayes@coginst.uwf.edu>; Fri, 7 Aug 1998 01:25:04 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id CAA03121;
Fri, 7 Aug 1998 02:21:32 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
id CAA28332; Fri, 7 Aug 1998 02:18:00 -0400
Date: Fri, 7 Aug 1998 02:18:00 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199808070618.CAA28332@west>
To: axf@ksl.stanford.edu, chezewiz@erols.com, cmenzel@philebus.tamu.edu,
cmenzel@tamu.edu, doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu,
fritz@cyc.com, guarino@ladseb.pd.cnr.it, jmc@cs.stanford.edu,
p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu, skydog@pacbell.net,
sowa@west.poly.edu
Subject: Re: Ontology Document Comments
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 2983
Status:

Peter,

Thanks for the comments. As usual, they are quite helpful.
Just a few comments on your comments:

>... Sensibly it should record what consensus we
>have achieved plus indications of the alternatives where dissent exists.

I don't believe that we have any real dissent. It is rather that
we have a fair number of agreements and a large number of issues
where there are quite a few alternatives that have not been explored
in enough depth for us to take a firm stand. The document we prepare
should endorse whatever consensus we have achieved and then indicate
alternatives.

And many alternatives are not matters of dispute. Instead they are
true options where one might choose one version for one application
and another version for a different application. An example is the
option of discrete, continuous, or lumpy axioms for mereology. Any
of those axioms might be appropriate for some application, but we
have to warn people that no two of them can be used together without
generating a contradiction.

>If one does not believe in necessary existents (as I for one do not) then
>"a depends on b iff necessarily if a exists then b exists and ..." as in
>PARTS is not so dusty. But this shows that even how one defines something
>is not theoretically neutral.

Many such axioms can be restated in terms of impossibility: "a depends on b iff a cannot exist without b...." I find it easier to believe that certain things are impossible than that certain things are necessary.

>To my knowledge the first person to use 'trope' to mean what realists about
>universals might call an instance of a universal was Donald Carey Williams
>in his "The Elements of Being". ca. 1951. I believe he took the term from
>Santayana, but used it in this sense precisely because it was not in other
>use in philosophy. The term is now very popular in Australia and is more or
>less established in this usage.

When I suggested that 'trope' should be put on the deprecated list, I did not mean that we could not use it in quotations or discussions about some topic where it is commonly used. But it is an example of a large number of metalevel terms that do not actually appear in the ontology itself, but in discussions about the ontology. If we are using logic as the formal notation, the word 'trope' would not occur. If we want to talk about the topic in English, we can do so by talking about the logical constructions, using phrases like "instance of".

Sometimes the clearest way to talk about a metalevel issue in English is just to put a logical expression like $P(x)$ in the middle of a sentence. It is easier to read and write " $2+2$ " than "the sum of two and two".

That doesn't solve all the problems, since people argue about what a statement in first-order logic "really means". But as long as we can agree on how something is represented in logic, our computer systems can interoperate even though the users have different opinions about what the symbols denote.

John

From ???@??? Sun Aug 09 16:27:52 1998

Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id EAA14151
for <phayes@coginst.uwf.edu>; Sat, 8 Aug 1998 04:46:11 -0500 (CDT)

Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id FAA06301;
Sat, 8 Aug 1998 05:42:37 -0400 (EDT)

Received: by west (SMI-8.6/SMI-SVR4)
id FAA04561; Sat, 8 Aug 1998 05:39:13 -0400

Date: Sat, 8 Aug 1998 05:39:13 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199808080939.FAA04561@west>
To: fritz@cyc.com
Subject: Re: Ontology Document Comments
Cc: axf@ksl.stanford.edu, chezewiz@erols.com, cmenzel@philebus.tamu.edu,
cmenzel@tamu.edu, doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu,
guarino@ladseb.pd.cnr.it, jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk,
phayes@coginst.uwf.edu, skydog@pacbell.net, sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 922
Status:

Fritz,

The term 'deprecated' is standard in the standards arena (where I've been spending a fair amount of time) in the sense of a term, feature, or construct that should be avoided in particular system. It does not carry any negative connotations or implications about the value of the feature in other contexts.

I don't disagree with your comments. But I do believe that we should adopt a small vocabulary of terms that we use for the presentation of the ontology and its rationale. To avoid pejorative connotations, we could call them the 'core terms' and put other words in the 'non-core' category. The non-core terms could be used in a survey or comparison of various positions, but they would not be used in the basic definitions, examples, and tutorials.

The word 'trope' is not a term that I have strong feelings about one way or another. I just mentioned it as a typical example of a non-core term.

John

From ???@??? Wed Jul 15 00:13:16 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id BAA23166
for <phayes@nuts.coginst.uwf.edu>; Wed, 15 Jul 1998 01:14:47 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id CAA16818;
Wed, 15 Jul 1998 02:11:28 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
id CAA19519; Wed, 15 Jul 1998 02:08:30 -0400
Date: Wed, 15 Jul 1998 02:08:30 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199807150608.CAA19519@west>

To: axf@KSL.Stanford.edu, phayes@coginst.uwf.edu
Subject: Continuants and occurrents
Cc: cmenzel@tamu.edu, doug@csi.uottawa.ca, fritz@cyc.com,
guarino@ladseb.pd.cnr.it, p.m.simons@leeds.ac.uk,
phayes@nuts.coginst.uwf.edu, sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 8389
Status:

At Heidelberg, there was some criticism of the distinction between continuants and occurrents based on temporal vs. spatial parts because it was not generalizable to a four-dimensional coordinate system. As a result of that discussion, I added two more subsections to my Ch. 2 on ontology, the first on time's arrow and the second on a generalization of the definitions to a four-dimensional space-time.

Following are the two sections (some of the mathematical notation does not print in this font, but it should be more or less intelligible). For most applications limited to the earth and the nearby solar system, Peter Simons' definitions in terms of time-dependent mereology seem to be preferable to the more complex definitions in terms of entropy, etc. But it is an interesting exercise to show the possible generalizations.

I fully agree that we should design ontologies that are formally axiomatizable, but I don't believe that any fixed set of axioms will be adequate for all applications. I believe that the framework of concept types will be more stable than any axioms we attach to any of the nodes in the framework.

John

TIME'S ARROW. In Einstein's theory of relativity, three-dimensional space and one-dimensional time are combined in a four-dimensional space-time continuum. That treatment simplifies the equations of physics, but it masks a fundamental asymmetry in time that does not affect space: the difference between past and future. To characterize that difference, Arthur Stanley Eddington (1928) imagined an arrow drawn somewhere in the space-time continuum:

Let us draw an arrow arbitrarily. If as we follow the arrow we find more and more of the random element in the world, then the arrow is pointing towards the future; if the random element decreases the arrow points towards the past.... I shall use the phrase "time's arrow" to express this one-way property of time

which has no analogue in space.

The random element is only apparent when there are enough objects in a situation to be statistically significant. If a movie of two billiard balls colliding were played forward or backward, both directions would represent physically possible events. But if any movie showed sixteen billiard balls coalescing into a triangle of fifteen while spitting out the cue ball in the direction of the player, it would be safe to assume that the film was running in reverse. When left to themselves, situations become more randomized with the passage of time, and the effects that are noticeable with a few billiard balls become more pronounced with larger numbers of interacting atoms and molecules.

For physical entities, the measure of randomness is called entropy; for abstractions, the measure of randomness is called information. Both entropy and information increase with time, and their increase is governed by the same mathematical laws. When atoms or billiard balls scatter, the increase in entropy is proportional to the increase in the number of bits required to encode their configuration. Like entropy and information, causality also affects the future, and not the past. In his lectures on cause and chance in physics, Max Born (1949) stated three assumptions that dominated the classical view:

- o "Causality postulates that there are laws by which the occurrence of an entity B of a certain class depends on the occurrence of an entity A of another class, where the word entity means any physical object, phenomenon, situation, or event. A is called the cause, B the effect."
- o "Antecedence postulates that the cause must be prior to, or at least simultaneous with, the effect."
- o "Contiguity postulates that cause and effect must be in spatial contact or connected by a chain of intermediate things in contact."

Relativity and quantum mechanics have forced physicists to abandon these assumptions as exact statements of what happens at the most fundamental levels, but they remain valid at the level of human experience. After analyzing them in terms of modern physics, Born concluded "chance has become the primary notion, mechanics an expression of its quantitative laws, and the overwhelming evidence of causality with all its attributes in the realm of ordinary experience is satisfactorily explained by the statistical laws of large numbers."

The arrow of time with its implications for entropy, information, and causality is well defined only because the universe is still evolving

from the big bang, when entropy was extremely low. Huw Price (1996) maintained that in a much older universe near thermal equilibrium, the arrow would be undefined or randomly fluctuating. In exceptional circumstances, such as matter and energy falling into a black hole, entropy and information might be destroyed, and the arrow of time could even be reversed. A universal ontology should accommodate any novel phenomena discovered by physicists and astronomers. But to describe ordinary events on the earth and the surrounding solar system, the ontology must also include a well developed stock of concepts for representing the familiar direction of time and its implications for causality and information flow in everyday life. Those concepts are discussed further in Chapter 4.

CONTINUANTS AND OCCURRENTS. The distinction between continuants and occurrents, which was informally discussed in Section 2.3, can only be formalized in terms of some axiomatization for time. Simons (1987) used time-dependent mereology to state the definitions:

- o A continuant has spatially distinguishable parts, but no temporal parts. A human being, for example, is constantly gaining and losing molecules, but at any time t when a person x exists, all of x 's parts exist at the same time t . Even if x loses some part, such as a tooth, that tooth (or at least its atoms) would continue to exist at the same time as x .
- o An occurrent has both spatial and temporal parts. Examples of occurrents include concerts, sports events, journeys, storms, and earthquakes. A complete specification of an occurrent must include all the parts (called *_stages_*) from beginning to end. Although a human being is a continuant, the life of a human being is an occurrent whose stages are spread out over the interval from birth to death.

These definitions apply equally well to physical entities and to the abstractions that encode their structure, but without the accompanying matter or energy.

Simons' definitions depend on a representation that separates one-dimensional time from three-dimensional space. They are adequate for ordinary human experience, but they would have to be generalized to four-dimensional space-time to accommodate modern theories of physics. A generalization could be defined in terms of entropy and the arrow of time:

- o For any simply connected four-dimensional space-time region R , determine whether there exists a continuous path I , called a *_time*

line_, that satisfies the following axioms:

The path I extends from one boundary point $t_{\text{sub}0}$ of R to another boundary point $t_{\text{sub}1}$ of R , where the entropy $S(t_{\text{sub}0})$ is a minimum for the region R , and $S(t_{\text{sub}1})$ is a maximum for R .

Every point t of the path I is tangent to a _time arrow_, which is defined as the direction of maximum increase in entropy ΔS at the point t .

- o If a time line I can be found for the region R , then at each point t of I , define a _snapshot_ $S(t)$ as the intersection of R with a three-dimensional hyperplane perpendicular to I at t .
- o A _spatial form_ is defined as a predicate $P(s)$ that makes a true or false statement about the configuration of matter and energy in a snapshot s .
- o The region R is said to contain the successive stages of a continuant if there exists a spatial form P such that for every time $t_{\text{sub}0} \leq t \leq t_{\text{sub}1}$, P is true of the snapshot $S(t)$.
- o If any of the above conditions is false, R is said to contain an occurrent.

For most human experience, these definitions coincide with Simons' definitions, but they may have to be revised when physicists make new discoveries. Simons' definitions are easier to state, and they do not change with current trends in physics.

From ???@??? Mon Jul 27 10:53:52 1998

Received: from [143.88.7.104] (lanrover6.coginst.uwf.edu [143.88.7.106])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id PAA17428;
Thu, 23 Jul 1998 15:29:40 -0500 (CDT)

X-Sender: phayes@mail.coginst.uwf.edu

Message-Id: <v04003a00b1dd5531a123@[143.88.7.104]>

In-Reply-To: <3.0.32.19980722125419.00ecb758@catbert.cyc.com>

Mime-Version: 1.0

Date: Thu, 23 Jul 1998 13:21:51 -0800

To: Fritz Lehmann <fritz@cyc.com>

From: Pat Hayes <phayes@coginst.uwf.edu>

Subject: Re: Continuants and occurrents

Cc: sowa@west.poly.edu, axf@ksl.stanford.edu, phayes@coginst.uwf.edu,
doug@csi.uottawa.ca, guarino@ladseb.pd.cnr.it, p.m.simons@leeds.ac.uk,
chezewiz@erols.com, e6nl001@coe.coppin.umd.edu, skydog@pacbell.net,
cmenzel@tamu.edu

Content-Type: text/plain; charset="us-ascii"
Content-Length: 2504
Status: O

>At 09:27 AM 7/22/98 -0700, Christopher Menzel wrote:
>>>[...]
>>> 4. Minimize the use, and discussion, of logical notations. Use English.
>>> This is ontology, not logic class.
>>
>>For the general presentation of the ontology, I strongly agree. But if
>>you mean anything more than that, I strongly disagree. While we do not
>>>want to *present* the ontology as a logical theory, the ontology should
>>still be *completely* formalized, to the last detail, and included in an
>>appendix.
>>[...]
>
>Dear Chris,
>
>Agreed. I meant that the general reader should not have to be familiar
>with any logical notation to know whether a particular distinction applies
>to a new term. Certainly it should be formalized, to the extent possible,
>in the appendix. I suspect that most of the precise characterization in
>logic can also be stated earlier in crisp, precise English.
>

Of course we have to write readably. However we also need to be conscious that English carries already within it all kinds of ontological assumptions and prejudices, and we must take care to be explicit in drawing attention to these and indicating which of them we are assuming, and which we are calling into question. Its not logic class but it is ontology class, inevitably. So while I agree about the use of crisp, precise, English, I think that there is no way to avoid the need for a certain amount of introductory material which the reader must be prepared to wade through and understand. We can't write this as though it were a car repair manual.

But there's little point in arguing about these issues. It would be better to get on with the actual writing. I hope to be able to get down to this seriously in about a week: right now I'm hitting deadlines on other projects.

Pat

PS. Fritz, your CC list now haas me on it twice again; could you delete one of them? Thanks. And I notice that it contains some new entries. Should we include these into the general Villa-Bosch team?

PPS. I also realise that we have not included John McCarthy in our email discussions. As he was present and took part in some of the discussions, and may even have a useful thing to say now and again, maybe we should add him to the list. Does anyone disagree?

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Sat Aug 01 13:20:41 1998
Received: from LADSEB.LADSEB.PD.CNR.IT (ladseb.ladseb.pd.cnr.it [150.178.2.3])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id NAA03363
for <phayes@coginst.uwf.edu>; Wed, 29 Jul 1998 13:10:21 -0500 (CDT)
Received: from [150.178.2.93] by LADSEB.LADSEB.PD.CNR.IT with ESMTP;
Wed, 29 Jul 1998 20:06:45 +0200
X-Sender: guarino@ladseb.ladseb.pd.cnr.it
Message-Id: <v0310280db1e51142e03d@[150.178.2.93]>
In-Reply-To: <199807150608.CAA19519@west>
Mime-Version: 1.0
Date: Wed, 29 Jul 1998 20:08:22 +0200
To: sowa@west.poly.edu (John F. Sowa), axf@KSL.Stanford.edu,
phayes@coginst.uwf.edu
From: Nicola Guarino <guarino@ladseb.pd.cnr.it>
Subject: Re: Continuants and occurrents
Cc: cmenzel@tamu.edu, doug@csi.uottawa.ca, fritz@cyc.com,
p.m.simons@leeds.ac.uk, phayes@nuts.coginst.uwf.edu,
sowa@west.poly.edu
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2457
Status:

At 2:08 AM -0400 7/15/98, John F. Sowa wrote:
>CONTINUANTS AND OCCURRENTS. The distinction between continuants and
>occurrents, which was informally discussed in Section 2.3, can only be
>formalized in terms of some axiomatization for time. Simons (1987) used
>time-dependent mereology to state the definitions:
>
>o A continuant has spatially distinguishable parts, but no temporal
> parts. A human being, for example, is constantly gaining and losing
> molecules, but at any time t when a person x exists, all of x's parts
> exist at the same time t. Even if x loses some part, such as a

- > tooth, that tooth (or at least its atoms) would continue to exist at
- > the same time as x.
- >
- >o An occurrent has both spatial and temporal parts. Examples of
- > occurrents include concerts, sports events, journeys, storms, and
- > earthquakes. A complete specification of an occurrent must include
- > all the parts (called _stages_) from beginning to end. Although
- > a human being is a continuant, the life of a human being is an
- > occurrent whose stages are spread out over the interval from birth
- > to death.
- >

Are those above Peter's words? I don't think so! I don't find these statements in his book. I believe you should avoid to refer to these words as "Simons' definitions".

By the way, the second part of the first "definition" is misleading. The point is not so much that the tooth continues to exist after being detached from the body, but rather that all the parts of a continuant exists at any time of its life: before losing the tooth, x is wholly there; after losing the tooth, x is still wholly there. On the other hand, we can't properly say that parts of occurrents *exist* at certain time: they just "occur"... When my current writing is at the end, its beginning (which is part of it) has already gone; this means that in this moment my writing this message is not *wholly* there...

Another aspect is that continuants may have contingent parts (like the tooth), while the (temporal) parts of occurrents are all essential....

Nicola Guarino
National Research Council phone: +39 49 8295751
LADSEB-CNR fax: +39 49 8295763
Corso Stati Uniti, 4 email: guarino@ladseb.pd.cnr.it
I-35127 Padova
Italy

Home page: ** updated 4/5/98 **
<http://www.ladseb.pd.cnr.it/infor/ontology/ontology.html>

FOIS'98 home page:

<http://mnemosyne.itc.it:1024/fois98/>

From ???@??? Sat Aug 01 13:20:47 1998

Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id FAA16747
for <phayes@nuts.coginst.uwf.edu>; Thu, 30 Jul 1998 05:52:58 -0500 (CDT)

Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id GAA05230;
Thu, 30 Jul 1998 06:49:27 -0400 (EDT)

Received: by west (SMI-8.6/SMI-SVR4)
id GAA14977; Thu, 30 Jul 1998 06:46:22 -0400

Date: Thu, 30 Jul 1998 06:46:22 -0400

From: sowa@west.poly.edu (John F. Sowa)

Message-Id: <199807301046.GAA14977@west>

To: axf@KSL.Stanford.edu, guarino@ladseb.pd.cnr.it, phayes@coginst.uwf.edu

Subject: Re: Continuants and occurrents

Cc: cmenzel@tamu.edu, doug@csi.uottawa.ca, fritz@cyc.com,
p.m.simons@leeds.ac.uk, phayes@nuts.coginst.uwf.edu

Mime-Version: 1.0

Content-Type: text/plain; charset=us-ascii

Content-Length: 717

Status:

Nicola,

I derived the ideas from Peter Simons' book, but those points are not direct quotes. Following are his actual words:

- Page 129: "Occurrents comprise what are variously called events, processes, happenings, occurrences, and states. They are, like continuants, in time, but unlike continuants they have temporal parts."
- Page 175: "A continuant is an object which is in time, but of which it makes no sense to say that it has temporal parts or phases. At any time at which it exists, a continuant is wholly present."

I think that I will quote these two passages, and then continue with my own discussion. That will avoid any question about what Simons said and what I am saying.

John

From ???@??? Sat Aug 01 13:20:49 1998

Received: from LADSEB.LADSEB.PD.CNR.IT (ladseb.ladseb.pd.cnr.it [150.178.2.3])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id IAA21833
for <phayes@coginst.uwf.edu>; Thu, 30 Jul 1998 08:03:37 -0500 (CDT)

Received: from [150.178.2.93] by LADSEB.LADSEB.PD.CNR.IT with ESMTP;
Thu, 30 Jul 1998 14:59:35 +0200
X-Sender: guarino@ladseb.ladseb.pd.cnr.it
Message-Id: <v03102803b1e6085fae27@[150.178.2.93]>
In-Reply-To: <199807150608.CAA19519@west>
Mime-Version: 1.0
Date: Thu, 30 Jul 1998 14:51:09 +0200
To: sowa@west.poly.edu (John F. Sowa), axf@KSL.Stanford.edu,
phayes@coginst.uwf.edu
From: Nicola Guarino <guarino@ladseb.pd.cnr.it>
Subject: Re: Continuants and occurrents
Cc: cmenzel@tamu.edu, doug@csi.uottawa.ca, fritz@cyc.com,
p.m.simons@leeds.ac.uk, phayes@nuts.coginst.uwf.edu
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2421
Status:

At 2:08 AM -0400 7/15/98, John F. Sowa wrote:

- > Simons' definitions depend on a representation that separates
- > one-dimensional time from three-dimensional space. They are adequate
- > for ordinary human experience, but they would have to be generalized
- > to four-dimensional space-time to accommodate modern theories of
- > physics. A generalization could be defined in terms of entropy and
- > the arrow of time:
- >
- >o For any simply connected four-dimensional space-time region R,
- > determine whether there exists a continuous path I, called a _time
- > line_, that satisfies the following axioms:
- >
- > The path I extends from one boundary point $t_{\text{sub}0}$. of R to another
- > boundary point $t_{\text{sub}1}$. of R, where the entropy $S(t_{\text{sub}0})$ is a
- > minimum for the region R, and $S(t_{\text{sub}1})$ is a maximum for R.
- >
- > Every point t of the path I is tangent to a _time arrow_, which is
- > defined as the direction of maximum increase in entropy ΔS at
- > the point t.
- >
- >o If a time line I can be found for the region R, then at each point t
- > of I, define a _snapshot_ $S(t)$ as the intersection of R with a
- > three-dimensional hyperplane perpendicular to I at t.
- >
- >o A _spatial form_ is defined as a predicate $P(s)$ that makes a true or
- > false statement about the configuration of matter and energy in a
- > snapshot s.
- >
- >o The region R is said to contain the successive stages of a continuant

- > if there exists a spatial form P such that for every time
- > t sub 0 le t le t sub 1, P is true of the snapshot S(t).
- >
- >o If any of the above conditions is false, R is said to contain an
- > occurrent.

Sorry, I do not understand this formulation. I (sort of) understand the reconstruction of the time arrow, but I can't see how to distinguish continuants from occurrents on the basis of the truth-persistence of a predicate P. For instance, P could be "is a snapshot" or "is a configuration of matter and energy", and in this case it would be always true both for continuants and occurrents...

-- Nicola

Nicola Guarino
National Research Council phone: +39 49 8295751
LADSEB-CNR fax: +39 49 8295763
Corso Stati Uniti, 4 email: guarino@ladseb.pd.cnr.it
I-35127 Padova
Italy

Home page: ** updated 4/5/98 **
<http://www.ladseb.pd.cnr.it/infor/ontology/ontology.html>

FOIS'98 home page:
<http://mnemosyne.itc.it:1024/fois98/>

From ???@??? Sat Aug 01 13:20:57 1998
Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id NAA07936
for <phayes@coginst.uwf.edu>; Thu, 30 Jul 1998 13:52:19 -0500 (CDT)
Received: from scratchy (scratchy [207.207.8.118])
by catbert.cyc.com (8.8.8/8.8.8) with SMTP id NAA14112;
Thu, 30 Jul 1998 13:45:44 -0500 (CDT)
Message-Id: <3.0.32.19980730134550.00a88dd8@catbert.cyc.com>
X-Sender: fritz@catbert.cyc.com
X-Mailer: Windows Eudora Pro Version 3.0 (32)
Date: Thu, 30 Jul 1998 13:46:51 -0500
To: sowa@west.poly.edu (John F. Sowa)
From: Fritz Lehmann <fritz@cyc.com>

Subject: Re: Continuants and occurrents

Cc: guarino@ladseb.pd.cnr.it, jmc@cs.stanford.edu, chezewiz@erols.com,
phayes@coginst.uwf.edu, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@nuts.coginst.uwf.edu,
e6nl001@coe.coppin.umd.edu, axf@KSL.Stanford.edu

Mime-Version: 1.0

Content-Type: text/plain; charset="us-ascii"

Content-Length: 1928

Status:

At 06:46 AM 7/30/98 -0400, John F. Sowa wrote:

>Nicola,

>I derived the ideas from Peter Simons' book, but those points are

>not direct quotes. Following are his actual words:

> - Page 129: "Occurrents comprise what are variously called events,

> processes, happenings, occurrences, and states. They are, like

> continuants, in time, but unlike continuants they have temporal parts."

> - Page 175: "A continuant is an object which is in time, but of which

> it makes no sense to say that it has temporal parts or phases. At any

> time at which it exists, a continuant is wholly present."

>I think that I will quote these two passages, and then continue with

>my own discussion. That will avoid any question about what Simons said

>and what I am saying.

>John

Dear John,

This is a definition of continuant (being "wholly present at present").

I haven't read Peter's book. Can anyone explain why it is a good
definition of continuant?

I continue to believe what I said in my article on "Big Posets of Participations" (in Conceptual Structures-96 in Sydney) basically: an occurrent is a change, i.e. a difference in something correlated with a difference in time (it may be simple or composite, discrete or continuous); a continuant is something that persists (does not change with respect to its identity criteria) during a difference in time. So "occurrent" is to "continuant" as "difference" is to "sameness". Although the difference/sameness is with respect to time, it has nothing to do with McTaggart's A-Time (tensed, indexical) versus B-Time (God's-eye, spacetime) notion.

This may or may not conflict with Peter's notion. It doesn't look identical at all.

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

=====
=====

From ???@??? Mon Aug 17 12:21:49 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA15243;
Mon, 17 Aug 1998 11:27:50 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a0eb1fd7a6f7218@[143.88.7.103]>
In-Reply-To: <v03102803b1e6085fae27@[150.178.2.93]>
References: <199807150608.CAA19519@west>
Mime-Version: 1.0
Date: Mon, 17 Aug 1998 01:29:44 -0600
To: Nicola Guarino <guarino@ladseb.pd.cnr.it>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Continuants and occurrents
Cc: cmenzel@tamu.edu, doug@csi.uottawa.ca, fritz@cyc.com,
p.m.simons@leeds.ac.uk, phayes@nuts.coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>
Content-Type: text/plain; charset="us-ascii"
Content-Length: 1330
Status:

>At 2:08 AM -0400 7/15/98, John F. Sowa wrote:
>> Simons' definitions depend on a representation that separates
>>one-dimensional time from three-dimensional space. They are adequate
>>for ordinary human experience, but they would have to be generalized
>>to four-dimensional space-time to accommodate modern theories of
>>physics. A generalization could be defined in terms of entropy and
>>the arrow of time:
[[definitions deleted]]
>
>Sorry, I do not understand this formulation. I (sort of) understand the
>reconstruction of the time arrow, but I can't see how to distinguish
>continuants from occurrents on the basis of the truth-persistence of a
>predicate P. For instance, P could be "is a snapshot" or "is a
>configuration of matter and energy", and in this case it would be always
>>true both for continuants and occurrents...

Yes, I agree. I think that the entire business of defining time's arrow in terms of entropy is, while fascinating, best left aside for ontological

purposes. We can just take it as given that time has an arrow.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Mon Aug 17 11:23:04 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
 by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id DAA21662
 for <phayes@coginst.uwf.edu>; Mon, 17 Aug 1998 03:01:01 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
 by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id DAA26411;
 Mon, 17 Aug 1998 03:57:01 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
 id DAA21073; Mon, 17 Aug 1998 03:53:07 -0400
Date: Mon, 17 Aug 1998 03:53:07 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199808170753.DAA21073@west>
To: skydog@pacbell.net
Subject: Similarity measure
Cc: E6NL001@coe.coppin.umd.edu, axf@KSL.Stanford.EDU, chezewiz@erols.com,
 cmenzel@turing.stanford.edu, doug@site.uottawa.ca, fritz@cyc.com,
 guarino@ladseb.pd.cnr.it, jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk,
 phayes@coginst.uwf.edu, sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 2114
Status:

Bob,

Your point is well taken:

>Is this measure of similarity, $m(x,y)$, a consistent one? Should one
>attempt to classify a cow in the same way one would classify a contract
>or a lamp? If an entity can be classified in more than one way - a dog
>is a pet and a canine - how does one know when it is properly and/or
>exhaustively classified?

I revised that definition to add a third argument to the similarity measure: $m(x,y,c)$ compares two entities x and y according to some standard for category c . For example, if x is a black cat and y is an orange cat, they would register high on a similarity measure by a standard for the category Cat, but they would not measure high by a standard for BlackEntity.

>Does a measure of similarity imply a measure of distance? Since
>ontologies can be of arbitrary size and uneven in their granularity, by
>distance I mean something more than simply the shortest path length. Are
>these measures related? Would two entities similarity or distance be
>different if the size or granularity of the ontology were different?

A measure of semantic distance would be the inverse of similarity. But defining either one is definitely nontrivial. People have been working on measures of similarity or semantic distance for years with only moderate success: there are a lot of simple measures that are OK for some purposes, but there are no really good ones that are good or even adequate for general consumption.

The problem of evaluating similarity or semantic distance is the major problem that info retrieval engines face. No one has yet developed a really good way to reduce the enormous number of hits that are produced for Internet searches.

The primary purpose of my note was to answer Adam's question about the kind of logic necessary to define 'natural kind' and/or related notions, such as prototype-based ontologies. All such approaches depend on having a similarity or semantic distance measure. And they all run aground in essentially the same way: defining a good measure that works for a wide range of cases is still a major research problem.

John

From ???@??? Mon Aug 17 16:10:56 1998
Received: from vapor.stanford.edu (vapor.Stanford.EDU [171.64.71.11])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id OAA22520
for <phayes@coginst.uwf.edu>; Mon, 17 Aug 1998 14:03:13 -0500 (CDT)
Received: (from jmc@localhost)
by vapor.stanford.edu (8.8.8/8.8.8) id LAA25629;
Mon, 17 Aug 1998 11:59:29 -0700 (PDT)
Date: Mon, 17 Aug 1998 11:59:29 -0700 (PDT)
Message-Id: <199808171859.LAA25629@vapor.stanford.edu>
From: John McCarthy <jmc@Steam.Stanford.EDU>
To: sowa@west.poly.edu
CC: skydog@pacbell.net, E6NL001@coe.coppin.umd.edu, axf@KSL.Stanford.EDU,

chezewiz@erols.com, cmenzel@turing.stanford.edu, doug@site.uottawa.ca,
fritz@cyc.com, guarino@ladseb.pd.cnr.it, p.m.simons@leeds.ac.uk,
phayes@coginst.uwf.edu, sowa@west.poly.edu

In-reply-to: <199808170753.DAA21073@west> (sowa@west.poly.edu)

Subject: Re: Similarity measure

Reply-to: jmc@cs.Stanford.EDU

Mime-Version: 1.0

Content-Type: text/plain; charset=us-ascii

Content-Length: 957

Status:

John Sowa includes:

The primary purpose of my note was to answer Adam's question about the kind of logic necessary to define 'natural kind' and/or related notions, such as prototype-based ontologies. All such approaches depend on having a similarity or semantic distance measure. And they all run aground in essentially the same way: defining a good measure that works for a wide range of cases is still a major research problem.

I don't think the human notions of natural kinds are based on similarity or semantic distance measures, and I don't think AI should depend on them either.

Here's an idea.

NaturalKind(ϕ) \Rightarrow (exists set-of-properties)((forall property in set-of-properties)(Interesting(property) and (forall x)($\phi(x) \Rightarrow$ property(x))) and BigEnough(set-of-propertiesZ))

However Interesting and BigEnough are left unspecified and seem to be context dependent.

Similarity measures are not involved.

From ???@??? Mon Aug 17 16:10:56 1998

Received: from vapor.stanford.edu (vapor.Stanford.EDU [171.64.71.11])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id OAA22520

for <phayes@coginst.uwf.edu>; Mon, 17 Aug 1998 14:03:13 -0500 (CDT)

Received: (from jmc@localhost)

by vapor.stanford.edu (8.8.8/8.8.8) id LAA25629;

Mon, 17 Aug 1998 11:59:29 -0700 (PDT)

Date: Mon, 17 Aug 1998 11:59:29 -0700 (PDT)

Message-Id: <199808171859.LAA25629@vapor.stanford.edu>

From: John McCarthy <jmc@Steam.Stanford.EDU>

To: sowa@west.poly.edu

CC: skydog@pacbell.net, E6NL001@coe.coppin.umd.edu, axf@KSL.Stanford.EDU,
chezewiz@erols.com, cmenzel@turing.stanford.edu, doug@site.uottawa.ca,
fritz@cyc.com, guarino@ladseb.pd.cnr.it, p.m.simons@leeds.ac.uk,
phayes@coginst.uwf.edu, sowa@west.poly.edu

In-reply-to: <199808170753.DAA21073@west> (sowa@west.poly.edu)

Subject: Re: Similarity measure

Reply-to: jmc@cs.Stanford.EDU

Mime-Version: 1.0

Content-Type: text/plain; charset=us-ascii

Content-Length: 957

Status:

John Sowa includes:

The primary purpose of my note was to answer Adam's question about the kind of logic necessary to define 'natural kind' and/or related notions, such as prototype-based ontologies. All such approaches depend on having a similarity or semantic distance measure. And they all run aground in essentially the same way: defining a good measure that works for a wide range of cases is still a major research problem.

I don't think the human notions of natural kinds are based on similarity or semantic distance measures, and I don't think AI should depend on them either.

Here's an idea.

NaturalKind(phi) => (exists set-of-properties)((forall property in set-of-properties)(Interesting(property) and (forall x)(phi(x) => property(x))) and BigEnough(set-of-propertiesZ))

However Interesting and BigEnough are left unspecified and seem to be context dependent.

Similarity measures are not involved.

From ???@??? Mon Aug 17 16:10:58 1998

Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id PAA26731;
Mon, 17 Aug 1998 15:29:35 -0500 (CDT)

X-Sender: phayes@mail.coginst.uwf.edu

Message-Id: <v04003a0eb1fe4e6e0f79@[143.88.7.118]>

In-Reply-To: <199808161810.LAA23914@vapor.stanford.edu>

References: <199808161651.MAA18092@west> (sowa@west.poly.edu)

Mime-Version: 1.0

Date: Mon, 17 Aug 1998 15:25:40 -0600

To: jmc@cs.Stanford.EDU
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: FW: Re: Rigid, Nonrigid, Antirigid
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@nuts.coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.Stanford.EDU>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 1558
Status:

...
>> Adam wrote
>>
>> >I still have a question regarding natural kinds. What sort of logic is
>> >required to provide a coherent definition of 'natural kind'. Is the
>> >standard definition something like: C is a natural kind if there is no
>> >sentence S, free in X such that C(X) is true iff S?
>>
JS:
>> The term "natural kind" is commonly used in certain philosophical
>> discussions,
>> but no one has ever given a coherent definition. I suggest that we put it
>> in the class of "noncore" terms that are used in discussions of other
>> positions, but that we should avoid using it in the core vocabulary.
>
>I don't agree with either Adam Farquhar's proposed definition or with
>or John Sowa's proposal to omit it from the core vocabulary.

I tend to agree with JMC that we need the concept of 'natural kind', but
also with JS that there may not be a useful *definition* of it. OK, lets
use it and not have a definition. Most concepts don't have definitions.
There still might be something useful to say about it.

However, that said, let me ask JMC to put up:

JMC:
>2. That there are natural kinds in the world is an important fact, and
>the term should not be avoided.
>

If it is important, what can be inferred from knowing that something is a
natural kind term?

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Thu Aug 20 11:17:31 1998
Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])
 by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA08521
 for <phayes@coginst.uwf.edu>; Thu, 20 Aug 1998 11:10:44 -0500 (CDT)
Received: from scratchy (scratchy [207.207.8.118])
 by catbert.cyc.com (8.8.8/8.8.8) with SMTP id LAA18347;
 Thu, 20 Aug 1998 11:06:37 -0500 (CDT)
Message-Id: <3.0.32.19980820110638.008edc18@catbert.cyc.com>
X-Sender: fritz@catbert.cyc.com
X-Mailer: Windows Eudora Pro Version 3.0 (32)
Date: Thu, 20 Aug 1998 11:06:40 -0500
To: Pat Hayes <phayes@coginst.uwf.edu>
From: Fritz Lehmann <fritz@cyc.com>
Subject: Re: United Nations, was: Natural kinds
Cc: sowa@west.poly.edu (John F. Sowa), Piek.Vossen@let.uva.nl,
 chezewiz@erols.com, cmenzel@tamu.edu, doug@csi.uottawa.ca,
 e6nl001@coe.coppin.umd.edu, fritz@cyc.com, guarino@ladseb.pd.cnr.it,
 p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu, skydog@pacbell.net,
 sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset="us-ascii"
Content-Length: 571
Status:

At 10:54 AM 8/20/98 -0600, Pat Hayes wrote:

>[...]

>I think everyone should be assigned a unique integer at birth by the UN and
>have it firmly tattooed on their forehead. That would save a great deal of
>trouble, and probably be less painful than circumcision.

>Pat

Keep an eye out for the person with the the integer DCLXVI.

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

====

From ???@??? Thu Aug 20 11:17:31 1998
Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA08521
for <phayes@coginst.uwf.edu>; Thu, 20 Aug 1998 11:10:44 -0500 (CDT)
Received: from scratchy (scratchy [207.207.8.118])
by catbert.cyc.com (8.8.8/8.8.8) with SMTP id LAA18347;
Thu, 20 Aug 1998 11:06:37 -0500 (CDT)
Message-Id: <3.0.32.19980820110638.008edc18@catbert.cyc.com>
X-Sender: fritz@catbert.cyc.com
X-Mailer: Windows Eudora Pro Version 3.0 (32)
Date: Thu, 20 Aug 1998 11:06:40 -0500
To: Pat Hayes <phayes@coginst.uwf.edu>
From: Fritz Lehmann <fritz@cyc.com>
Subject: Re: United Nations, was: Natural kinds
Cc: sowa@west.poly.edu (John F. Sowa), Piek.Vossen@let.uva.nl,
chezewiz@erols.com, cmenzel@tamu.edu, doug@csi.uottawa.ca,
e6nl001@coe.coppin.umd.edu, fritz@cyc.com, guarino@ladseb.pd.cnr.it,
p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu, skydog@pacbell.net,
sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset="us-ascii"
Content-Length: 571
Status:

At 10:54 AM 8/20/98 -0600, Pat Hayes wrote:

>[...]

>I think everyone should be assigned a unique integer at birth by the UN and

>have it firmly tattooed on their forehead. That would save a great deal of

>trouble, and probably be less painful than circumcision.

>Pat

Keep an eye out for the person with the the integer DCLXVI.

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

====

From ???@??? Thu Aug 20 12:48:28 1998
Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA10310

for <phayes@coginst.uwf.edu>; Thu, 20 Aug 1998 12:05:46 -0500 (CDT)
Received: from scratchy (scratchy [207.207.8.118])
by catbert.cyc.com (8.8.8/8.8.8) with SMTP id LAA19167;
Thu, 20 Aug 1998 11:57:47 -0500 (CDT)

Message-Id: <3.0.32.19980820115746.00a12940@catbert.cyc.com>

X-Sender: fritz@catbert.cyc.com

X-Mailer: Windows Eudora Pro Version 3.0 (32)

Date: Thu, 20 Aug 1998 11:58:10 -0500

To: Robert Spillers <skydog@pacbell.net>

From: Fritz Lehmann <fritz@cyc.com>

Subject: Re: Bob Spillers Objections to Path Length

Cc: John Sowa <sowa@west.poly.edu>, Bill Andersen <chezewiz@erols.com>,
Chris Menzel <cmenzel@turing.stanford.edu>,
Doug Skuce <doug@site.uottawa.ca>, Fritz Lehmann <fritz@cyc.com>,
John McCarthy <jmc@cs.stanford.edu>,
Nancy Lawler <E6NL001@coe.coppin.umd.edu>,
Nicola Guarino <guarino@ladseb.pd.cnr.it>,
Pat Hayes <phayes@coginst.uwf.edu>,
Peter Simons <p.m.simons@leeds.ac.uk>,
Adam Farquhar <axf@KSL.Stanford.EDU>,
Robert Spillers <skydog@pacbell.net>

Mime-Version: 1.0

Content-Type: text/enriched

Content-Length: 2681

Status:

<x-rich>At 10:38 PM 8/16/98 -0700, Robert Spillers wrote:

>>>>

<excerpt>

Does a measure of similarity imply a measure of distance?

</excerpt><<<<<<<

I guess yes, since distance is supposed to mean dissimilarity in these systems.

>>>>

<excerpt> Since ontologies can be of arbitrary size and uneven in their granularity, by distance I mean something more than simply the shortest path length. Are these measures related?

</excerpt><<<<<<<<

Very good point. As Bob has pointed out to me, the different amount of work in different subject areas of the knowledge base means that a short path over in the "Animals" department is about the same "dissimilarity" as a long path over in "Furniture", if "Furniture" has been ontologized in much more detail than "Animals" has. The fact that I can freely introduce intermediate nodes in a knowledge base tells me that there is something wrong with these path-length-based "semantic distance" measures (as I've been repeating for years). Suppose "Armchair" has a certain semantic distance from "Lemon" based on some path that goes through furniture, physical object, fruit, etc. Then I add two new classes "Furniture With Four Legs" above "Armchair" and "Yellow Citrus Fruit" above "Lemon". Can we then rightly say that the notions of armchair and lemon are now more distantly related than before? I could put in a thousand more intermediate concepts, including some very artificial and uninteresting ones.

I think there CAN be a genuine measure of distance in such a structure, but it has to recognize some sorts of genuine levels or distinctions that are preserved "up to homeomorphism" between different knowledge bases --- that is, any number of intermediate nodes will not affect the real "semantic distance". Most proposals for path-length-based semantic distance fail this test.

The measures are probably "related" in the weak sense that a genuine increase in semantic distance would never _decrease_ the path-length-based distance.

>>>>

<excerpt> Would two entities similarity or distance be different if the size or granularity of the ontology were different?

</excerpt><<<<<<<<

They shouldn't be, but in most proposals I've seen, they are.

<excerpt>

</excerpt><<<<<<<<

Bob

>>>>

<excerpt>

</excerpt>

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA

email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

=====

====

</x-rich>

From ???@??? Thu Aug 20 12:48:28 1998

Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA10310

for <phayes@coginst.uwf.edu>; Thu, 20 Aug 1998 12:05:46 -0500 (CDT)

Received: from scratchy (scratchy [207.207.8.118])

by catbert.cyc.com (8.8.8/8.8.8) with SMTP id LAA19167;

Thu, 20 Aug 1998 11:57:47 -0500 (CDT)

Message-Id: <3.0.32.19980820115746.00a12940@catbert.cyc.com>

X-Sender: fritz@catbert.cyc.com

X-Mailer: Windows Eudora Pro Version 3.0 (32)

Date: Thu, 20 Aug 1998 11:58:10 -0500

To: Robert Spillers <skydog@pacbell.net>

From: Fritz Lehmann <fritz@cyc.com>

Subject: Re: Bob Spillers Objections to Path Length

Cc: John Sowa <sowa@west.poly.edu>, Bill Andersen <chezewiz@erols.com>,

Chris Menzel <cmenzel@turing.stanford.edu>,

Doug Skuce <doug@site.uottawa.ca>, Fritz Lehmann <fritz@cyc.com>,

John McCarthy <jmc@cs.stanford.edu>,

Nancy Lawler <E6NL001@coe.coppin.umd.edu>,

Nicola Guarino <guarino@ladseb.pd.cnr.it>,

Pat Hayes <phayes@coginst.uwf.edu>,
Peter Simons <p.m.simons@leeds.ac.uk>,
Adam Farquhar <axf@KSL.Stanford.EDU>,
Robert Spillers <skydog@pacbell.net>

Mime-Version: 1.0
Content-Type: text/enriched
Content-Length: 2681
Status:

<x-rich>At 10:38 PM 8/16/98 -0700, Robert Spillers wrote:

>>>>

<excerpt>

Does a measure of similarity imply a measure of distance?

</excerpt><<<<<<<<

I guess yes, since distance is supposed to mean dissimilarity in these systems.

>>>>

<excerpt> Since ontologies can be of arbitrary size and uneven in their granularity, by distance I mean something more than simply the shortest path length. Are these measures related?

</excerpt><<<<<<<<

Very good point. As Bob has pointed out to me, the different amount of work in different subject areas of the knowledge base means that a short path over in the "Animals" department is about the same "dissimilarity" as a long path over in "Furniture", if "Furniture" has been ontologized in much more detail than "Animals" has. The fact that I can freely introduce intermediate nodes in a knowledge base tells me that there is something wrong with these path-length-based "semantic distance" measures (as I've been repeating for years). Suppose "Armchair" has a certain semantic distance from "Lemon" based on some path that goes through furniture, physical object, fruit, etc. Then I add two new classes "Furniture With Four Legs" above "Armchair" and "Yellow Citrus Fruit" above "Lemon". Can we then rightly say that the notions of armchair and lemon are now more distantly related than before? I could put in a

thousand more intermediate concepts, including some very artificial and uninteresting ones.

I think there CAN be a genuine measure of distance in such a structure, but it has to recognize some sorts of genuine levels or distinctions that are preserved "up to homeomorphism" between different knowledge bases --- that is, any number of intermediate nodes will not affect the real "semantic distance". Most proposals for path-length-based semantic distance fail this test.

The measures are probably "related" in the weak sense that a genuine increase in semantic distance would never _decrease_ the path-length-based distance.

>>>>

<excerpt> Would two entities similarity or distance be different if the size or granularity of the ontology were different?

</excerpt><<<<<<<<

They shouldn't be, but in most proposals I've seen, they are.

<excerpt>

</excerpt><<<<<<<<

Bob

>>>>

<excerpt>

</excerpt>

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA

email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

=====
=====

</x-rich>

From ???@??? Thu Aug 20 13:28:43 1998

Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id NAA14697;
Thu, 20 Aug 1998 13:12:35 -0500 (CDT)

X-Sender: phayes@mail.coginst.uwf.edu

Message-Id: <v04003a0db202205f02c5@[143.88.7.118]>

In-Reply-To: <3.0.32.19980820115746.00a12940@catbert.cyc.com>

Mime-Version: 1.0

Date: Thu, 20 Aug 1998 13:08:45 -0600

To: Fritz Lehmann <fritz@cyc.com>

From: Pat Hayes <phayes@coginst.uwf.edu>

Subject: Re: Bob Spillers Objections to Path Length

Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net

Content-Type: text/plain; charset="us-ascii"

Content-Length: 2026

Status:

Fritz-

I hadnt realized that you guys are talking about path distance in the ontology itself. That's obviously a disaster.

Theres a strange historical mixture of ideas involved here. What we now call semantic networks (and which everyone now knows are just a graphical notation for logic :-)) have two distinct intellectual ancestors: tree-like heirarchical classification systems, and 'associative' concept networks, which were much prized by psychologists when I was a grad student. Associative networks were supposed to encode 'association', ie two concepts (words, maybe) were linked if people tended to think of one when you said the other: cats and dogs, fishes and oceans, meat and potatoes, etc. . In these networks, path distance might be a plausible measure of how far "apart" two concepts/words are; but in a classification tree, path distance

makes no sense at all, as your examples help to illustrate.

Heirarchies arent compatible with distance measures; they trivialize them. To take a notorious common-sense example, consider 'meeting distance', ie path-distance in a graph on humans with links whenever two of them have actually met. One would think this would be quite large, but in fact it is pretty small because social heirachies mean that there are a few people who have met an incredible number of other people. Almost everyone in England is only three or four steps from the Queen, for example (in the US, use recent presidents and consider that they are all only one or two meets from the others.) All one needs is a small percentage of richly connected nodes to provide short-cuts, and path distance becomes meaningless. This wasnt a problem for the original word-association work because there are no shortcuts there.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Thu Aug 20 13:28:43 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id NAA14697;
Thu, 20 Aug 1998 13:12:35 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a0db202205f02c5@[143.88.7.118]>
In-Reply-To: <3.0.32.19980820115746.00a12940@catbert.cyc.com>
Mime-Version: 1.0
Date: Thu, 20 Aug 1998 13:08:45 -0600
To: Fritz Lehmann <fritz@cyc.com>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Bob Spillers Objections to Path Length
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2026

Status:

Fritz-

I hadn't realized that you guys are talking about path distance in the ontology itself. That's obviously a disaster.

There's a strange historical mixture of ideas involved here. What we now call semantic networks (and which everyone now knows are just a graphical notation for logic :-)) have two distinct intellectual ancestors: tree-like hierarchical classification systems, and 'associative' concept networks, which were much prized by psychologists when I was a grad student. Associative networks were supposed to encode 'association', ie two concepts (words, maybe) were linked if people tended to think of one when you said the other: cats and dogs, fishes and oceans, meat and potatoes, etc. . In these networks, path distance might be a plausible measure of how far "apart" two concepts/words are; but in a classification tree, path distance makes no sense at all, as your examples help to illustrate.

Hierarchies aren't compatible with distance measures; they trivialize them. To take a notorious common-sense example, consider 'meeting distance', ie path-distance in a graph on humans with links whenever two of them have actually met. One would think this would be quite large, but in fact it is pretty small because social hierarchies mean that there are a few people who have met an incredible number of other people. Almost everyone in England is only three or four steps from the Queen, for example (in the US, use recent presidents and consider that they are all only one or two meets from the others.) All one needs is a small percentage of richly connected nodes to provide short-cuts, and path distance becomes meaningless. This wasn't a problem for the original word-association work because there are no shortcuts there.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Tue Aug 25 09:55:50 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id AAA18863

for <phayes@coginst.uwf.edu>; Tue, 25 Aug 1998 00:16:07 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id BAA18082;
Tue, 25 Aug 1998 01:12:04 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
id BAA06884; Tue, 25 Aug 1998 01:07:57 -0400
Date: Tue, 25 Aug 1998 01:07:57 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199808250507.BAA06884@west>
To: chezewiz@erols.com, fritz@cyc.com
Subject: Re: Heidelberg: continuants etc.
Cc: cmenzel@tamu.edu, doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu,
guarino@ladseb.pd.cnr.it, hovy@isi.edu, jmc@cs.stanford.edu,
p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu, skydog@pacbell.net,
sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 549
Status:

I want to put a bit more emphasis on Bill's points:

>It may be the case that the upper level is underutilized for a
>number of reasons:

- > 1) It isn't axiomatized richly enough
- > 2) It isn't axiomatized properly
- > 3) It isn't tied in with the "lower level" axioms properly

These are definitely my impressions of the Cyc upper levels and their relationships to the lower levels. I must admit that I haven't looked at Cyc in sufficient detail to be certain, but I'm glad that someone who has worked with it has voiced similar concerns.

John

From ???@??? Tue Aug 25 14:06:27 1998
Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA28433
for <phayes@coginst.uwf.edu>; Tue, 25 Aug 1998 12:41:49 -0500 (CDT)
Received: from scratchy (scratchy [207.207.8.118])
by catbert.cyc.com (8.8.8/8.8.8) with SMTP id MAA12987;
Tue, 25 Aug 1998 12:30:27 -0500 (CDT)
Message-Id: <3.0.32.19980825123019.00a9ea70@catbert.cyc.com>
X-Sender: fritz@catbert.cyc.com
X-Mailer: Windows Eudora Pro Version 3.0 (32)
Date: Tue, 25 Aug 1998 12:30:28 -0500

To: sowa@west.poly.edu (John F. Sowa)
From: Fritz Lehmann <fritz@cyc.com>
Subject: Re: Heidelberg: continuants etc.
Cc: chezewiz@erols.com, fritz@cyc.com, cmenzel@tamu.edu, doug@csi.uottawa.ca,
e6nl001@coe.coppin.umd.edu, guarino@ladseb.pd.cnr.it, hovy@isi.edu,
jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
skydog@pacbell.net, sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset="us-ascii"
Content-Length: 783
Status:

Let me emphasize that I was describing Doug Lenat's view that the upper levels aren't as important for reasoning as the intermediate levels --- not my own view. I think that although there may be only few axioms that characterize highest level concepts, they are important and are of course very widely inherited.

For disambiguating natural language, my standard example is "The legislature repealed the statue." That is known to be an error ("statue" for "statute"), due to a very high-level distinction between concrete and abstract.

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

=====
=====

From ???@??? Wed Aug 26 09:55:41 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id SAA17042;
Tue, 25 Aug 1998 18:35:36 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a0fb20906bd4d0d@[143.88.7.118]>
In-Reply-To: <3.0.32.19980825150113.00a9e858@catbert.cyc.com>
Mime-Version: 1.0
Date: Tue, 25 Aug 1998 18:31:52 -0600
To: Fritz Lehmann <fritz@cyc.com>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Heidelberg: continuants etc.
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,

Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 863
Status:

Fritz, what is your point? Why should we give a damn about Yahoo and yer average bookstore? Your original comment was about philosophy, not web browsers and bookstores; and in reply to that, Chris is exactly right. Ive been in philosophy departments in Rochester, Illinois and (indirectly) Stanford, and moved among philosophers (as it were) for about a decade, and Ive never come across the nonsense that you seem to be so worried about, even at the APA. (Even Heideggerian scholars often have something coherent and sensible to say, if one listens to them carefully enough.)

Pat Hayes

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Aug 26 09:55:42 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id SAA18347;
Tue, 25 Aug 1998 18:41:08 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a10b2090840a80f@[143.88.7.118]>
In-Reply-To: <3.0.32.19980825163715.00a9e858@catbert.cyc.com>
Mime-Version: 1.0
Date: Tue, 25 Aug 1998 18:37:22 -0600
To: Fritz Lehmann <fritz@cyc.com>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Heidelberg: continuants etc.
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 855

Status:

>Yes, that's why I included "gavagai" in my silly list --- Quine's example.
> I wouldn't hire Quine as my chauffeur. He might decline to stop at a sign
>saying "BRIDGE OUT", based on his "gavagai" argument.
>

No, you havnt got the argument right. If it said GAVAGAI, Quine would expect to see rabbits. The point is not the correspondence of language to the world, but the claim (controversial) that the way that world is conceptualised cannot be determined from language. (Here be rabbits, or Rabbitness is locally manifested.) Its controversial, but not entirely silly.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Aug 26 09:55:46 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id DAA14735
for <phayes@coginst.uwf.edu>; Wed, 26 Aug 1998 03:21:20 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id EAA21405;
Wed, 26 Aug 1998 04:17:38 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
id EAA13515; Wed, 26 Aug 1998 04:13:51 -0400
Date: Wed, 26 Aug 1998 04:13:51 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199808260813.EAA13515@west>
To: fritz@cyc.com, jmc@cs.Stanford.EDU
Subject: Re: Heidelberg: continuants etc.
Cc: chezewiz@erols.com, cmenzel@csl.stanford.edu, cmenzel@tamu.edu,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu,
guarino@ladseb.pd.cnr.it, hovy@isi.edu, p.m.simons@leeds.ac.uk,
phayes@coginst.uwf.edu, skydog@pacbell.net, sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii

Content-Length: 1294

Status:

>Should you have an opportunity to hire Quine in any capacity, don't
>pass it up. If I understood what he said, I think he's wrong about
>"gavagai", but he's right about a lot of other things.

I agree with that point. On almost any issue, Quine usually has a quotable nugget that captures the essence of the matter in two sentences.

Even his "gavagai" point is hard to refute. Quine's example illustrates the point that even when you translate a sentence from A to B in such a way that you capture the truth conditions, you might still miss the connotations. As a real-life example, I like to compare the English word "stump", which is usually translated into French as "souche".

Although that translation works for many cases, "stump" has the connotation of something that has been cut off and is now dead. But "souche" has the connotation of the nourishing source. I once saw a French article in linguistics that was talking about the "souche" of the Indo-European languages. But if you translated that as "the stump of the Indo-European languages", you would get the wrong image, in fact, almost the same kind of gap in meaning as "rabbit" and "undetached rabbit parts."

But if you are doing cryptography, the truth conditions are usually sufficient to decide what to bomb.

John

From ???@??? Tue Sep 01 10:09:25 1998

Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id TAA03377

for <phayes@coginst.uwf.edu>; Mon, 31 Aug 1998 19:16:22 -0500 (CDT)

Received: from west (west.poly.edu [128.238.20.21])

by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id UAA03776;

Mon, 31 Aug 1998 20:12:41 -0400 (EDT)

Received: by west (SMI-8.6/SMI-SVR4)

id UAA14275; Mon, 31 Aug 1998 20:08:25 -0400

Date: Mon, 31 Aug 1998 20:08:25 -0400

From: sowa@west.poly.edu (John F. Sowa)

Message-Id: <199809010008.UAA14275@west>

To: fritz@cyc.com, phayes@coginst.uwf.edu

Subject: Machine translation

Cc: chezewiz@erols.com, cmenzel@tamu.edu, doug@csi.uottawa.ca,

e6nl001@coe.coppin.umd.edu, guarino@ladseb.pd.cnr.it,

jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk, skydog@pacbell.net,

sowa@west.poly.edu

Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 1187
Status:

Our recent discussion of Quine's "gavagai" example reminded me of the following story by the linguist Guy Cardan, who had been involved in an MT project during the Vietnam-war era. It indicates how elusive the "meaning" of a text can be and how difficult it can be to tell whether the essential points are getting across.

John

The US Air Force had two machine translation projects going for English to Vietnamese manual translating. Here are the results:

1. Bilingual experts judged output A and output B to be approximately equally good.
2. The Vietnamese technicians said they preferred the original English manuals to either A or B translations.
3. However, the Vietnamese made twice as many errors with the English manuals and worked at half the speed as American technicians.
4. With translation A, the Vietnamese worked just as well as the American technicians.
5. With translation B, the Vietnamese worked worse than they did with the English manuals.

Apparently, there are human factors in translations that neither "experts" nor "users" can reliably predict or detect.

From ???@??? Thu Aug 20 11:17:29 1998
Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id KAA07319
for <phayes@coginst.uwf.edu>; Thu, 20 Aug 1998 10:47:43 -0500 (CDT)
Received: from scratchy (scratchy [207.207.8.118])
by catbert.cyc.com (8.8.8/8.8.8) with SMTP id KAA17910;
Thu, 20 Aug 1998 10:42:33 -0500 (CDT)
Message-Id: <3.0.32.19980820104233.009a0e20@catbert.cyc.com>
X-Sender: fritz@catbert.cyc.com
X-Mailer: Windows Eudora Pro Version 3.0 (32)
Date: Thu, 20 Aug 1998 10:42:36 -0500

To: jmc@cs.Stanford.EDU

From: Fritz Lehmann <fritz@cyc.com>

Subject: Re: Natural kinds

Cc: sowa@west.poly.edu, phayes@coginst.uwf.edu, Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@tamu.edu, doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu, fritz@cyc.com, guarino@ladseb.pd.cnr.it, p.m.simons@leeds.ac.uk, skydog@pacbell.net, sowa@west.poly.edu

Mime-Version: 1.0

Content-Type: text/plain; charset="us-ascii"

Content-Length: 2366

Status:

At 07:05 PM 8/19/98 -0700, John McCarthy wrote:

>Hmm. I see that what I took to be the notion of natural kind does not
>correspond to the literature John Sowa cites. My excuse is that I
>once told Putnam what I had in mind, and he said it was the same
>notion. The essence of my notion is that objects belonging to the
>natural kind may be expected to have presently unknown common
>properties, e.g. as lemons do and sea mammals do (distinguishing the
>latter from fish). I don't need to say what the similarities are, and
>they may be discovered (if at all) much later. If one suspects that
>most of some objects distinguished by presently noticed properties
>form a natural kind, the reaction is to study the objects further.
>The study may result in excluding some presently included objects and
>including some superficially different objects.

>

>I think mine is the right notion, but I'll lay off for now.

I too think McCarthy's is the right notion about this, and I further think that there has been an unvoiced transition from the subject of natural kinds to the subject of prototype theory. They need have nothing to do with each other. It is in prototype theory that new entities are thrown in with the "nearest" prototype, according to some (often inherently bogus) distance measure or vector of weighted distances. Whether the selected prototypes are natural kinds doesn't matter; the prototypes could be as artificial as you like.

I've been using "striped cup" as my standard NON-natural kind --- something not worth representing with a Cyc concept. It is "analytic" in Kant-speak: all we know is that it is striped and that it is a cup, and nothing new or more about it. That's the difference between white cat and black cat: the black cat has the additional non-analytic feature of being unlucky.

Natural kinds to me are exactly as McCarthy says, and they exist because of regularities in our world about classes of entities of which we have only imperfect knowledge. New things remain to be discovered about lemons that

will, lo (because of the aforesaid regularity), be common to all lemons.

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

=====
=====

From ???@??? Thu Aug 20 11:17:30 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id KAA07348;
Thu, 20 Aug 1998 10:58:46 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a05b20201becd56@[143.88.7.118]>
In-Reply-To: <199808201214.IAA09615@west>
Mime-Version: 1.0
Date: Thu, 20 Aug 1998 10:54:53 -0600
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Natural kinds
Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@tamu.edu,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu, fritz@cyc.com,
guarino@ladseb.pd.cnr.it, p.m.simons@leeds.ac.uk,
phayes@coginst.uwf.edu, skydog@pacbell.net, sowa@west.poly.edu
Content-Type: text/plain; charset="us-ascii"
Content-Length: 1494
Status:

>
>PS: As someone named John, I am very well aware of the context-dependent
>nature of proper names and the need for some supplementary rigid identifiers
>like social-security numbers or URL's. "McCarthy" helps to distinguish it a
>bit, but not much. "Sowa" narrows it down more, but there is a John R. Sowa
>who works in Organic Chemistry, and in the Science Citation Index we tend
>to get credited with each other's works.

Its not just the Johns. I was amazed to discover how many Patrick J. Hayes's there are, including a physicist in Texas somewhere. One is a Justice of the Peace in an English county, another an unsuccessful Republican candidate in New Jersey. If one retreats to "Pat Hayes", AltaVista gets several thousand hits. There were three Pat Hayes' at the University of Illinois while I was there. Once I got accidentally enrolled in someone else's health insurance scheme; another time I was sent someone other PJH's employment contract to sign. (His salary was lower, so I didn't.)

I think everyone should be assigned a unique integer at birth by the UN and have it firmly tattooed on their forehead. That would save a great deal of trouble, and probably be less painful than circumcision.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Thu Aug 20 11:33:34 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
 by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA08761;
 Thu, 20 Aug 1998 11:20:43 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a06b2020568aa1d@[143.88.7.118]>
In-Reply-To: <199808201214.IAA09615@west>
Mime-Version: 1.0
Date: Thu, 20 Aug 1998 11:16:52 -0600
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Natural kinds
Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@tamu.edu,
 doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu, fritz@cyc.com,
 guarino@ladseb.pd.cnr.it, p.m.simons@leeds.ac.uk,
 phayes@coginst.uwf.edu, skydog@pacbell.net, sowa@west.poly.edu
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2653
Status:

John, (er...Dr. Sowa?) :

More seriously, on the matter of similarity measures. The account you give only has bite when you specify what the measure domain is like. Are measures placed in a metric space of some kind? (Continuous? Obeys triangle inequality? *Almost* obeys triangle inequality? Etc.) If the set of measures can be anything, then the account is trivial: just declare things to be their own measure (from the prototype), and use the identity mapping. (A Herbrand measure, one might call it.)

Also, this 'distance-from-prototype' account has been criticised by

psychologists on the grounds that it doesn't fit the empirical facts (except in some very special cases, such as color judgements.) Fodor has written a number of things on this, and there have been surveys in Psych Review (dont have the refs handy, sorry).

One variation that has been suggested is the 'family measure' idea, where what makes things into a kind is that they share some significant number of a set of attributes, but without any presumption that this sharing can be mapped into a metric distance space in any significant way (since there is no ordering on attributes.) This also has been criticised, and alternatives include allowing *negative* 'sharing', ie concepts are related if one could have an attribute that the other does, but in fact doesnt have it. All logically adept folk will see that this is now essentially vacuous, since it allows arbitrary propositional connections between concepts (just use disjunctive normal form) to be the reason for them being in the same 'kind'.

I repeat my (genuine, not rhetorical) question to all natural-kindlers: what *follows* from knowing that two things are of the same 'natural kind' ? If you tell me that oranges and lemons are the same natural kind, what can I then conclude that I couldnt conclude before? (Or maybe, can't conclude that I previously could?) For example, how about this: there is then a default presumption that oranges have all the same properties that lemons have (of a certain kind, ie the 'biological' properties...or maybe that the biotropes are..what?...similar??) Or how about this: that it is worth (a heuristic meta-judgement) creating or finding a category of which these shall both be instances, and attaching any shared property, by meta-default, to this category rather than to them both separately?

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Fri Aug 21 13:05:17 1998

Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id GAA03676

for <phayes@coginst.uwf.edu>; Fri, 21 Aug 1998 06:18:09 -0500 (CDT)

Received: from west (west.poly.edu [128.238.20.21])

by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id HAA09409;

Fri, 21 Aug 1998 07:14:31 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
id HAA15170; Fri, 21 Aug 1998 07:09:01 -0400
Date: Fri, 21 Aug 1998 07:09:01 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199808211109.HAA15170@west>
To: phayes@coginst.uwf.edu
Subject: Re: Natural kinds
Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@tamu.edu,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu, fritz@cyc.com,
guarino@ladseb.pd.cnr.it, p.m.simons@leeds.ac.uk, skydog@pacbell.net,
sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 4796
Status:

Fritz writes

>I too think McCarthy's is the right notion about this, and I further think
>that there has been an unvoiced transition from the subject of natural
>kinds to the subject of prototype theory.

It wasn't unvoiced. I was very explicitly trying to voice Quine's
position that there is a natural progression in the history of science:

intuitive similarity -> prototype + similarity measure -> theory

This is all discussed quite nicely in his 20 page article on natural kinds.

> It is in prototype theory that new entities are thrown in
>with the "nearest" prototype, according to some (often inherently bogus)
>distance measure or vector of weighted distances. Whether the selected
>prototypes are natural kinds doesn't matter; the prototypes could be as
>artificial as you like.

Yes, that is the issue that the biologists have been working on for
two thousand years and the neural network people have been working on
for two dozen years. The biologists have developed guidelines that are
reasonably good for their domain. The neural network people vary from
naive pot-pourri types to somewhat more sophisticated versions, but they
tend to get hung up on the technology of how you wire up the networks.
The statisticians have been around long enough that they appreciate
the need for models (approaching the theory end of the scale).

>Natural kinds to me are exactly as McCarthy says, and they exist because of

>regularities in our world about classes of entities of which we have only
>imperfect knowledge. New things remain to be discovered about lemons that
>will, lo (because of the aforesaid regularity), be common to all lemons.

Yes, and that is the point of Quine's article as well. When the new things are discovered, we approach the theory end of the scale. But that is the point of my criticism of natural kinds: we can't give reasonable guidelines to our ontology consumers/developers that tell them to look out for "undiscovered things". They want a cookbook, not a research proposal.

Pat writes

>More seriously, on the matter of similarity measures. The account you give
>only has bite when you specify what the measure domain is like....

I deliberately left the measure unspecified (as Quine did in his article). As he said, it is extremely domain dependent. The various examples you mentioned have been tried, and as I said in an earlier note, the matter of finding a good measure is still an unresolved research issue, and I believe that it always will be.

>I repeat my (genuine, not rhetorical) question to all natural-kinders: what
>*follows* from knowing that two things are of the same 'natural kind' ?

I have exactly the same questions, and I have never seen any answer to them in the natural-kind literature. That is why I believe that the term "natural kind" should NOT be placed in the core vocabulary because it doesn't give the ontology user/consumer/developer any guidelines about what to do.

>I hadn't realized that you guys are talking about path distance in the
>ontology itself. That's obviously a disaster.

I most definitely did not mean path length in my discussion. I meant something closer to what the neural network guys use. But as I said above, I only think of that as a temporary expedient -- or hack.

The definitions I put in Appendix B are based on our two years of ontology workshops in 1996 and 1997. By placing definitions in the appendix, I am trying to state the positions as clearly as I can, but I am not necessarily advocating any one of them.

Following are the kinds of ontologies:

1. Terminological ontology, as in WordNet with a hierarchy and minimal

definitions and almost no axioms (other than the partial ordering).

2. Axiomatized ontology, as in Cyc with lots of definitions and axioms.
3. Prototype-based ontology, as in the biological domain before DNA and in any domain where neural networks are being applied.
4. Mixed ontology, which uses some combination of the above.

To relate this to Quine's position, the terminological ontology based on intuitive similarity is the naive starting point. The prototype ontology with a domain-dependent similarity measure is the next step in the attempt to systematize any field (and for many fields, it is difficult to go any further at the present stage of knowledge). And Quine's theory-based approach corresponds to the axiomatized ontology.

John McCarthy's definition of natural kind is in the same spirit as Quine's discussion, but it doesn't tell the user what to do. I was trying to capture the various stages of development in the definitions. The ultimate goal is to have a fully axiomatized or theory-based ontology for every word in the English language. But I don't think that we're going to get there very soon.

John

From ???@??? Tue Sep 08 10:16:17 1998

Received: from LADSEB.LADSEB.PD.CNR.IT (ladseb.ladseb.pd.cnr.it [150.178.2.3])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id DAA25006
for <phayes@coginst.uwf.edu>; Mon, 7 Sep 1998 03:43:23 -0500 (CDT)

Received: from [150.178.2.93] by LADSEB.LADSEB.PD.CNR.IT with ESMTP;
Mon, 7 Sep 1998 10:37:44 +0200

"Wed, 12 Aug 1998 18:09:21 +0200." <v0310280bb1f76c72aca2@[150.178.2.93]>

X-Sender: guarino@ladseb.ladseb.pd.cnr.it

Message-Id: <v03102801b218a85b22b8@[150.178.2.93]>

In-Reply-To: <199808151506.KAA08897@correo.austin.apc.slb.com>

References: <199808121759.KAA03801@Turing.Stanford.EDU> <Your message of
"Wed, 12 Aug 1998 18:09:21 +0200." <v0310280bb1f76c72aca2@[150.178.2.93]>

Mime-Version: 1.0

Date: Mon, 7 Sep 1998 10:41:06 +0200

To: Adam Farquhar <farquhar@austin.apc.slb.com>,
Christopher Menzel <cmenzel@csl.stanford.edu>

From: Nicola Guarino <guarino@ladseb.pd.cnr.it>

Subject: Natural kinds

Cc: Christopher Menzel <cmenzel@csl.stanford.edu>, cmenzel@tamu.edu,
sowa@west.poly.edu, axf@KSL.Stanford.EDU, phayes@coginst.uwf.edu,
doug@csi.uottawa.ca, fritz@cyc.com, p.m.simons@leeds.ac.uk,
e6nl001@coe.coppin.umd.edu, chezewiz@erols.com, jmc@cs.stanford.edu,

Piek Vossen <Piek.Vossen@let.uva.nl>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2437
Status:

At 2:43 PM -0500 8/14/98, Adam Farquhar wrote:

>Chris, Nicola,

>

>I think that Chris's recent characterization and vocabulary for properties
>(strongly essential, non strongly essential, strongly non-essential and
>equivalent phrasing in terms of contingent) are clear and cover the ground
>fairly well. It seems best to avoid the possible confusion with the common
>(in certain circles!) meaning of rigid.

>I still have a question regarding natural kinds. What sort of logic is
>required to provide a coherent definition of 'natural kind'. Is the
>standard definition something like: C is a natural kind if there is no
>sentence S, free in X such that C(X) is true iff S?

>

Hi Adam, sorry for this late reply. The above definition holds for any primitive property. But a property can be primitive and still it may not correspond to a natural kind, in the usual sense of this expression. The notion of natural kind is much stronger. Natural kinds, according to my intuition, must be both primitive and "strongly essential" (using the terminology proposed by Chris). A further condition is that they each type must carry its own identity criterion for their instances. These three conditions actually characterize what I have called *types*. I don't know whether this notion of types actually coincides with the usual notion of natural kind, or maybe it is a bit more generic. It seems however good enough for many practical purposes.

Pat correctly asks what *follows* from explicitly introducing such a notion:

1. from the fact that they are strongly essential, it follows that they are not fluents: something cannot belong to a type now and to another type tomorrow.

2. assuming that classes with different identity criterias are disjoint, it follows that types form a tree, where all nodes at the same level are mutually disjoint.

These properties are especially relevant when compared with that of roles and attributions. As acknowledged by Piek Vossen at the workshop, these distinctions between kinds of property turn out to be quite useful also for linguistic purposes.

Cheers,

-- Nicola

Nicola Guarino
National Research Council phone: +39 49 8295751
LADSEB-CNR fax: +39 49 8295763
Corso Stati Uniti, 4 email: guarino@ladseb.pd.cnr.it
I-35127 Padova
Italy

Home page: ** updated 4/5/98 **
<http://www.ladseb.pd.cnr.it/infor/ontology/ontology.html>

From ???@??? Tue Sep 08 15:15:18 1998
Received: from LADSEB.LADSEB.PD.CNR.IT (ladseb.ladseb.pd.cnr.it [150.178.2.3])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id MAA14620
for <phayes@coginst.uwf.edu>; Tue, 8 Sep 1998 12:06:54 -0500 (CDT)
Received: from [150.178.2.93] by LADSEB.LADSEB.PD.CNR.IT with ESMTP;
Tue, 8 Sep 1998 19:00:07 +0200
<199808151506.KAA08897@correo.austin.apc.slb.com>
<199808121759.KAA03801@Turing.Stanford.EDU>
X-Sender: guarino@ladseb.ladseb.pd.cnr.it
Message-Id: <v0310280cb21b10199c2b@[150.178.2.93]>
In-Reply-To: <v04003a0ab21b16071152@[143.88.7.118]>
References: <v03102801b218a85b22b8@[150.178.2.93]>
<199808151506.KAA08897@correo.austin.apc.slb.com>
<199808121759.KAA03801@Turing.Stanford.EDU>
Mime-Version: 1.0
Date: Tue, 8 Sep 1998 19:03:40 +0200
To: Pat Hayes <phayes@coginst.uwf.edu>
From: Nicola Guarino <guarino@ladseb.pd.cnr.it>
Subject: Re: Natural kinds
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Transfer-Encoding: 8bit
X-MIME-Autoconverted: from quoted-printable to 8bit by nuts.coginst.uwf.edu id
MAA14620

Content-Type: text/plain
Content-Length: 2622
Status:

At 11:33 AM -0600 9/8/98, Pat Hayes wrote:

>>At 2:43 PM -0500 8/14/98, Adam Farquhar wrote:

>....

>>>I still have a question regarding natural kinds. What sort of logic is
>>>required to provide a coherent definition of 'natural kind'. Is the
>>>standard definition something like: C is a natural kind if there is no
>>>sentence S, free in X such that C(X) is true iff S?

>>>

>>

>>Hi Adam, sorry for this late reply. The above definition holds for any
>>primitive property.

>

>Wait a gosh-darned minute here.

Thanks for giving me an opportunity to improve my english...

>As given, Adam's definition is provably

>everywhere false, since for any kind that can be named in the language (by
>a predicate 'C', say), the sentence 'C(x)' is itself a suitable S. OK, so
>we want to say that S doesn't include 'C' (thus ruling out recursion, by the
>way.)

I agree, let's assume it doesn't include 'C'.

>But what this illustrates is that 'primitive' in this sense isn't a

>property of a predicate as such, but a relationship between a predicate and
>a theory (or maybe between a predicate and the signature of a theory?)

>Primitiveness - not having a definition - depends on the available
>vocabulary for writing definitions.

Of course, when we speak of a property of a predicate, it implicitly refers
to a certain available vocabulary and a certain axiomatization. In other
words, we speak of a property of a predicate *within a certain ontology*.

For instance, my definition of "rigidity" (or "strong essentiality") refer
to a property of a predicate as axiomatized within a certain ontology.

Loosely speaking, if we sometimes omit to specify the ontology it is
because we assume to implicitly refer to the same ontology, namely the one
we are trying to build...

So, within a certain ontology, a property can be primitive or not. Being
primitive within a certain ontology does not mean to be a natural kind
within that ontology.

When we say that "Horse" is a natural kind *per se* we are committing to an implicit common ontology (better, conceptualization...) constraining somehow the meaning of such predicate. When Adam says that natural kinds are primitives, he just says that they are not defined in this common ontology (or maybe they cannot be defined in any ontology...)

-- Nicola

Nicola Guarino
National Research Council phone: +39 49 8295751
LADSEB-CNR fax: +39 49 8295763
Corso Stati Uniti, 4 email: guarino@ladseb.pd.cnr.it
I-35127 Padova
Italy

Home page: ** updated 4/5/98 **
<http://www.ladseb.pd.cnr.it/infor/ontology/ontology.html>

From ???@??? Tue Sep 08 15:15:17 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA12838;
Tue, 8 Sep 1998 11:36:28 -0500 (CDT)
<199808121759.KAA03801@Turing.Stanford.EDU> <Your message of "Wed, 12 Aug
1998 18:09:21 +0200." <v0310280bb1f76c72aca2@[150.178.2.93]>
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a0ab21b16071152@[143.88.7.118]>
In-Reply-To: <v03102801b218a85b22b8@[150.178.2.93]>
References: <199808151506.KAA08897@correo.austin.apc.slb.com>
<199808121759.KAA03801@Turing.Stanford.EDU> <Your message of "Wed, 12 Aug
1998 18:09:21 +0200." <v0310280bb1f76c72aca2@[150.178.2.93]>
Mime-Version: 1.0
Date: Tue, 8 Sep 1998 11:33:00 -0600
To: Nicola Guarino <guarino@ladseb.pd.cnr.it>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Natural kinds
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,

JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 1330
Status:

>At 2:43 PM -0500 8/14/98, Adam Farquhar wrote:

....

>>I still have a question regarding natural kinds. What sort of logic is
>>required to provide a coherent definition of 'natural kind'. Is the
>>standard definition something like: C is a natural kind if there is no
>>sentence S, free in X such that C(X) is true iff S?

>>

>

>Hi Adam, sorry for this late reply. The above definition holds for any
>primitive property.

Wait a gosh-darned minute here. As given, Adam's definition is provably everywhere false, since for any kind that can be named in the language (by a predicate 'C', say), the sentence 'C(x)' is itself a suitable S. OK, so we want to say that S doesn't include 'C' (thus ruling out recursion, by the way.) But what this illustrates is that 'primitive' in this sense isn't a property of a predicate as such, but a relationship between a predicate and a theory (or maybe between a predicate and the signature of a theory?) Primitiveness - not having a definition - depends on the available vocabulary for writing definitions.

Pat Hayes

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Sep 16 11:14:00 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id KAA09128;
Wed, 16 Sep 1998 10:28:15 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a08b21b0f7084d1@[143.88.7.118]>
In-Reply-To: <199809080745.DAA25700@west>
Mime-Version: 1.0
Date: Tue, 15 Sep 1998 09:41:18 -0600

To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Necessity of "necessity"
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2291
Status:

Hi John and Nicola

This particular debate about the best way to approach "necessity" has been raging for many long years now, and we are unlikely to come to a philosophical consensus on it here. If we have to talk of necessity at all, then let us agree to differ on whether is best described using Other Worlds or by asking for the Necessary Rules. In this post-Kripkean age we should be able to translate back and forth between them, or at least acknowledge that such translation ought to be possible one day. (I think almost nobody agrees with David Lewis's cosmic super-realism, but I agree with J.S. that there is something very odd about a 'possible world' outside of Kripke semantics. I have trouble imagining what it would be to imagine *this* world, let alone an alternative one.)

Do we need necessity at all? I'm unconvinced by Nicola on this issue, but since all these wierd distinctions betwen different kinds of necessity seem to be a perfect example of ragged philosophy wandering around looking for a good home, I am willing to sit back and nod quietly while those who enjoy this arcane business amuse themselves, just as long as it doesnt get in the way of any actual substantive work. As an example of the latter, from Nicola's message:

[Nicola]

>

>2. assuming that classes with different identity criterias are disjoint, it
>follows that types form a tree, where all nodes at the same level are
>mutually disjoint.

>

Something may simultaneously belong to two different types which have different identity criteria across time. For example, a thunderstorm is both a continuant and an occurrent (instantaneously), and a river is both a piece of liquid (instantaneously) and a liquid object. Any proposed way of

building ontologies which rules out examples like these is simply unacceptable, as it will grossly interfere with practical reasoning. I'm not sure if Nicola's tree-criteria would forbid it, but if so, then point 2 is a reductio rather than a positive argument, in my view.

Pat Hayes

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Sep 16 13:45:06 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA16843
for <phayes@coginst.uwf.edu>; Wed, 16 Sep 1998 12:36:14 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id NAA22938;
Wed, 16 Sep 1998 13:31:30 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
id NAA24854; Wed, 16 Sep 1998 13:26:27 -0400
Date: Wed, 16 Sep 1998 13:26:27 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199809161726.NAA24854@west>
To: phayes@coginst.uwf.edu
Subject: Re: Necessity of "necessity"
Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@tamu.edu,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu, fritz@cyc.com,
guarino@ladseb.pd.cnr.it, jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk,
skydog@pacbell.net
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 569
Status:

Pat,

I agree with John McCarthy "Modality si! Modal logic, no!"

The modal auxiliaries like 'can' and 'must' are essential whenever you are trying to state a standard.

Short point: A formulation of modality in terms of laws is finer grained than a formulation in terms of possible worlds. Everything you can do with modal operators can be done by Dunning (a la Michael D). But there are things that you can state and reason about in terms of laws (or axioms) that cannot be stated with only the modal operators. John M. makes similar points in his paper.

John

From ???@??? Thu Sep 17 17:37:45 1998

Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id OAA23494

for <phayes@coginst.uwf.edu>; Thu, 17 Sep 1998 14:19:43 -0500 (CDT)

Received: from scratchy (scratchy [207.207.8.118])

by catbert.cyc.com (8.8.8/8.8.8) with SMTP id OAA23141;

Thu, 17 Sep 1998 14:15:10 -0500 (CDT)

Message-Id: <3.0.32.19980917141509.00f21eb8@catbert.cyc.com>

X-Sender: fritz@catbert.cyc.com

X-Mailer: Windows Eudora Pro Version 3.0 (32)

Date: Thu, 17 Sep 1998 14:15:21 -0500

To: Pat Hayes <phayes@coginst.uwf.edu>

From: Fritz Lehmann <fritz@cyc.com>

Subject: Models intermediate or not? Was: Re: tropes

Cc: sowa@west.poly.edu (John F. Sowa), Piek.Vossen@let.uva.nl,

chezewiz@erols.com, cmenzel@tamu.edu, doug@csi.uottawa.ca,

e6nl001@coe.coppin.umd.edu, fritz@cyc.com, guarino@ladseb.pd.cnr.it,

p.m.simons@leeds.ac.uk, skydog@pacbell.net

Mime-Version: 1.0

Content-Type: text/plain; charset="us-ascii"

Content-Length: 1244

Status:

At 12:25 PM 9/17/98 -0600, Pat Hayes wrote:

>[...]

>However, I disagee (along with Tarski, Quine and others.) Models may be

>built from such things as integers, but they can also be built from such

>things as tables, events or pieces of cheese. Model theory requires a

>domain which is a set, but that doesnt mean that models must be 'formal',

>and hence lacking that rich, loamy texture that real worlds have. A set can

>be a set of anything, including parts of the real world.

>[...]

>Pat

Dear Pat,

Do I understand your view correctly if I understand it to imply that a

model can never be wrong? A sentence can be false with respect to the model in both views. In Sowa's view a model could also fail to correctly model the real world. Your view, taking the elements of the model to be entities in the real world, seems to make the model incapable of error. Can the model even be capable of underspecification of the real world (even in a technically inaccessible way), in your approach? How?

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

=====
=====

From ???@??? Fri Sep 18 17:55:58 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA12669;
Fri, 18 Sep 1998 12:24:54 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a06b22851071726@[143.88.7.118]>
In-Reply-To: <3.0.32.19980917141509.00f21eb8@catbert.cyc.com>
Mime-Version: 1.0
Date: Fri, 18 Sep 1998 12:21:38 -0600
To: Fritz Lehmann <fritz@cyc.com>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Models intermediate or not? Was: Re: tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 3181
Status:

>At 12:25 PM 9/17/98 -0600, Pat Hayes wrote:

>>[...]

>>However, I disagee (along with Tarski, Quine and others.) Models may be
>>built from such things as integers, but they can also be built from such
>>things as tables, events or pieces of cheese. Model theory requires a
>>domain which is a set, but that doesnt mean that models must be 'formal',
>>and hence lacking that rich, loamy texture that real worlds have. A set can
>>be a set of anything, including parts of the real world.

>>[...]

>>Pat

>

>

>Dear Pat,

>

>Do I understand your view correctly if I understand it to imply that a
>model can never be wrong? A sentence can be false with respect to the
>model in both views.

Im not sure what you mean by a model being 'wrong'. Models, or interpretations, assign truthvalues to closed sentences; so a model (actually one ought to say an interpretation, since in model theory "model" means an interpretation which makes a sentence true, but lets continue using "model" to be synonymous with "interpretation") is, in a sense, a world. So of course it can't be 'wrong' about itself, although it can certainly make some sentences false, ie those sentences may be false in it.

My problem with John's account of 'models' is that it seems to need two different ways for a sentence to be false of the world: it might be false in a model, or that model may be inaccurate in some way. But I can only see one way in which a sentence can be false, which involves giving an interpretation of it in some world. Model theory talks of the relationship between sentences and interpretations, and thats all.

In Sowa's view a model could also fail to correctly
>model the real world. Your view, taking the elements of the model to be
>entities in the real world, seems to make the model incapable of error. Can
>the model even be capable of underspecification of the real world (even in
>a technically inaccessible way), in your approach? How?

Well, models *need* not be made of real-world stuff; my point is only that they *can* be. There are also, for example, Herbrand models, which are made entirely of symbolic expressions.

Im not sure what you mean by 'underspecification'. A sentence can surely underspecify the world. Heres an example: $(\text{forall } x)(P \ x)$. If we interpret 'P' to mean the property of having zero rest mass and the domain to be all leptons, then this sentence is true in the actual world. But if we take P to mean the property of having charge, then it isnt true. Two different interpretations of the same sentence, both in the real world, give different values. Thats why the sentence doesnt tell us much about - underspecifies - the world. But its the *sentence* that 'underspecifies', not a model. The first interpretation isnt some abstract thing that needs to be fitted onto the world: it just is (part of) the world, plain and simple: all the leptons, just as God made 'em.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Fri Sep 18 17:55:58 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA12669;
Fri, 18 Sep 1998 12:24:54 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a06b22851071726@[143.88.7.118]>
In-Reply-To: <3.0.32.19980917141509.00f21eb8@catbert.cyc.com>
Mime-Version: 1.0
Date: Fri, 18 Sep 1998 12:21:38 -0600
To: Fritz Lehmann <fritz@cyc.com>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Models intermediate or not? Was: Re: tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 3181
Status:

>At 12:25 PM 9/17/98 -0600, Pat Hayes wrote:
>>[...]
>>However, I disagee (along with Tarski, Quine and others.) Models may be
>>built from such things as integers, but they can also be built from such
>>things as tables, events or pieces of cheese. Model theory requires a
>>domain which is a set, but that doesnt mean that models must be 'formal',
>>and hence lacking that rich, loamy texture that real worlds have. A set can
>>be a set of anything, including parts of the real world.
>>[...]
>>Pat
>
>
>Dear Pat,
>

>Do I understand your view correctly if I understand it to imply that a
>model can never be wrong? A sentence can be false with respect to the
>model in both views.

Im not sure what you mean by a model being 'wrong'. Models, or interpretations, assign truthvalues to closed sentences; so a model (actually one ought to say an interpretation, since in model theory "model" means an interpretation which makes a sentence true, but lets continue using "model" to be synonymous with "interpretation") is, in a sense, a world. So of course it can't be 'wrong' about itself, although it can certainly make some sentences false, ie those sentences may be false in it.

My problem with John's account of 'models' is that it seems to need two different ways for a sentence to be false of the world: it might be false in a model, or that model may be inaccurate in some way. But I can only see one way in which a sentence can be false, which involves giving an interpretation of it in some world. Model theory talks of the relationship between sentences and interpretations, and thats all.

In Sowa's view a model could also fail to correctly
>model the real world. Your view, taking the elements of the model to be
>entities in the real world, seems to make the model incapable of error. Can
>the model even be capable of underspecification of the real world (even in
>a technically inaccessible way), in your approach? How?

Well, models *need* not be made of real-world stuff; my point is only that they *can* be. There are also, for example, Herbrand models, which are made entirely of symbolic expressions.

Im not sure what you mean by 'underspecification'. A sentence can surely underspecify the world. Heres an example: $(\text{forall } x)(P \ x)$. If we interpret 'P' to mean the property of having zero rest mass and the domain to be all leptons, then this sentence is true in the actual world. But if we take P to mean the property of having charge, then it isnt true. Two different interpretations of the same sentence, both in the real world, give different values. Thats why the sentence doesnt tell us much about - underspecifies - the world. But its the *sentence* that 'underspecifies', not a model. The first interpretation isnt some abstract thing that needs to be fitted onto the world: it just is (part of) the world, plain and simple: all the leptons, just as God made 'em.

Pat

11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Sep 30 11:06:49 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id JAA11919
for <phayes@coginst.uwf.edu>; Sat, 26 Sep 1998 09:01:49 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id JAA06369;
Sat, 26 Sep 1998 09:57:44 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
id JAA29390; Sat, 26 Sep 1998 09:52:42 -0400
Date: Sat, 26 Sep 1998 09:52:42 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199809261352.JAA29390@west>
To: fritz@cyc.com, phayes@coginst.uwf.edu
Subject: Re: Models intermediate or not? Was: Re: tropes
Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@tamu.edu,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu,
guarino@ladseb.pd.cnr.it, jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk,
skydog@pacbell.net, sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 5478
Status:

Pat,

There are several reasons for drawing a sharp distinction between models of the world and the actual world we all live in:

1. It replaces many ill-defined notions with more conventional mathematical constructions. To borrow one of your favorite phrases, "I don't know what it means to say that" a model can be "built from tables, events, and pieces of cheese." I understand very well how to build a mathematical model from symbolic expressions, and I understand how to build an engineering model of an airplane at one-tenth scale to test in a wind tunnel. The first is abstract, and the second is physical. But I do not understand how you can "build" a mathematical model from cheese.
2. Engineers have a long history of developing two kinds of models: mathematical models stated entirely in terms of symbolic expressions,

and physical models that are instrumented with sensors for measuring the quantities represented by the symbolic expressions. Their techniques involve exactly the kinds of empirical tests, approximations, and criteria for success that give meaning to what they are doing. Unlike the unnamed "logicians" you have not cited, they are using words in a very precise, testable way. (As I have said many times, Tarski does not belong in that list of logicians, because he defined his terms so precisely that he deliberately excluded tables, events, and cheese. He said that his models were intended only for formal languages, and to the end of his days, he objected to the loose application of his approach to natural language -- or to pseudo-formal languages with predicates like `pieceOfCheese(x)`.)

3. I realize that you have a long background in physics as well as mathematics, so I don't believe that you are really as confused as some of the unnamed "logicians" you refer to. I believe that your terminology and my terminology can be reconciled by a global change of the term 'possible world' to 'possible model of the world'. That change has several consequences:

- a) It eliminates the need for adjectives like "possible" or "real" in front of "world". There is only one world.
- b) It makes all the models of model theory into mathematical models "built" from symbolic entities.
- c) It allows you to say that some privileged model happens to be isomorphic to the world or some aspect of it. You can define the truth of a sentence s about the world as truth of s about a model that is isomorphic to the world. Once you have given such a definition, then you can adopt convenient abbreviations that would allow you to drop the phrase "isomorphic model of" whenever you can be sure that your readers understand the point.
- d) But it also allows you to adopt an agnostic attitude towards your models in those cases (actually the overwhelming majority of cases) for which you don't know whether an exact isomorphism exists.
- e) Finally (and for many purposes this is the most important point), this distinction gives you a convenient place to introduce measures of approximation, granularity, and experimental error. We have all expressed serious reservations (or worse) about Lotfi Z's fuzziness. This separation of models from the world gives us a precise two-valued logic for relating sentences to models and a continuously variable range of tolerances for measuring the adequacy of a model.

Some comments on your comments:

>My problem with John's account of 'models' is that it seems to need two
>different ways for a sentence to be false of the world: it might be false
>in a model, or that model may be inaccurate in some way. But I can only see
>one way in which a sentence can be false, which involves giving an
>interpretation of it in some world.

If you replace that last word 'world' with 'model', I would be much happier. I don't like to put quantifiers in front of 'world' because that suggests that there might be other possible worlds, which no one has ever observed or ever suggested how they might be observed. I agree with you that there is only one way that a sentence can be false. The inaccuracy of a model does not make a sentence false -- it gives you another metalevel for talking about your interpretation. In effect, we have multiple metalevels: truth and falsity are metalevel terms that are not part of the first-order language. Talking about the level of approximation goes to the metametalevel:

1. "It is true that p." is a metalevel statement about p, which at the metalevel happens to imply p.
2. "p is true in model M" is another metalevel statement about p.
3. "M is an accurate model of the world within the tolerance of our measuring instruments" is a statement about M, which can be combined with the previous statement to make a metametastatement about p:

"p is true in a model that accurately represents the world within the tolerance of our measuring instruments."

I claim that it is very important to be able to make statements like this. It meets the lofty requirements in a much clearer way than fuzzy logic.

>Model theory talks of the relationship
>between sentences and interpretations, and that's all.

I agree with this point. That's why I also want to make metametastatements about the models.

John

From ???@??? Wed Sep 30 17:31:09 1998

Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id PAA25289;
Wed, 30 Sep 1998 15:51:20 -0500 (CDT)

X-Sender: phayes@mail.coginst.uwf.edu

Message-Id: <v04003a03b2381b8853b5@[143.88.7.118]>

In-Reply-To: <199809261352.JAA29390@west>
Mime-Version: 1.0
Date: Wed, 30 Sep 1998 15:48:13 -0500
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Models intermediate or not? Was: Re: tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 15622
Status:

John, it is foolish for us to repeat these old arguments. I am never going to agree with you about this matter, and I think your position is profoundly mistaken on both technical and historical grounds. You are of course free to adopt any philosophical position you like, but your extreme nominalism is, shall we say, not to my taste; and while I would not attempt to ever change your mind or question your faith, I register my disagreement. Since we do disagree, let us instead agree that a 'standard' ontology should as far as possible allow either position, ie be as catholic as possible. Towards that end, therefore, we should allow the notion of interpretation (in the formal sense) to range as widely as possible. If anyone wishes then to restrict their attention to some particular class of interpretations (eg only 'mathematical' ones, or those containing no cheese) then they are free to do so, but we also allow people with a broader view of what constitutes an interpretation to not feel excluded from using the results of our deliberations.

Pat

PS. In response to your message, for the benefit of those who havnt seen us arguing about this before.

[John Sowa:]

>There are several reasons for drawing a sharp distinction between

>models of the world and the actual world we all live in:

>

> 1. It replaces many ill-defined notions with more conventional mathematical
> constructions. To borrow one of your favorite phrases, "I don't know
> what it means to say that" a model can be "built from tables, events,
> and pieces of cheese." I understand very well how to build a mathematical

- > model from symbolic expressions, and I understand how to build an
- > engineering model of an airplane at one-tenth scale to test in a wind
- > tunnel. The first is abstract, and the second is physical. But I do not
- > understand how you can "build" a mathematical model from cheese.

As I explained in my previous message, you are using 'model' in two different senses. Using the distinction I introduced there (model-1 is an engineering simplification, model-2 is a Tarkian interpretation of a sentence) , of course we cannot 'build' a model-2 in the same sense that we can build a model-1. But if you allow linguistic interpretations - models-2 - which are physical enough to be placed in a windtunnel (ie taking your confused use of 'model' at face value), then I see no reason to exclude models made of cheese or indeed of anything else. (Unless perhaps you want to raise certain materials - stainless steel and plexiglass, etc. - to a different semantic plane than the one inhabited by such lowly stuff as dairy produce?)

Notice that when you talk of a 'mathematical' model, to me this can only mean a description in a mathematical language. But such a description may well refer to models-2 containing cheese (as when I add up the total weight of cheese I purchased, multiply it by the price per pound and discover that the supermarket overcharged me.) I think that you take a different view; to you, the world is sharply divided into solid physical stuff on the one hand and ethereal mathematical abstractions on the other, and any math-sounding term like 'number' or 'set' necessarily refers to the ethereal. Most platonists are quite happy to say that these abstractions are real, but you seem to want to combine two views which are usually taken to be in opposition: a platonist view of mathematics and a sharply nominalist view of physical talk. It seems simpler to me to just say that there are languages and ways they can be interpreted. If we say it this baldly you can evidently agree with it, so lets agree to stop at a point where we have not yet diverged.

- > 2. Engineers have a long history of developing two kinds of models:
 - > mathematical models stated entirely in terms of symbolic expressions,
 - > and physical models that are instrumented with sensors for measuring the
 - > quantities represented by the symbolic expressions. Their techniques
 - > involve exactly the kinds of empirical tests, approximations, and
 - > criteria for success that give meaning to what they are doing.

Well theres all kinds of things to say to this. Just two will have to do for now. First, most engineering 'models' involve both physical and symbolic aspects. Certainly anything modelled on a computer does, for example; but think also of slide rules. Second, the idea that all these tests are what "give meaning" to the model in the same sense as a theory of semantics, is highly debateable. At any rate, its an extreme view of

'meaning'.

- > Unlike
- > the unnamed "logicians" you have not cited, they are using words in
- > a very precise, testable way. (As I have said many times, Tarski does
- > not belong in that list of logicians, because he defined his terms
- > so precisely that he deliberately excluded tables, events, and cheese.
- > He said that his models were intended only for formal languages,

Another confusion. Of course we are here talking about formal languages; but the formality of the *language* is a separate matter from the *semantics* of that language. Here we all are trying to create a formal general-purpose ontology, for goodness sake. If we are going to rule out the possibility of our formalism referring to parts of the real world then we have given up before even beginning.

As to the "unnamed logicians"; as you know, I have given you ample citations previously, but if you insist:

Quine: (Shoes, From "Set Theory and its Logic", Introduction:) "The notion of class is so fundamental to thought that we cannot hope to define it in more fundamental terms. We can say that a class is any aggregate, any collection, any combination of objects of any sort....[but] the aggregating or collecting or combining is to connote no actual displacement of the objects, and further that the aggregation or collection or combination of say seven pairs of shoes is not to be identified with the aggregation or collection or combination of those fourteen shoes...."

Carnap: (Astronomy, From "Introduction to Symbolic Logic and its Applications", p.4) "Eg in a certain application 'P' might designate the property Spherical....Now suppose that ..we take 'a' to designate the sun and 'b' to designate the moon. Then in our symbolic language we write the sentence 'P(a)' for the "The sun is spherical"." [et.seq; the entire book is full of physical examples like this.]

In fact, let me make a counterchallenge. Find *any* logician or logic text which talks of Tarski's theory of truth in such a way as to exclude domains of any particular kind, eg restricts the theory to apply only to 'mathematical' domains. I rule out Peirce and Barwise & Perry.

- > to the end of his days, he objected to the loose application of his
- > approach to natural language -- or to pseudo-formal languages with
- > predicates like pieceOfCheese(x).

I challenge you to back this claim up with citations. Tarski may have been more interested in mathematics than ontology, but he was always

crystal-clear about what counted as a member of a set, and it certainly allowed sets of physical things. (And by the way, there's nothing 'pseudo' about the formality of 'pieceOfCheese'.)

- >
- > 3. I realize that you have a long background in physics as well as
- > mathematics, so I don't believe that you are really as confused as
- > some of the unnamed "logicians" you refer to. I believe that your
- > terminology and my terminology can be reconciled by a global change
- > of the term 'possible world' to 'possible model of the world'.

Sorry, that is exactly what I object to, and my own view cannot be reconciled to this change. To do so would deny almost everything I've ever thought and written about the importance of a semantic theory in considering representations in AI. The semantic theory relates a sentence to a way the world might be, not to a model (ie a model-1) of the world (unless that model *is*, a way the world might be, of course.) In fact it's not even coherent to speak of model-2's "of the world": these are models of sentences.

That

- > change has several consequences:
- >
- > a) It eliminates the need for adjectives like "possible" or "real"
- > in front of "world". There is only one world.

But even if we accept this as a piece of metaphysics, it's not a good basis for a theory of meaning, for several reasons. First, there are many *ways* that the world might be used to interpret a sentence: that is, there are many interpretations whose domain of discourse is part of the real world. Second, we often talk and reason counterfactually, ie about ways the world might have been (but isn't). Third, we seem to need to be able to talk about such alternative worlds to account for (mis)perception reports. (None of this talk amounts to Lewis-like 'strong claim' about the real existence of these alternative worlds; but our theory of truth needs to be able to discuss how our sentences *would* be interpreted *if* they were interpreted in worlds different from this one, since neither we nor our languages have access to all the details of this world.)

- >
- > b) It makes all the models of model theory into mathematical models
- > "built" from symbolic entities.

This sense of "build", as your use of scare-quotes maybe illustrates, is quite different from the engineers-model wind-tunnel sense you have used earlier. And the difference is crucially important, because the models-2 that matter are often physically impossible (maybe transfinite, for

example) ; which is precisely why they *are* important, since they are often useful as a way to establish a flaw in a physical axiomatisation.

- > c) It allows you to say that some privileged model happens to be
- > isomorphic to the world or some aspect of it. You can define
- > the truth of a sentence s about the world as truth of s about
- > a model that is isomorphic to the world. Once you have given
- > such a definition, then you can adopt convenient abbreviations
- > that would allow you to drop the phrase "isomorphic model of"
- > whenever you can be sure that your readers understand the point.

This argument is rather like saying that one should wear a mask because people will see you face better when you take it off.

- > d) But it also allows you to adopt an agnostic attitude towards your
- > models in those cases (actually the overwhelming majority of cases)
- > for which you don't know whether an exact isomorphism exists.

It always exists, by definition. That is, every model-2 is isomorphic to a "world", ie itself.

- > e) Finally (and for many purposes this is the most important point),
- > this distinction gives you a convenient place to introduce measures
- > of approximation, granularity, and experimental error. We have all
- > expressed serious reservations (or worse) about Lotfi Z's fuzziness.
- > This separation of models from the world gives us a precise two-valued
- > logic for relating sentences to models and a continuously variable
- > range of tolerances for measuring the adequacy of a model.

Rubbish. Ahem, please forgive my rudeness. Let me rephrase that: I'll believe it when I see it. I've been trying, on and off, for almost 15 years now to create a useful nontrivial account of approximation and granularity (especially one that can be used to account for vernier effects in perceptual judgements, the usefulness of dithering, etc..) and I've not been able to and havnt seen anyone else being able to. People are constantly claiming to have some magic way to handle it, and now you seem to think that driving an truck over Tarski's theory of truth is somehow going to solve this particular hard problem. I can't type the sound I am inclined to make.

Again, as a challenge, just for a start: show us how you propose to get past the first classical snag, the heap paradox. Remember, we need to be able to *reason* about those heaps and grains of sand; its not enough to give us a theory of heaps and then talk of grains in a metatheory.

>Some comments on your comments:

>
>>My problem with John's account of 'models' is that it seems to need two
>>different ways for a sentence to be false of the world: it might be false
>>in a model, or that model may be inaccurate in some way. But I can only see
>>one way in which a sentence can be false, which involves giving an
>>interpretation of it in some world.

>
>If you replace that last word 'world' with 'model', I would be much happier.
>I don't like to put quantifiers in front of 'world' because that suggests
>that there might be other possible worlds, which no one has ever observed
>or ever suggested how they might be observed.

Read Kripke "Naming and Necessity" for a nice analysis of what is this wrong with this way of thinking. Possible worlds aren't other places, like infinitely distant galaxies. They are alternative ways this world could be.

> I agree with you that there
>is only one way that a sentence can be false. The inaccuracy of a model
>does not make a sentence false -- it gives you another metalevel for talking
>about your interpretation. In effect, we have multiple metalevels: truth
>and falsity are metalevel terms that are not part of the first-order language.
>Talking about the level of approximation goes to the metametalevel:

>
> 1. "It is true that p." is a metalevel statement about p, which at the
> metalevel happens to imply p.
>
> 2. "p is true in model M" is another metalevel statement about p.

OK, lets look at this. Where in Tarskian (or any other) semantic theory is there a requirement that the name 'M', when interpreted in (a suitable metatheory extension of) M itself, must denote M? If you are here proposing a formal language (as opposed to just doing mathematics) then you can't just stipulate what a name shall denote: you must provide a theory of (meta)truth and then see what interpretations it provides. Good luck with the Godel sentences.

> 3. "M is an accurate model of the world within the tolerance of our
> measuring instruments" is a statement about M, which can be combined
> with the previous statement to make a metametastatement about p:

I'd like to see a sketch of the instrument that can detect, say, a domain with cardinality \aleph_3 , or the first inaccessible ordinal. Sorry to sound sarcastic, but I'm reduced to it here. Your entire discussion is just completely confused by using an ambiguous word ('model') first in one sense and then another.

Pat

PS. On a more conciliatory note. If what worries you is referring (even in our metaln-discussions) to any 'other world' (which is ridiculous since there is only one world, etc.) or still worse to 'possible worlds', then there is an alternative way to talk. Instead of saying, (sentence) P is true in (possible) world M, say rather: M is a (possible) *way* for the sentence P to be true. This focusses on the (important!) fact that there is lot more to an interpretation than just its domain, and also that aspects of the world which are invisible to the signature of the sentence are irrelevant to the truth-conditions. This way, you can be content with the unity of the single World, and I can think to myself: this is another way to say, 'in a possible world'.

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Sep 30 17:31:09 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id PAA25289;
Wed, 30 Sep 1998 15:51:20 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a03b2381b8853b5@[143.88.7.118]>
In-Reply-To: <199809261352.JAA29390@west>
Mime-Version: 1.0
Date: Wed, 30 Sep 1998 15:48:13 -0500
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Models intermediate or not? Was: Re: tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 15622
Status:

John, it is foolish for us to repeat these old arguments. I am never going

to agree with you about this matter, and I think your position is profoundly mistaken on both technical and historical grounds. You are of course free to adopt any philosophical position you like, but your extreme nominalism is, shall we say, not to my taste; and while I would not attempt to ever change your mind or question your faith, I register my disagreement. Since we do disagree, let us instead agree that a 'standard' ontology should as far as possible allow either position, ie be as catholic as possible. Towards that end, therefore, we should allow the notion of interpretation (in the formal sense) to range as widely as possible. If anyone wishes then to restrict their attention to some particular class of interpretations (eg only 'mathematical' ones, or those containing no cheese) then they are free to do so, but we also allow people with a broader view of what constitutes an interpretation to not feel excluded from using the results of our deliberations.

Pat

PS. In response to your message, for the benefit of those who havnt seen us arguing about this before.

[John Sowa:]

>There are several reasons for drawing a sharp distinction between
>models of the world and the actual world we all live in:

>

> 1. It replaces many ill-defined notions with more conventional mathematical
> constructions. To borrow one of your favorite phrases, "I don't know
> what it means to say that" a model can be "built from tables, events,
> and pieces of cheese." I understand very well how to build a mathematical
> model from symbolic expressions, and I understand how to build an
> engineering model of an airplane at one-tenth scale to test in a wind
> tunnel. The first is abstract, and the second is physical. But I do not
> understand how you can "build" a mathematical model from cheese.

As I explained in my previous message, you are using 'model' in two different senses. Using the distinction I introduced there (model-1 is an engineering simplification, model-2 is a Tarkian interpretation of a sentence) , of course we cannot 'build' a model-2 in the same sense that we can build a model-1. But if you allow linguistic interpretations - models-2 - which are physical enough to be placed in a windtunnel (ie taking your confused use of 'model' at face value), then I see no reason to exclude models made of cheese or indeed of anything else. (Unless perhaps you want to raise certain materials - stainless steel and plexiglass, etc. - to a different semantic plane than the one inhabited by such lowly stuff as dairy produce?)

Notice that when you talk of a 'mathematical' model, to me this can only mean a description in a mathematical language. But such a description may well refer to models-2 containing cheese (as when I add up the total weight of cheese I purchased, multiply it by the price per pound and discover that the supermarket overcharged me.) I think that you take a different view; to you, the world is sharply divided into solid physical stuff on the one hand and ethereal mathematical abstractions on the other, and any math-sounding term like 'number' or 'set' necessarily refers to the ethereal. Most platonists are quite happy to say that these abstractions are real, but you seem to want to combine two views which are usually taken to be in opposition: a platonist view of mathematics and a sharply nominalist view of physical talk. It seems simpler to me to just say that there are languages and ways they can be interpreted. If we say it this baldly you can evidently agree with it, so lets agree to stop at a point where we have not yet diverged.

- > 2. Engineers have a long history of developing two kinds of models:
- > mathematical models stated entirely in terms of symbolic expressions,
- > and physical models that are instrumented with sensors for measuring the
- > quantities represented by the symbolic expressions. Their techniques
- > involve exactly the kinds of empirical tests, approximations, and
- > criteria for success that give meaning to what they are doing.

Well theres all kinds of things to say to this. Just two will have to do for now. First, most engineering 'models' involve both physical and symbolic aspects. Certainly anything modelled on a computer does, for example; but think also of slide rules. Second, the idea that all these tests are what "give meaning" to the model in the same sense as a theory of semantics, is highly debateable. At any rate, its an extreme view of 'meaning'.

- > Unlike
- > the unnamed "logicians" you have not cited, they are using words in
- > a very precise, testable way. (As I have said many times, Tarski does
- > not belong in that list of logicians, because he defined his terms
- > so precisely that he deliberately excluded tables, events, and cheese.
- > He said that his models were intended only for formal languages,

Another confusion. Of course we are here talking about formal languages; but the formality of the *language* is a seperate matter from the *semantics* of that language. Here we all are trying to create a formal general-purpose ontology, for goodness sake. If we are going to rule out the possibility of our formalism referring to parts of the real world then we have given up before even beginning.

As to the "unnamed logicians"; as you know, I have given you ample

citations previously, but if you insist:

Quine: (Shoes, From "Set Theory and its Logic", Introduction:) "The notion of class is so fundamental to thought that we cannot hope to define it in more fundamental terms. We can say that a class is any aggregate, any collection, any combination of objects of any sort.....[but] the aggregating or collecting or combining is to connote no actual displacement of the objects, and further that the aggregation or collection or combination of say seven pairs of shoes is not to be identified with the aggregation or collection or combination of those fourteen shoes...."

Carnap:(Astronomy, From "Introduction to Symbolic Logic and its Applications", p.4) "Eg in a certain application 'P' might designate the property Spherical....Now suppose that ..we take 'a' to designate the sun and 'b' to designate the moon. Then in our symbolic language we write the sentence 'P(a)' for the "The sun is spherical"." [et.seq; the entire book is full of physical examples like this.]

In fact, let me make a counterchallenge. Find *any* logician or logic text which talks of Tarki's theory of truth in such a way as to exclude domains of any particular kind, eg restricts the theory to apply only to 'mathematical' domains. I rule out Peirce and Barwise & Perry.

- > to the end of his days, he objected to the loose application of his
- > approach to natural language -- or to pseudo-formal languages with
- > predicates like pieceOfCheese(x.)

I challenge you to back this claim up with citations. Tarski may have been more interested in mathematics than ontology, but he was always crystal-clear about what counted as a member of a set, and it certainly allowed sets of physical things. (And by the way, theres nothing 'pseudo' about the formality of 'pieceOfCheese'.)

>

- > 3. I realize that you have a long background in physics as well as
- > mathematics, so I don't believe that you are really as confused as
- > some of the unnamed "logicians" you refer to. I believe that your
- > terminology and my terminology can be reconciled by a global change
- > of the term 'possible world' to 'possible model of the world'.

Sorry, that is exactly what I object to, and my own view cannot be reconciled to this change. To do so would deny almost everything Ive ever thought and written about the importance of a semantic theory in considering representations in AI. The semantic theory relates a sentence to a way the world might be, not to a model (ie a model-1) of the world (unless that model *is*, a way the world might be, of course.) In fact its

not even coherent to speak of model-2's "of the world": these are models of sentences.

That

- > change has several consequences:
- >
- > a) It eliminates the need for adjectives like "possible" or "real"
- > in front of "world". There is only one world.

But even if we accept this as a piece of metaphysics, its not a good basis for a theory of meaning, for several reasons. First, there are many *ways* that the world might be used to interpret a sentence: that is, there are many interpretations whose domain of discourse is part of the real world. Second, we often talk and reason counterfactually, ie about ways the world might have been (but isnt). Third, we seem to need to be able to talk about such alternative worlds to account for (mis)perception reports. (None of this talk amounts to Lewis-like 'strong claim' about the real existence of these alternative worlds; but our theory of truth needs to be able discuss how our sentences *would* be interpreted *if* they were interpreted in worlds different from this one, since neither we nor our languages have access to all the details of this world.)

- >
- > b) It makes all the models of model theory into mathematical models
- > "built" from symbolic entities.

This sense of "build", as your use of scare-quotes maybe illustrates, is quite different from the engineers-model wind-tunnel sense you have used earlier. And the difference is crucially important, because the models-2 that matter are often physically impossible (maybe transfinite, for example) ; which is precisely why they *are* important, since they are often useful as a way to establish a flaw in a physical axiomatisation.

- > c) It allows you to say that some privileged model happens to be
- > isomorphic to the world or some aspect of it. You can define
- > the truth of a sentence s about the world as truth of s about
- > a model that is isomorphic to the world. Once you have given
- > such a definition, then you can adopt convenient abbreviations
- > that would allow you to drop the phrase "isomorphic model of"
- > whenever you can be sure that your readers understand the point.

This argument is rather like saying that one should wear a mask because people will see you face better when you take it off.

- > d) But it also allows you to adopt an agnostic attitude towards your
- > models in those cases (actually the overwhelming majority of cases)
- > for which you don't know whether an exact isomorphism exists.

It always exists, by definition. That is, every model-2 is isomorphic to a "world", ie itself.

- > e) Finally (and for many purposes this is the most important point),
- > this distinction gives you a convenient place to introduce measures
- > of approximation, granularity, and experimental error. We have all
- > expressed serious reservations (or worse) about Lotfi Z's fuzziness.
- > This separation of models from the world gives us a precise two-valued
- > logic for relating sentences to models and a continuously variable
- > range of tolerances for measuring the adequacy of a model.

Rubbish. Ahem, please forgive my rudeness. Let me rephrase that: I'll believe it when I see it. I've been trying, on and off, for almost 15 years now to create a useful nontrivial account of approximation and granularity (especially one that can be used to account for vernier effects in perceptual judgements, the usefulness of dithering, etc..) and I've not been able to and havnt seen anyone else being able to. People are constantly claiming to have some magic way to handle it, and now you seem to think that driving an truck over Tarski's theory of truth is somehow going to solve this particular hard problem. I can't type the sound I am inclined to make.

Again, as a challenge, just for a start: show us how you propose to get past the first classical snag, the heap paradox. Remember, we need to be able to *reason* about those heaps and grains of sand; its not enough to give us a theory of heaps and then talk of grains in a metatheory.

>Some comments on your comments:

>

>>My problem with John's account of 'models' is that it seems to need two
>>different ways for a sentence to be false of the world: it might be false
>>in a model, or that model may be inaccurate in some way. But I can only see
>>one way in which a sentence can be false, which involves giving an
>>interpretation of it in some world.

>

>If you replace that last word 'world' with 'model', I would be much happier.
>I don't like to put quantifiers in front of 'world' because that suggests
>that there might be other possible worlds, which no one has ever observed
>or ever suggested how they might be observed.

Read Kripke "Naming and Necessity" for a nice analysis of what is this wrong with this way of thinking. Possible worlds arent other places, like infinitely distant galaxies. They are alternative ways this world could be.

> I agree with you that there

- >is only one way that a sentence can be false. The inaccuracy of a model
- >does not make a sentence false -- it gives you another metalevel for talking
- >about your interpretation. In effect, we have multiple metalevels: truth
- >and falsity are metalevel terms that are not part of the first-order language.
- >Talking about the level of approximation goes to the metametalevel:
- >
- > 1. "It is true that p." is a metalevel statement about p, which at the
- > metalevel happens to imply p.
- >
- > 2. "p is true in model M" is another metalevel statement about p.

OK, lets look at this. Where in Tarskian (or any other) semantic theory is there a requirement that the name 'M', when interpreted in (a suitable metatheory extension of) M itself, must denote M? If you are here proposing a formal language (as opposed to just doing mathematics) then you can't just stipulate what a name shall denote: you must provide a theory of (meta)truth and then see what interpretations it provides. Good luck with the Godel sentences.

- > 3. "M is an accurate model of the world within the tolerance of our
- > measuring instruments" is a statement about M, which can be combined
- > with the previous statement to make a metametastatement about p:

I'd like to see a sketch of the instrument that can detect, say, a domain with cardinality \aleph_3 , or the first inaccessible ordinal. Sorry to sound sarcastic, but I'm reduced to it here. Your entire discussion is just completely confused by using an ambiguous word ('model') first in one sense and then another.

Pat

PS. On a more conciliatory note. If what worries you is referring (even in our metaIn-discussions) to any 'other world' (which is ridiculous since there is only one world, etc.) or still worse to 'possible worlds', then there is an alternative way to talk. Instead of saying, (sentence) P is true in (possible) world M, say rather: M is a (possible) *way* for the sentence P to be true. This focusses on the (important!) fact that there is lot more to an interpretation than just its domain, and also that aspects of the world which are invisible to the signature of the sentence are irrelevant to the truth-conditions. This way, you can be content with the unity of the single World, and I can think to myself: this is another way to say, 'in a possible world'.

Pensacola, FL 32514
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

(850)474 3023 fax

From ???@??? Wed Sep 30 17:31:09 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id PAA25289;
Wed, 30 Sep 1998 15:51:20 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a03b2381b8853b5@[143.88.7.118]>
In-Reply-To: <199809261352.JAA29390@west>
Mime-Version: 1.0
Date: Wed, 30 Sep 1998 15:48:13 -0500
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Models intermediate or not? Was: Re: tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 15622
Status:

John, it is foolish for us to repeat these old arguments. I am never going to agree with you about this matter, and I think your position is profoundly mistaken on both technical and historical grounds. You are of course free to adopt any philosophical position you like, but your extreme nominalism is, shall we say, not to my taste; and while I would not attempt to ever change your mind or question your faith, I register my disagreement. Since we do disagree, let us instead agree that a 'standard' ontology should as far as possible allow either position, ie be as catholic as possible. Towards that end, therefore, we should allow the notion of interpretation (in the formal sense) to range as widely as possible. If anyone wishes then to restrict their attention to some particular class of interpretations (eg only 'mathematical' ones, or those containing no cheese) then they are free to do so, but we also allow people with a broader view of what constitutes an interpretation to not feel excluded from using the results of our deliberations.

Pat

PS. In response to your message, for the benefit of those who havnt seen us arguing about this before.

[John Sowa:]

>There are several reasons for drawing a sharp distinction between
>models of the world and the actual world we all live in:

>

> 1. It replaces many ill-defined notions with more conventional mathematical
> constructions. To borrow one of your favorite phrases, "I don't know
> what it means to say that" a model can be "built from tables, events,
> and pieces of cheese." I understand very well how to build a mathematical
> model from symbolic expressions, and I understand how to build an
> engineering model of an airplane at one-tenth scale to test in a wind
> tunnel. The first is abstract, and the second is physical. But I do not
> understand how you can "build" a mathematical model from cheese.

As I explained in my previous message, you are using 'model' in two different senses. Using the distinction I introduced there (model-1 is an engineering simplification, model-2 is a Tarkian interpretation of a sentence) , of course we cannot 'build' a model-2 in the same sense that we can build a model-1. But if you allow linguistic interpretations - models-2 - which are physical enough to be placed in a windtunnel (ie taking your confused use of 'model' at face value), then I see no reason to exclude models made of cheese or indeed of anything else. (Unless perhaps you want to raise certain materials - stainless steel and plexiglass, etc. - to a different semantic plane than the one inhabited by such lowly stuff as dairy produce?)

Notice that when you talk of a 'mathematical' model, to me this can only mean a description in a mathematical language. But such a description may well refer to models-2 containing cheese (as when I add up the total weight of cheese I purchased, multiply it by the price per pound and discover that the supermarket overcharged me.) I think that you take a different view; to you, the world is sharply divided into solid physical stuff on the one hand and ethereal mathematical abstractions on the other, and any math-sounding term like 'number' or 'set' necessarily refers to the ethereal. Most platonists are quite happy to say that these abstractions are real, but you seem to want to combine two views which are usually taken to be in opposition: a platonist view of mathematics and a sharply nominalist view of physical talk. It seems simpler to me to just say that there are languages and ways they can be interpreted. If we say it this baldly you can evidently agree with it, so lets agree to stop at a point where we have not yet diverged.

> 2. Engineers have a long history of developing two kinds of models:

- > mathematical models stated entirely in terms of symbolic expressions,
- > and physical models that are instrumented with sensors for measuring the
- > quantities represented by the symbolic expressions. Their techniques
- > involve exactly the kinds of empirical tests, approximations, and
- > criteria for success that give meaning to what they are doing.

Well theres all kinds of things to say to this. Just two will have to do for now. First, most engineering 'models' involve both physical and symbolic aspects. Certainly anything modelled on a computer does, for example; but think also of slide rules. Second, the idea that all these tests are what "give meaning" to the model in the same sense as a theory of semantics, is highly debateable. At any rate, its an extreme view of 'meaning'.

- > Unlike
- > the unnamed "logicians" you have not cited, they are using words in
- > a very precise, testable way. (As I have said many times, Tarski does
- > not belong in that list of logicians, because he defined his terms
- > so precisely that he deliberately excluded tables, events, and cheese.
- > He said that his models were intended only for formal languages,

Another confusion. Of course we are here talking about formal languages; but the formality of the *language* is a seperate matter from the *semantics* of that language. Here we all are trying to create a formal general-purpose ontology, for goodness sake. If we are going to rule out the possibility of our formalism referring to parts of the real world then we have given up before even beginning.

As to the "unnamed logicians"; as you know, I have given you ample citations previously, but if you insist:

Quine: (Shoes, From "Set Theory and its Logic", Introduction:) "The notion of class is so fundamental to thought that we cannot hope to define it in more fundamental terms. We can say that a class is any aggregate, any collection, any combination of objects of any sort.....[but] the aggregating or collecting or combining is to connote no actual displacement of the objects, and further that the aggregation or collection or combination of say seven pairs of shoes is not to be identified with the aggregation or collection or combination of those fourteen shoes...."

Carnap:(Astronomy, From "Introduction to Symbolic Logic and its Applications", p.4) "Eg in a certain application 'P' might designate the property Spherical....Now suppose that ..we take 'a' to designate the sun and 'b' to designate the moon. Then in our symbolic language we write the sentence 'P(a)' for the "The sun is spherical"." [et.seq; the entire book is full of physical examples like this.]

In fact, let me make a counterchallenge. Find *any* logician or logic text which talks of Tarki's theory of truth in such a way as to exclude domains of any particular kind, eg restricts the theory to apply only to 'mathematical' domains. I rule out Peirce and Barwise & Perry.

- > to the end of his days, he objected to the loose application of his
- > approach to natural language -- or to pseudo-formal languages with
- > predicates like pieceOfCheese(x).

I challenge you to back this claim up with citations. Tarski may have been more interested in mathematics than ontology, but he was always crystal-clear about what counted as a member of a set, and it certainly allowed sets of physical things. (And by the way, theres nothing 'pseudo' about the formality of 'pieceOfCheese'.)

- >
- > 3. I realize that you have a long background in physics as well as
- > mathematics, so I don't believe that you are really as confused as
- > some of the unnamed "logicians" you refer to. I believe that your
- > terminology and my terminology can be reconciled by a global change
- > of the term 'possible world' to 'possible model of the world'.

Sorry, that is exactly what I object to, and my own view cannot be reconciled to this change. To do so would deny almost everything Ive ever thought and written about the importance of a semantic theory in considering representations in AI. The semantic theory relates a sentence to a way the world might be, not to a model (ie a model-1) of the world (unless that model *is*, a way the world might be, of course.) In fact its not even coherent to speak of model-2's "of the world": these are models of sentences.

That

- > change has several consequences:
- >
- > a) It eliminates the need for adjectives like "possible" or "real"
- > in front of "world". There is only one world.

But even if we accept this as a piece of metaphysics, its not a good basis for a theory of meaning, for several reasons. First, there are many *ways* that the world might be used to interpret a sentence: that is, there are many interpretations whose domain of discourse is part of the real world. Second, we often talk and reason counterfactually, ie about ways the world might have been (but isnt). Third, we seem to need to be able to talk about such alternative worlds to account for (mis)perception reports. (None of this talk amounts to Lewis-like 'strong claim' about the real existence of

these alternative worlds; but our theory of truth needs to be able discuss how our sentences *would* be interpreted *if* they were interpreted in worlds different from this one, since neither we nor our languages have access to all the details of this world.)

- >
- > b) It makes all the models of model theory into mathematical models
- > "built" from symbolic entities.

This sense of "build", as your use of scare-quotes maybe illustrates, is quite different from the engineers-model wind-tunnel sense you have used earlier. And the difference is crucially important, because the models-2 that matter are often physically impossible (maybe transfinite, for example) ; which is precisely why they *are* important, since they are often useful as a way to establish a flaw in a physical axiomatisation.

- > c) It allows you to say that some privileged model happens to be
- > isomorphic to the world or some aspect of it. You can define
- > the truth of a sentence s about the world as truth of s about
- > a model that is isomorphic to the world. Once you have given
- > such a definition, then you can adopt convenient abbreviations
- > that would allow you to drop the phrase "isomorphic model of"
- > whenever you can be sure that your readers understand the point.

This argument is rather like saying that one should wear a mask because people will see you face better when you take it off.

- > d) But it also allows you to adopt an agnostic attitude towards your
- > models in those cases (actually the overwhelming majority of cases)
- > for which you don't know whether an exact isomorphism exists.

It always exists, by definition. That is, every model-2 is isomorphic to a "world", ie itself.

- > e) Finally (and for many purposes this is the most important point),
- > this distinction gives you a convenient place to introduce measures
- > of approximation, granularity, and experimental error. We have all
- > expressed serious reservations (or worse) about Lotfi Z's fuzziness.
- > This separation of models from the world gives us a precise two-valued
- > logic for relating sentences to models and a continuously variable
- > range of tolerances for measuring the adequacy of a model.

Rubbish. Ahem, please forgive my rudeness. Let me rephrase that: I'll believe it when I see it. I've been trying, on and off, for almost 15 years now to create a useful nontrivial account of approximation and granularity (especially one that can be used to account for vernier effects in perceptual judgements, the usefulness of dithering, etc..) and I've not been

able to and havnt seen anyone else being able to. People are constantly claiming to have some magic way to handle it, and now you seem to think that driving an truck over Tarski's theory of truth is somehow going to solve this particular hard problem. I can't type the sound I am inclined to make.

Again, as a challenge, just for a start: show us how you propose to get past the first classical snag, the heap paradox. Remember, we need to be able to *reason* about those heaps and grains of sand; its not enough to give us a theory of heaps and then talk of grains in a metatheory.

>Some comments on your comments:

>

>>My problem with John's account of 'models' is that it seems to need two >>different ways for a sentence to be false of the world: it might be false >>in a model, or that model may be inaccurate in some way. But I can only see >>one way in which a sentence can be false, which involves giving an >>interpretation of it in some world.

>

>If you replace that last word 'world' with 'model', I would be much happier.
>I don't like to put quantifiers in front of 'world' because that suggests
>that there might be other possible worlds, which no one has ever observed
>or ever suggested how they might be observed.

Read Kripke "Naming and Necessity" for a nice analysis of what is this wrong with this way of thinking. Possible worlds arent other places, like infinitely distant galaxies. They are alternative ways this world could be.

> I agree with you that there

>is only one way that a sentence can be false. The inaccuracy of a model
>does not make a sentence false -- it gives you another metalevel for talking
>about your interpretation. In effect, we have multiple metalevels: truth
>and falsity are metalevel terms that are not part of the first-order language.
>Talking about the level of approximation goes to the metametalevel:

>

> 1. "It is true that p." is a metalevel statement about p, which at the
> metalevel happens to imply p.

>

> 2. "p is true in model M" is another metalevel statement about p.

OK, lets look at this. Where in Tarskian (or any other) semantic theory is there a requirement that the name 'M', when interpreted in (a suitable metatheory extension of) M itself, must denote M? If you are here proposing a formal language (as opposed to just doing mathematics) then you can't just stipulate what a name shall denote: you must provide a theory of (meta)truth and then see what interpretations it provides. Good luck with

the Godel sentences.

- > 3. "M is an accurate model of the world within the tolerance of our
- > measuring instruments" is a statement about M, which can be combined
- > with the previous statement to make a metametastatement about p:

I'd like to see a sketch of the instrument that can detect, say, a domain with cardinality aleph-3, or the first inaccessible ordinal. Sorry to sound sarcastic, but I'm reduced to it here. Your entire discussion is just completely confused by using an ambiguous word ('model') first in one sense and then another.

Pat

PS. On a more conciliatory note. If what worries you is referring (even in our metaln-discussions) to any 'other world' (which is ridiculous since there is only one world, etc.) or still worse to 'possible worlds', then there is an alternative way to talk. Instead of saying, (sentence) P is true in (possible) world M, say rather: M is a (possible) *way* for the sentence P to be true. This focusses on the (important!) fact that there is lot more to an interpretation than just its domain, and also that aspects of the world which are invisible to the signature of the sentence are irrelevant to the truth-conditions. This way, you can be content with the unity of the single World, and I can think to myself: this is another way to say, 'in a possible world'.

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Tue Oct 06 09:54:04 1998
Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id RAA26247
for <phayes@coginst.uwf.edu>; Mon, 5 Oct 1998 17:23:59 -0500 (CDT)
Received: from scratchy (scratchy [207.207.8.118])
by catbert.cyc.com (8.8.8/8.8.8) with SMTP id RAA15392;
Mon, 5 Oct 1998 17:19:14 -0500 (CDT)
Message-Id: <3.0.32.19981005171949.00a39818@catbert.cyc.com>
X-Sender: fritz@catbert.cyc.com
X-Mailer: Windows Eudora Pro Version 3.0 (32)
Date: Mon, 05 Oct 1998 17:19:55 -0500

To: Pat Hayes <phayes@coginst.uwf.edu>
From: Fritz Lehmann <fritz@cyc.com>
Subject: Re: Models Intermediate or not. Was Re: Tropes
Cc: Fritz Lehmann <fritz@cyc.com>, Piek.Vossen@let.uva.nl, cmenzel@tamu.edu,
doug@csi.uottawa.ca, fritz@cyc.com, p.m.simons@leeds.ac.uk,
phayes@coginst.uwf.edu, sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net

Mime-Version: 1.0

Content-Type: text/plain; charset="us-ascii"

Content-Length: 5572

Status:

At 03:03 PM 10/5/98 -0500, Pat Hayes wrote:

>Hi Fritz

>

>[John Sowa, to Pat:]

>>> You cannot solve

>>>the symbol grounding problem in a Cantorian transfinite universe.

>

>>Right. It also bothers me that the entire apparatus of that universe

>>depends on Cantor's Main Theorem (that there is no surjection of an

>>infinite set onto its powerset), whose proof is isomorphic to the Barber

>>Paradox, and whose proof fails in perfectly good set theories with

>>universal sets, like Quine's New Foundations (NF) set theory.

>

>Are you proposing that we create a universal high-level ontology which is

>based on a denial of Cantor's theorem and a bold declaration that (contrary

>to what most people have been led to believe) the set of real numbers is

>countable after all? If so, come out clearly and let's debate that

>proposal.

No, I do not propose that. But I would prefer that the Upper Ontology be as agnostic as possible about transfinite stuff, and in particular it should be agnostic about controversial set-theoretic axioms that in concert are deemed to support Cantorism. In fact the ontology doesn't need much set-theory at all, and would be safer without it. There's no need, for practical purposes, for us to choose among ZF, ZFC, VGB, NF, RTT (Ramified Types), Church, etc. When I say "it bothers me" above, that is merely the whine of an amateur who hasn't taken the time to fully understand the basics, but strongly suspects that the mainstream is misguided; different people accord different weights to such amateur reservations. (If I ever come up with a cogent refutation or alternative, I'll attempt to publish it, but meanwhile ...) Let's have axioms that are correct for finite sets, and be silent where the different systems diverge for infinite, uncountable

(if any) and universal (if any) sets. Also, let's not unduly emphasize set theory over (simpler) mereology.

>[...]

>>I generally agree with John Sowa in this (renewed) debate. I think a >>"model" can err. A model (it's annoying to have to call it "a Tarskian >>structure" especially since half of Tarski's work was the Hayes-disparaged >>algebraic approach which dispensed with them entirely) could contain the >>ground atomic assertion that Dusseldorf is in Wales. A sentence like (NOT >>(NOT (IN-REGION Dusseldorf Wales))) is true in the model, but false.

>

>(Surely you are joking, Fritz?) Of course a Tarskian model can be "false" >in this sense. Did anyone ever claim not?

Good; I mistakenly had thought you'd disagreed earlier in an email message questioning what a model being "wrong" could mean.

>[...]

>However, I'd be interested to know what semantic theory justifies your claim >that '(IN-REGION Dusseldorf Wales)' is false. I presume you mean the >brackets to indicate a first-order sentence, and that it is false because >the English sentence "Dusseldorf is in Wales" is false? That is so in >virtue of it containing two proper names, among other things. What is your >justification for claiming that a first-order constant symbol carries the >meaning of an English proper name? (Or, if you like, for claiming that the >FO constant symbol "Wales" denotes the land whose natives call Cymri?)

The formula is meant to be false because Dusseldorf is not in Wales. That fact also happens to make the English sentence "Dusseldorf is in Wales" false, but I was not referring to the English version any more than the same sentence in any of about 5000 other languages in which it could be asserted. I thought I could skip including sentences like "'Dusseldorf' denotes Dusseldorf.", etc. in my email. As you rightly point out, model theory is incapable of capturing (and is not not intended to address) what I call the "Fiction Bit", which determines whether a conjunctive-existential description of the world in fact "holds". That's why I support Sowa et al. in their view that there is a two-step relation of a formal sentence to the world: 1. Is the sentence true with respect to the model? 2. Is the model itself factual?

>Pat Hayes

>PS. I've never had the temerity to disparage any of Tarski's work.

>[...]

Maybe only by implication. Such temerity has begun to accumulate in your critique of Robert Burch's "A Peircean Reduction Thesis" because your

(somewhat persuasive, alas) attack on the "tinker-toy" or chemical-valence treatment of relations also seems to be an attack on all work in the algebra of relations: those syntactic, equational theories of relations (including the Cylindric Algebra of Henkin, Monk & Tarski) that make no reference at all to domains or "individuals" being related. The gist of your critique, as I remember it, is that it's simply artificial and pointless not to allow individuals themselves to be first-class nodes and junction-points in the directed hypergraph of a relational structure, and to quantify over such nodes. You see no need for limiting a system to equational substitution of relations; given that, I can't see what merit you would find in Tarski's last book (with S. Givant) "A Formalization of Set Theory without Variables", which explores the power and limits of such systems.

Incidentally, I wonder what such equational theories (which lack quantified variables) do to Quine's oft-quoted-by-Sowa statement: "To be is to be the value of a quantified variable."?

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

```
=====
=====
From ???@??? Tue Oct 06 09:54:07 1998
Received: from Steam.Stanford.EDU (steam.Stanford.EDU [171.64.71.20])
    by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id RAA28244
    for <phayes@coginst.uwf.edu>; Mon, 5 Oct 1998 17:42:07 -0500 (CDT)
Received: (from jmc@localhost)
    by Steam.Stanford.EDU (8.8.8/8.8.8) id PAA13847;
    Mon, 5 Oct 1998 15:37:43 -0700 (PDT)
Date: Mon, 5 Oct 1998 15:37:43 -0700 (PDT)
Message-Id: <199810052237.PAA13847@Steam.Stanford.EDU>
From: John McCarthy <jmc@Steam.Stanford.EDU>
To: fritz@cyc.com
CC: phayes@coginst.uwf.edu, fritz@cyc.com, Piek.Vossen@let.uva.nl,
    cmenzel@tamu.edu, doug@csi.uottawa.ca, fritz@cyc.com,
    p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu, sowa@west.poly.edu,
    guarino@ladseb.pd.cnr.it, chezewiz@erols.com,
    e6n1001@coe.coppin.umd.edu, skydog@pacbell.net
In-reply-to: <3.0.32.19981005171949.00a39818@catbert.cyc.com> (message from
    Fritz Lehmann on Mon, 05 Oct 1998 17:19:55 -0500)
Subject: Re: Models Intermediate or not. Was Re: Tropes
Reply-to: jmc@cs.Stanford.EDU
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
```

Content-Length: 50

Status:

Please remove me from this particular discussion.

From ???@??? Tue Oct 06 09:54:14 1998

Received: from philebus.tamu.edu (root@PHILEBUS.TAMU.EDU [165.91.161.22])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id UAA06432
for <phayes@coginst.uwf.edu>; Mon, 5 Oct 1998 20:08:54 -0500 (CDT)

Received: from philebus.tamu.edu (localhost [127.0.0.1])
by philebus.tamu.edu (8.8.5/8.8.5) with ESMTP id UAA00731;
Mon, 5 Oct 1998 20:04:59 -0500

Message-Id: <199810060104.UAA00731@philebus.tamu.edu>

X-Mailer: exmh version 2.0.2 2/24/98

To: Fritz Lehmann <fritz@cyc.com>

cc: Pat Hayes <phayes@coginst.uwf.edu>, Piek.Vossen@let.uva.nl,
doug@csi.uottawa.ca, p.m.simons@leeds.ac.uk, sowa@west.poly.edu,
guarino@ladseb.pd.cnr.it, Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>, skydog@pacbell.net

Subject: Re: Models Intermediate or not. Was Re: Tropes

In-reply-to: Your message of "Mon, 05 Oct 1998 17:19:55 CDT."
<3.0.32.19981005171949.00a39818@catbert.cyc.com>

Mime-Version: 1.0

Date: Mon, 05 Oct 1998 20:04:59 -0500

From: Chris Menzel <cmenzel@philebus.tamu.edu>

Content-Type: text/plain; charset=us-ascii

Content-Length: 923

Status:

Hi again, Fritz.

> ...In fact the ontology doesn't need much
> set-theory at all,...

That is so.

> Let's have axioms that are correct for finite sets,
> and be silent where the different systems diverge for infinite, uncountable
> (if any) and universal (if any) sets.

I'd suggest a look at Kripke-Platek set theory with urelements (KPU) that Barwise uses to such great effect in Admissible Sets and Structures. It is *much* weaker than ZF et al, but can do plenty. It is in particular missing an infinity axiom and -- the real culprit -- the power set axiom. Power sets what cranks us into the uncountable, after all. No power set, no Cantor's theorem.

> Also, let's not unduly emphasize set
> theory over (simpler) mereology.

Well, now, have a look at Lesniewski's full system and the work of his heirs like Sobocinski before you start making serious claims about the simplicity of mereology!

-chris

From ???@??? Tue Oct 06 12:52:35 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA12246;
Tue, 6 Oct 1998 12:04:19 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a02b23fe3326540@[143.88.7.118]>
In-Reply-To: <3.0.32.19981005171949.00a39818@catbert.cyc.com>
Mime-Version: 1.0
Date: Tue, 6 Oct 1998 12:01:17 -0500
To: Fritz Lehmann <fritz@cyc.com>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Models Intermediate or not. Was Re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>, skydog@pacbell.net
Content-Type: text/enriched; charset="us-ascii"
Content-Length: 7857
Status:

<x-rich>Hi Fritz

>>

>>[John Sowa, to Pat:]

>>>> You cannot solve

>>>>the symbol grounding problem in a Cantorian transfinite universe.

>>[Fritz:]

>>>Right. It also bothers me that the entire apparatus of that universe

>>>depends on Cantor's Main Theorem (that there is no surjection of an

>>>infinite set onto its powerset), whose proof is isomorphic to the Barber

>>>Paradox, and whose proof fails in perfectly good set theories with

>>>universal sets, like Quine's New Foundations (NF) set theory.

>>[Pat:]

>>Are you proposing that we create a universal high-level ontology which is

>>based on a denial of Cantor's theorem and a bold declaration that (contrary

>>to what most people have been led to believe) the set of real numbers is

>>countable after all? If so, come out clearly and let's debate that

>>proposal.

>[Fritz:]

>No, I do not propose that. But I would prefer that the Upper Ontology be

>as agnostic as possible about transfinite stuff, and in particular it

>should be agnostic about controversial set-theoretic axioms.....

By all means; but (just to try to keep the issues clear) you have now moved to a different topic. John and I were arguing about the proper way to understand a metatheory of truth, not about how to formalize this (or parts of it) in an ontology. (I tend to agree with Chris that NF is a rather peculiar formal set theory and that ZF is more widely accepted for good reasons, but that's a completely different area of discussion.)

I entirely agree with you that the upper ontology should be as agnostic as possible. One way to make that agnosticism very plain is to show

that our ontological axioms have interpretations which are finite and also interpretations which are infinite, perhaps to various degrees. That might almost be a *definition* of 'being agnostic', in fact; but its a way that would be ruled out a priori if we were to adopt John's metatheory, since all interpretations would then be finite, by definition.

>

>>[...]

>>However, I'd be interested to know what semantic theory justifies your claim

>>that '(IN-REGION Dusseldorf Wales)' is false. I presume you mean the

>>brackets to indicate a first-order sentence, and that it is false because

>>the English sentence "Dusseldorf is in Wales" is false? That is so in

>>virtue of it containing two proper names, among other things. What is your

>>justification for claiming that a first-order constant symbol carries the

>>meaning of an English proper name? (Or, if you like, for claiming that the

>>FO constant symbol "Wales" denotes the land whose natives call Cymri?)

>

>The formula is meant to be false because Dusseldorf is not in Wales.

That's just not a good enough answer. Why do the symbols 'Dusseldorf' and 'Wales' in the formula '(IN-REGION Dusseldorf Wales)' denote Dusseldorf and Wales respectively? This really does seem like an example of the mistake that John was accusing me of: just declaring that one's symbolic names shall refer to some particular thing by a kind

of armchair ostension. Notice this is different from saying 'I shall call the set of European countries D..', since (unlike things like 'D'), proper names are not ours to define.

>.... I thought I could skip including sentences like "'Dusseldorf'
>denotes Dusseldorf.", etc. in my email.

No, you can't, because such claims are at the heart of what makes a name into something more than a mere constant symbol. You are using the naming power of the metatheory (English, in this case) and just assuming you can reflect the meaning of its names down into a formalism; but FOL (even CYCL) isn't English, and it doesn't have names in this (essentially social) sense.

As you rightly point out, model

>theory is incapable of capturing (and is not not intended to address) what

>I call the "Fiction Bit", which determines whether a

>conjunctive-existential description of the world in fact "holds".
That's

>why I support Sowa et al. in their view that there is a two-step relation

>of a formal sentence to the world: 1. Is the sentence true with respect to

>the model? 2. Is the model itself factual?

Your use of the singular here ("the model") makes me wonder if we are communicating properly. If you mean 'model' in Tarski's sense, which one are you referring to? Consistent sentences have many such models.

Let me suggest another way to phrase the 'fictive bit' question: 1.
Does the formal sentence have any interpretations which make it true?
2. Is the real world one of them?

Pat Hayes

>>PS. Ive never had the temerity to disparage any of Tarski's work.

>>[...]

>

>Maybe only by implication. Such temerity has begun to accumulate in your

>critique of Robert Burch's "A Peircean Reduction Thesis" because your

>(somewhat persuasive, alas) attack on the "tinker-toy" or chemical-valence

>treatment of relations also seems to be an attack on all work in the

>algebra of relations: those syntactic, equational theories of relations

>(including the Cylindric Algebra of Henkin, Monk & Tarski) that make no

>reference at all to domains or "individuals" being related.

Well, that was a review of Burch, not Tarski; but in any case, it wasnt an attack on relational algebra (how can anyone attack an algebra? Thats like attacking the natural numbers) but on the daft conclusions about 'essential thirdness' that some people, including Burch, wanted to extract from it. (For the record, I also noted how a purely algebraic extension - add the inverse operator for Peirce's 'comma' - also destroyed the 'thirdness' arguments.)

.

The gist of

>your critique, as I remember it, is that it's simply artificial and

>pointless not to allow individuals themselves to be first-class nodes and

>junction-points in the directed hypergraph of a relational structure,
and

>to quantify over such nodes.

Not exactly. I showed that given any such graph (not hypergraph: if we allow hypergraphs then the 'essential thirdness' proof fails immediately; that's my 'algebraic extension' point), we can *interpret* its identity nodes as existentially quantified variables. The very same graph can be looked at in either way: as expressing quantification over a language of (binary) predicates, or as a purely relational expression involving (at least one ternary) identity. The quantification isn't over the nodes, but over the individuals that the identities are supposed to hold between.

>..... You see no need for limiting a system to

>equational substitution of relations; given that, I can't see what merit

>you would find in Tarski's last book (with S. Givant) "A Formalization of

>Set Theory without Variables", which explores the power and limits of such

>systems.

Well, I view it as any other piece of pure mathematics. I don't find it very interesting; but that's not a disparagement of it. However, lest people get too excited about the merits of doing away with bound variables, it's worth noting that such formalisations have been tried for machine use already. See for example J. A. Robinson, 'A note on mechanising higher-order logic', Machine Intelligence 5, Edinburgh U.P. 1969. They tend to generate rather opaque expressions. For example, the sentence

$(\exists y z)(R(y z z) \ \& \ (\text{not}(\exists x)(S(y x z) \ \& \ R(x y z))))$

has the following translation in Burch's algebra:

J12(J14(J16(
I3oJ13(J17(J17(J16(I4o(R3oNEG(J13(J16(I2o(S3oR3)))))))))))))

The only way I can follow this is to think of it as kind of a knitting pattern which produces a graph isomorphic to the 1PC expression.

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax

phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

</x-rich>
From ???@??? Tue Oct 06 12:52:35 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA12246;
Tue, 6 Oct 1998 12:04:19 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a02b23fe3326540@[143.88.7.118]>
In-Reply-To: <3.0.32.19981005171949.00a39818@catbert.cyc.com>
Mime-Version: 1.0
Date: Tue, 6 Oct 1998 12:01:17 -0500
To: Fritz Lehmann <fritz@cyc.com>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Models Intermediate or not. Was Re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>, skydog@pacbell.net
Content-Type: text/enriched; charset="us-ascii"
Content-Length: 7857
Status:

<x-rich>Hi Fritz

>>

>>[John Sowa, to Pat:]

>>>> You cannot solve

>>>>the symbol grounding problem in a Cantorian transfinite universe.

>>[Fritz:]

>>>Right. It also bothers me that the entire apparatus of that universe

>>>depends on Cantor's Main Theorem (that there is no surjection of an

>>>infinite set onto its powerset), whose proof is isomorphic to the Barber

>>>Paradox, and whose proof fails in perfectly good set theories with

>>>universal sets, like Quine's New Foundations (NF) set theory.

>>[Pat:]

>>Are you proposing that we create a universal high-level ontology which is

>>based on a denial of Cantor's theorem and a bold declaration that (contrary

>>to what most people have been led to believe) the set of real numbers is

>>countable after all? If so, come out clearly and let's debate that

>>proposal.

>[Fritz:]

>No, I do not propose that. But I would prefer that the Upper Ontology be

>as agnostic as possible about transfinite stuff, and in particular it

>should be agnostic about controversial set-theoretic axioms.....

By all means; but (just to try to keep the issues clear) you have now moved to a different topic. John and I were arguing about the proper way to understand a metatheory of truth, not about how to formalize this (or parts of it) in an ontology. (I tend to agree with Chris that NF is a rather peculiar formal set theory and that ZF is more widely accepted for good reasons, but that's a completely different area of discussion.)

I entirely agree with you that the upper ontology should be as agnostic as possible. One way to make that agnosticism very plain is to show that our ontological axioms have interpretations which are finite and also interpretations which are infinite, perhaps to various degrees. That might almost be a *definition* of 'being agnostic', in fact; but it's a way that would be ruled out a priori if we were to adopt John's metatheory, since all interpretations would then be finite, by definition.

>

>>[...]

>>However, I'd be interested to know what semantic theory justifies your claim

>>that '(IN-REGION Dusseldorf Wales)' is false. I presume you mean the

>>brackets to indicate a first-order sentence, and that it is false because

>>the English sentence "Dusseldorf is in Wales" is false? That is so in

>>virtue of it containing two proper names, among other things. What is your

>>justification for claiming that a first-order constant symbol carries the

>>meaning of an English proper name? (Or, if you like, for claiming that the

>>FO constant symbol "Wales" denotes the land whose natives call Cymri?)

>

>The formula is meant to be false because Dusseldorf is not in Wales.

That's just not a good enough answer. Why do the symbols 'Dusseldorf' and 'Wales' in the formula '(IN-REGION Dusseldorf Wales)' denote Dusseldorf and Wales respectively? This really does seem like an example of the mistake that John was accusing me of: just declaring that one's symbolic names shall refer to some particular thing by a kind of armchair ostension. Notice this is different from saying 'I shall call the set of European countries D..', since (unlike things like 'D'), proper names are not ours to define.

>... I thought I could skip including sentences like "'Dusseldorf'

>denotes Dusseldorf.", etc. in my email.

No, you can't, because such claims are at the heart of what makes a name into something more than a mere constant symbol. You are using the naming power of the metatheory (English, in this case) and just assuming you can reflect the meaning of its names down into a formalism; but FOL (even CYCL) isn't English, and it doesn't have names in this (essentially social) sense.

As you rightly point out, model

>theory is incapable of capturing (and is not not intended to address) what

>I call the "Fiction Bit", which determines whether a

>conjunctive-existential description of the world in fact "holds". That's

>why I support Sowa et al. in their view that there is a two-step relation

>of a formal sentence to the world: 1. Is the sentence true with

respect to

>the model? 2. Is the model itself factual?

Your use of the singular here ("the model") makes me wonder if we are communicating properly. If you mean 'model' in Tarski's sense, which one are you referring to? Consistent sentences have many such models.

Let me suggest another way to phrase the 'fictive bit' question: 1. Does the formal sentence have any interpretations which make it true? 2. Is the real world one of them?

Pat Hayes

>>PS. Ive never had the temerity to disparage any of Tarski's work.

>>[...]

>

>Maybe only by implication. Such temerity has begun to accumulate in your

>critique of Robert Burch's "A Peircean Reduction Thesis" because your

>(somewhat persuasive, alas) attack on the "tinker-toy" or chemical-valence

>treatment of relations also seems to be an attack on all work in the

>algebra of relations: those syntactic, equational theories of relations

>(including the Cylindric Algebra of Henkin, Monk & Tarski) that make no

>reference at all to domains or "individuals" being related.

Well, that was a review of Burch, not Tarski; but in any case, it wasn't an attack on relational algebra (how can anyone attack an algebra? That's like attacking the natural numbers) but on the daft conclusions

about 'essential thirdness' that some people, including Burch, wanted to extract from it. (For the record, I also noted how a purely algebraic extension - add the inverse operator for Peirce's 'comma' - also destroyed the 'thirdness' arguments.)

.

The gist of

>your critique, as I remember it, is that it's simply artificial and

>pointless not to allow individuals themselves to be first-class nodes and

>junction-points in the directed hypergraph of a relational structure, and

>to quantify over such nodes.

Not exactly. I showed that given any such graph (not hypergraph: if we allow hypergraphs then the 'essential thirdness' proof fails immediately; that's my 'algebraic extension' point), we can *interpret* its identity nodes as existentially quantified variables. The very same graph can be looked at in either way: as expressing quantification over a language of (binary) predicates, or as a purely relational expression involving (at least one ternary) identity. The quantification isn't over the nodes, but over the individuals that the identities are supposed to hold between.

>..... You see no need for limiting a system to

>equational substitution of relations; given that, I can't see what merit

>you would find in Tarski's last book (with S. Givant) "A Formalization of

>Set Theory without Variables", which explores the power and limits of such

>systems.

Well, I view it as any other piece of pure mathematics. I don't find it

very interesting; but that's not a disparagement of it. However, lest people get too excited about the merits of doing away with bound variables, it's worth noting that such formalisations have been tried for machine use already. See for example J. A. Robinson, 'A note on mechanising higher-order logic', Machine Intelligence 5, Edinburgh U.P. 1969. They tend to generate rather opaque expressions. For example, the sentence

$(\exists x y z) (R(y z z) \ \& \ (\text{not}(\exists x)(S(y x z) \ \& \ R(x y z))))$

has the following translation in Burch's algebra:

J12(J14(J16(I3oJ13(J17(J17(J16(I4o(R3oNEG(J13(J16(I2o(S3oR3))))))))))))))

The only way I can follow this is to think of it as kind of a knitting pattern which produces a graph isomorphic to the 1PC expression.

IHMC, University of West Florida	(850)434 8903	home
11000 University Parkway	(850)474 2091	office
Pensacola, FL 32514	(850)474 3023	fax

phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

</x-rich>
From ???@??? Tue Oct 06 12:52:36 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA16477;
Tue, 6 Oct 1998 12:54:07 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a04b23ffdc5a657@[143.88.7.118]>
In-Reply-To: <3.0.32.19981006102559.00a369c0@catbert.cyc.com>
Mime-Version: 1.0

Date: Tue, 6 Oct 1998 12:51:07 -0500

To: Fritz Lehmann <fritz@cyc.com>

From: Pat Hayes <phayes@coginst.uwf.edu>

Subject: Re: Ontology of sets (was: re: Tropes)

Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net

Content-Type: text/plain; charset="us-ascii"

Content-Length: 2947

Status:

>

>I do not propose NF as the basis for set theory for our ontology; I
>recommend a practical, agnostic, non-committal ontology which allows people
>to pick any (or none) of the many, divergent-at-infinity axiomatizations of
>set theory that work OK for finite sets.

OK, lets try to construct this cottage-industry set theory. We presumably want some of the ZF axioms (maybe modified a bit) but certainly not all of them. Lets trash Choice, for a start, since it is kind of provably agnosticisable from the rest. If you want to allow finite models we should get rid of Infinity. How about Foundation? (Do we want to allow 'infinitely deep' subset chains? It has some advantages for everyday use, according to Barwise.) If we get rid of those, that leaves Pair, Sum and Extensionality (which all seem innocuous (?)) , Powerset and Separation. Powerset is the thing that generates all that Cantorian slashing, so maybe you want to leave that one out (it seems fine for finite sets, but beyond that it gets a bit murky.) Its the last that is really the difficult one, because we obviously need *some* kind of comprehension principle to connect sets with predicates, but this is where all the Russel-paradox trouble comes in; and its a schema, not an axiom. Maybe we should just not have a general comprehension schema at all, but just allow particular cases as the need arises. That doesnt keep out curiously circular definitions which produce Russelish contradictions, but it does mean that if anyone puts one in, its their fault, not ours. (We provide the naive set theory, but with "danger!" labels, and then its up to the customer not to misuse it.)

Or, we could insist that instances use a Separation-type restriction to keep the theory safe. (I can see a possible snag. It would be tempting to allow a 'universe' class which everything was an element of, to serve as a general-purpose set to seprate anything from; but this would lead to Cantor's paradox. So there wouldnt be any straightforward general strategy for using Separation.)

Any comments or other ideas?

BTW, Quine is the master of idiosyncratic set theories, but his later version (in 'set theory and its logic') seems better than the NF system. Its more elegant, and it is almost meticulous in its 'agnosticism'. As Quine says: "the weak axioms that govern the main body of the work are such as to imply the existence of none but finite classes." He even develops arithmetic without needing to postulate infinity, a trick I've always admired. It has most of the other major systems as axiomatic extensions, as well, as he shows. It needs rather a lot of care with quantifier substitution, is all.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Tue Oct 06 12:52:36 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA16477;
Tue, 6 Oct 1998 12:54:07 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a04b23ffdc5a657@[143.88.7.118]>
In-Reply-To: <3.0.32.19981006102559.00a369c0@catbert.cyc.com>
Mime-Version: 1.0
Date: Tue, 6 Oct 1998 12:51:07 -0500
To: Fritz Lehmann <fritz@cyc.com>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Ontology of sets (was: re: Tropes)
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2947
Status:

>

>I do not propose NF as the basis for set theory for our ontology; I
>recommend a practical, agnostic, non-committal ontology which allows people
>to pick any (or none) of the many, divergent-at-infinity axiomatizations of
>set theory that work OK for finite sets.

OK, lets try to construct this cottage-industry set theory. We presumably want some of the ZF axioms (maybe modified a bit) but certainly not all of them. Lets trash Choice, for a start, since it is kind of provably agnosticisable from the rest. If you want to allow finite models we should get rid of Infinity. How about Foundation? (Do we want to allow 'infinitely deep' subset chains? It has some advantages for everyday use, according to Barwise.) If we get rid of those, that leaves Pair, Sum and Extensionality (which all seem innocuous (?)) , Powerset and Separation. Powerset is the thing that generates all that Cantorian slashing, so maybe you want to leave that one out (it seems fine for finite sets, but beyond that it gets a bit murky.) Its the last that is really the difficult one, because we obviously need *some* kind of comprehension principle to connect sets with predicates, but this is where all the Russel-paradox trouble comes in; and its a schema, not an axiom. Maybe we should just not have a general comprehension schema at all, but just allow particular cases as the need arises. That doesnt keep out curiously circular definitions which produce Russelish contradictions, but it does mean that if anyone puts one in, its their fault, not ours. (We provide the naive set theory, but with "danger!" labels, and then its up to the customer not to misuse it.)

Or, we could insist that instances use a Separation-type restriction to keep the theory safe. (I can see a possible snag. It would be tempting to allow a 'universe' class which everything was an element of, to serve as a general-purpose set to separate anything from; but this would lead to Cantor's paradox. So there wouldnt be any straightforward general strategy for using Separation.)

Any comments or other ideas?

BTW, Quine is the master of idiosyncratic set theories, but his later version (in 'set theory and its logic') seems better than the NF system. Its more elegant, and it is almost meticulous in its 'agnosticism'. As Quine says: "the weak axioms that govern the main body of the work are such as to imply the existence of none but finite classes." He even develops arithmetic without needing to postulate infinity, a trick Ive always admired. It has most of the other major systems as axiomatic extensions, as well, as he shows. It needs rather a lot of care with quantifier substitution, is all.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Oct 07 10:09:46 1998
Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id SAA10050
for <phayes@coginst.uwf.edu>; Tue, 6 Oct 1998 18:29:21 -0500 (CDT)
Received: from scratchy (scratchy [207.207.8.118])
by catbert.cyc.com (8.8.8/8.8.8) with SMTP id SAA01288;
Tue, 6 Oct 1998 18:24:35 -0500 (CDT)
Message-Id: <3.0.32.19981006182420.00a124a8@catbert.cyc.com>
X-Sender: fritz@catbert.cyc.com
X-Mailer: Windows Eudora Pro Version 3.0 (32)
Date: Tue, 06 Oct 1998 18:24:27 -0500
To: Chris Menzel <cmenzel@philebus.tamu.edu>
From: Fritz Lehmann <fritz@cyc.com>
Subject: Re: Sets, was re: Tropes
Cc: Fritz Lehmann <fritz@cyc.com>, phayes@coginst.uwf.edu,
Piek.Vossen@let.uva.nl, chezewiz@erols.com, doug@csi.uottawa.ca,
e6nl001@coe.coppin.umd.edu, guarino@ladseb.pd.cnr.it,
p.m.simons@leeds.ac.uk, skydog@pacbell.net, sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2164
Status:

Dear Chris,

At 03:35 PM 10/6/98 -0500, you wrote:

>Fritz wrote:

>> >> Right. It also bothers me that the entire apparatus of that universe
>> >> depends on Cantor's Main Theorem (that there is no surjection of an
>> >> infinite set onto its powerset), whose proof is isomorphic to the Barber
>> >> Paradox, and whose proof fails in perfectly good set theories with
>> >> universal sets, like Quine's New Foundations (NF) set theory.

>> >

>> >I don't know what what criteria you use for counting theories as perfectly
>> >good, Fritz,

>>
>> At least the reference is appreciated ...
>
>No disrespect intended, Fritz!

My sense of humor is too obscure for my own good; I thought you'd been successfully amused by my handy invocation of NF -- no disrespect suspected!

>[...]
>> >[...] Furthermore, NF has never been proved consistent relative
>> >to ZF [...]
>>
>> Will that be a flaw if ZF proves to be inconsistent?

Although legitimate, this question was also supposed to be at least mildly funny.

>Certainly not, since then every theory will be (trivially) consistent
>relative to ZF. (And of course the last 90 years strongly suggests that
>the likelihood of finding a contradiction in ZF is about as high as that
>of finding one in Peano Arithmetic.)

Much effort was spent in those 90 years unsuccessfully trying to prove ZF consistent.

>[... finite sets ...]
>Right. All reasonable set theories should agree on those. Problem is,
>we want at least the natural numbers around, don't we? If so, it is
>difficult not to confront the issues of infinite sets that quickly arise
>out of our number ontology.

I thought we agreed in Heidelberg to let numbers be primitives, and not "ground" them at all in anything else. I don't know that we'll be forced us into transfiniteness. I'd prefer the ontology "not to confront the issues of infinite sets", if at all possible.

>[...]
>Highest regards,
>-chris

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

=====

====

From ???@??? Wed Oct 07 10:09:46 1998
Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id SAA10050
for <phayes@coginst.uwf.edu>; Tue, 6 Oct 1998 18:29:21 -0500 (CDT)
Received: from scratchy (scratchy [207.207.8.118])
by catbert.cyc.com (8.8.8/8.8.8) with SMTP id SAA01288;
Tue, 6 Oct 1998 18:24:35 -0500 (CDT)
Message-Id: <3.0.32.19981006182420.00a124a8@catbert.cyc.com>
X-Sender: fritz@catbert.cyc.com
X-Mailer: Windows Eudora Pro Version 3.0 (32)
Date: Tue, 06 Oct 1998 18:24:27 -0500
To: Chris Menzel <cmenzel@philebus.tamu.edu>
From: Fritz Lehmann <fritz@cyc.com>
Subject: Re: Sets, was re: Tropes
Cc: Fritz Lehmann <fritz@cyc.com>, phayes@coginst.uwf.edu,
Piek.Vossen@let.uva.nl, chezewiz@erols.com, doug@csi.uottawa.ca,
e6nl001@coe.coppin.umd.edu, guarino@ladseb.pd.cnr.it,
p.m.simons@leeds.ac.uk, skydog@pacbell.net, sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2164
Status:

Dear Chris,

At 03:35 PM 10/6/98 -0500, you wrote:

>Fritz wrote:

>> >> Right. It also bothers me that the entire apparatus of that universe
>> >> depends on Cantor's Main Theorem (that there is no surjection of an
>> >> infinite set onto its powerset), whose proof is isomorphic to the Barber
>> >> Paradox, and whose proof fails in perfectly good set theories with
>> >> universal sets, like Quine's New Foundations (NF) set theory.

>> >

>> >I don't know what what criteria you use for counting theories as perfectly
>> >good, Fritz,

>>

>> At least the reference is appreciated ...

>

>No disrespect intended, Fritz!

My sense of humor is too obscure for my own good; I thought you'd been
successfully amused by my handy invocation of NF -- no disrespect suspected!

>[...]

>> >[...] Furthermore, NF has never been proved consistent relative

>> >to ZF [...]
>>
>> Will that be a flaw if ZF proves to be inconsistent?

Although legitimate, this question was also supposed to be at least mildly funny.

>Certainly not, since then every theory will be (trivially) consistent
>relative to ZF. (And of course the last 90 years strongly suggests that
>the likelihood of finding a contradiction in ZF is about as high as that
>of finding one in Peano Arithmetic.)

Much effort was spent in those 90 years unsuccessfully trying to prove ZF consistent.

>[... finite sets ...]
>Right. All reasonable set theories should agree on those. Problem is,
>we want at least the natural numbers around, don't we? If so, it is
>difficult not to confront the issues of infinite sets that quickly arise
>out of our number ontology.

I thought we agreed in Heidelberg to let numbers be primitives, and not "ground" them at all in anything else. I don't know that we'll be forced us into transfiniteness. I'd prefer the ontology "not to confront the issues of infinite sets", if at all possible.

>[...]
>Highest regards,
>-chris

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

=====
=====

From ???@??? Wed Oct 07 13:39:57 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id KAA19168;
Wed, 7 Oct 1998 10:56:13 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a02b2413c65fab8@[143.88.7.118]>
In-Reply-To: <3.0.32.19981006182420.00a124a8@catbert.cyc.com>
Mime-Version: 1.0
Date: Wed, 7 Oct 1998 10:53:16 -0500
To: Fritz Lehmann <fritz@cyc.com>

From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Sets, was re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 1923
Status:

[Chris:]

>>[... finite sets ...]

>>Right. All reasonable set theories should agree on those. Problem is,
>>we want at least the natural numbers around, don't we? If so, it is
>>difficult not to confront the issues of infinite sets that quickly arise
>>out of our number ontology.

{Fritz:}

>I thought we agreed in Heidelberg to let numbers be primitives, and not
>"ground" them at all in anything else. I don't know that we'll be forced
>us into transfiniteness. I'd prefer the ontology "not to confront the
>issues of infinite sets", if at all possible.

>

Lets agree that we dont want to define away the natural numbers as a
set-theoretic construction. Most formal set theories all show how they can
do this as a kind of rite of passage for a foundation of matmematics, but
we arent looking for such a foundation. (Right? Anyone object?)

However, that said, Chris does have a good point. If

1. we have the natural numbers around, and
2. we allow ourselves to talk about the set of them all, and
3. we assume Powerset,

then the issue of infinite sets is right there in our faces, whether we
like it or not. Chris suggests that we should give up 3., but I think its
worth taking a look at Quine's STAIL system, which (kind of) gives up 2.
instead. We can have all the numbers, and quantify over them (and even use
what looks like set-talk about them all, but isn't really) without making
an ontological committment to the existence of that set, or indeed to
anything infinite. This seems like it might give us the best of both
worlds, since Powerset is a very valuable tool to throw away carelessly and
is quite harmless for finite sets.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Oct 07 13:39:57 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id KAA19168;
Wed, 7 Oct 1998 10:56:13 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a02b2413c65fab8@[143.88.7.118]>
In-Reply-To: <3.0.32.19981006182420.00a124a8@catbert.cyc.com>
Mime-Version: 1.0
Date: Wed, 7 Oct 1998 10:53:16 -0500
To: Fritz Lehmann <fritz@cyc.com>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Sets, was re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 1923
Status:

[Chris:]

>>[... finite sets ...]

>>Right. All reasonable set theories should agree on those. Problem is,
>>we want at least the natural numbers around, don't we? If so, it is
>>difficult not to confront the issues of infinite sets that quickly arise
>>out of our number ontology.

{Fritz:]

>I thought we agreed in Heidelberg to let numbers be primitives, and not
>"ground" them at all in anything else. I don't know that we'll be forced
>us into transfiniteness. I'd prefer the ontology "not to confront the
>issues of infinite sets", if at all possible.
>

Lets agree that we dont want to define away the natural numbers as a set-theoretic construction. Most formal set theories all show how they can do this as a kind of rite of passage for a foundation of matmematics, but we arent looking for such a foundation. (Right? Anyone object?)

However, that said, Chris does have a good point. If

1. we have the natural numbers around, and
2. we allow ourselves to talk about the set of them all, and
3. we assume Powerset,

then the issue of infinite sets is right there in our faces, whether we like it or not. Chris suggests that we should give up 3., but I think its worth taking a look at Quine's STAIL system, which (kind of) gives up 2. instead. We can have all the numbers, and quantify over them (and even use what looks like set-talk about them all, but isn't really) without making an ontological committment to the existence of that set, or indeed to anything infinite. This seems like it might give us the best of both worlds, since Powerset is a very valuable tool to throw away carelessly and is quite harmless for finite sets.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Oct 07 13:39:57 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id KAA19168;
Wed, 7 Oct 1998 10:56:13 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a02b2413c65fab8@[143.88.7.118]>
In-Reply-To: <3.0.32.19981006182420.00a124a8@catbert.cyc.com>
Mime-Version: 1.0
Date: Wed, 7 Oct 1998 10:53:16 -0500
To: Fritz Lehmann <fritz@cyc.com>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Sets, was re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,

Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 1923
Status:

[Chris:]

>>[... finite sets ...]

>>Right. All reasonable set theories should agree on those. Problem is,
>>we want at least the natural numbers around, don't we? If so, it is
>>difficult not to confront the issues of infinite sets that quickly arise
>>out of our number ontology.

{Fritz:]

>I thought we agreed in Heidelberg to let numbers be primitives, and not
>"ground" them at all in anything else. I don't know that we'll be forced
>us into transfiniteness. I'd prefer the ontology "not to confront the
>issues of infinite sets", if at all possible.
>

Lets agree that we dont want to define away the natural numbers as a
set-theoretic construction. Most formal set theories all show how they can
do this as a kind of rite of passage for a foundation of matmematics, but
we arent looking for such a foundation. (Right? Anyone object?)

However, that said, Chris does have a good point. If

1. we have the natural numbers around, and
2. we allow ourselves to talk about the set of them all, and
3. we assume Powerset,

then the issue of infinite sets is right there in our faces, whether we
like it or not. Chris suggests that we should give up 3., but I think its
worth taking a look at Quine's STAIL system, which (kind of) gives up 2.
instead. We can have all the numbers, and quantify over them (and even use
what looks like set-talk about them all, but isn't really) without making
an ontological committment to the existence of that set, or indeed to
anything infinite. This seems like it might give us the best of both
worlds, since Powerset is a very valuable tool to throw away carelessly and
is quite harmless for finite sets.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu

<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Oct 07 10:09:46 1998
Received: from mail-gw.pacbell.net (mail-gw.pacbell.net [206.13.28.25])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id TAA16178
for <phayes@coginst.uwf.edu>; Tue, 6 Oct 1998 19:56:19 -0500 (CDT)
Received: from pacbell.net (ppp-206-170-6-40.rdcy01.pacbell.net [206.170.6.40]) by
mail-gw.pacbell.net (8.8.8/8.7.1+antispam) with ESMTP id RAA21023; Tue, 6 Oct 1998
17:51:03 -0700 (PDT)
Message-ID: <361ABAE7.73789CB0@pacbell.net>
Date: Tue, 06 Oct 1998 17:50:47 -0700
From: Robert Spillers <skydog@pacbell.net>
X-Mailer: Mozilla 4.04 [en] (Win95; I)
MIME-Version: 1.0
To: Chris Menzel <cmenzel@tamu.edu>, Doug Skuce <doug@site.uottawa.ca>,
Fritz Lehmann <fritz@cyc.com>, John McCarthy <jmc@cs.stanford.edu>,
John Sowa <sowa@west.poly.edu>,
Nancy Lawler <E6NL001@coe.coppin.umd.edu>,
Pat Hayes <phayes@coginst.uwf.edu>,
Peter Simons <p.m.simons@leeds.ac.uk>,
Piek Vossen <piek.vossen@let.uva.nl>
Subject: Short paper on theory topics
Content-Type: multipart/mixed; boundary="-----DD73AC43257FA7FCAB39A054"
Content-Length: 117990
Status:

Attached are both MS Word and ps.prn files for this message.

Content-Type: application/x-unknown-content-type-prn_auto_file;
name="Shortpaper.prn"
Content-Disposition: inline; filename="Shortpaper.prn"

Attachment converted: lonestar:Shortpaper.prn (????/----) (00007441)
Attachment converted: lonestar:Topic Assignments for Short Pap (WDBN/MSWD)
(00007442)
From ???@??? Wed Oct 07 10:09:56 1998
Received: from philebus.tamu.edu (root@PHILEBUS.TAMU.EDU [165.91.161.22])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id JAA14696
for <phayes@coginst.uwf.edu>; Wed, 7 Oct 1998 09:49:02 -0500 (CDT)
Received: from philebus.tamu.edu (localhost [127.0.0.1])
by philebus.tamu.edu (8.8.5/8.8.5) with ESMTP id JAA09501;
Wed, 7 Oct 1998 09:44:42 -0500
Message-Id: <199810071444.JAA09501@philebus.tamu.edu>
X-Mailer: exmh version 2.0.2 2/24/98

To: sowa@west.poly.edu (John F. Sowa)
cc: E6NL001@coe.coppin.umd.edu, doug@site.uottawa.ca, fritz@cyc.com,
jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
piek.vossen@let.uva.nl, skydog@pacbell.net
Subject: Re: Short paper on theory topics
In-reply-to: Your message of "Tue, 06 Oct 1998 22:15:30 EDT."
<199810070215.WAA02045@west>
Mime-Version: 1.0
Date: Wed, 07 Oct 1998 09:44:41 -0500
From: Chris Menzel <cmenzel@philebus.tamu.edu>
Content-Type: text/plain; charset=us-ascii
Content-Length: 617
Status:

> Both the ps.prn and the MS word file were received as unreadable
> gibberish. I can't deal with MS word, but I can print Postscript.
>
> However, ps.prn seems to be some kind of encoded Postscript.
> Can you resend it in either plain ASCII or in a plain Postscript?

Just another datapoint, John -- Bob's attachment came to me as perfectly readable postscript. I can view it and print it with gv or ghostview (under linux), so perhaps the attachment just got munged in your case for some reason. I'll send you the raw postscript in a separate message that you can just copy and paste into a file.

-chris

From ???@??? Wed Oct 07 13:39:57 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA19298;
Wed, 7 Oct 1998 11:08:22 -0500 (CDT)
<199810070215.WAA02045@west>
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a03b2413fcc77e@[143.88.7.118]>
In-Reply-To: <199810071444.JAA09501@philebus.tamu.edu>
References: Your message of "Tue, 06 Oct 1998 22:15:30 EDT."
<199810070215.WAA02045@west>
Mime-Version: 1.0
Date: Wed, 7 Oct 1998 11:05:25 -0500
To: Chris Menzel <cmenzel@philebus.tamu.edu>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Short paper on theory topics
Cc: E6NL001@coe.coppin.umd.edu, doug@site.uottawa.ca, fritz@cyc.com,
jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,

piek.vossen@let.uva.nl, skydog@pacbell.net,
sowa@west.poly.edu (John F. Sowa)
Content-Type: text/plain; charset="us-ascii"
Content-Length: 1598
Status:

>> Both the ps.prn and the MS word file were received as unreadable
>> gibberish. I can't deal with MS word, but I can print Postscript.
>>
>> However, ps.prn seems to be some kind of encoded Postscript.
>> Can you resend it in either plain ASCII or in a plain Postscript?
>
>Just another datapoint, John -- Bob's attachment came to me as perfectly
>readable postscript. I can view it and print it with gv or ghostview
>(under linux), so perhaps the attachment just got munged in your case for
>some reason. I'll send you the raw postscript in a separate message that
>you can just copy and paste into a file.
>

Plain ASCII however is a lot more use than Postscript for a working document. I can print out postscript, but thats all (Ghostscript is just a printer-to-the-screen.) Sending a postscript file to someone is exactly as useful as sending them a piece of paper: it is readable, but has to be re-typed into ones computer in order to become a working document. It is suitable for a finished publication, not soemthing one is supposed to contribute to. So let me agree with John , and suggest that the best 'universal' format at this stage might be simply a cut-and-paste of the plain text into a mailer window.

Also, Chris, while I admire your committment to Linux, its hardly an international standard!

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Oct 07 13:39:57 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA19298;

Wed, 7 Oct 1998 11:08:22 -0500 (CDT)
<199810070215.WAA02045@west>
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a03b2413fccc77e@[143.88.7.118]>
In-Reply-To: <199810071444.JAA09501@philebus.tamu.edu>
References: Your message of "Tue, 06 Oct 1998 22:15:30 EDT."
<199810070215.WAA02045@west>
Mime-Version: 1.0
Date: Wed, 7 Oct 1998 11:05:25 -0500
To: Chris Menzel <cmenzel@philebus.tamu.edu>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Short paper on theory topics
Cc: E6NL001@coe.coppin.umd.edu, doug@site.uottawa.ca, fritz@cyc.com,
jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
piek.vossen@let.uva.nl, skydog@pacbell.net,
sowa@west.poly.edu (John F. Sowa)
Content-Type: text/plain; charset="us-ascii"
Content-Length: 1598
Status:

>> Both the ps.prn and the MS word file were received as unreadable
>> gibberish. I can't deal with MS word, but I can print Postscript.
>>
>> However, ps.prn seems to be some kind of encoded Postscript.
>> Can you resend it in either plain ASCII or in a plain Postscript?
>
>Just another datapoint, John -- Bob's attachment came to me as perfectly
>readable postscript. I can view it and print it with gv or ghostview
>(under linux), so perhaps the attachment just got munged in your case for
>some reason. I'll send you the raw postscript in a separate message that
>you can just copy and paste into a file.
>

Plain ASCII however is a lot more use than Postscript for a working document. I can print out postscript, but thats all (Ghostscript is just a printer-to-the-screen.) Sending a postscript file to someone is exactly as useful as sending them a piece of paper: it is readable, but has to be re-typed into ones computer in order to become a working document. It is suitable for a finished publication, not soemthing one is supposed to contribute to. So let me agree with John , and suggest that the best 'universal' format at this stage might be simply a cut-and-paste of the plain text into a mailer window.

Also, Chris, while I admire your committment to Linux, its hardly an international standard!

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Oct 07 13:39:58 1998
Received: from philebus.tamu.edu (root@PHILEBUS.TAMU.EDU [165.91.161.22])
 by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA21610
 for <phayes@coginst.uwf.edu>; Wed, 7 Oct 1998 11:31:43 -0500 (CDT)
Received: from philebus.tamu.edu (localhost [127.0.0.1])
 by philebus.tamu.edu (8.8.5/8.8.5) with ESMTP id LAA09951;
 Wed, 7 Oct 1998 11:27:47 -0500
Message-Id: <199810071627.LAA09951@philebus.tamu.edu>
X-Mailer: exmh version 2.0.2 2/24/98
To: Pat Hayes <phayes@coginst.uwf.edu>
Cc: E6NL001@coe.coppin.umd.edu, doug@site.uottawa.ca, fritz@cyc.com,
 p.m.simons@leeds.ac.uk, piek.vossen@let.uva.nl, skydog@pacbell.net,
 sowa@west.poly.edu (John F. Sowa)
Subject: Re: Short paper on theory topics
In-reply-to: Your message of "Wed, 07 Oct 1998 11:05:25 CDT."
 <v04003a03b2413fcc77e@[143.88.7.118]>
Mime-Version: 1.0
Date: Wed, 07 Oct 1998 11:27:46 -0500
From: Chris Menzel <cmenzel@philebus.tamu.edu>
Content-Type: text/plain; charset=us-ascii
Content-Length: 936
Status:

> Also, Chris, while I admire your committment to Linux, its hardly an
> international standard!

Well, I am an admitted linux fanatic, Pat, but of course the emphasis
there was on the use of ghostview -- I assume that if I can display the
file on my platform it will display on others as well. The reference to
linux was for John's benefit, as he is himself a linux buff.

<geek>

As to Linux's not being a standard, that is good -- no operating system
should be! Micro\$soft is doing its best, however, to get that idea into
the heads of corporate types everywhere vis-a-vis Windows. They have had

good success with this evil scheme, but the 7 million current users and > 100% growth per year of a free OS -- which now has solid corporate support from the likes of Oracle, Sybase, IBM, Intel, Netscape, and Corel -- that runs circles around their own allegedly "industrial strength" OS NT has MS scared shitless!

</geek>

-chris

From ???@??? Mon Sep 14 11:08:19 1998
Received: from LADSEB.LADSEB.PD.CNR.IT (ladseb.ladseb.pd.cnr.it [150.178.2.3])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with SMTP id NAA29391
for <phayes@coginst.uwf.edu>; Sat, 12 Sep 1998 13:04:12 -0500 (CDT)
Received: from [150.178.99.15] by LADSEB.LADSEB.PD.CNR.IT with ESMTP;
Sat, 12 Sep 1998 20:00:01 +0200
X-Sender: guarino@ladseb.ladseb.pd.cnr.it
Message-Id: <v03102805b219ade7f256@[150.178.2.93]>
Mime-Version: 1.0
Date: Sat, 12 Sep 1998 20:01:52 +0200
To: Pat Hayes <phayes@coginst.uwf.edu>
From: Nicola Guarino <guarino@ladseb.pd.cnr.it>
Subject: Tropes
Cc: sowa@west.poly.edu, axf@ksl.stanford.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, e6nl001@coe.coppin.umd.edu,
chezewiz@erols.com, jmc@cs.stanford.edu, cmenzel@tamu.edu,
skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 14204
Status:

Pat, I have found your comments on tropes very appropriate and stimulating. The issues you raised are extremely important. My feeling is that, for our practical applications, we can adopt a liberal aptitude, borrowing some hints from the trope view without addressing all the intricacies of the current philosophical debate (in particular, without committing to the assumption that tropes are the ultimate constituents of the universe and remaining rather neutral on the problem of universals). As I will discuss below, I believe that the main advantage of the theory of tropes is that it can help to understand the determinable/determinate distinction, which is of utmost importance for practical cases. It would have deserved a whole session at Heidelberg, unfortunately we had no time to discuss it.

>I don't recall this being mentioned at Heidelberg. Maybe I was away just >then. As for the similarity of the reds, this seems to be a matter for an

>ontology of colors, not tropes. According to Chris' explanation, as I
>follow it, the color of the apple is one trope, and the color of the cheek
>is another. They are *necessarily* distinct, being individuated by the
>things they are tropes of. So one can't just say that the apple's trope is
>identical to the cheek's trope. One would have to say, presumably, that the
>*color* of the apple's trope was identical to the *color* of the cheek's
>trope (?? Or maybe that they were, while not identical, similar with
>respect to color...but since they consist of nothing more than being
>instances of the color-property, its hard to see how else they could be
>similar...)

the apple's trope (i.e., the color of that apple) does not *have* a color;
it just *is* a color. The cheek's trope, even if qualitatively identical to
the apple's trope, is still numerically distinct from it. The relation of
"exact similarity" is very much like geometrical congruence: two spatial
shapes can be congruent and yet not identical.

>But then why bother with these tropes at all: one can simply
>talk about the color of the apple and the color of the cheek, just like the
>rest of the human race does.

Independently of tropes, I understand that you agree on the utility of
considering these individual colors as distinct elements of our domain. You
argue that "this seems to be a matter for an ontology of colors": I would
say, more generally, that this is a matter for an ontology of *qualities*.
The color of the cheek and the color of the apple are what I (and you) call
qualities ("individual qualities", if we want to stress their
particularity). My feeling is that this is the most relevant kind of trope
we should deal with. There are other kinds of trope, most notably events
(occurents), but their ontological status seems to be not a problem after
Davidson, and we have already included them in our ontology.

The spectrum of various possible kinds of tropes can be shown by the
following linguistic examples, which all can be seen as cases of
"particularized properties" [I deliberately avoid using "instances of
properties" - see discussion below]:

1. the color of this apple
2. the redness of this apple, intended as:
 - 2a. the kind of redness (or: the kind of red) of this apple
 - 2b. the event of this apple having been red for a certain time
 - 2c. the state of affairs corresponding to the being red of this apple
 - 2d. the fact that this apple has been red for a certain time

Just for clarification, the difference between events, states of affairs
and facts is the following, in my mind:

- an event only occurs at a specific time
- the same state of affairs can occur at different times
- a fact does not occur, it can only be either true or false.

According to my knowledge, all of the examples above occur in the literature. It is not clear whether they correspond to alternative interpretations of the notion of trope: some philosophers only see tropes as "properties at a time at a place" (opting therefore for interpretation 2b), while others (especially Campbell) seem to leave space for the other interpretations (in particular the first one), admitting for instance that the apple's trope maintains its identity while the apple moves, behaving as a continuant.

This interpretation is for me the most interesting one, since facts, events and states of affairs have already gained their decent ontological status without a specific appeal to the theory of tropes. Maybe, an advantage of the tropes theory could be the generalization of all these entities under a common label, but its real novelty lies in my opinion in giving a peculiar ontological status to entities described by the expressions 1 and 2a.

Such expressions have the form "the X of Y", where X is a determinable like color, shape, size, etc., or "the Z-ness of Y", where Z is a determinate like red, round, etc., and the "ness" suffix is a way to construct a sort of "sub-determinable" by means of a determinate. The theory of tropes states that these expressions denote particulars. The color of this apple is an instance of the class of colors; the redness of this apple right now is identical to the color of this apple right now, they are the same trope. Therefore the class "redness", i.e. the class of all the instances of a red color, is a subclass of the class "color".

Now, the crucial point raised by Pat in the statement below is that the instances of "redness" (which are qualities, namely colors) are very different from the instances of "red", which are everyday things like apples and cheeks. To make this point evident, we can label these two classes "red-color" and "red-thing": a red thing (like an apple) *is* not a red color, it just *has* a red color.

>The best sense I can make of all this, derived from thinking about the
 >example of the wisdom of Socrates, is this. Socrates is wise. OK. Now, let
 >us speak of wisdom. There is wisdom, pure and simple. Then there is a
 >particular kind, or instance (?) of wisdom, which is the particular wisdom
 >that Socrates has by virtue of being wise. (This is the trope that 'glues'
 >Socrates to Wisdom?) Similarly, whenever something, a, has a property, P,
 >ie when P(a), there is a particular instance of P which is the particular P
 >that a has.
 >Is that more or less right? Because if so, the Socrates example suggests

>that there is something wrong with it. We say that Socrates is WISE, but
>that what Socrates has is WISDOM, not WISE.

According to the discussion above, tropes are not instances of properties (at least not in the logical sense we are used to, such that Socrates is an instance of Wise). Indeed, tropes are sometimes defined in this way, but, in what is considered as the first paper on this subject, Williams is well aware of the possible confusion behind the notion of instantiation, and makes clear that we may distinguish two senses of instantiation, one for tropes and one for ordinary things. I suspect that the ambiguity is also due to the limitations of the usual logical predication, which allows for two different readings of the proposition Red(a): either "a is red" or "a is a red" (this linguistic observation does not work very much with "wise", however).

Because of this ambiguity, I prefer to regard tropes as *particularized properties* (i.e., properties particularized within things) rather than instances of properties. In other words, tropes are not instances or properties, but rather something that things must have in order to be instances of properties.

So, whenever P(a) is true, there is another particular, say b, different from a and dependent on it, that is responsible of the fact that P(a) is true. Socrates is wise because of its wisdom. Such a wisdom is a trope (a quality), but it is NOT an instance of the property "wise". This trope, glued with many other, contributes to form Socrates. The "glue" here does not so much link together the particular "Socrates" with the universal "Wisdom", but rather the particular "Socrates' wisdom" with the (much more complex) particular "Socrates". Under this view, individuals can be seen as "bundles of tropes".

>Heres another problem. We are used to talking of particular pieces of
>wisdom, and particular colors of things. But if this is a general notion
>then, presumably, all properties have tropes. How about (one of my favorite
>examples for testing intuitions) the property of being further north than
>the oldest plumber born in Philadelphia? Does this have tropes? Is *my*
>particular way of being further north different from *your* particular way
>of being further north? In general, how are the tropes of ($\lambda x. \phi$)
>related to the tropes of the things denoted by the syntactic constituents
>of ϕ ?

I should add that the account of predication I described above does not hold for arbitrary properties, not at least in Campbell's view: his theory of tropes is "sparse", in the sense that not necessarily, whenever P(a) is true, there is a P-trope that makes it true. For instance, if Socrates is a

human, this is so not because of a "humanity" trope, but because of a compresence of various tropes corresponding to the *qualities* of Socrates. I understand therefore that the correspondence only holds for those properties that are *determinates* of some *determinable*: so something is red (determinate) because of its color (determinable), and wise because of its (degree of) wisdom.

However, if we take a more general view of tropes which includes events, fact and states of affairs, we can observe that the position above is not true, in the sense that, whenever P(a) holds, we always have the corresponding event (if a is concrete), fact, or state of affairs. For instance, if we take a relational predication predication like To-the-right-of-this-table(a), there is a corresponding event, a fact, and a state of affairs, but there is no specific quality involved. I would say therefore that Campbell's view of "sparse" tropes only holds for tropes which are qualities (the most important ones, as I said).

In conclusion, I completely agree with the following:

>One must distinguish the property P from what something has when P is true
>of it. This is orthogonal to the type-instance distinction: there is
>redness (the quality), for example, just as there is red (the property),
>and the redness of this apple is an instance of redness (not of red). If a
>trope is a particular redness or wiseness or rigidity, etc. , then what
>are redness, wisdom and rigidity? I tend to think that most of these tropes
>(if that is what they are) are best thought of in particular terms with
>their own ontologies: colors are one thing, physical rigidity another, and
>wisdom something else altogether. But in any case, whatever these things
>are, they arent properties. One doesnt say "Socrates is wisdom", or write
>Wisdom(Socrates).

I hope to have shown however that this argument does not undermine the trope view. Simply, the original assumption (that tropes are instances of properties) was wrong or at least inaccurate.

>Apparently a trope can't be properly described in English. Can you,
>then, give us an example of an *inference* that involves tropes? That is,
>can you show us something that couldnt be inferred unless we have tropes in
>our ontology? (Because if you can't describe them in English and you can't
>point to any inference in which they are involved, what possible motivation
>can there be for invoking them, whatever they are?)

I agree that a general notion of trope which includes all the various interpretations I discussed above can be in fact difficult to describe in English. The single notions of quality/event/fact/state of affairs are however amenable to be described. Focusing on qualities, a very useful

inference made very simple by their explicit introduction in our ontology regards the determinable/determinate distinction: from (1) "x is a wise" and (2) "wisdom is a moral quality" we would like to infer "x has a moral quality" as well as "x has a quality". To perform such inference, we need first to link the *attribution* "wise" with the corresponding quality "wisdom"; then, we have to reformulate (2) as "any wisdom is a moral quality", since wisdom is actually a class of individual qualities. This move leaves space to other sub-qualities of wisdom (say political wisdom, etc...) with no difficulty. For instance, from the following program written is pseudo-Prolog:

```
Socrates is wise.  
wise is an attribution of wisdom.  
any wisdom is a moral quality.  
any moral quality is a quality.  
X has a Q if  
  X is Y,  
  Y is an attribution of Z,  
  any Z is a Q.
```

we can derive:

```
Socrates has a wisdom.  
Socrates has a moral quality  
Socrates has a quality.
```

>(Things get even murkier in Nicola's messages, where tropes seem to have
>all kinds of remarkable structure: temporal extent, parts, etc.. I have no
>idea how this can be reconciled with the account that Chris gives. For
>details see my response to Nicola.)

In fact, in my mind (and I believe also in Campbell's mind, but this is not clear to me) many qualities are continuants (at least, the qualities of continuants): they have a spatial location, they can move, they keep their identity while admitting changes, they *can* have parts. Consider these examples:

```
The color of a car has a specific spatial location (namely the external  
surface of the body of the car).  
The color of this apple changed from red to brown in a couple of days.  
The volume of my bedroom is part of the volume of my house.  
*The color of the door of my car is part of the color of my car.
```

Well, more on this aspect (which is a delicate one) on a next message. I have been long enough this time.

Bottom line: I suggest to avoid to use the word "trope" in our document and to speak just of qualities (as well as of events, facts and [possibly] states of affairs).

-- Nicola

Nicola Guarino
National Research Council phone: +39 49 8295751
LADSEB-CNR fax: +39 49 8295763
Corso Stati Uniti, 4 email: guarino@ladseb.pd.cnr.it
I-35127 Padova
Italy

Home page: ** updated 4/5/98 **
<http://www.ladseb.pd.cnr.it/infor/ontology/ontology.html>

From ???@??? Mon Sep 14 11:08:20 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
 by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id UAA26010
 for <phayes@coginst.uwf.edu>; Sat, 12 Sep 1998 20:23:15 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
 by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id VAA18375;
 Sat, 12 Sep 1998 21:19:23 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
 id VAA02093; Sat, 12 Sep 1998 21:14:57 -0400
Date: Sat, 12 Sep 1998 21:14:57 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199809130114.VAA02093@west>
To: guarino@ladseb.pd.cnr.it, phayes@coginst.uwf.edu
Subject: Re: Tropes
Cc: axf@ksl.stanford.edu, chezewiz@erols.com, cmenzel@tamu.edu,
 doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu, fritz@cyc.com,
 jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk, skydog@pacbell.net
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 2847
Status:

Nicola,

I think that we are beginning to converge on a useful core vocabulary

that will be acceptable to most of our customers, which includes computer scientists, philosophers, linguists, and a variety of others.

I agree that words like 'trope' are known to a small subset of those people and the benefits of using it are far outweighed by the difficulty of explaining it to the rest (and even to ourselves). But we can use those terms when referring to or quoting an author who does use them.

There are a couple of other points that have to be resolved. You mentioned that "the same state of affairs can occur at different times." However, we (at least some of us) have agreed that a state is an occurrent, which is limited to a specific region of space-time. Another occurrence of the same type of state would therefore be a different instance.

I would prefer to use the category Situation as a supertype of both Event and State, since I believe that we can give a definition of Situation that would be acceptable to both the situation semantics crowd (i.e. Barwise & Perry et al.) and to the situation calculus crowd (i.e. those who cite McCarthy & Hayes, 1969). Therefore, I propose that we adopt the word 'situation' instead of 'state of affairs' and define the category Situation as a subtype of Occurrent. We can, of course, use the term 'state of affairs' when discussing authors who use that term, but I would prefer to leave it outside the core vocabulary.

The word 'fact' raises a lot of further questions, as does the word 'proposition'. Can we define a fact as a situation that happens to exist at some place and time in the real world? If so, then we can eliminate the word as a primitive (although we can use it as a defined term, if convenient).

The word 'proposition' is very convenient for talking about the meaning of sentences in different languages (e.g. English, Italian, KIF, and CGs). We frequently say that two different sentences mean the same "thing", and that thing cannot be a fact, since we want to talk about hypothetical and false propositions as well.

My favorite definition of proposition is "an equivalence class of sentences that are derivable from one another by some agreed-upon rule". That rule would be the translation rules from one language to another (which can be formally defined in the case of formal languages, such as CGs and KIF). For natural languages, we run into the gavagai problems, but I think that we can give a precise definition of 'proposition' for formal languages, and leave the term as a useful informal term for natural languages.

This definition of 'proposition' in terms of some rule of equivalence corresponds to Church's definition of 'intension' for functions and

relations -- as indeed it should, since a proposition can be considered a 0-adic relation.

John

From ???@??? Wed Sep 16 15:49:00 1998

Received: from vapor.stanford.edu (vapor.Stanford.EDU [171.64.71.11])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id PAA29714
for <phayes@coginst.uwf.edu>; Wed, 16 Sep 1998 15:50:13 -0500 (CDT)

Received: (from jmc@localhost)
by vapor.stanford.edu (8.8.8/8.8.8) id NAA04840;
Wed, 16 Sep 1998 13:46:19 -0700 (PDT)

Date: Wed, 16 Sep 1998 13:46:19 -0700 (PDT)

Message-Id: <199809162046.NAA04840@vapor.stanford.edu>

From: John McCarthy <jmc@Steam.Stanford.EDU>

To: phayes@coginst.uwf.edu

CC: sowa@west.poly.edu, Piek.Vossen@let.uva.nl, cmenzel@tamu.edu,
doug@csi.uottawa.ca, fritz@cyc.com, p.m.simons@leeds.ac.uk,
phayes@coginst.uwf.edu, sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
chezewiz@erols.com, e6nl001@coe.coppin.umd.edu, skydog@pacbell.net

In-reply-to: <v04003a01b225dafc624b@[143.88.7.118]> (message from Pat Hayes on
Wed, 16 Sep 1998 15:32:46 -0600)

Subject: Re: Tropes

Reply-to: jmc@cs.Stanford.EDU

Mime-Version: 1.0

Content-Type: text/plain; charset=us-ascii

Content-Length: 872

Status:

My original view of situations was as states of the universe. As such, they can never be described fully, but you can say things about particular situations or situations in general. Situations are therefore **rich** entities. The use of theories with rich entities is that a person or robot can always find out more about them.

Most, maybe all, of the use that has been made of situation calculus is compatible with situations being **poor** entities, i.e. fully describable.

Pat's idea that situations must contain their pasts makes me nervous.

Another important way of looking at a particular situation calculus formalism is as an axiomatic theory applicable to whatever collections of entities, real or abstract, can benefit from it. I prefer this way of looking at situation calculus.

Whatever definition the ontologists come up with, I will probably

violate.

From ???@??? Wed Sep 16 15:49:00 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id PAA27916;
Wed, 16 Sep 1998 15:36:02 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu (Unverified)
Message-Id: <v04003a01b225dafc624b@[143.88.7.118]>
In-Reply-To: <199809161735.NAA24889@west>
Mime-Version: 1.0
Date: Wed, 16 Sep 1998 15:32:46 -0600
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2871
Status:

[John Sowa]

>I agree that the original statement of the situation calculus talked
>about a state of the entire world (or even universe), but no actual
>implementation could ever (or has ever) used more than a tiny
>(or even toy) world. In that sense, I don't believe that there
>is any practical difference between the sit. calc. view and the B&P view.
>If you know of any example where that distinction made a difference,
>I would be very interested in hearing about it.

John, as you know, this reflects an old, old source of disagreement between our views of ontology and representation. The things in the (intended) models, and the things in the implementations, are not in the same category (in my view, shared by many.) (Not all, and no doubt not even all the best thinkers; but enough to make the simple identification very confusing, which was my original point.) So while sitcalc was indeed a rather simple account of some very large things, it did nevertheless define its chosen ontology (in the old philosophical sense) to be those large things, just as Kripke semantics talks about alternative possible worlds. And this did have some practical consequences: for example, it follows from the presumed completeness of a situation that it 'contains' all its past, which is why situations always have a tree-like temporal structure. Without this assumption there could be 'backward' alternative pasts as well, and many of

the sitcalc axiomatizations would become immediately false. B&P situations make no such assumption.

In the other direction, a crucial aspect of B&P's theory is that there is no denotation relation between a descriptive language and the thing described, so that among other things there are no quantified variables. Instead, a Situation is supposed to contain "indeterminates", which are first-class entities which just *are* indeterminate, like clouds of ghostly ur-things hovering around waiting to be born (but not by being bound to a free variable, because there aren't any variables: the things that seem like variables are *names* of indeterminates, not indeterminate names.) B&P Situation theory, in fact, is really quite amazingly peculiar, and can't be given a semantics in the ordinary Tarski sense. I'm not sure if it can even be given a mathematical description within set theory.

>I agree that if we use the term 'situation' formally, we must give
>our own definition, but we can and should give some informal discussion
>about how it relates to other common uses.

I agree, as long as we don't casually say things that are in fact deeply misleading.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Thu Sep 17 10:36:18 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id EAA15763
for <phayes@coginst.uwf.edu>; Thu, 17 Sep 1998 04:03:22 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id EAA23960;
Thu, 17 Sep 1998 04:59:13 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
id EAA29734; Thu, 17 Sep 1998 04:54:46 -0400
Date: Thu, 17 Sep 1998 04:54:46 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199809170854.EAA29734@west>
To: jmc@cs.Stanford.EDU, phayes@coginst.uwf.edu

Subject: Re: Tropes

Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@tamu.edu,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu, fritz@cyc.com,
guarino@ladseb.pd.cnr.it, p.m.simons@leeds.ac.uk, skydog@pacbell.net

Mime-Version: 1.0

Content-Type: text/plain; charset=us-ascii

Content-Length: 1281

Status:

John and Pat,

In manz of our previous discussions, I have consistentlz made a sharp distinction between the model and the aspect of the world to which that model is applied. (Sorrz about the Z's -- I'm using a German kezboard, in which the Z's and Y's are interchanged.)

I prefer to identify the engineering sense of 'model' with the Tarskian sense of 'model' and I distinguish the model from the real world. As my motto, I adopt the engineers' slogan, "All models are wrong, but some are useful." In that sense, the formal language (i.e. logic) makes true or false statements about some formal model, which is at best an approximation to some aspect of the present world or of some past, future, or hypothetical aspect.

I agree with John that 'history' and 'situation' should be kept separate.

In any case, this discussion emphasizes my point: there is no universally accepted definition of 'situation', but there are many different defintions that all share some family resemblances -- to use a deprecated term, 'sitation' is a natural kind term.

I believe that it is a useful word, which we should define in a technical sense for the purpose of the report we are writing. But we should recogniye that we have no authoritz to enforce our definition on anzone else.

John

From ???@??? Fri Sep 18 10:33:40 1998

Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id FAA15991

for <phayes@coginst.uwf.edu>; Fri, 18 Sep 1998 05:09:07 -0500 (CDT)

Received: from west (west.poly.edu [128.238.20.21])

by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id GAA25403;

Fri, 18 Sep 1998 06:05:15 -0400 (EDT)

Received: by west (SMI-8.6/SMI-SVR4)

id GAA06879; Fri, 18 Sep 1998 06:00:41 -0400
Date: Fri, 18 Sep 1998 06:00:41 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199809181000.GAA06879@west>
To: phayes@coginst.uwf.edu
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@tamu.edu,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu, fritz@cyc.com,
guarino@ladseb.pd.cnr.it, p.m.simons@leeds.ac.uk, skydog@pacbell.net
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 1190
Status:

Pat,

Bz all means, I want to deal with that "rich loamy texture of the real world", and I would love to be able to capture that texture in our models. But the simple fact is that no one -- no logician, no philosopher, no engineer, and no ontologist has ever come close.

You may want to "postulate" that some model M *is* the real world. But that postulation is an empty claim -- it has no operational translation to actual measurements and actions that can be performed on that rich loam. The engineers are the most pragmatic of all the groups cited above, and they make a clear distinction between models and the world for a very good reason: there is no operational meaning to the claim that some model M *is* the world. The best you can say is that some model M is a useful approximation to the world for some purpose. To claim anything else confuses the issue more than it helps.

Sorry, I was being sarcastic when I said that 'situation' is a natural kind term. What I meant is that if we want to use it effectively, we must give a formal definition that takes it out of the natural kind category. And I would like that definition to have some clear operational meaning.

John

From ???@??? Mon Sep 21 15:30:28 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id OAA06102;
Mon, 21 Sep 1998 14:18:44 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a07b22856515584@[143.88.7.118]>
In-Reply-To: <199809181000.GAA06879@west>
Mime-Version: 1.0

Date: Mon, 21 Sep 1998 14:15:34 -0600
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2658
Status:

John, we are rehearsing an old dispute, but let me give my old reply to your engineering worries. You are here confusing two senses of 'model'. As used by engineers (and almost everyone outside of model theory) this means a simplified version or description of something else, often used to provide understanding or enable predictions. A model of a bridge might be a small-scale physical replica, or it might be a simplified description of the bridge in terms of, say, rigid bars hinged together, used to predict stresses at joints. In this sense of model - call it model-1 - a model is presumed to be physically or computationally accessible, as it were: it gives answers to the modeller who designed it.

In Tarskian model theory, however, the meaning is almost the opposite. TMT is a theory of truth: of how a sentence might be true of something. Here, models (interpretations, strictly, or call them model-2s) are ways the world might be, and the central issue is, given a description, what worlds would make it true or false? or, put another way, what worlds does it rule out, if claimed to be true? There is no assumption that these 'ways the world might be' are in any sense accessible or computable or manipulable (they might have uncomputable or infinite domains, for example) and they are not themselves subject to further interpretation: they *are* the interpretations, the realities against which the sentences are measured, as it were. One would not, in general, expect that a model-2 be accessible or computable (although some may be): to claim this would, among other things, deny Godel's theorems. Similarly, model-2s are not designed by anyone, and they don't deliver answers. They are just ways the world might be. They can be anything that obeys the semantic constraints of the truth theory, and all it requires the universe to be is a set. Any set will do.

This confusion of nomenclature is particularly unfortunate since the meanings are not just different, but often opposite. If - as often in our field - a model-1 is an axiomatic description, then when we apply Tarskian semantics, the world being modelled-1 is a model-2 of that model-1. The

relationships 'being a model-1 of' and 'being a model-2 of' are then exact inverses, if the model-1 is accurate, ie true, and if we establish this by using logical semantics, ie Tarskian, model-2-theory.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Mon Sep 21 15:30:28 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id OAA06102;
Mon, 21 Sep 1998 14:18:44 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a07b22856515584@[143.88.7.118]>
In-Reply-To: <199809181000.GAA06879@west>
Mime-Version: 1.0
Date: Mon, 21 Sep 1998 14:15:34 -0600
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2658
Status:

John, we are rehearsing an old dispute, but let me give my old reply to your engineering worries. You are here confusing two senses of 'model'. As used by engineers (and almost everyone outside of model theory) this means a simplified version or description of something else, often used to provide understanding or enable predictions. A model of a bridge might be a small-scale physical replica, or it might be a simplified description of the

bridge in terms of, say, rigid bars hinged together, used to predict stresses at joints. In this sense of model - call it model-1 - a model is presumed to be physically or computationally accessible, as it were: it gives answers to the modeller who designed it.

In Tarskian model theory, however, the meaning is almost the opposite. TMT is a theory of truth: of how a sentence might be true of something. Here, models (interpretations, strictly, or call them model-2 s) are ways the world might be, and the central issue is, given a description, what worlds would make it true or false? or, put another way, what worlds does it rule out, if claimed to be true? There is no assumption that these 'ways the world might be' are in any sense accessible or computable or manipulable (they might have uncomputable or infinite domains, for example) and they are not themselves subject to further interpretation: they *are* the interpretations, the realities against which the sentences are measured, as it were. One would not, in general, expect that a model-2 be accessible or computable (although some may be): to claim this would, among other things, deny Godel's theorems. Similarly, model-2's are not designed by anyone, and they don't deliver answers. They are just ways the world might be. They can be anything that obeys the semantic constraints of the truth theory, and all it requires the universe to be is a set. Any set will do.

This confusion of nomenclature is particularly unfortunate since the meanings are not just different, but often opposite. If - as often in our field - a model-1 is an axiomatic description, then when we apply Tarskian semantics, the world being modelled-1 is a model-2 of that model-1. The relationships 'being a model-1 of' and 'being a model-2 of' are then exact inverses, if the model-1 is accurate, ie true, and if we establish this by using logical semantics, ie Tarskian, model-2-theory.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Sep 30 11:06:56 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA18743
for <phayes@coginst.uwf.edu>; Sun, 27 Sep 1998 11:18:15 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id MAA07375;
Sun, 27 Sep 1998 12:14:23 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
id MAA04962; Sun, 27 Sep 1998 12:09:22 -0400
Date: Sun, 27 Sep 1998 12:09:22 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199809271609.MAA04962@west>
To: phayes@coginst.uwf.edu
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@tamu.edu,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu, fritz@cyc.com,
guarino@ladseb.pd.cnr.it, jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk,
skydog@pacbell.net
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 696
Status:

Pat,

I am not *confusing* the two senses of 'model'. I am deliberately *equating* them. In Ch 1 of my 1984 book, I quoted Carl Adam Petri, of net fame, who said that there was a deep equivalence between the two. He was right, and I built on that point in my 1984 book, my forthcoming 1999 book, and my publications in between.

Model theory as the logicians have defined it includes only the mathematical part. The engineers have a better idea: they formulate both mathematical models (which are equivalent to the logicians' versions) and physical models (which are far superior to the logicians' versions by including precise methods of measurement and error evaluation and control).

John

From ???@??? Wed Sep 30 17:31:09 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id QAA29538;
Wed, 30 Sep 1998 16:50:57 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a06b2384d2fff6f@[143.88.7.118]>
In-Reply-To: <199809271609.MAA04962@west>
Mime-Version: 1.0
Date: Wed, 30 Sep 1998 16:47:50 -0500
To: sowa@west.poly.edu (John F. Sowa)

From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 4059
Status:

>Pat,
>
>I am not *confusing* the two senses of 'model'. I am deliberately
>*equating* them.

John,

Confusion is no less confused when it is deliberate. You can't 'equate' them because they refer to different things. A Tarskian 'model' relates a *sentence* to something - lets call it a structure - and defines truthconditions for the sentence relative to the structure. We might debate endlessly about the real nature of these structures, but its just a fact that they are not sentences (they have no syntax, for example) and have no truthvalues and hence no truth-conditions. So the relation of these structures to something else *cannot* be the same as the relation between sentences and them.

Just declaring that you propose to equate fish with wildebeeste doesnt alter the facts of biology, even if you write a book about it.

Look, consider what you have written:

>Model theory as the logicians have defined it includes only the mathematical
>part. The engineers have a better idea: they formulate both mathematical
>models (which are equivalent to the logicians' versions) and physical models
>(which are far superior to the logicians' versions by including precise
>methods of measurement and error evaluation and control).

Take a real example, a modern windtunnel where scale models of whole aircraft are used to measure lift/drag ratios. (Built by NACA in the thirties.) This sense of 'model' (my model-2) is a *physical* thing which is *built* to be *similar* to the physical thing being modelled in some crucial way that enables one to *extrapolate* from *measurements* made during *experiments* *using* the model to *values* for the thing being

modelled. Now consider the highlighted words in the previous sentence and ask how many of them apply to the Tarskian case. Answer: none.

It's not that this is "superior" to the 'logician's version' ; it is utterly dissimilar to it, in almost every respect. This kind of modelling relies on physical similarity between model and thing modelled; it makes sense only in the context of measurements and experimentation (on the model), and it can succeed only when some overarching theory (of both model and thing modelled) can account for how the results are to be extrapolated. (Many of the early problems with flight arose from the lack of a proper understanding of Reynolds number effects, and hence improper extrapolations of what measurements were available.)

None of this applies to a theory of truth. There is no similarity between a sentence and what it describes, no experiments are performed on sentences, no measurements are taken, no extrapolations are made, and hence no overarching physical theory is required to explain how to make it. We don't experiment on sentences to discover anything about the world they describe. We don't wire up sentences with measuring instruments and put them in windtunnels to find out what happens to them. Sentences are used to convey meaning, not to discover facts. Ideas like precision and statistical error simply don't apply to the relationship between sentences and what they refer to (unless they actually talk of such matters, of course.) If I look up and say "The sky is clear tonight", the truth or falsity of what I say depends on what I mean by 'clear'; it can't be determined by *measuring* anything, and it certainly can't be determined by measuring the sentence itself.

You may disagree; but the main point I want to make (which, unlike the above, is relevant to our ontological task) is that your position is, at best, idiosyncratic. We shouldn't incorporate such a very particular view of meaning into the basic fabric of a high-level ontology which should be useable, as far as possible, by people with all kinds of metaphysical prejudices.

Pat Hayes

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Sep 30 17:31:09 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id QAA29538;
Wed, 30 Sep 1998 16:50:57 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a06b2384d2fff6f@[143.88.7.118]>
In-Reply-To: <199809271609.MAA04962@west>
Mime-Version: 1.0
Date: Wed, 30 Sep 1998 16:47:50 -0500
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 4059
Status:

>Pat,
>
>I am not *confusing* the two senses of 'model'. I am deliberately
>*equating* them.

John,

Confusion is no less confused when it is deliberate. You can't 'equate'
them because they refer to different things. A Tarskian 'model' relates a
sentence to something - lets call it a structure - and defines
truthconditions for the sentence relative to the structure. We might debate
endlessly about the real nature of these structures, but its just a fact
that they are not sentences (they have no syntax, for example) and have no
truthvalues and hence no truth-conditions. So the relation of these
structures to something else *cannot* be the same as the relation between
sentences and them.

Just declaring that you propose to equate fish with wildebeeste doesnt
alter the facts of biology, even if you write a book about it.

Look, consider what you have written:

>Model theory as the logicians have defined it includes only the mathematical

>part. The engineers have a better idea: they formulate both mathematical
>models (which are equivalent to the logicians' versions) and physical models
>(which are far superior to the logicians' versions by including precise
>methods of measurement and error evaluation and control).

Take a real example, a modern windtunnel where scale models of whole aircraft are used to measure lift/drag ratios. (Built by NACA in the thirties.) This sense of 'model' (my model-2) is a *physical* thing which is *built* to be *similar* to the physical thing being modelled in some crucial way that enables one to *extrapolate* from *measurements* made during *experiments* *using* the model to *values* for the thing being modelled. Now consider the highlighted words in the previous sentence and ask how many of them apply to the Tarskian case. Answer: none.

It's not that this is "superior" to the 'logician's version' ; it is utterly dissimilar to it, in almost every respect. This kind of modelling relies on physical similarity between model and thing modelled; it makes sense only in the context of measurements and experimentation (on the model), and it can succeed only when some overarching theory (of both model and thing modelled) can account for how the results are to be extrapolated. (Many of the early problems with flight arose from the lack of a proper understanding of Reynolds number effects, and hence improper extrapolations of what measurements were available.)

None of this applies to a theory of truth. There is no similarity between a sentence and what it describes, no experiments are performed on sentences, no measurements are taken, no extrapolations are made, and hence no overarching physical theory is required to explain how to make it. We don't experiment on sentences to discover anything about the world they describe. We don't wire up sentences with measuring instruments and put them in windtunnels to find out what happens to them. Sentences are used to convey meaning, not to discover facts. Ideas like precision and statistical error simply don't apply to the relationship between sentences and what they refer to (unless they actually talk of such matters, of course.) If I look up and say "The sky is clear tonight", the truth or falsity of what I say depends on what I mean by 'clear'; it can't be determined by *measuring* anything, and it certainly can't be determined by measuring the sentence itself.

You may disagree; but the main point I want to make (which, unlike the above, is relevant to our ontological task) is that your position is, at best, idiosyncratic. We shouldn't incorporate such a very particular view of meaning into the basic fabric of a high-level ontology which should be useable, as far as possible, by people with all kinds of metaphysical prejudices.

Pat Hayes

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Thu Oct 01 11:41:34 1998

Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id VAA16346
for <phayes@coginst.uwf.edu>; Wed, 30 Sep 1998 21:13:20 -0500 (CDT)

Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id WAA12153;
Wed, 30 Sep 1998 22:09:21 -0400 (EDT)

Received: by west (SMI-8.6/SMI-SVR4)
id WAA26032; Wed, 30 Sep 1998 22:04:15 -0400

Date: Wed, 30 Sep 1998 22:04:15 -0400

From: sowa@west.poly.edu (John F. Sowa)

Message-Id: <199810010204.WAA26032@west>

To: phayes@coginst.uwf.edu

Subject: Re: Tropes

Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@tamu.edu,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu, fritz@cyc.com,
guarino@ladseb.pd.cnr.it, jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk,
skydog@pacbell.net

Mime-Version: 1.0

Content-Type: text/plain; charset=us-ascii

Content-Length: 4250

Status:

Pat,

I think that we understand each other very well, and there is no need to continue the fruitless attempt to persuade one another of the "correctness" of some metaphysical position.

The main issue that should concern us is the adoption of a suitable core vocabulary in which we can define the ontology and discuss its implications. The use of terms outside the core is appropriate in the informal discussions, especially when quoting or discussing positions taken by other authors.

>.... A Tarskian 'model' relates a
>*sentence* to something - lets call it a structure - and defines
>truthconditions for the sentence relative to the structure. We might debate
>endlessly about the real nature of these structures, but its just a fact
>that they are not sentences (they have no syntax, for example) and have no
>truthvalues and hence no truth-conditions. So the relation of these
>structures to something else *cannot* be the same as the relation between
>sentences and them.

I agree completely with this statement. The term 'model' by the way was not used by Tarski, and it is also avoided by many contemporary logicians, who prefer the word 'interpretation' for the relationship between a sentence and one of these structures (cf. Genesereth & Nilsson for a widely-used book in AI that avoids the word 'model').

>Take a real example, a modern windtunnel where scale models of whole
>aircraft are used to measure lift/drag ratios. (Built by NACA in the
>thirties.) This sense of 'model' (my model-2) is a *physical* thing which
>is *built* to be *similar* to the physical thing being modelled in some
>crucial way that enables one to *extrapolate* from *measurements* made
>during *experiments* *using* the model to *values* for the thing being
>modelled. Now consider the highlighted words in the previous sentence and
>ask how many of them apply to the Tarskian case. Answer: none.

Precisely. This is an example of a physical model. The engineers use the term 'model' for two different, but related things: a mathematical model consisting of symbolic expressions, which is quite suitable for serving as one of the "structures" used to interpret the truth values of a sentence, and a physical model whose measurements conform to the mathematical model as closely as the construction techniques and measuring instruments permit. That physical model is related by similarity to the existing or planned physical system of interest.

Even in computer science, the engineering terminology is the most widely used. See the article by Mark Fox & Co. in the latest issue of *_AI Magazine_* on "Enterprise Modeling". The ANSI organization which started this ontology project has been working on Conceptual Schema Modeling Facilities (CSMF), which uses the term 'model' in the same sense as Fox & Co. And recently, the object-oriented community has been coalescing around a notation called UML (Universal Modeling Language), which consists of diagrams and symbolic expressions used to specify mathematical models in the engineering sense.

>None of this applies to a theory of truth. There is no similarity between a
>sentence and what it describes, no experiments are performed on sentences,

>no measurements are taken, no extrapolations are made, and hence no
>overarching physical theory is required to explain how to make it. We don't
>experiment on sentences to discover anything about the world they describe...

I agree completely with this statement and the continuation, which I deleted.
The relation of interpretation between a sentence and the mathematical
structure we have discussed above is very different from the relationship
of measurement and approximation between the mathematical structure and
the physical structure.

Suggestion: Most of the "customers" for our ontology efforts will be
familiar with the term 'model' in the engineering sense or the computer
systems sense, as in Mark Fox's usage or the ANSI CSMF usage. Therefore,
I suggest that we use the term 'interpretation' for the relation between
a sentence and a Tarskian-style mathematical structure. If we use the
word 'model', it will be necessary to remember that many, if not most
of our readers will not be familiar with the model-theoretic sense.

John

From ???@??? Thu Oct 01 13:57:46 1998

Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA11232;
Thu, 1 Oct 1998 12:17:02 -0500 (CDT)

X-Sender: phayes@mail.coginst.uwf.edu

Message-Id: <v04003a02b23963169962@[143.88.7.118]>

In-Reply-To: <199810010204.WAA26032@west>

Mime-Version: 1.0

Date: Thu, 1 Oct 1998 12:13:57 -0500

To: sowa@west.poly.edu (John F. Sowa)

From: Pat Hayes <phayes@coginst.uwf.edu>

Subject: Re: Tropes

Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net

Content-Type: text/plain; charset="us-ascii"

Content-Length: 3699

Status:

Hi John.

You wrote:

>I think that we understand each other very well, and there is no need
>to continue the fruitless attempt to persuade one another of the

>"correctness" of some metaphysical position.

It seems we can agree on this.

>.... The term 'model' by the way

>was not used by Tarski, and it is also avoided by many contemporary

>logicians, who prefer the word 'interpretation' for the relationship

>between a sentence and one of these structures (cf. Genesereth & Nilsson

>for a widely-used book in AI that avoids the word 'model').

Yes, I prefer 'interpretation'. We may need a word for 'an interpretation which makes the sentence in question true', and 'model' is widely used for that purpose in model theory. But our correspondence itself illustrates how dangerous it is to use such a word carelessly, so by all means let us agree to make a careful distinction between (engineering) models and (Tarsk/Kripk/Monatgu-ian semantic) interpretations.

>..... The engineers use

>the term 'model' for two different, but related things: a mathematical

>model consisting of symbolic expressions, which is quite suitable for

>serving as one of the "structures" used to interpret the truth values

>of a sentence, and a physical model whose measurements conform to the

>mathematical model as closely as the construction techniques and

>measuring instruments permit. That physical model is related by

>similarity to the existing or planned physical system of interest.

At first glance you seem here to be simply repeating what I have said, but just for clarification, I don't agree with the claim implicit in the clause " which is quite suitable for

>serving as one of the "structures" used to interpret the truth values

>of a sentence".

Mathematical models do not denote sentences, they *consist* of sentences, often in the form of equations, and the process of 'using' them consists in drawing conclusions from these equations, ie of performing inferences (albeit in some particularly efficient way using special-purpose circuitry.) Similarly, the sense used by Fox & co refers to modelling *languages* (my emphasis), so clearly we are here talking about lexical items with a syntax. Indeed, these languages are often syntactic variations of predicate logic, as you know. So in this sense of 'model', the 'models' are sentences, rather than the structures used to interpret those sentences. (Unless of course they are Herbrand structures.)

>Suggestion: Most of the "customers" for our ontology efforts will be

>familiar with the term 'model' in the engineering sense or the computer

>systems sense, as in Mark Fox's usage or the ANSI CSMF usage. Therefore,

>I suggest that we use the term 'interpretation' for the relation between

>a sentence and a Tarskian-style mathematical structure. If we use the
>word 'model', it will be necessary to remember that many, if not most
>of our readers will not be familiar with the model-theoretic sense.

Fine, I can go along with that, with one caveat. Your use of 'mathematical'
here suggests you may intend it to be a classifying adjective, to contrast
with 'nonmathematical structures'. If so, I register a disagreement.
However this is an internal dispute, and we can still agree to maintain the
external vocabulary to try to avoid confusion. At any rate, we should
strive to be clear about what we mean at all times. (Hey, whats wrong with
ecumenicalism?)

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Thu Oct 01 13:57:46 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA11232;
Thu, 1 Oct 1998 12:17:02 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a02b23963169962@[143.88.7.118]>
In-Reply-To: <199810010204.WAA26032@west>
Mime-Version: 1.0
Date: Thu, 1 Oct 1998 12:13:57 -0500
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 3699
Status:

Hi John.

You wrote:

>I think that we understand each other very well, and there is no need
>to continue the fruitless attempt to persuade one another of the
>"correctness" of some metaphysical position.

It seems we can agree on this.

>.... The term 'model' by the way
>was not used by Tarski, and it is also avoided by many contemporary
>logicians, who prefer the word 'interpretation' for the relationship
>between a sentence and one of these structures (cf. Genesereth & Nilsson
>for a widely-used book in AI that avoids the word 'model').

Yes, I prefer 'interpretation'. We may need a word for 'an interpretation which makes the sentence in question true', and 'model' is widely used for that purpose in model theory. But our correspondence itself illustrates how dangerous it is to use such a word carelessly, so by all means let us agree to make a careful distinction between (engineering) models and (Tarsk/Kripk/Monatgu-ian semantic) interpretations.

>..... The engineers use
>the term 'model' for two different, but related things: a mathematical
>model consisting of symbolic expressions, which is quite suitable for
>serving as one of the "structures" used to interpret the truth values
>of a sentence, and a physical model whose measurements conform to the
>mathematical model as closely as the construction techniques and
>measuring instruments permit. That physical model is related by
>similarity to the existing or planned physical system of interest.

At first glance you seem here to be simply repeating what I have said, but just for clarification, I don't agree with the claim implicit in the clause "which is quite suitable for
>serving as one of the "structures" used to interpret the truth values
>of a sentence".

Mathematical models do not denote sentences, they *consist* of sentences, often in the form of equations, and the process of 'using' them consists in drawing conclusions from these equations, ie of performing inferences (albeit in some particularly efficient way using special-purpose circuitry.) Similarly, the sense used by Fox & co refers to modelling *languages* (my emphasis), so clearly we are here talking about lexical items with a syntax. Indeed, these languages are often syntactic variations of predicate logic, as you know. So in this sense of 'model', the 'models' are sentences, rather than the structures used to interpret those

sentences. (Unless of course they are Herbrand structures.)

>Suggestion: Most of the "customers" for our ontology efforts will be
>familiar with the term 'model' in the engineering sense or the computer
>systems sense, as in Mark Fox's usage or the ANSI CSMF usage. Therefore,
>I suggest that we use the term 'interpretation' for the relation between
>a sentence and a Tarskian-style mathematical structure. If we use the
>word 'model', it will be necessary to remember that many, if not most
>of our readers will not be familiar with the model-theoretic sense.

Fine, I can go along with that, with one caveat. Your use of 'mathematical'
here suggests you may intend it to be a classifying adjective, to contrast
with 'nonmathematical structures'. If so, I register a disagreement.
However this is an internal dispute, and we can still agree to maintain the
external vocabulary to try to avoid confusion. At any rate, we should
strive to be clear about what we mean at all times. (Hey, whats wrong with
ecumenicalism?)

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Thu Oct 01 17:08:09 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id PAA26667;
Thu, 1 Oct 1998 15:56:28 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a0eb2398ffd5a10@[143.88.7.118]>
In-Reply-To: <199810011925.PAA00743@west>
Mime-Version: 1.0
Date: Thu, 1 Oct 1998 15:53:22 -0500
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,

Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 1831
Status:

>Re models & sentences: A Tarski-style "interpretation" consists of a set
>of "individuals" and a set of relationships that are true of those
>individuals.
>Such a structure is isomorphic to a conjunction of sentences, each of which
>consists of a single predicate symbol applied to a list of constants.
>Hintikka replaced the notion of possible world with such conjunctions of
>ground-level atoms, which he called 'model sets'. Hintikka's model sets
>are equivalent to Kripke's worlds, but without the metaphysical baggage.

They are a special case. True, the completeness theorem (actually the usual way of proving it) shows that any sentence with a true interpretation has a Herbrand interpretation (essentially an infinite conjunction) which makes exactly the same first-order sentences true. So one could argue that in any 'effective' sense, all of logic could be understood to refer only to its own syntax, in a kind of grand anti-Econian denial of the fundamental axioms of semiotics; but that seems unproductive to me, both from the philosophical and technical points of view. Also, note that it applies to any 'mathematical' model, so from this perspective, Russell's famous definition of mathematics should be altered to be "that subject where we are talking about our own words, and we can compute whether or not what we say is true." Your 'baggage' is, to me, the muscle on what would otherwise be a bare skeleton. But, every man to his taste, as they say. I'm happy for us to agree to disagree, as long as nobody frightens the horses.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Thu Oct 01 17:08:09 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id PAA26667;

Thu, 1 Oct 1998 15:56:28 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a0eb2398ffd5a10@[143.88.7.118]>
In-Reply-To: <199810011925.PAA00743@west>
Mime-Version: 1.0
Date: Thu, 1 Oct 1998 15:53:22 -0500
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 1831
Status:

>Re models & sentences: A Tarski-style "interpretation" consists of a set
>of "individuals" and a set of relationships that are true of those
>individuals.
>Such a structure is isomorphic to a conjunction of sentences, each of which
>consists of a single predicate symbol applied to a list of constants.
>Hintikka replaced the notion of possible world with such conjunctions of
>ground-level atoms, which he called 'model sets'. Hintikka's model sets
>are equivalent to Kripke's worlds, but without the metaphysical baggage.

They are a special case. True, the completeness theorem (actually the usual way of proving it) shows that any sentence with a true interpretation has a Herbrand interpretation (essentially an infinite conjunction) which makes exactly the same first-order sentences true. So one could argue that in any 'effective' sense, all of logic could be understood to refer only to its own syntax, in a kind of grand anti-Econian denial of the fundamental axioms of semiotics; but that seems unproductive to me, both from the philosophical and technical points of view. Also, note that it applies to any 'mathematical' model, so from this perspective, Russell's famous definition of mathematics should be altered to be "that subject where we are talking about our own words, and we can compute whether or not what we say is true." Your 'baggage' is, to me, the muscle on what would otherwise be a bare skeleton. But, every man to his taste, as they say. I'm happy for us to agree to disagree, as long as nobody frightens the horses.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Thu Oct 01 17:08:14 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id QAA28883
for <phayes@coginst.uwf.edu>; Thu, 1 Oct 1998 16:35:08 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id RAA13614;
Thu, 1 Oct 1998 17:30:54 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
id RAA01379; Thu, 1 Oct 1998 17:25:56 -0400
Date: Thu, 1 Oct 1998 17:25:56 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199810012125.RAA01379@west>
To: phayes@coginst.uwf.edu
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@tamu.edu,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu, fritz@cyc.com,
guarino@ladseb.pd.cnr.it, jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk,
skydog@pacbell.net, sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 1341
Status:

Pat,

I think that we both have the same goal of addressing the symbol grounding problem by relating language (formal and/or natural) to the real world:

>... Your 'baggage' is, to me, the muscle on what would otherwise
>be a bare skeleton....

That is why I put so much emphasis on the problems of measurement and approximation. That is where the real difficulties arise. I acknowledge that the "structures" that serve as the basis for determining the truth values of sentences are "bare skeletons". The muscle comes from the mechanisms of perception, measurement, etc., that map the bones to the outside world. If you just "postulate" that your "individuals" are actual

physical entities, you haven't done anything to put muscle on the bone.

You can't claim that your symbols are grounded in reality unless you supplement your Tarski-style models or interpretations with a theory of perception, measurement, experiment, or something similar. Only two logicians seriously set out to address that problem: Carnap with his *Logische Aufbau der Welt* and Nelson Goodman with his *Structure of Appearance*. And both of them found that the amount of work involved in mapping symbols to the world is much, much bigger than a Tarski-style denotation function that computes the truth value of a sentence in terms of a set of relations.

John

From ???@??? Thu Oct 01 18:06:07 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id SAA05316;
Thu, 1 Oct 1998 18:05:36 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a12b239b1eec832@[143.88.7.118]>
In-Reply-To: <199810012125.RAA01379@west>
Mime-Version: 1.0
Date: Thu, 1 Oct 1998 18:02:32 -0500
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 3451
Status:

John -

.....

>

>You can't claim that your symbols are grounded in reality unless you
>supplement your Tarski-style models or interpretations with a theory
>of perception, measurement, experiment, or something similar.

Yes, I can. To demonstrate, I will now do so. (Watch my cursor!) I hereby define D to be the set containing the books in my office (now, ie at the

time of typing.) I may be wrong about what is in the set D, and you will probably never know what is in it; but that's quite beside the point. I don't have to measure or check anything in the world, or to perceive anything, in order to simply define D to be that set. Of course, I can't compute the extension of D without doing some observing or measuring (though none of it involving a sophisticated engineering concern with niceties of vernier gauges); but that's a matter of epistemics, not definition. If I were suffering from an elaborate hallucination and there never had been such an office, then D would be the empty set, but it would still be the set I've defined it to be.

I rather enjoyed defining D. I think I'll do it again. E is the set of all the stars which will be visible to me tonight; F is the set of dishrags that I have ever used to clean up spilled Xylene; G is the set of train whistles which I ever heard while driving through Illinois. It's easy to define sets of physical things. You ought to try it some time: it's an oddly liberating experience. The secret to bear in mind is that the real world **really is there**! Carnap and Goodman are quite beside the point; try Lao Tzu.

Only

>two logicians seriously set out to address that problem: Carnap with
>his *Logische Aufbau der Welt* and Nelson Goodman with his *Structure
>of Appearance*. And both of them found that the amount of work involved
>in mapping symbols to the world is much, much bigger than a Tarski-style
>denotation function that computes the truth value of a sentence in terms
>of a set of relations.

Sigh. I've explained this to you SO many times, John, and you STILL make the same mistake. Tarski-style denotation functions are not things that you **compute**. They aren't LISP code. They are just mappings from symbols to things; not something that can (in general) be implemented. Consider: were this not so, then to even **talk** of anything being noncomputable would be a oxymoron by necessity.

I know that there is a philosophical position - a kind of militant antiplatonism - which insists that, indeed, this is right; that there is something wrong with Cantor's slash, that there are no uncountable things, that the whole business of computability is vacuous since everything real is finite and, in some sense, computable. I even have some sympathy myself for this radical position. Maybe, with work, it could be made internally coherent. But this is a very, very minority view, to put it mildly. If we are going to be able to talk to almost anyone else in the entire world and expect them to take us seriously, we had better keep quiet about this idea until someone gets the bugs worked out. In the meantime, let's be very careful to keep all traces of this extreme finitism out of our ontological

writing.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Thu Oct 01 18:06:07 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id SAA05316;
Thu, 1 Oct 1998 18:05:36 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a12b239b1eec832@[143.88.7.118]>
In-Reply-To: <199810012125.RAA01379@west>
Mime-Version: 1.0
Date: Thu, 1 Oct 1998 18:02:32 -0500
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 3451
Status:

John -

.....
>

>You can't claim that your symbols are grounded in reality unless you
>supplement your Tarski-style models or interpretations with a theory
>of perception, measurement, experiment, or something similar.

Yes, I can. To demonstrate, I will now do so. (Watch my cursor!) I hereby
define D to be the set containing the books in my office (now, ie at the
time of typing.) I may be wrong about what is in the set D, and you will

probably never know what is in it; but that's quite beside the point. I don't have to measure or check anything in the world, or to perceive anything, in order to simply define D to be that set. Of course, I can't compute the extension of D without doing some observing or measuring (though none of it involving a sophisticated engineering concern with niceties of vernier gauges); but that's a matter of epistemics, not definition. If I were suffering from an elaborate hallucination and there never had been such an office, then D would be the empty set, but it would still be the set I've defined it to be.

I rather enjoyed defining D. I think I'll do it again. E is the set of all the stars which will be visible to me tonight; F is the set of dishrags that I have ever used to clean up spilled Xylene; G is the set of train whistles which I ever heard while driving through Illinois. It's easy to define sets of physical things. You ought to try it some time: it's an oddly liberating experience. The secret to bear in mind is that the real world *really is there*! Carnap and Goodman are quite beside the point; try Lao Tsu.

Only

>two logicians seriously set out to address that problem: Carnap with
>his Logische Aufbau der Welt and Nelson Goodman with his Structure
>of Appearance. And both of them found that the amount of work involved
>in mapping symbols to the world is much, much bigger than a Tarski-style
>denotation function that computes the truth value of a sentence in terms
>of a set of relations.

Sigh. I've explained this to you SO many times, John, and you STILL make the same mistake. Tarski-style denotation functions are not things that you *compute*. They aren't LISP code. They are just mappings from symbols to things; not something that can (in general) be implemented. Consider: were this not so, then to even *talk* of anything being noncomputable would be a oxymoron by necessity.

I know that there is a philosophical position - a kind of militant antiplatonism - which insists that, indeed, this is right; that there is something wrong with Cantor's slash, that there are no uncountable things, that the whole business of computability is vacuous since everything real is finite and, in some sense, computable. I even have some sympathy myself for this radical position. Maybe, with work, it could be made internally coherent. But this is a very, very minority view, to put it mildly. If we are going to be able to talk to almost anyone else in the entire world and expect them to take us seriously, we had better keep quiet about this idea until someone gets the bugs worked out. In the meantime, let's be very careful to keep all traces of this extreme finitism out of our ontological writing.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Fri Oct 02 10:04:18 1998

Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id TAA11323
for <phayes@coginst.uwf.edu>; Thu, 1 Oct 1998 19:15:16 -0500 (CDT)

Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id UAA13720;
Thu, 1 Oct 1998 20:11:18 -0400 (EDT)

Received: by west (SMI-8.6/SMI-SVR4)
id UAA01948; Thu, 1 Oct 1998 20:06:25 -0400

Date: Thu, 1 Oct 1998 20:06:25 -0400

From: sowa@west.poly.edu (John F. Sowa)

Message-Id: <199810020006.UAA01948@west>

To: phayes@coginst.uwf.edu

Subject: Re: Tropes

Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@tamu.edu,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu, fritz@cyc.com,
guarino@ladseb.pd.cnr.it, jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk,
skydog@pacbell.net, sowa@west.poly.edu

Mime-Version: 1.0

Content-Type: text/plain; charset=us-ascii

Content-Length: 2041

Status:

Pat,

Your methods of "defining" sets have exactly the same flaw as the mathematician's method for hunting lions by "postulating" repeated partitions of the Sahara desert. You cannot build an AI system that is grounded in reality by "postulating" or "defining" variables that magically represent objects in the world.

>Sigh. Ive explained this to you SO many times, John ,and you STILL make the
>same mistake. Tarski-style denotation functions are not things that you
>*compute*. They aren't LISP code. They are just mappings from symbols to

>things; not something that can (in general) be implemented. Consider: were
>this not so, then to even *talk* of anything being noncomputable would be a
>oxymoron by necessity.

Sigh. I've explained this to you SO many times, Pat, and you STILL make the same mistake. Computer science in general and AI in particular cannot even deal with large finite sets. "Intractable" computations are finite, but exponentially growing ones. Countably infinite sets are so far beyond the capabilities of anything computable that philosophical questions about uncountable ones are beyond consideration.

>I know that there is a philosophical position - a kind of militant
>antiplatonism....

That is beside the point. In my youth, I enjoyed studying Cantor's theories, and I still feel quite comfortable with a Platonistic philosophy of mathematics. But if we want to have grounded symbols in our AI systems, we have to do more than "postulate" or "define" a mapping.

The beauty of Tarski's denotation function is that it can be computed in polynomial time by an ordinary SQL query to any relational database. To solve the symbol grounding problem, we have to show how the symbols in that database can be related to the inputs from the sensors and the outputs to the effectors.

I realize that you and I are not prepared to go out and build those sensors and effectors by ourselves, but I believe that it is our responsibility to show the engineers where to plug their devices into our framework.

John

From ???@??? Mon Oct 05 10:07:43 1998

Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id TAA20024
for <phayes@coginst.uwf.edu>; Fri, 2 Oct 1998 19:17:22 -0500 (CDT)

Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id UAA15186;
Fri, 2 Oct 1998 20:12:42 -0400 (EDT)

Received: by west (SMI-8.6/SMI-SVR4)
id UAA08254; Fri, 2 Oct 1998 20:07:32 -0400

Date: Fri, 2 Oct 1998 20:07:32 -0400

From: sowa@west.poly.edu (John F. Sowa)

Message-Id: <199810030007.UAA08254@west>

To: phayes@coginst.uwf.edu

Subject: Re: Tropes

Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@tamu.edu,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu, fritz@cyc.com,
guarino@ladseb.pd.cnr.it, jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk,
skydog@pacbell.net, sowa@west.poly.edu

Mime-Version: 1.0

Content-Type: text/plain; charset=us-ascii

Content-Length: 6616

Status:

Pat,

>They dont suffer from this flaw because, unlike your rhetorical example of
>the lions, they are not METHODS for DOING anything....

Exactly my point! This ontology project is not a project in pure mathematics.
It is a project in AI and computer science, both of which are engineering
disciplines that use mathematics to compute. Even for the pure scientific
side of AI, the only functions of interest are ones that are computable.

>Your final conclusion follows only if one equates "computable" with "within
>the scope of philosophical consideration"....

There are many issues worthy of philosophical consideration in this ontology
effort. When we are dealing with the ontology of mathematics, then I would
be happy to adopt a Platonistic stance, which is usually the most fruitful
approach to the nature of mathematical objects. But when we are trying to
define relationships between language and the world that are either
(a) psychologically realistic or (b) computationally tractable, then
the finitistic stance is the only one that makes sense. You cannot solve
the symbol grounding problem in a Cantorian transfinite universe.

>... We are arguing

>here not about those computational properties (which have to do with things
>like branching rates in search spaces) but about a *theory* of *truth*....

As I said before, I have no objection to uncountable sets for a theory
of mathematical truth. But all available evidence indicates that the
universe is finite, and what we can perceive, compute, think about, and
talk about is certainly finite. The really difficult problems of a *theory*
of *truth* are finite.

>(A side comment on something you said in an earlier message. You said that
>any Tarskian interpretation is 'isomorphic' to a countable set of ground
>atomic sentences. This isnt quite correct, technically, since....

I think that I just said "a conjunction of ground atoms" without using

the word "countable". If you want to generalize the point to an uncountable conjunction, I have no objection.

>>The beauty of Tarski's denotation function is that it can be computed
>>in polynomial time by an ordinary SQL query to any relational database.
>
>This is just a plain error. It isn't even coherent to say that Tarski's
>denotation function is computable, since it may target a noncomputable set.

As I have said before, the mathematical problem of reasoning about Platonic structures is the easy part. For that part, you don't have to dig around in the "rich loamy soil" to build your database of relations. You can just sit back in your easy chair and imagine uncountable infinities.

>What is the appropriate definition of 'computable' for, say, the set of
>galaxies?

Since the universe is finite, the number of galaxies is finite. The main issue here is not the denotation side of the problem, but the perception side, which puts the data into the database.

>... In any case, why would the process of answering a SQL query
>involve computing a denotation function? That process is one of inference;
>it manipulates sentences (of which the relational database entries are a
>simple sort, ie ground atoms.)

A Tarski-style structure is isomorphic to a relational DB in the finite case. In the infinite case, which never arises in databases obtained from the perception of physical situations, it is still isomorphic to a conjunction of possibly uncountably many ground atoms.

The definition of the denotation function is *identical* to the algorithm for evaluating an SQL query, which is identical to the algorithm for evaluating a Prolog expression that contains predicates defined only by ground-level assertions. All three of these things can be evaluated in polynomial time, where the degree of the polynomial is equal to the number of quantifiers in the expression (which may be implicit in SQL & Prolog). In fact, Tarski's definition can be optimized in the same way as SQL and Prolog by building an index so that many of the existential quantifiers can be evaluated in logarithmic time.

>The symbol grounding problem is indeed a difficult and interesting problem,
>but it is quite different from what we are talking about. Grounding has to
>do with how a system which is both reasoning about a world and causally
>embedded in it could *establish* that a name must refer to something
>'outside' in the physical world. It's an issue in robot epistemology,

>concerned with ways of knowing that a name refers. Model theory is
>concerned with ways of referring, not ways of knowing that one refers. Put
>another way, MT is concerned with how a world *could* be, grounding is
>concerned with how the perceived world *is*, and how we might know that.
>The difference is crucial.

I agree with your definitions. But they get to the crux of our dispute.
For a theory of mathematical truth, I have no quarrel with you. But
my primary argument against model theory, as a theory of truth about the
physical world, is that it fails to address the symbol grounding problem.
Every criticism that has been hurled against AI projects that relate
language to symbols in a computer can be leveled against the claim that
Tarski's denotation function "solves" the problem of defining truth.

Your claim that the "individuals" in Tarski's structures *are* the
actual physical objects begs every question that Roger Schank begs
when he claims that his GENSYMs *stand for* the actual physical objects.
Using the verb 'are' instead of 'stand for' in no way solves the symbol
grounding problem. Unlike Roger Schank, I am willing to agree that
Tarski's approach solves one half of the problem, but you can't solve
the symbol grounding problem by postulating it away.

>I can't agree that we have a responsibility to provide grounding sockets
>for the engineers to plug their sensors into, because (1) I don't think this
>is central to our task - we aren't in the robot-building game here - and (2)
>it's much too heavy a burden, since we can't be expected to provide
>ready-made solutions to problems that haven't even been adequately formulated
>yet. Grounding is an open problem, and we don't even know if the current -
>very partial - ideas about it really work.

I agree with you completely on this point. The point I was trying to make
is that the database is the boundary between the data that Tarski's
algorithm applies to and the data that the robot builders generate.
The engineers are responsible for filling up the database with ground
atoms about what is in the world, and the denotation function uses SQL
to determine whether a particular sentence is true or false.

John

From ???@??? Mon Oct 05 10:07:46 1998

Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id TAA20092
for <phayes@coginst.uwf.edu>; Fri, 2 Oct 1998 19:23:30 -0500 (CDT)

Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id UAA15191;
Fri, 2 Oct 1998 20:19:28 -0400 (EDT)

Received: by west (SMI-8.6/SMI-SVR4)
id UAA08284; Fri, 2 Oct 1998 20:14:37 -0400
Date: Fri, 2 Oct 1998 20:14:37 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199810030014.UAA08284@west>
To: phayes@coginst.uwf.edu
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@tamu.edu,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu, fritz@cyc.com,
guarino@ladseb.pd.cnr.it, jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk,
skydog@pacbell.net
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 143
Status:

Pat,

I just read your latest note after I responded to the one before it.
I think it looks promising, and I'll send another note later.

John

From ???@??? Mon Oct 05 10:07:57 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA25876
for <phayes@coginst.uwf.edu>; Sat, 3 Oct 1998 11:41:18 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id MAA15796;
Sat, 3 Oct 1998 12:37:14 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
id MAA11970; Sat, 3 Oct 1998 12:32:20 -0400
Date: Sat, 3 Oct 1998 12:32:20 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199810031632.MAA11970@west>
To: phayes@coginst.uwf.edu
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@tamu.edu,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu, fritz@cyc.com,
guarino@ladseb.pd.cnr.it, jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk,
skydog@pacbell.net, sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 6065
Status:

Pat,

I think that your symbolic way of stating things captures what both of us were trying to say in different words. I have some comments about what words I would attach to the symbols.

>S is a FOL sentence (using a vocabulary V)
>R is a RDB (ie a set of ground literals using V)
>R+ is R together with appropriate closed-world axioms
>I is a Tarskian interpretation I (of V) over a domain D
>(where D is any set whatever, including cheese.)

I like this very much. Just one comment about the cheese:
In this formalism, the cheese has to be broken up into a set of lumps, each individually indexed, before you can talk about it formally in S. That is the main difficulty we face in trying to formalize NL: the cheese doesn't get broken up until the end of the discussion. Until then, you can't turn the crank on the formalism.

>Such RDB's provide, in a sense, a sufficient 'surrogate' for the
>interpretations:

I like the metaphor of relational DBs, because in my experience of teaching programmers, they get the main points very quickly when I state them in terms of RDBs. But possible worlds lead to glassy-eyed stares. I think that will also be true of most of our customers. I also agree with the theorems & other stuff I deleted without comment.

>S <---> R <---> I

>So far this is just theorems. How do we interpret it? To me, R belongs in
>the same category as S: both are lexical, both are describable in terms of
>computability, and both can be interpreted by a semantic theory. The
>relationship between them is one of derivability (I-) rather than
>interpretation (I=).

When you get to ground-level atoms, as in a relational DB, there is no structural difference between a model and a set of sentences: they are isomorphic. When you add quantified axioms to form R+, I agree that the relationship between S and R+ is provability (I-). But if you have just the raw DB, the SQL algorithm you apply to R is equivalent to Tarski's definition of interpretation (I=).

>RDBs are made is essentially lexical, or at any rate symbolic; and the
>notion of a 'physical' RDB seems incoherent or at best wrong-headed.
>Similarly, with this view of 'interpretation', the things I'm calling 'I'

>above may seem so metaphysically scruffy as to be beyond the pale of mere
>semantic theory.

That's partly my objection. But the more general problem is that the world doesn't come nicely packaged into clearly defined lumps. Animals, being mobile, are an exception because they cannot function unless they are detachable. But most of the natural world, including plants, is a continuum where you cannot find clear boundaries. Examples include the aspen trees, where a whole forest may consist of two or three individuals, each with dozens or even hundreds of individual trunks, each of which looks like a separate tree from the surface. In the autumn, however, you can see each of the three individuals turning color on different days, with all the trunks of each individual in synch.

Therefore, if you have selected a predefined convention for breaking the world into pieces, there is an isomorphism between R and I. But if you still have messy details to attend to, the most you can claim is an approximation between R and I. In that case, Tarski's definition still works with R, but it can't be applied to I.

>All of which (if correct) suggests that we might communicate better if we
>just agreed that you want to be talking about RDB's (with appropriate
>closed-world assumptions), but I don't. If we can agree that when you say
>'interpretation' you mean the $S \leftrightarrow R$ relationship, and when I say it I
>mean the $S \leftrightarrow I$ one, then I suspect that we will get along much more
>smoothly.

Yes, but.... If we have a well-behaved domain in which your I breaks up into clear individuals, then R and I are isomorphic, and we can both perform equivalent operations on them. But most domains aren't well behaved.

>.... Most applications of ontologies are in situations where the
>physical relationship between the reasoner and its immediate circumstances
>play no role in how it uses its ontology to reason. There may be no sensors
>involved in, say, an 'intelligent agent' or a smart GPS system, and still
>less in, say, something that reasons about medicine or military logistics.
>Grounding simply isn't an issue here...

Au contraire.... That is where all the thorny knowledge sharing problems arise. One of my favorite examples is of the oil company that couldn't reconcile their geological DB with their financial DB: the GDB defined 'oil well' as any hole in the ground dug or drilled for the purpose of getting oil whether or not it happened to be dry; the FDB defined 'oil well' as any pipe connected to one or more holes that produced oil. As a result, they couldn't correlate the production results with the geological data.

In medical applications, the data that goes into the various databases come from an incredible number of sources: sensors attached to patients, scribbling by a doctor or nurse about some observation, prescribed amounts of drugs, which may have little correlation with the amounts taken, etc.

Tarski's definition applies very well to what's in R and it can be extended to I *only if* I happens to be isomorphic to R. But the reasons for the discrepancies are so varied and so significant that you have to record and consider for each piece of data the conditions under which it was derived.

>... Whatever; but please don't tell me that I'm somehow forbidden
>to refer to them, or that doing so goes beyond the bounds of precise
>discussion. It doesn't, since to talk about something doesn't imply that one
>is able to compute it.

I agree with this final point completely. The great power of NL is that it allows us to talk about things before we have completely correlated our conventions for referring to things in the world. What goes into a relational DB must be sanitized for the protection of SQL, but that does not mean it is an accurate reflection of what's really in the messy world outside.

John

From ???@??? Mon Oct 05 14:29:03 1998

Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id KAA06830
for <phayes@coginst.uwf.edu>; Mon, 5 Oct 1998 10:56:03 -0500 (CDT)

Received: from scratchy (scratchy [207.207.8.118])

by catbert.cyc.com (8.8.8/8.8.8) with SMTP id KAA01530;
Mon, 5 Oct 1998 10:50:47 -0500 (CDT)

Message-Id: <3.0.32.19981005105120.030bb488@catbert.cyc.com>

X-Sender: fritz@catbert.cyc.com

X-Mailer: Windows Eudora Pro Version 3.0 (32)

Date: Mon, 05 Oct 1998 10:51:46 -0500

To: sowa@west.poly.edu (John F. Sowa)

From: Fritz Lehmann <fritz@cyc.com>

Subject: Re: Tropes

Cc: phayes@coginst.uwf.edu, Piek.Vossen@let.uva.nl, chezewiz@erols.com,
cmenzel@tamu.edu, doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu,
fritz@cyc.com, guarino@ladseb.pd.cnr.it, jmc@cs.stanford.edu,
p.m.simons@leeds.ac.uk, skydog@pacbell.net, sowa@west.poly.edu

Mime-Version: 1.0

Content-Type: text/plain; charset="us-ascii"

Content-Length: 1323

Status:

At 08:07 PM 10/2/98 -0400, John F. Sowa wrote:

>Pat,

>[...]

> You cannot solve

>the symbol grounding problem in a Cantorian transfinite universe.

Right. It also bothers me that the entire apparatus of that universe depends on Cantor's Main Theorem (that there is no surjection of an infinite set onto its powerset), whose proof is isomorphic to the Barber Paradox, and whose proof fails in perfectly good set theories with universal sets, like Quine's New Foundations (NF) set theory.

>[...] But all available evidence indicates that the

>universe is finite [...]

I am unaware of any important evidence one way or the other.

>[...]

>John

I generally agree with John Sowa in this (renewed) debate. I think a "model" can err. A model (it's annoying to have to call it "a Tarskian structure" especially since half of Tarski's work was the Hayes-disparaged algebraic approach which dispensed with them entirely) could contain the ground atomic assertion that Dusseldorf is in Wales. A sentence like (NOT (NOT (IN-REGION Dusseldorf Wales))) is true in the model, but false.

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

=====

====

From ???@??? Mon Oct 05 14:29:06 1998

Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA10685;

Mon, 5 Oct 1998 11:40:58 -0500 (CDT)

X-Sender: phayes@mail.coginst.uwf.edu

Message-Id: <v04003a02b23e93fd49@[143.88.7.118]>

In-Reply-To: <199810031632.MAA11970@west>

Mime-Version: 1.0

Date: Mon, 5 Oct 1998 11:37:59 -0500

To: sowa@west.poly.edu (John F. Sowa)

From: Pat Hayes <phayes@coginst.uwf.edu>

Subject: Re: Tropes

Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,

fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net

Content-Type: text/plain; charset="us-ascii"

Content-Length: 11344

Status:

Hi John

>>S is a FOL sentence (using a vocabulary V)
>>R is a RDB (ie a set of ground literals using V)
>>R+ is R together with appropriate closed-world axioms
>>I is a Tarskian interpretation I (of V) over a domain D
>>(where D is any set whatever, including cheese.)

>

>I like this very much. Just one comment about the cheese:
>In this formalism, the cheese has to be broken up into a set of
>lumps, each individually indexed, before you can talk about it
>formally in S. That is the main difficulty we face in trying
>to formalize NL: the cheese doesn't get broken up until the
>end of the discussion. Until then, you can't turn the crank
>on the formalism.

FOL requires us to have available some notion of individual in order to make sense of quantification. That's usually expressed simply by saying that the domain is a set. But 'individuals' in this sense might be all kinds of things, not just separated lumps, and they needn't be 'indexed'. Logical individuals can be physical objects, times, substances, sentences,... you name it. In particular, they can be kinds of cheese, or substances, or the denotata of mass terms. To keep ourselves clear, we maybe need to have what Nicola calls 'individuation criteria' for the things in the domain; but notice that the theory of truth is a metatheory. There is no assumption that these criteria of individuation be expressed (or even expressible) in the FOL axioms being interpreted.

>>Such RDB's provide, in a sense, a sufficient 'surrogate' for the
>>interpretations:

>

>I like the metaphor of relational DBs, because in my experience of
>teaching programmers, they get the main points very quickly when I
>state them in terms of RDBs. But possible worlds lead to glassy-eyed
>stares.

All depends on your audience. My students have no idea what an RDB is.

>I think that will also be true of most of our customers.

What worries me is that if we are too free with RDB talk, our customers will *think* they understand us, but in fact will be confused.

>I also agree with the theorems & other stuff I deleted without comment.

>

>>S <---> R <---> I

>

>>So far this is just theorems. How do we interpret it? To me, R belongs in
>>the same category as S: both are lexical, both are describable in terms of
>>computability, and both can be interpreted by a semantic theory. The
>>relationship between them is one of derivability (I-) rather than
>>interpretation (I=).

>

>When you get to ground-level atoms, as in a relational DB, there is no
>structural difference between a model and a set of sentences: they are
>isomorphic.

I presume you mean 'interpretation' rather than 'model', given our earlier discussion? If so, the above just isn't generally true, since interpretations may be uncountable. In fact, it's almost never true, since a set of ground atoms (with no closed world axioms) *cannot* uniquely specify a single interpretation: it's always possible to add elements to a domain which aren't names in the RDB vocabulary. I stated the exact relationship in my earlier message, so let's not get it all muddied up by talking about 'isomorphisms' that provably aren't isomorphic. (You may wish to declare uncountable interpretations to be computationally indecent, but there's no gainsaying the fact that they exist.)

Your use of 'isomorphism' here is also potentially dangerous for a different reason. If we consider an RDB to be an algebraic structure, then it may well have self-isomorphisms which do not conserve truth. It is guaranteed to be an interpretation of itself only when interpreted in the Herbrand style, so that each symbol denotes its own lexical extension in the RDB; but other interpretations are possible. An interpretation includes the mapping from the vocabulary of the interpreted theory to the abstract structure (domain plus relations), and the 'same' - ie isomorphic- structure can be interpreted differently, giving different truth values. Again, to interpret 'R' to denote the relation listed by 'S' and vice versa, say, may seem perverse, but there's no mathematical reason why it couldn't be done. (Which is another way of saying: if you want to stop it, you need to invoke some other mechanism with some semantic bite.)

>..... the more general problem is that the

>world doesn't come nicely packaged into clearly defined lumps. Animals,
>being mobile, are an exception because they cannot function unless they
>are detachable. But most of the natural world, including plants, is
>a continuum where you cannot find clear boundaries. Examples include the
>aspen trees, where a whole forest may consist of two or three individuals,
>each with dozens or even hundreds of individual trunks, each of which looks
>like a separate tree from the surface. In the autumn, however, you can see
>each of the three individuals turning color on different days, with all
>the trunks of each individual in synch.

>

>Therefore, if you have selected a predefined convention for breaking
>the world into pieces, there is an isomorphism between R and I. But
>if you still have messy details to attend to, the most you can claim
>is an approximation between R and I. In that case, Tarski's definition
>still works with R, but it can't be applied to I.

John, you are just plain confused. Of course Tarski's definition works with I, since I is *defined* to be a Tarskian interpretation! I really don't have the time (or the energy) to give you a basic correspondence course in model theory. In brief: any particular way to divide up the physical world (separate trees, trees apart from root systems, separate cells, you name it) provides the domain of *an* interpretation. Of course this doesn't show that any one of these is *the* interpretation, but that's not how model theory works. (The fact that you think that talking physical interpretations must somehow involve solving the symbol grounding problem suggests that you haven't understood this properly. Perhaps there is only one world, but there are many ways to make Tarskian interpretations out of it.) The relationship you call 'approximation' is properly stated in model theory by saying that the theory R is true in many possible interpretations I (of a certain kind, perhaps), ie it fails to completely specify what it is describing. Model theory gives you just the theoretical tools you need to talk about this kind of 'approximation'.

>>All of which (if correct) suggests that we might communicate better if we
>>just agreed that you want to be talking about RDB's (with appropriate
>>closed-world assumptions), but I don't. If we can agree that when you say
>>'interpretation' you mean the $S \leftrightarrow R$ relationship, and when I say it I
>>mean the $S \leftrightarrow I$ one, then I suspect that we will get along much more
>>smoothly.

>

>Yes, but.... If we have a well-behaved domain in which your I breaks
>up into clear individuals, then R and I are isomorphic, and we can both
>perform equivalent operations on them. But most domains aren't well >behaved.

In this sense, *all* domains are well-behaved, by definition, since a domain is a set.

>
>>.... Most applications of ontologies are in situations where the
>>physical relationship between the reasoner and its immediate circumstances
>>play no role in how it uses its ontology to reason. There may be no sensors
>>involved in, say, an 'intelligent agent' or a smart GPS system, and still
>>less in, say, something that reasons about medicine or military logistics.
>>Grounding simply isn't an issue here...

>
>Au contraire.... That is where all the thorny knowledge sharing problems
>arise. One of my favorite examples is of the oil company that couldn't
>reconcile their geological DB with their financial DB: the GDB defined
>'oil well' as any hole in the ground dug or drilled for the purpose of
>getting oil whether or not it happened to be dry; the FDB defined 'oil well'
>as any pipe connected to one or more holes that produced oil. As a result,
>they couldn't correlate the production results with the geological data.

That's a problem of mutual definition and mutual consistency, not the symbol grounding problem (as that term is used elsewhere, anyway). This would be the symbol grounding problem if one or both of these companies were trying to incorporate their Kbase into a robot which was expected to recognise an oilwell when it saw one.

I agree with you that this problem of unrecognised mutually incompatible definitions is also hard and that we need to consider it. (As I understand it, this is one of the main motivations for this entire project of trying to form a mutually acceptable, reasonably useful and hopefully universal high-level ontology in the first place.) However, none of this involves us refusing to consider Tarskian interpretations with physical domains. Those two definitions you cite above are **textual** differences (and maybe axiomatic differences) between two theories, and presumably a way to reconcile them would be to find a way to express oil-well-1 in theory-2 and vice versa. If we try to do this, we will be writing axioms which I'm sure both parties would prefer to have at least **some** physical interpretations.

>In medical applications, the data that goes into the various databases
>come from an incredible number of sources: sensors attached to patients,
>scribbling by a doctor or nurse about some observation, prescribed amounts
>of drugs, which may have little correlation with the amounts taken, etc.

>
>Tarski's definition applies very well to what's in R and it can be extended
>to I **only if** I happens to be isomorphic to R.

Again, this is just flat wrong.

But the reasons for the

>discrepancies are so varied and so significant that you have to record
>and consider for each piece of data the conditions under which it was derived.
>
>>... Whatever; but please don't tell me that I'm somehow forbidden
>>to refer to them, or that doing so goes beyond the bounds of precise
>>discussion. It doesn't, since to talk about something doesn't imply that one
>>is able to compute it.
>
>I agree with this final point completely. The great power of NL is that
>it allows us to talk about things before we have completely correlated
>our conventions for referring to things in the world.

So does FOL! In fact, if anything, FOL is even looser than NL in this regard: it makes *no* assumptions about what its names mean.

> What goes into
>a relational DB must be sanitized for the protection of SQL, but that
>does not mean it is an accurate reflection of what's really in the messy
>world outside.

To speak of interpretations over domains containing physical things is not to claim that the theory being interpreted must be the final and sole correct Theory of the Universe. It doesn't even assume that it is true at all, in fact. All it says is that when talking within a metatheory of truth, we can at least consider the *possibility* that our axioms *might* be true of reality. That seems like a very small and modest assumption on which to base a theory of truth, since to deny it is to give up before we even begin.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Mon Oct 05 14:29:06 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA10685;
Mon, 5 Oct 1998 11:40:58 -0500 (CDT)

X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a02b23e93dfd49@[143.88.7.118]>
In-Reply-To: <199810031632.MAA11970@west>
Mime-Version: 1.0
Date: Mon, 5 Oct 1998 11:37:59 -0500
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 11344
Status:

Hi John

>>S is a FOL sentence (using a vocabulary V)
>>R is a RDB (ie a set of ground literals using V)
>>R+ is R together with appropriate closed-world axioms
>>I is a Tarskian interpretation I (of V) over a domain D
>>(where D is any set whatever, including cheese.)
>
>I like this very much. Just one comment about the cheese:
>In this formalism, the cheese has to be broken up into a set of
>lumps, each individually indexed, before you can talk about it
>formally in S. That is the main difficulty we face in trying
>to formalize NL: the cheese doesn't get broken up until the
>end of the discussion. Until then, you can't turn the crank
>on the formalism.

FOL requires us to have available some notion of individual in order to make sense of quantification. That's usually expressed simply by saying that the domain is a set. But 'individuals' in this sense might be all kinds of things, not just separated lumps, and they needn't be 'indexed'. Logical individuals can be physical objects, times, substances, sentences,... you name it. In particular, they can be kinds of cheese, or substances, or the denotata of mass terms. To keep ourselves clear, we maybe need to have what Nicola calls 'individuation criteria' for the things in the domain; but notice that the theory of truth is a metatheory. There is no assumption that these criteria of individuation be expressed (or even expressible) in the FOL axioms being interpreted.

>>Such RDB's provide, in a sense, a sufficient 'surrogate' for the
>>interpretations:

>

>I like the metaphor of relational DBs, because in my experience of
>teaching programmers, they get the main points very quickly when I
>state them in terms of RDBs. But possible worlds lead to glassy-eyed
>stares.

All depends on your audience. My students have no idea what an RDB is.

>I think that will also be true of most of our customers.

What worries me is that if we are too free with RDB talk, our customers
will *think* they understand us, but in fact will be confused.

>I also agree with the theorems & other stuff I deleted without comment.

>

>>S <---> R <---> I

>

>>So far this is just theorems. How do we interpret it? To me, R belongs in
>>the same category as S: both are lexical, both are describable in terms of
>>computability, and both can be interpreted by a semantic theory. The
>>relationship between them is one of derivability (I-) rather than
>>interpretation (I=).

>

>When you get to ground-level atoms, as in a relational DB, there is no
>structural difference between a model and a set of sentences: they are
>isomorphic.

I presume you mean 'interpretation' rather than 'model', given our earlier
discussion? If so, the above just isn't generally true, since
interpretations may be uncountable. In fact, it's almost never true, since a
set of ground atoms (with no closed world axioms) *cannot* uniquely specify
a single interpretation: it's always possible to add elements to a domain
which aren't names in the RDB vocabulary. I stated the exact relationship in
my earlier message, so let's not get it all muddied up by talking about
'isomorphisms' that provably aren't isomorphic. (You may wish to declare
uncountable interpretations to be computationally indecent, but there's no
gainsaying the fact that they exist.)

Your use of 'isomorphism' here is also potentially dangerous for a
different reason. If we consider an RDB to be an algebraic structure, then
it may well have self-isomorphisms which do not conserve truth. It is
guaranteed to be an interpretation of itself only when interpreted in the
Herbrand style, so that each symbol denotes its own lexical extension in
the RDB; but other interpretations are possible. An interpretation includes

the mapping from the vocabulary of the interpreted theory to the abstract structure (domain plus relations), and the 'same' - ie isomorphic- structure can be interpreted differently, giving different truthvalues. Again, to interpret 'R' to denote the relation listed by 'S' and vice versa, say, may seem perverse, but there's no mathematical reason why it couldn't be done. (Which is another way of saying: if you want to stop it, you need to invoke some other mechanism with some semantic bite.)

>..... the more general problem is that the
>world doesn't come nicely packaged into clearly defined lumps. Animals,
>being mobile, are an exception because they cannot function unless they
>are detachable. But most of the natural world, including plants, is
>a continuum where you cannot find clear boundaries. Examples include the
>aspen trees, where a whole forest may consist of two or three individuals,
>each with dozens or even hundreds of individual trunks, each of which looks
>like a separate tree from the surface. In the autumn, however, you can see
>each of the three individuals turning color on different days, with all
>the trunks of each individual in synch.
>
>Therefore, if you have selected a predefined convention for breaking
>the world into pieces, there is an isomorphism between R and I. But
>if you still have messy details to attend to, the most you can claim
>is an approximation between R and I. In that case, Tarski's definition
>still works with R, but it can't be applied to I.

John, you are just plain confused. Of course Tarski's definition works with I, since I is *defined* to be a Tarskian interpretation! I really don't have the time (or the energy) to give you a basic correspondence course in model theory. In brief: any particular way to divide up the physical world (separate trees, trees apart from root systems, separate cells, you name it) provides the domain of *an* interpretation. Of course this doesn't show that any one of these is *the* interpretation, but that's not how model theory works. (The fact that you think that talking physical interpretations must somehow involve solving the symbol grounding problem suggests that you haven't understood this properly. Perhaps there is only one world, but there are many ways to make Tarskian interpretations out of it.) The relationship you call 'approximation' is properly stated in model theory by saying that the theory R is true in many possible interpretations I (of a certain kind, perhaps), ie it fails to completely specify what it is describing. Model theory gives you just the theoretical tools you need to talk about this kind of 'approximation'.

>>All of which (if correct) suggests that we might communicate better if we
>>just agreed that you want to be talking about RDB's (with appropriate
>>closed-world assumptions), but I don't. If we can agree that when you say
>>'interpretation' you mean the $S \leftrightarrow R$ relationship, and when I say it I

>>mean the S <-> I one, then I suspect that we will get along much more
>>smoothly.

>

>Yes, but.... If we have a well-behaved domain in which your I breaks
>up into clear individuals, then R and I are isomorphic, and we can both
>perform equivalent operations on them. But most domains aren't well >behaved.

In this sense, *all* domains are well-behaved, by definition, since a domain is a set.

>

>>.... Most applications of ontologies are in situations where the
>>physical relationship between the reasoner and its immediate circumstances
>>play no role in how it uses its ontology to reason. There may be no sensors
>>involved in, say, an 'intelligent agent' or a smart GPS system, and still
>>less in, say, something that reasons about medicine or military logistics.
>>Grounding simply isn't an issue here...

>

>Au contraire.... That is where all the thorny knowledge sharing problems
>arise. One of my favorite examples is of the oil company that couldn't
>reconcile their geological DB with their financial DB: the GDB defined
>'oil well' as any hole in the ground dug or drilled for the purpose of
>getting oil whether or not it happened to be dry; the FDB defined 'oil well'
>as any pipe connected to one or more holes that produced oil. As a result,
>they couldn't correlate the production results with the geological data.

That's a problem of mutual definition and mutual consistency, not the symbol grounding problem (as that term is used elsewhere, anyway). This would be the symbol grounding problem if one or both of these companies were trying to incorporate their Kbase into a robot which was expected to recognise an oilwell when it saw one.

I agree with you that this problem of unrecognised mutually incompatible definitions is also hard and that we need to consider it. (As I understand it, this is one of the main motivations for this entire project of trying to form a mutually acceptable, reasonably useful and hopefully universal high-level ontology in the first place.) However, none of this involves us refusing to consider Tarskian interpretations with physical domains. Those two definitions you cite above are *textual* differences (and maybe axiomatic differences) between two theories, and presumably a way to reconcile them would be to find a way to express oil-well-1 in theory-2 and vice versa. If we try to do this, we will be writing axioms which I'm sure both parties would prefer to have at least *some* physical interpretations.

>In medical applications, the data that goes into the various databases
>come from an incredible number of sources: sensors attached to patients,

>scribbling by a doctor or nurse about some observation, prescribed amounts
>of drugs, which may have little correlation with the amounts taken, etc.
>
>Tarski's definition applies very well to what's in R and it can be extended
>to I *only if* I happens to be isomorphic to R.

Again, this is just flat wrong.

But the reasons for the
>discrepancies are so varied and so significant that you have to record
>and consider for each piece of data the conditions under which it was derived.
>
>>... Whatever; but please don't tell me that I'm somehow forbidden
>>to refer to them, or that doing so goes beyond the bounds of precise
>>discussion. It doesn't, since to talk about something doesn't imply that one
>>is able to compute it.
>
>I agree with this final point completely. The great power of NL is that
>it allows us to talk about things before we have completely correlated
>our conventions for referring to things in the world.

So does FOL! In fact, if anything, FOL is even looser than NL in this regard: it makes *no* assumptions about what its names mean.

> What goes into
>a relational DB must be sanitized for the protection of SQL, but that
>does not mean it is an accurate reflection of what's really in the messy
>world outside.

To speak of interpretations over domains containing physical things is not to claim that the theory being interpreted must be the final and sole correct Theory of the Universe. It doesn't even assume that it is true at all, in fact. All it says is that when talking within a metatheory of truth, we can at least consider the *possibility* that our axioms *might* be true of reality. That seems like a very small and modest assumption on which to base a theory of truth, since to deny it is to give up before we even begin.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax

phayes@ai.uwf.edu
http://www.coginst.uwf.edu/~phayes

From ???@??? Mon Oct 05 15:20:24 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id OAA13387;
Mon, 5 Oct 1998 14:31:45 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a05b23eb7966341@[143.88.7.118]>
In-Reply-To: <199810030007.UAA08254@west>
Mime-Version: 1.0
Date: Mon, 5 Oct 1998 14:28:37 -0500
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 14074
Status:

John, greetings

>
>>They dont suffer from this flaw because, unlike your rhetorical example of
>>the lions, they are not METHODS for DOING anything. ...
>
>Exactly my point! This ontology project is not a project in pure mathematics.
>It is a project in AI and computer science, both of which are engineering
>disciplines that use mathematics to compute.

?? Then Im at a loss to see what your point is. Tarkian truth isn't
concerned with methods (you just agreed) ...but we are.... So what? So we
must reject Tarskian model theory? (Can you suggest an alternative?) I
agree that we arent doing pure matheamtics; but Ive never said or thought
we were, so that seems to have nothing to do with what we have been arguing
about.

> Even for the pure scientific
>side of AI, the only functions of interest are ones that are computable.

Rubbish. Our business here is constructing first-order formalisations (or maybe deciding suitable vocabularies for use by first-order axioms, etc.; but in any case:) Our subjectmatter is the vocabularies of such formal theories, and much of what we discuss concerns what these vocabularies are supposed to mean. Now, the processes (of database retrieval, searching for inferences, etc.) that use these formalisms are, of course, computable. But a metatheory of truth - the language that we use when talking about what the formal sentences mean - needs to consider *all* possible interpretations of these formalisms; it - the metatheory - is not confined to meaning being computable. So to claim that 'the only functions which are of interest are ones that are computable' is wrong. Or at any rate, *I* am interested in other functions, even if you aren't; and the (meta)theory of truth that Tarski has given us also allows for such interpretations; and we ignore them at our peril, in my view. Finiteness isn't even first-order expressible, for example, so (far from being on solid computational ground) for us to assume that all our semantic domains were finite would be to assume that our theories already have expressive powers that are beyond recursive enumerability, let alone recursion.

>>Your final conclusion follows only if one equates "computable" with "within the scope of philosophical consideration"....

>

>There are many issues worthy of philosophical consideration in this ontology >effort. When we are dealing with the ontology of mathematics, then I would >be happy to adopt a Platonistic stance, which is usually the most fruitful >approach to the nature of mathematical objects. But when we are trying to >define relationships between language and the world that are either >(a) psychologically realistic or (b) computationally tractable, then >the finitistic stance is the only one that makes sense. You cannot solve >the symbol grounding problem in a Cantorian transfinite universe.

The relationship that has to be computationally tractable is not that between the formalism and the world, but between some sentences and others. Many applications of our ontologies have no computational access to the worlds they are about. (See my earlier message for comments on symbol grounding.)

>>.... We are arguing

>>here not about those computational properties (which have to do with things >>like branching rates in search spaces) but about a *theory* of *truth*....

>

>As I said before, I have no objection to uncountable sets for a theory >of mathematical truth. But all available evidence indicates that the >universe is finite, and what we can perceive, compute, think about, and >talk about is certainly finite. The really difficult problems of a *theory*

>of *truth* are finite.

Look, John, I respect this extreme finitist perspective of yours, and I even have some intellectual sympathy with it. But you must recognise that it is a very marginal and extreme position, quite at odds with the way that almost everyone else thinks, and that we cannot possibly expect to build it into the very fabric of our assumptions as though it were received opinion. It is way out in left field, and has many internal philosophical problems waiting to be worked out in detail.

>>(A side comment on something you said in an earlier message. You said that
>>any Tarskian interpretation is 'isomorphic' to a countable set of ground
>>atomic sentences. This isn't quite correct, technically, since....

>

>I think that I just said "a conjunction of ground atoms" without using
>the word "countable". If you want to generalize the point to an
>uncountable conjunction, I have no objection.

All collections of atomic sentences are countable.

>>>The beauty of Tarski's denotation function is that it can be computed
>>>in polynomial time by an ordinary SQL query to any relational database.

>>

>>This is just a plain error. It isn't even coherent to say that Tarski's
>>denotation function is computable, since it may target a noncomputable set.

>

>As I have said before, the mathematical problem of reasoning about
>Platonic structures is the easy part. For that part, you don't have to
>dig around in the "rich loamy soil" to build your database of relations.
>You can just sit back in your easy chair and imagine uncountable infinities.

That's not the point. How do you propose to eliminate these uncomfortably large interpretations from being true interpretations of your axioms? The way any semantics theory works is, you specify the rules of interpretation and then we will see what interpretations we find. If you use Tarskian rules, then these uncountable interpretations exist. You are the one refusing to face reality here, by assuming that we can somehow legislate them out of our metatheory just by saying they are computationally indecent, or that we don't need to consider them, or something. Give us a semantic theory which rules them out!

>>What is the appropriate definition of 'computable' for, say, the set of
>>galaxies?

>

>Since the universe is finite, the number of galaxies is finite. The main
>issue here is not the denotation side of the problem, but the perception

>side, which puts the data into the database.

I disagree. Perception, and the issues that go with it (such as grounding), are not in our ontological province. They might be if we were doing robotics, but we aren't.

>>... In any case, why would the process of answering a SQL query
>>involve computing a denotation function? That process is one of inference;
>>it manipulates sentences (of which the relational database entries are a
>>simple sort, ie ground atoms.)

>

>A Tarski-style structure is isomorphic to a relational DB in the
>finite case. In the infinite case, which never arises in databases
>obtained from the perception of physical situations, it is still
>isomorphic to a conjunction of possibly uncountably many ground atoms.

That last sentence doesn't make sense; see earlier message for comments about 'isomorphic' here.

>The definition of the denotation function is *identical* to the algorithm
>for evaluating an SQL query, which is identical to the algorithm for
>evaluating a Prolog expression that contains predicates defined only by
>ground-level assertions. All three of these things can be evaluated in
>polynomial time, where the degree of the polynomial is equal to the number
>of quantifiers in the expression (which may be implicit in SQL & Prolog).
>In fact, Tarski's definition can be optimized in the same way as SQL and
>Prolog by building an index so that many of the existential quantifiers
>can be evaluated in logarithmic time.

All of the above is entirely to do with efficient ways to compute unsatisfiability, which of course is the same problem as computing consequence. Now, in a sense, one detects unsatisfiability by showing that a systematic process of trying to build a (Herbrand) interpretation must fail. But this applies only to Herbrand interpretations, not to all interpretations. For purely computational purposes, indeed, Herbrand interpretations can stand in for arbitrary ones; but when we are arguing about the nature of the distinction between continuants and occurrents, say, our intuitions are going to become very impoverished if we are only allowed to consider models made of ground expressions. The metatheory of a computable theory need not itself be computable.

>>The symbol grounding problem is indeed a difficult and interesting problem,
>>but it is quite different from what we are talking about. Grounding has to
>>do with how a system which is both reasoning about a world and causally
>>embedded in it could *establish* that a name must refer to something
>>'outside' in the physical world. It's an issue in robot epistemology,

>>concerned with ways of knowing that a name refers. Model theory is is
>>concerned with ways of referring, not ways of knowing that one refers. Put
>>another way, MT is concerned with how a world *could* be, grounding is
>>concerned with how the perceived world *is*, and how we might know that.
>>The difference is crucial.

>

>I agree with your definitions. But they get to the crux of our dispute.
>For a theory of mathematical truth, I have no quarrel with you. But
>my primary argument against model theory, as a theory of truth about the
>physical world, is that it fails to address the symbol grounding problem.

The above illustrates why communicating with you is often so difficult. You say you agree, then go on to immediately contradict what you just agreed with. Read what I wrote again: model theory *isnt concerned* with symbol grounding. Of course it 'fails to address' the symbol grounding problem, because symbol grounding is an entirely separate issue. It isnt concerned with making better ballbearings, either, so its hardly a valid criticism to say that it should be rejected because bearings dont run smoothly.

(Even if symbol grounding were in our ambit, this is hardly a good argument *against* model theory. It might, at best, be an argument against a kind of over-complacent confidence that MT solved all problems, but that's not being claimed here. All we have been arguing about is allowing people to even consider interpretations over the physical world. You seem to think that until we have 'solved' symbol grounding (which would at the very least involve a realistic theory of perception), *all* talk of physical interpretations must be banned. This doesnt seem like a good way to make progress, especially when this theory, even if incomplete, seems to offer some ways to analyse the problems that arise in symbol grounding.)

>Every criticism that has been hurled against AI projects that relate
>language to symbols in a computer can be leveled against the claim that
>Tarski's denotation function "solves" the problem of defining truth.

>

>Your claim that the "individuals" in Tarski's structures *are* the
>actual physical objects begs every question that Roger Schank begs
>when he claims that his GENSYMs *stand for* the actual physical objects.

It doesnt beg any questions at all. If I can talk of a set of, say, planets, then I can talk of an interpretation over a domain of planets. (Not *the* interpretation, notice, but *an* interpretation.) If you reply: but you can't talk of a set of planets before specifying exactly what "planet" means; then I may agree; but we are now having a metalevel discussion about the meaning of English words. I might respond, "set of planets in the ordinary sense of 'planet' "; or maybe we will come to agree that a planet is something orbiting a star; or whatever. Each of these ways

to get exact about what "planet" means might result in different things being allowed as a set of planets. But no matter what agreement we come to, each of these sets is, in fact, a set (of planets or something else), and that's all Tarski's construction needs.

>Using the verb 'are' instead of 'stand for' in no way solves the symbol
>grounding problem. Unlike Roger Schank, I am willing to agree that
>Tarski's approach solves one half of the problem, but you can't solve
>the symbol grounding problem by postulating it away.

>

>>I can't agree that we have a responsibility to provide grounding sockets
>>for the engineers to plug their sensors into, because (1) I don't think this
>>is central to our task - we aren't in the robot-building game here - and (2)
>>it's much too heavy a burden, since we can't be expected to provide
>>ready-made solutions to problems that haven't even been adequately formulated
>>yet. Grounding is an open problem, and we don't even know if the current -
>>very partial - ideas about it really work.

>

>I agree with you completely on this point. The point I was trying to make
>is that the database is the boundary between the data that Tarski's
>algorithm applies to and the data that the robot builders generate.
>The engineers are responsible for filling up the database with ground
>atoms about what is in the world, and the denotation function uses SQL
>to determine whether a particular sentence is true or false.

Tarski didn't define an 'algorithm', and the denotation function doesn't 'use' anything, just as the exponential function (say) doesn't 'use' anything. It's a mathematical function, not a computational process. (John, you have an amazing ability to stretch historical truth by misusing technical vocabulary.)

You seem to have an image of a kind of ideal robot whose sensors deliver ground atoms to it which are a complete ground-atomic description of its world, and it then uses some systematic inference process to test its quantified beliefs against that database. We could argue about whether this picture is at all convincing (either as robotics or psychology), but (a) it is a very particular picture, not shared by most, and (b) it is irrelevant here in any case, since we aren't doing robotics but general-purpose ontology design.

Pat

IHMC, University of West Florida	(850)434 8903	home
11000 University Parkway	(850)474 2091	office
Pensacola, FL 32514	(850)474 3023	fax

phayes@ai.uwf.edu
http://www.coginst.uwf.edu/~phayes

From ???@??? Mon Oct 05 15:20:24 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id OAA13387;
Mon, 5 Oct 1998 14:31:45 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a05b23eb7966341@[143.88.7.118]>
In-Reply-To: <199810030007.UAA08254@west>
Mime-Version: 1.0
Date: Mon, 5 Oct 1998 14:28:37 -0500
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 14074
Status:

John, greetings

>
>>They dont suffer from this flaw because, unlike your rhetorical example of
>>the lions, they are not METHODS for DOING anything. ...
>
>Exactly my point! This ontology project is not a project in pure mathematics.
>It is a project in AI and computer science, both of which are engineering
>disciplines that use mathematics to compute.

?? Then Im at a loss to see what your point is. Tarkian truth isn't
concerned with methods (you just agreed) ...but we are.... So what? So we
must reject Tarskian model theory? (Can you suggest an alternative?) I
agree that we arent doing pure matheamtics; but Ive never said or thought
we were, so that seems to have nothing to do with what we have been arguing
about.

> Even for the pure scientific
>side of AI, the only functions of interest are ones that are computable.

Rubbish. Our business here is constructing first-order formalisations (or maybe deciding suitable vocabularies for use by first-order axioms, etc.; but in any case:) Our subjectmatter is the vocabularies of such formal theories, and much of what we discuss concerns what these vocabularies are supposed to mean. Now, the processes (of database retrieval, searching for inferences, etc.) that use these formalisms are, of course, computable. But a metatheory of truth - the language that we use when talking about what the formal sentences mean - needs to consider *all* possible interpretations of these formalisms; it - the metatheory - is not confined to meaning being computable. So to claim that 'the only functions which are of interest are ones that are computable' is wrong. Or at any rate, *I* am interested in other functions, even if you aren't; and the (meta)theory of truth that Tarski has given us also allows for such interpretations; and we ignore them at our peril, in my view. Finiteness isn't even first-order expressible, for example, so (far from being on solid computational ground) for us to assume that all our semantic domains were finite would be to assume that our theories already have expressive powers that are beyond recursive enumerability, let alone recursion.

>>Your final conclusion follows only if one equates "computable" with "within the scope of philosophical consideration"....

>

>There are many issues worthy of philosophical consideration in this ontology >effort. When we are dealing with the ontology of mathematics, then I would >be happy to adopt a Platonistic stance, which is usually the most fruitful >approach to the nature of mathematical objects. But when we are trying to >define relationships between language and the world that are either >(a) psychologically realistic or (b) computationally tractable, then >the finitistic stance is the only one that makes sense. You cannot solve >the symbol grounding problem in a Cantorian transfinite universe.

The relationship that has to be computationally tractable is not that between the formalism and the world, but between some sentences and others. Many applications of our ontologies have no computational access to the worlds they are about. (See my earlier message for comments on symbol grounding.)

>>.... We are arguing

>>here not about those computational properties (which have to do with things >>like branching rates in search spaces) but about a *theory* of *truth*....

>

>As I said before, I have no objection to uncountable sets for a theory >of mathematical truth. But all available evidence indicates that the >universe is finite, and what we can perceive, compute, think about, and >talk about is certainly finite. The really difficult problems of a *theory*

>of *truth* are finite.

Look, John, I respect this extreme finitist perspective of yours, and I even have some intellectual sympathy with it. But you must recognise that it is a very marginal and extreme position, quite at odds with the way that almost everyone else thinks, and that we cannot possibly expect to build it into the very fabric of our assumptions as though it were received opinion. It is way out in left field, and has many internal philosophical problems waiting to be worked out in detail.

>>(A side comment on something you said in an earlier message. You said that
>>any Tarskian interpretation is 'isomorphic' to a countable set of ground
>>atomic sentences. This isn't quite correct, technically, since....

>

>I think that I just said "a conjunction of ground atoms" without using
>the word "countable". If you want to generalize the point to an
>uncountable conjunction, I have no objection.

All collections of atomic sentences are countable.

>>>The beauty of Tarski's denotation function is that it can be computed
>>>in polynomial time by an ordinary SQL query to any relational database.

>>

>>This is just a plain error. It isn't even coherent to say that Tarski's
>>denotation function is computable, since it may target a noncomputable set.

>

>As I have said before, the mathematical problem of reasoning about
>Platonic structures is the easy part. For that part, you don't have to
>dig around in the "rich loamy soil" to build your database of relations.
>You can just sit back in your easy chair and imagine uncountable infinities.

That's not the point. How do you propose to eliminate these uncomfortably large interpretations from being true interpretations of your axioms? The way any semantics theory works is, you specify the rules of interpretation and then we will see what interpretations we find. If you use Tarskian rules, then these uncountable interpretations exist. You are the one refusing to face reality here, by assuming that we can somehow legislate them out of our metatheory just by saying they are computationally indecent, or that we don't need to consider them, or something. Give us a semantic theory which rules them out!

>>What is the appropriate definition of 'computable' for, say, the set of
>>galaxies?

>

>Since the universe is finite, the number of galaxies is finite. The main
>issue here is not the denotation side of the problem, but the perception

>side, which puts the data into the database.

I disagree. Perception, and the issues that go with it (such as grounding), are not in our ontological province. They might be if we were doing robotics, but we aren't.

>>... In any case, why would the process of answering a SQL query
>>involve computing a denotation function? That process is one of inference;
>>it manipulates sentences (of which the relational database entries are a
>>simple sort, ie ground atoms.)

>

>A Tarski-style structure is isomorphic to a relational DB in the
>finite case. In the infinite case, which never arises in databases
>obtained from the perception of physical situations, it is still
>isomorphic to a conjunction of possibly uncountably many ground atoms.

That last sentence doesn't make sense; see earlier message for comments about 'isomorphic' here.

>The definition of the denotation function is *identical* to the algorithm
>for evaluating an SQL query, which is identical to the algorithm for
>evaluating a Prolog expression that contains predicates defined only by
>ground-level assertions. All three of these things can be evaluated in
>polynomial time, where the degree of the polynomial is equal to the number
>of quantifiers in the expression (which may be implicit in SQL & Prolog).
>In fact, Tarski's definition can be optimized in the same way as SQL and
>Prolog by building an index so that many of the existential quantifiers
>can be evaluated in logarithmic time.

All of the above is entirely to do with efficient ways to compute unsatisfiability, which of course is the same problem as computing consequence. Now, in a sense, one detects unsatisfiability by showing that a systematic process of trying to build a (Herbrand) interpretation must fail. But this applies only to Herbrand interpretations, not to all interpretations. For purely computational purposes, indeed, Herbrand interpretations can stand in for arbitrary ones; but when we are arguing about the nature of the distinction between continuants and occurrents, say, our intuitions are going to become very impoverished if we are only allowed to consider models made of ground expressions. The metatheory of a computable theory need not itself be computable.

>>The symbol grounding problem is indeed a difficult and interesting problem,
>>but it is quite different from what we are talking about. Grounding has to
>>do with how a system which is both reasoning about a world and causally
>>embedded in it could *establish* that a name must refer to something
>>'outside' in the physical world. It's an issue in robot epistemology,

>>concerned with ways of knowing that a name refers. Model theory is is
>>concerned with ways of referring, not ways of knowing that one refers. Put
>>another way, MT is concerned with how a world *could* be, grounding is
>>concerned with how the perceived world *is*, and how we might know that.
>>The difference is crucial.

>

>I agree with your definitions. But they get to the crux of our dispute.
>For a theory of mathematical truth, I have no quarrel with you. But
>my primary argument against model theory, as a theory of truth about the
>physical world, is that it fails to address the symbol grounding problem.

The above illustrates why communicating with you is often so difficult. You say you agree, then go on to immediately contradict what you just agreed with. Read what I wrote again: model theory *isnt concerned* with symbol grounding. Of course it 'fails to address' the symbol grounding problem, because symbol grounding is an entirely separate issue. It isnt concerned with making better ballbearings, either, so its hardly a valid criticism to say that it should be rejected because bearings dont run smoothly.

(Even if symbol grounding were in our ambit, this is hardly a good argument *against* model theory. It might, at best, be an argument against a kind of over-complacent confidence that MT solved all problems, but that's not being claimed here. All we have been arguing about is allowing people to even consider interpretations over the physical world. You seem to think that until we have 'solved' symbol grounding (which would at the very least involve a realistic theory of perception), *all* talk of physical interpretations must be banned. This doesnt seem like a good way to make progress, especially when this theory, even if incomplete, seems to offer some ways to analyse the problems that arise in symbol grounding.)

>Every criticism that has been hurled against AI projects that relate
>language to symbols in a computer can be leveled against the claim that
>Tarski's denotation function "solves" the problem of defining truth.
>
>Your claim that the "individuals" in Tarski's structures *are* the
>actual physical objects begs every question that Roger Schank begs
>when he claims that his GENSYMs *stand for* the actual physical objects.

It doesnt beg any questions at all. If I can talk of a set of, say, planets, then I can talk of an interpretation over a domain of planets. (Not *the* interpretation, notice, but *an* interpretation.) If you reply: but you can't talk of a set of planets before specifying exactly what "planet" means; then I may agree; but we are now having a metalevel discussion about the meaning of English words. I might respond, "set of planets in the ordinary sense of 'planet' "; or maybe we will come to agree that a planet is something orbiting a star; or whatever. Each of these ways

to get exact about what "planet" means might result in different things being allowed as a set of planets. But no matter what agreement we come to, each of these sets is, in fact, a set (of planets or something else), and that's all Tarski's construction needs.

>Using the verb 'are' instead of 'stand for' in no way solves the symbol
>grounding problem. Unlike Roger Schank, I am willing to agree that
>Tarski's approach solves one half of the problem, but you can't solve
>the symbol grounding problem by postulating it away.

>

>>I can't agree that we have a responsibility to provide grounding sockets
>>for the engineers to plug their sensors into, because (1) I don't think this
>>is central to our task - we aren't in the robot-building game here - and (2)
>>it's much too heavy a burden, since we can't be expected to provide
>>ready-made solutions to problems that haven't even been adequately formulated
>>yet. Grounding is an open problem, and we don't even know if the current -
>>very partial - ideas about it really work.

>

>I agree with you completely on this point. The point I was trying to make
>is that the database is the boundary between the data that Tarski's
>algorithm applies to and the data that the robot builders generate.
>The engineers are responsible for filling up the database with ground
>atoms about what is in the world, and the denotation function uses SQL
>to determine whether a particular sentence is true or false.

Tarski didn't define an 'algorithm', and the denotation function doesn't 'use' anything, just as the exponential function (say) doesn't 'use' anything. It's a mathematical function, not a computational process. (John, you have an amazing ability to stretch historical truth by misusing technical vocabulary.)

You seem to have an image of a kind of ideal robot whose sensors deliver ground atoms to it which are a complete ground-atomic description of its world, and it then uses some systematic inference process to test its quantified beliefs against that database. We could argue about whether this picture is at all convincing (either as robotics or psychology), but (a) it is a very particular picture, not shared by most, and (b) it is irrelevant here in any case, since we aren't doing robotics but general-purpose ontology design.

Pat

IHMC, University of West Florida	(850)434 8903	home
11000 University Parkway	(850)474 2091	office
Pensacola, FL 32514	(850)474 3023	fax

phayes@ai.uwf.edu
http://www.coginst.uwf.edu/~phayes

From ???@??? Mon Oct 05 15:20:41 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id PAA15628;
Mon, 5 Oct 1998 15:06:20 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a06b23eceb9d4a0@[143.88.7.118]>
In-Reply-To: <3.0.32.19981005105120.030bb488@catbert.cyc.com>
Mime-Version: 1.0
Date: Mon, 5 Oct 1998 15:03:20 -0500
To: Fritz Lehmann <fritz@cyc.com>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 3475
Status:

Hi Fritz

[John Sowa, to Pat:]
>> You cannot solve
>>the symbol grounding problem in a Cantorian transfinite universe.
>
>Right. It also bothers me that the entire apparatus of that universe
>depends on Cantor's Main Theorem (that there is no surjection of an
>infinite set onto its powerset), whose proof is isomorphic to the Barber
>Paradox, and whose proof fails in perfectly good set theories with
>universal sets, like Quine's New Foundations (NF) set theory.

Are you proposing that we create a universal high-level ontology which is based on a denial of Cantor's theorem and a bold declaration that (contrary to what most people have been led to believe) the set of real numbers is countable after all? If so, come out clearly and let's debate that proposal.

(For the record, I never claimed you could solve the grounding problem in a

transfinite universe, by the way. This comment of John's is on some crazy tangent from our original discussion. The proposal strikes me as a category error.)

[JS]

>>[...] But all available evidence indicates that the
>>universe is finite [...]

>

>I am unaware of any important evidence one way or the other.

There is lots of evidence that the universe is finite; that's the only hypothesis consistent with the big-bang idea, which is pretty much established science now. That is why the idea of the 'total mass of the universe' makes sense, for example.

However, it's so amazingly big that it might as well be infinite for computational purposes, of course :-)

>I generally agree with John Sowa in this (renewed) debate. I think a
>"model" can err. A model (it's annoying to have to call it "a Tarskian
>structure" especially since half of Tarski's work was the Hayes-disparaged
>algebraic approach which dispensed with them entirely) could contain the
>ground atomic assertion that Dusseldorf is in Wales. A sentence like (NOT
>(NOT (IN-REGION Dusseldorf Wales))) is true in the model, but false.

(Surely you are joking, Fritz?) Of course a Tarskian model can be "false" in this sense. Did anyone ever claim not? Model theory isn't concerned with the way the world *is*, but the ways it *could* be. Of course semantics isn't a source of infallibility.

However, I'd be interested to know what semantic theory justifies your claim that '(IN-REGION Dusseldorf Wales)' is false. I presume you mean the brackets to indicate a first-order sentence, and that it is false because the English sentence "Dusseldorf is in Wales" is false? That is so in virtue of it containing two proper names, among other things. What is your justification for claiming that a first-order constant symbol carries the meaning of an English proper name? (Or, if you like, for claiming that the FO constant symbol "Wales" denotes the land whose natives call Cymri?)

Pat Hayes

PS. I've never had the temerity to disparage any of Tarski's work. The term 'Tarskian' is in widespread use, and has been for years, because Tarski invented the concept. And again, for the record, we had agreed to use the term "interpretation", and I've been trying to use it scrupulously ever since to avoid confusion with the other sense of (engineering) 'model'

which I presume you didnt intend.

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Mon Oct 05 15:20:41 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id PAA15628;
Mon, 5 Oct 1998 15:06:20 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a06b23eceb9d4a0@[143.88.7.118]>
In-Reply-To: <3.0.32.19981005105120.030bb488@catbert.cyc.com>
Mime-Version: 1.0
Date: Mon, 5 Oct 1998 15:03:20 -0500
To: Fritz Lehmann <fritz@cyc.com>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 3475
Status:

Hi Fritz

[John Sowa, to Pat:]
>> You cannot solve
>>the symbol grounding problem in a Cantorian transfinite universe.
>
>Right. It also bothers me that the entire apparatus of that universe
>depends on Cantor's Main Theorem (that there is no surjection of an
>infinite set onto its powerset), whose proof is isomorphic to the Barber
>Paradox, and whose proof fails in perfectly good set theories with
>universal sets, like Quine's New Foundations (NF) set theory.

Are you proposing that we create a universal high-level ontology which is

based on a denial of Cantor's theorem and a bold declaration that (contrary to what most people have been led to believe) the set of real numbers is countable after all? If so, come out clearly and let's debate that proposal.

(For the record, I never claimed you could solve the grounding problem in a transfinite universe, by the way. This comment of John's is on some crazy tangent from our original discussion. The proposal strikes me as a category error.)

[JS]

>>[...] But all available evidence indicates that the
>>universe is finite [...]

>

>I am unaware of any important evidence one way or the other.

There is lots of evidence that the universe is finite; that's the only hypothesis consistent with the big-bang idea, which is pretty much established science now. That is why the idea of the 'total mass of the universe' makes sense, for example.

However, it's so amazingly big that it might as well be infinite for computational purposes, of course :-)

>I generally agree with John Sowa in this (renewed) debate. I think a
>"model" can err. A model (it's annoying to have to call it "a Tarskian
>structure" especially since half of Tarski's work was the Hayes-disparaged
>algebraic approach which dispensed with them entirely) could contain the
>ground atomic assertion that Dusseldorf is in Wales. A sentence like (NOT
>(NOT (IN-REGION Dusseldorf Wales))) is true in the model, but false.

(Surely you are joking, Fritz?) Of course a Tarskian model can be "false" in this sense. Did anyone ever claim not? Model theory isn't concerned with the way the world *is*, but the ways it *could* be. Of course semantics isn't a source of infallibility.

However, I'd be interested to know what semantic theory justifies your claim that '(IN-REGION Dusseldorf Wales)' is false. I presume you mean the brackets to indicate a first-order sentence, and that it is false because the English sentence "Dusseldorf is in Wales" is false? That is so in virtue of it containing two proper names, among other things. What is your justification for claiming that a first-order constant symbol carries the meaning of an English proper name? (Or, if you like, for claiming that the FO constant symbol "Wales" denotes the land whose natives call Cymri?)

Pat Hayes

PS. Ive never had the temerity to disparage any of Tarski's work. The term 'Tarskian' is in widespread use, and has been for years, because Tarski invented the concept. And again, for the record, we had agreed to use the term "interpretation", and Ive been trying to use it scrupulously ever since to avoid confusion with the other sense of (engineering) 'model' which I presume you didnt intend.

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Tue Oct 06 09:54:14 1998
Received: from philebus.tamu.edu (root@PHILEBUS.TAMU.EDU [165.91.161.22])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id UAA06410
for <phayes@coginst.uwf.edu>; Mon, 5 Oct 1998 20:02:12 -0500 (CDT)
Received: from philebus.tamu.edu (localhost [127.0.0.1])
by philebus.tamu.edu (8.8.5/8.8.5) with ESMTP id TAA00679;
Mon, 5 Oct 1998 19:57:33 -0500
Message-Id: <199810060057.TAA00679@philebus.tamu.edu>
X-Mailer: exmh version 2.0.2 2/24/98
To: Fritz Lehmann <fritz@cyc.com>
Cc: phayes@coginst.uwf.edu, Piek.Vossen@let.uva.nl, chezewiz@erols.com,
cmenzel@tamu.edu, doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu,
guarino@ladseb.pd.cnr.it, p.m.simons@leeds.ac.uk, skydog@pacbell.net,
sowa@west.poly.edu
Subject: Re: Tropes
In-reply-to: Your message of "Mon, 05 Oct 1998 10:51:46 CDT."
<3.0.32.19981005105120.030bb488@catbert.cyc.com>
Mime-Version: 1.0
Date: Mon, 05 Oct 1998 19:57:33 -0500
From: Chris Menzel <cmenzel@philebus.tamu.edu>
Content-Type: text/plain; charset=us-ascii
Content-Length: 1942
Status:

John wrote:

>> You cannot solve
>>the symbol grounding problem in a Cantorian transfinite universe.

Fritz wrote:

> Right. It also bothers me that the entire apparatus of that universe
> depends on Cantor's Main Theorem (that there is no surjection of an
> infinite set onto its powerset), whose proof is isomorphic to the Barber
> Paradox, and whose proof fails in perfectly good set theories with
> universal sets, like Quine's New Foundations (NF) set theory.

I don't know what what criteria you use for counting theories as perfectly good, Fritz, but I count it against a theory (esp in the context of ontology) if it flouts strong intuitions or adopts principles with no intuitive justification. On both scores, NF fares poorly indeed. For instance, in NF, there is a set (the universal set, for example, though there's lots more) that is not equinumerous to the set of its singleton subsets. Huh? Furthermore, NF has never been proved consistent relative to ZF -- no surprise since no one has any clue what the universe of NF's sets looks like, in dramatic contrast to the clear (not to say *wholly* clear) image of the cumulative hierarchy one gets as the natural model of ZF. And what motivation is there for restricting Comprehension to stratified formulas? None, other than that the restriction seems to block the paradoxes. By contrast, the restriction in the Axiom of Separation is completely justified in the cumulative hierarchy -- some conditions apply to sets that occur arbitrarily high up in the hierarchy, and hence there is never any "level" of the hierarchy at which all the sets that satisfy the condition is "constructed", hence there is no *set* of things satisfying the condition.

NF is a kludge. An interesting, even mathematically rich kludge (witness the recent Oxford Logic Guides book), but a kludge. It is in particular not a viable candidate for a basic set theory for ontology.

Cheers!

-chris

From ???@??? Tue Oct 06 09:54:19 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id FAA17282
for <phayes@coginst.uwf.edu>; Tue, 6 Oct 1998 05:50:19 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id GAA18827;
Tue, 6 Oct 1998 06:45:58 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
id GAA27126; Tue, 6 Oct 1998 06:40:30 -0400
Date: Tue, 6 Oct 1998 06:40:30 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199810061040.GAA27126@west>

To: fritz@cyc.com

Subject: Re: Tropes

Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@tamu.edu,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu,
guarino@ladseb.pd.cnr.it, jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk,
phayes@coginst.uwf.edu, skydog@pacbell.net, sowa@west.poly.edu

Mime-Version: 1.0

Content-Type: text/plain; charset=us-ascii

Content-Length: 1595

Status:

Fritz,

>Right. It also bothers me that the entire apparatus of that universe
>depends on Cantor's Main Theorem (that there is no surjection of an
>infinite set onto its powerset), whose proof is isomorphic to the Barber
>Paradox, and whose proof fails in perfectly good set theories with
>universal sets, like Quine's New Foundations (NF) set theory.

Yes, mathematical systems that have had that much research pumped into them usually have multiple pathways of approach to their main theorems. But the unusual feature of all the uncountable stuff is that it hangs on the slender thread of a single proof by contradiction. As I said before, it was fun to study in my youth (like playing chess), but I adopt an agnostic stance towards it now. I won't say it's wrong, but I won't bother to waste any time on theories that critically depend on it.

>>[...] But all available evidence indicates that the
>>universe is finite [...]

>I am unaware of any important evidence one way or the other.

Astronomers today generally agree that the Big Bang theory is the most likely explanation of everything they have observed so far.

However, there are lots of theories about an infinite foam of bubbles, each of which is its own universe with different physical parameters. Our own universe may be just one big bubble that is inaccessible from the others.

In any case, all those possibilities are irrelevant to the main point that everything we or our computers can know or think about is finite. Even Cantor's theorems are metalevel thoughts about infinities, not actually infinite thoughts.

John

From ???@??? Tue Oct 06 12:01:35 1998
Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id KAA05877
for <phayes@coginst.uwf.edu>; Tue, 6 Oct 1998 10:29:58 -0500 (CDT)
Received: from scratchy (scratchy [207.207.8.118])
by catbert.cyc.com (8.8.8/8.8.8) with SMTP id KAA14047;
Tue, 6 Oct 1998 10:25:18 -0500 (CDT)
Message-Id: <3.0.32.19981006102559.00a369c0@catbert.cyc.com>
X-Sender: fritz@catbert.cyc.com
X-Mailer: Windows Eudora Pro Version 3.0 (32)
Date: Tue, 06 Oct 1998 10:26:02 -0500
To: Chris Menzel <cmenzel@philebus.tamu.edu>
From: Fritz Lehmann <fritz@cyc.com>
Subject: Re: Tropes
Cc: Fritz Lehmann <fritz@cyc.com>, phayes@coginst.uwf.edu,
Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@tamu.edu,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu,
guarino@ladseb.pd.cnr.it, p.m.simons@leeds.ac.uk, skydog@pacbell.net,
sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset="us-ascii"
Content-Length: 1525
Status:

At 07:57 PM 10/5/98 -0500, Chris Menzel wrote:

>John wrote:

>>> You cannot solve

>>> the symbol grounding problem in a Cantorian transfinite universe.

>

>Fritz wrote:

>> Right. It also bothers me that the entire apparatus of that universe
>> depends on Cantor's Main Theorem (that there is no surjection of an
>> infinite set onto its powerset), whose proof is isomorphic to the Barber
>> Paradox, and whose proof fails in perfectly good set theories with
>> universal sets, like Quine's New Foundations (NF) set theory.

>

>I don't know what what criteria you use for counting theories as perfectly
>good, Fritz,

At least the reference is appreciated ...

>[...] Furthermore, NF has never been proved consistent relative
>to ZF [...]

Will that be a flaw if ZF proves to be inconsistent?

>[...]

>NF is a kludge. An interesting, even mathematically rich kludge (witness
>the recent Oxford Logic Guides book), but a kludge. It is in particular
>not a viable candidate for a basic set theory for ontology.

I do not propose NF as the basis for set theory for our ontology; I
recommend a practical, agnostic, non-committal ontology which allows people
to pick any (or none) of the many, divergent-at-infinity axiomatizations of
set theory that work OK for finite sets.

>-chris

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

=====
=====

From ???@??? Wed Oct 07 10:09:44 1998

Received: from philebus.tamu.edu (root@PHILEBUS.TAMU.EDU [165.91.161.22])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id PAA27229
for <phayes@coginst.uwf.edu>; Tue, 6 Oct 1998 15:39:54 -0500 (CDT)

Received: from philebus.tamu.edu (localhost [127.0.0.1])
by philebus.tamu.edu (8.8.5/8.8.5) with ESMTP id PAA05189;
Tue, 6 Oct 1998 15:35:46 -0500

Message-Id: <199810062035.PAA05189@philebus.tamu.edu>

X-Mailer: exmh version 2.0.2 2/24/98

To: Fritz Lehmann <fritz@cyc.com>

cc: phayes@coginst.uwf.edu, Piek.Vossen@let.uva.nl, chezewiz@erols.com,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu,
guarino@ladseb.pd.cnr.it, p.m.simons@leeds.ac.uk, skydog@pacbell.net,
sowa@west.poly.edu

Subject: Re: Tropes

In-reply-to: Your message of "Tue, 06 Oct 1998 10:26:02 CDT."
<3.0.32.19981006102559.00a369c0@catbert.cyc.com>

Mime-Version: 1.0

Date: Tue, 06 Oct 1998 15:35:46 -0500

From: Chris Menzel <cmenzel@philebus.tamu.edu>

Content-Type: text/plain; charset=us-ascii

Content-Length: 3038

Status:

Fritz wrote:

>>> Right. It also bothers me that the entire apparatus of that universe

> >> depends on Cantor's Main Theorem (that there is no surjection of an
> >> infinite set onto its powerset), whose proof is isomorphic to the Barber
> >> Paradox, and whose proof fails in perfectly good set theories with
> >> universal sets, like Quine's New Foundations (NF) set theory.
> >
> > I don't know what what criteria you use for counting theories as perfectly
> > good, Fritz,
>
> At least the reference is appreciated ...

No disrespect intended, Fritz! It was just curious to me why you (whose judgment in matters ontological I hold in the highest esteem!) would think of NF as a perfectly good set theory -- especially given your interest in having a theory that is intuitive and natural. NF is neither.

> > [...] Furthermore, NF has never been proved consistent relative
> > to ZF [...]
>
> Will that be a flaw if ZF proves to be inconsistent?

Certainly not, since then every theory will be (trivially) consistent relative to ZF. (And of course the last 90 years strongly suggests that the likelihood of finding a contradiction in ZF is about as high as that of finding one in Peano Arithmetic.) The point I was trying to make is that pretty much every other set theory can be proved consistent relative to ZF or some extension thereof. These proofs generally consist in the construction of models of the theories that provide pretty vivid pictures of what their intended universes looks like. NF has resisted similar treatment, a fact underscoring its odd and unintuitive nature -- no one can provide any sort of picture of what the universe of sets looks like under NF. That is a bad thing for any theory, especially one as fundamental as set theory.

> > [...]
> > NF is a kludge. An interesting, even mathematically rich kludge (witness
> > the recent Oxford Logic Guides book), but a kludge. It is in particular
> > not a viable candidate for a basic set theory for ontology.
>
> I do not propose NF as the basis for set theory for our ontology;

Nor did I ever assert or imply that you did. I was only reacting to your assertion that NF is perfectly natural, and the implication that it would do as well as any other as a basis for talking about sets in the context of ontology. I think it is important to say why this is false. So I did.

> I recommend a practical, agnostic, non-committal ontology which allows people

> to pick any (or none) of the many, divergent-at-infinity axiomatizations of
> set theory that work OK for finite sets.

Right. All reasonable set theories should agree on those. Problem is, we want at least the natural numbers around, don't we? If so, it is difficult not to confront the issues of infinite sets that quickly arise out of our number ontology. I think it is clear how to deal with transfinities in ontology for most purposes, even those involving the numbers, viz., scotch the troublesome power set axiom (and perhaps more besides, as in KPU).

Highest regards,

-chris

From ???@??? Wed Oct 07 10:09:51 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id VAA20329
for <phayes@coginst.uwf.edu>; Tue, 6 Oct 1998 21:04:55 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id WAA20027;
Tue, 6 Oct 1998 22:00:56 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
id VAA01942; Tue, 6 Oct 1998 21:55:44 -0400
Date: Tue, 6 Oct 1998 21:55:44 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199810070155.VAA01942@west>
To: cmenzel@philebus.tamu.edu, fritz@cyc.com
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, doug@csi.uottawa.ca,
e6nl001@coe.coppin.umd.edu, guarino@ladseb.pd.cnr.it,
p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu, skydog@pacbell.net,
sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 4651
Status:

Pat, Chris, Fritz, et al.,

I think that these email exchanges have sufficiently established the point that we can all agree quickly (i.e. after one or two brief comments) about technical issues. But the issues we fight about endlessly are primarily terminological and turf wars (e.g. what is or is not included in what may or may not be called Tarskian model theory).

What I suggest is that we itemize a list of issues that should be considered for inclusion in the ontology report, and then decide what to name them after we have agreed on what should be included. Following are some issues we discussed in recent email exchanges. I do not claim that this is a definitive list, but only that it is a starting point:

1. Somewhere in the ontology there should be categories and theories for all the major mathematical structures, starting with integers and real numbers and going on to set theory, mereology, and variations thereof. Cantorian stuff and Lesniewskian stuff should all find a place somewhere in there, and we should not impose our preferences on what anyone else may select. We should also make room for various computer-oriented datatypes, including whatever standards ANSI, IEEE, ISO, and other standards bodies recommend.
2. I agree that Tarski presented his definition in formal terms, which did not constitute an "algorithm" in the modern sense. However, his definitions can be translated to an equivalent algorithm that is computable in polynomial time (unlike satisfiability, which is an NP-complete problem for which the known algorithms are exponential). The existence of efficient algorithms, although not directly relevant to the mathematical definitions, is of enormous practical importance for most, if not all of the potential users of our ontologies for domains other than the mathematical ones in point #1 above. (And by the way, I am much less hung up on computability than the KL-ONE crowd, who try to limit their languages to what is tractable.)
3. I agree that the term 'symbol grounding' in AI has become associated with the robotics issues, but the term 'grounding' was an important topic long before robots or computers were conceived. The people who have been working on such issues for the longest time are the engineers and scientists. They have developed sophisticated, but highly domain dependent techniques for dealing with experimental error, granularity, tolerance, etc. I am not suggesting that we delve into all that detail, but I believe we should recognize that they have done more significant work on the problem of how symbols relate to the physical world than any of the logicians. I am not claiming that this is or should be part of "Tarskian model theory", but I believe that the issues of how statements are related to the world must take into account the kinds of things that the engineers and physicists have been doing quite successfully.
4. I have often quoted Quine's dictum "To be is to be the value of a quantified variable" as a useful test for determining what are the implicit entities presupposed by a knowledge representation. It is

an example of the kind of thing that I believe we should be doing more consistently: developing well-defined formulas and tests that can help a knowledge engineer dig out the relevant aspects of a KRep that must be formalized. And to respond to Fritz's question: Quine did mention that his technique can be generalized to languages that do not have explicit quantifiers or explicit variables, including one of his own languages, which used combinatorics to get rid of explicit variables (cf. the paper "Variables Explained Away").

5. I cited the examples of aspen trees and pieces of cheese to emphasize the importance of identity criteria for individuals as a critical issue in relating symbols to physical things. Tarski did not talk about identity criteria, since for the kinds of mathematical domains he was working with, they were not necessary. But they are an essential aspect of a theory of reference, which is a precondition for any kind of theory of truth that relates sentences to the world. The issue of whether it belongs to model theory is irrelevant; it is something that is significant for the customers for our ontologies.

I'm sure that there are many other issues that are relevant to the ontology business, and I hope we can find ways of talking about them among ourselves and to our potential users.

John

From ???@??? Wed Oct 07 14:06:28 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id NAA02331;
Wed, 7 Oct 1998 13:42:27 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a06b24144d1f568@[143.88.7.118]>
In-Reply-To: <199810070155.VAA01942@west>
Mime-Version: 1.0
Date: Wed, 7 Oct 1998 13:39:28 -0500
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, doug@csi.uottawa.ca,
e6nl001@coe.coppin.umd.edu, guarino@ladseb.pd.cnr.it,
p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu, skydog@pacbell.net,
cmenzel@philebus.tamu.edu, fritz@cyc.com
Content-Type: text/plain; charset="us-ascii"
Content-Length: 9571
Status:

John,

I appreciate the conciliatory tone of your message and agree with much of it. And I'm sorry that these debates have left many people feeling tired and fed up. Nevertheless, there are still places where you have stated errors as though they were facts, and I worry that these errors might get woven into the document in ways which (history tells us) will continue to give rise to future confusions and communication problems. I've marked these below with triple asterisks so that those who aren't interested can ignore them more easily.

>What I suggest is that we itemize a list of issues that should be
>considered for inclusion in the ontology report, and then decide what
>to name them after we have agreed on what should be included. Following
>are some issues we discussed in recent email exchanges. I do not claim
>that this is a definitive list, but only that it is a starting point:

>

- > 1. Somewhere in the ontology there should be categories and theories
- > for all the major mathematical structures, starting with integers
- > and real numbers....

Why don't we simply state that certain well-known mathematical structures are being *assumed* by our ontology, rather than providing ontologies for them? I have in mind integers and reals (and finite groups, etc.; any mathematical structures which are well understood and adequately formalised.) This way we don't need to reinvent old wheels, and we can leave the users to invoke whatever formalisations they want to use (even NF set theory :-), and we keep carefully away from all pre-existing turf battlefields. And this isn't any kind of abdication, since there's no shortage of such formalisations. (I did this in one temporal ontology, for example, where I simply *assumed* the integers as having been provided, and defined time in terms of them, being careful to say that the arithmetic had to come from somewhere else.)

- > 2. I agree that Tarski presented his definition in formal terms, which
- > did not constitute an "algorithm" in the modern sense. However, his
- > definitions can be translated to an equivalent algorithm that is
- > computable in polynomial time (unlike satisfiability, which is an
- > NP-complete problem for which the known algorithms are exponential).

This only makes sense relative to an RDB, ie a Herbrand interpretation. Do you mean that there is a polynomial algorithm which follows the Tarski recursion (substituting a complete RDB for an interpretation)?

- > The existence of efficient algorithms, although not directly relevant
- > to the mathematical definitions, is of enormous practical importance

- > for most, if not all of the potential users of our ontologies for
- > domains other than the mathematical ones in point #1 above. (And
- > by the way, I am much less hung up on computability than the KL-ONE
- > crowd, who try to limit their languages to what is tractable.)

That's an interesting claim. Why is it of such practical importance? I can see one reason: if the ontology is to be used as a data model for a database, then this algorithm could be used to check that the DB conforms to the data model, ie that the ontological axioms are true in the Herbrand interpretation defined by the RDB. Is this the application you have in mind?

- > 3. I agree that the term 'symbol grounding' in AI has become associated
- > with the robotics issues, but the term 'grounding' was an important
- > topic long before robots or computers were conceived. The people who
- > have been working on such issues for the longest time are the engineers
- > and scientists. They have developed sophisticated, but highly domain
- > dependent techniques for dealing with experimental error, granularity,
- > tolerance, etc. I am not suggesting that we delve into all that detail,
- > but I believe we should recognize that they have done more significant
- > work on the problem of how symbols relate to the physical world than
- > any of the logicians. I am not claiming that this is or should be
- > part of "Tarskian model theory", but I believe that the issues of
- > how statements are related to the world must take into account the
- > kinds of things that the engineers and physicists have been doing
- > quite successfully.

*** I'm fairly aware of what 'engineers and physicists' have done in these areas. As a carpenter and clock-restorer, I use the engineering notions of tolerance constantly myself, and spend almost as much time in email correspondence with physicists (about the meaning of theoretical terms) and psychologists (about experimental design in psycholinguistics) as I do with you, John. And I think you are wrong: the philosophical issues of meaning aren't addressed by these people at all. They simply use language to refer, just like everyone else does, without being concerned about how it is that, say, "hospitals" refers to hospitals (or even how it is that "within .001 cm" means within 0.001 cm.) In fact, like almost everyone else, they are rather puzzled by what on earth it is that bothers anyone who worries about semantics and meaning. They don't have any way to deal with such classical problems as the paradox of the heap (which is the reef on which all attempts to formalize granularity still founder); like everyone else, they rely on robust common sense to not get caught in such difficulties. For example, house framers use a reference pole on which all the stud lengths, window beam heights, etc. are marked and from which all other measurements must be taken. It is easy to explain why this is a good idea: without it, acceptably small errors of transcription from one stud to the stock from which the next is cut, especially if systematic, can add up to an

unacceptably large error. That's kind of obvious; but I don't know any good way to *formalise* it. Scientists, engineers and carpenters just aren't in the formalisation business. All of which is not to disparage this stuff, of course, or even to say it is wholly irrelevant to us. But I do think that we should mention these issues (of experimental error, granularity, tolerance, etc.) only where they are relevant, rather than insisting that they must be acknowledged before even allowing ourselves to speak of the world at all. ***

> 4. I have often quoted Quine's dictum "To be is to be the value of a
> quantified variable" as a useful test for determining what are the
> implicit entities presupposed by a knowledge representation. It is
> an example of the kind of thing that I believe we should be doing
> more consistently: developing well-defined formulas and tests that
> can help a knowledge engineer dig out the relevant aspects of a KRep
> that must be formalized. And to respond to Fritz's question: Quine
> did mention that his technique can be generalized to languages that
> do not have explicit quantifiers or explicit variables, including
> one of his own languages, which used combinatorics to get rid of
> explicit variables (cf. the paper "Variables Explained Away").

I agree with the above.

> 5. I cited the examples of aspen trees and pieces of cheese to emphasize
> the importance of identity criteria for individuals as a critical issue
> in relating symbols to physical things. Tarski did not talk about
> identity criteria, since for the kinds of mathematical domains he was
> working with, they were not necessary. But they are an essential
> aspect of a theory of reference, which is a precondition for any kind
> of theory of truth that relates sentences to the world. The issue of
> whether it belongs to model theory is irrelevant; it is something that
> is significant for the customers for our ontologies.

I think we all agree that identity criteria are central. *** But I can't agree with your implication here that Tarskian semantics is concerned only with 'mathematical domains', where this is an ontological classification to be contrasted with 'physical domains'. This usage begs the question which we have been arguing about. First, it embodies what seems to me to be a philosophical error; in my view, mathematics is simply a language for talking about things, so there are no 'mathematical domains'. I concede that this is rather a strong anti-Platonist view and many don't share it; but my main objection is that even if your distinction is accepted by someone who thinks that mathematics is about abstract stuff, model theory applies just as well to physical domains as to any other kind.***

***This isn't just a turf war about what is to be called 'tarskian'; the

point is rather whether it is appropriate to use this useful conceptual tool - model theory - in our own deliberations. Your constantly reiterated view, re-embodied in the above paragraph, implies that we cannot; or at any rate that its used must be restricted in ways that are artificial and counter to the way the theory is used almost everywhere else in the academic community.***

BTW, getting identity criteria clear isnt always trivial even in 'mathematical domains' , as some of our own discussions at Heidelberg should have testified! (How many isomorphic graphs are there?)

Pat

PS. You know, maybe if we just got on with the actual work, we might find something to agree about. :-)

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Oct 07 14:06:28 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id NAA02331;
Wed, 7 Oct 1998 13:42:27 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a06b24144d1f568@[143.88.7.118]>
In-Reply-To: <199810070155.VAA01942@west>
Mime-Version: 1.0
Date: Wed, 7 Oct 1998 13:39:28 -0500
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, doug@csi.uottawa.ca,
e6nl001@coe.coppin.umd.edu, guarino@ladseb.pd.cnr.it,
p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu, skydog@pacbell.net,
cmenzel@philebus.tamu.edu, fritz@cyc.com
Content-Type: text/plain; charset="us-ascii"
Content-Length: 9571
Status:

John,

I appreciate the conciliatory tone of your message and agree with much of it. And I'm sorry that these debates have left many people feeling tired and fed up. Nevertheless, there are still places where you have stated errors as though they were facts, and I worry that these errors might get woven into the document in ways which (history tells us) will continue to give rise to future confusions and communication problems. I've marked these below with triple asterisks so that those who aren't interested can ignore them more easily.

>What I suggest is that we itemize a list of issues that should be
>considered for inclusion in the ontology report, and then decide what
>to name them after we have agreed on what should be included. Following
>are some issues we discussed in recent email exchanges. I do not claim
>that this is a definitive list, but only that it is a starting point:
>
> 1. Somewhere in the ontology there should be categories and theories
> for all the major mathematical structures, starting with integers
> and real numbers....

Why don't we simply state that certain well-known mathematical structures are being *assumed* by our ontology, rather than providing ontologies for them? I have in mind integers and reals (and finite groups, etc.; any mathematical structures which are well understood and adequately formalised.) This way we don't need to reinvent old wheels, and we can leave the users to invoke whatever formalisations they want to use (even NF set theory :-), and we keep carefully away from all pre-existing turf battlefields. And this isn't any kind of abdication, since there's no shortage of such formalisations. (I did this in one temporal ontology, for example, where I simply *assumed* the integers as having been provided, and defined time in terms of them, being careful to say that the arithmetic had to come from somewhere else.)

> 2. I agree that Tarski presented his definition in formal terms, which
> did not constitute an "algorithm" in the modern sense. However, his
> definitions can be translated to an equivalent algorithm that is
> computable in polynomial time (unlike satisfiability, which is an
> NP-complete problem for which the known algorithms are exponential).

This only makes sense relative to an RDB, ie a Herbrand interpretation. Do you mean that there is a polynomial algorithm which follows the Tarski recursion (substituting a complete RDB for an interpretation)?

- > The existence of efficient algorithms, although not directly relevant
- > to the mathematical definitions, is of enormous practical importance
- > for most, if not all of the potential users of our ontologies for
- > domains other than the mathematical ones in point #1 above. (And
- > by the way, I am much less hung up on computability than the KL-ONE
- > crowd, who try to limit their languages to what is tractable.)

That's an interesting claim. Why is it of such practical importance? I can see one reason: if the ontology is to be used as a data model for a database, then this algorithm could be used to check that the DB conforms to the data model, ie that the ontological axioms are true in the Herbrand interpretation defined by the RDB. Is this the application you have in mind?

- > 3. I agree that the term 'symbol grounding' in AI has become associated
- > with the robotics issues, but the term 'grounding' was an important
- > topic long before robots or computers were conceived. The people who
- > have been working on such issues for the longest time are the engineers
- > and scientists. They have developed sophisticated, but highly domain
- > dependent techniques for dealing with experimental error, granularity,
- > tolerance, etc. I am not suggesting that we delve into all that detail,
- > but I believe we should recognize that they have done more significant
- > work on the problem of how symbols relate to the physical world than
- > any of the logicians. I am not claiming that this is or should be
- > part of "Tarskian model theory", but I believe that the issues of
- > how statements are related to the world must take into account the
- > kinds of things that the engineers and physicists have been doing
- > quite successfully.

*** I'm fairly aware of what 'engineers and physicists' have done in these areas. As a carpenter and clock-restorer, I use the engineering notions of tolerance constantly myself, and spend almost as much time in email correspondence with physicists (about the meaning of theoretical terms) and psychologists (about experimental design in psycholinguistics) as I do with you, John. And I think you are wrong: the philosophical issues of meaning aren't addressed by these people at all. They simply use language to refer, just like everyone else does, without being concerned about how it is that, say, "hospitals" refers to hospitals (or even how it is that "within .001 cm" means within 0.001 cm.) In fact, like almost everyone else, they are rather puzzled by what on earth it is that bothers anyone who worries about semantics and meaning. They don't have any way to deal with such classical problems as the paradox of the heap (which is the reef on which all attempts to formalize granularity still founder); like everyone else, they rely on robust common sense to not get caught in such difficulties. For example, house framers use a reference pole on which all the stud lengths, window beam heights, etc. are marked and from which all other measurements must be taken. It is easy to explain why this is a good idea: without it,

acceptably small errors of transcription from one stud to the stock from which the next is cut, especially if systematic, can add up to an unacceptably large error. That's kind of obvious; but I don't know any good way to *formalise* it. Scientists, engineers and carpenters just aren't in the formalisation business. All of which is not to disparage this stuff, of course, or even to say it is wholly irrelevant to us. But I do think that we should mention these issues (of experimental error, granularity, tolerance, etc.) only where they are relevant, rather than insisting that they must be acknowledged before even allowing ourselves to speak of the world at all. ***

> 4. I have often quoted Quine's dictum "To be is to be the value of a
> quantified variable" as a useful test for determining what are the
> implicit entities presupposed by a knowledge representation. It is
> an example of the kind of thing that I believe we should be doing
> more consistently: developing well-defined formulas and tests that
> can help a knowledge engineer dig out the relevant aspects of a KRep
> that must be formalized. And to respond to Fritz's question: Quine
> did mention that his technique can be generalized to languages that
> do not have explicit quantifiers or explicit variables, including
> one of his own languages, which used combinatorics to get rid of
> explicit variables (cf. the paper "Variables Explained Away").

I agree with the above.

> 5. I cited the examples of aspen trees and pieces of cheese to emphasize
> the importance of identity criteria for individuals as a critical issue
> in relating symbols to physical things. Tarski did not talk about
> identity criteria, since for the kinds of mathematical domains he was
> working with, they were not necessary. But they are an essential
> aspect of a theory of reference, which is a precondition for any kind
> of theory of truth that relates sentences to the world. The issue of
> whether it belongs to model theory is irrelevant; it is something that
> is significant for the customers for our ontologies.

I think we all agree that identity criteria are central. *** But I can't agree with your implication here that Tarskian semantics is concerned only with 'mathematical domains', where this is an ontological classification to be contrasted with 'physical domains'. This usage begs the question which we have been arguing about. First, it embodies what seems to me to be a philosophical error; in my view, mathematics is simply a language for talking about things, so there are no 'mathematical domains'. I concede that this is rather a strong anti-Platonist view and many don't share it; but my main objection is that even if your distinction is accepted by someone who thinks that mathematics is about abstract stuff, model theory applies just as well to physical domains as to any other kind. ***

This isn't just a turf war about what is to be called 'tarskian'; the point is rather whether it is appropriate to use this useful conceptual tool - model theory - in our own deliberations. Your constantly reiterated view, re-embodied in the above paragraph, implies that we cannot; or at any rate that its used must be restricted in ways that are artificial and counter to the way the theory is used almost everywhere else in the academic community.

BTW, getting identity criteria clear isnt always trivial even in 'mathematical domains' , as some of our own discussions at Heidelberg should have testified! (How many isomorphic graphs are there?)

Pat

PS. You know, maybe if we just got on with the actual work, we might find something to agree about. :-)

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Thu Oct 08 10:16:54 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id WAA06638
for <phayes@coginst.uwf.edu>; Wed, 7 Oct 1998 22:00:16 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id WAA21668;
Wed, 7 Oct 1998 22:56:05 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
id WAA10870; Wed, 7 Oct 1998 22:50:44 -0400
Date: Wed, 7 Oct 1998 22:50:44 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199810080250.WAA10870@west>
To: phayes@coginst.uwf.edu
Subject: Re: Tropes
Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@philebus.tamu.edu,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu, fritz@cyc.com,
guarino@ladseb.pd.cnr.it, p.m.simons@leeds.ac.uk, skydog@pacbell.net,

sowa@west.poly.edu

Mime-Version: 1.0

Content-Type: text/plain; charset=us-ascii

Content-Length: 5553

Status:

Pat,

The major disagreement between us is not caused by my "errors", but by my conscious decision about how the multiple uses of the word 'model' can be reconciled. Section 1.1 of my 1984 book, which I still find quite persuasive, is entitled "Knowledge and Models". In that section and throughout the remainder of the book, I build a strong case for the point that the multiple uses of the word 'model' in logic, science, engineering, and computer science are not unrelated homonyms, but different aspects of a common central meaning.

As you know, the world is divided into "splitters" and "lumpers". When it comes to terminology, I have strong preferences for the lumping approach together with a strong historical sense in which I try to look for commonalities in the way the different senses have evolved.

In that earlier book, I cited with approval C. A. Petri's observation that there is an affinity between the engineers' and the logicians' use of the term. I also cited Kenneth Craik (1943) who stated that the brain was a "machine for making models", a comment which Minsky picked up and promoted quite convincingly, in my opinion. I believe one can put together strong arguments for the point that mental models in the head, structures of GENSYMs in AI programs, abstract relational structures consisting of a set of individuals together with a collection of relations over those individuals, architectural blueprints, engineering plans and drawings, and physical devices built according to those plans all serve a common purpose for which the word 'model' is the best label available in the English language.

I realize that you are likely to call my usage "idiosyncratic", but I would like to ask who in AI, besides yourself, you would consider *not* to be idiosyncratic in underlying philosophy and terminology.

In any case, I realize that for a consensus document, such as the one we are planning to produce, we should be fairly conservative in our choice of terminology. That is, unfortunately, quite a challenge, since we have to reconcile contributions from multiple disciplines, most of which have evolved different and usually conflicting terminology for similar concepts. That was one of my motivations for my "lumping" in my previous book, and I don't see any way to develop a satisfactory

terminology for this proposed ontology document without doing some degree of lumping.

When I sent around those 5 points, I deliberately avoided suggesting any names for them (other than the words used in their description, which were not intended to be candidates for the primary names).

A few comments about your recent comments:

>Why dont we simply state that certain well-known mathematical structure are
>being *assumed* by our ontology, rather than providing ontologies for them?

I agree. I think that we should list them and either copy the basic axioms (e.g. Peano's axioms for arithmetic) or cite the best source we can find. But we do have a responsibility for showing where they fit in relationship to other related categories.

>This only makes sense relative to an RDB, ie a Herbrand intepretation. Do
>you mean that there is a polynomial algorithm which follows the tarski
>recursion (substituting a complete RDB for an interpretation)?

Yes, for finite relational structures, the RDB is isomorphic to the Tarski-style interpretation (given the closed-world assumption).

>Thats an interesting claim. Why is it of such practical importance? I can
>see one reason: if the ontology is to be used as a data model for a
>database, then this algorithm could be used to check that the DB conforms
>to the data model, ie that the ontological axioms are true in the Herbrand
>interpretation defined by the RDB. Is this the application you have in mind?

Yes. I consider the development of resources for building knowledge bases and databases to be one of the primary goals of this ontology project.

>,,, Scientists, engineers and carpenters just arent in
>the formalisation business. All of which is not to disparage this stuff, of
>course, or even to say it is wholly irrelevant to us. But I do think that
>we should mention these issues (of experimental error, granularity,
>tolerance, etc.) only where they are relevant, rather than insisting that
>they must be acknowledged before even allowing ourselves to speak of the
>world at all. ***

Yes, I agree. But there is a lot of work by philosophers of science and by the more philosophically inclined scientists that is quite relevant. Again, I don't mean that we should reinvent the wheel. Much of it can be "included by reference" when appropriate.

>someone who thinks that mathematics is about abstract stuff, model theory
>applies just as well to physical domains as to any other kind.***

Yes, I agree that it is applicable. And as I said before, I prefer not to continue the arguments about where model theory starts or stops or where different developers of the formalism happened to draw the lines that distinguish the theory from its applications.

>BTW, getting identity criteria clear isnt always trivial even in
>'mathematical domains', as some of our own discussions at Heidelberg
>should have testified! (How many isomorphic graphs are there?)

No, I wouldn't say it's trivial, but at least the definitions and axioms are under the control of the mathematician -- unlike the physical world which has a "mind" of its own.

>PS. You know, maybe if we just got on with the actual work, we might find
>something to agree about. :-)

I hope so.

John

From ???@??? Thu Oct 08 15:22:55 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id PAA18554
for <phayes@coginst.uwf.edu>; Thu, 8 Oct 1998 15:12:55 -0500 (CDT)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id QAA22616;
Thu, 8 Oct 1998 16:08:51 -0400 (EDT)
Received: by west (SMI-8.6/SMI-SVR4)
id QAA15441; Thu, 8 Oct 1998 16:03:36 -0400
Date: Thu, 8 Oct 1998 16:03:36 -0400
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199810082003.QAA15441@west>
To: phayes@coginst.uwf.edu
Subject: Re: More **** (was: Re: Tropes)
Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@philebus.tamu.edu,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu, fritz@cyc.com,
guarino@ladseb.pd.cnr.it, p.m.simons@leeds.ac.uk, skydog@pacbell.net
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 1282
Status:

Pat,

In your response to my last note, you said that you appreciated my "conciliatory" tone. That was because the phone line dropped while I was composing my first response.

Your list of nonidiosyncratic logicians is pure fantasy on your part.

Rather than respond in detail, which is looking more and more like a waste of time, I suggest you read Quine's book Roots of Reference.

I don't always agree with Quine -- he voted for Nixon in 1972 and he is too nominalistic in temperament. But that book is something that expresses my concerns about the mapping from language (formal or otherwise) to the world in very clear terms. In that book, Quine also mentions models a bit -- he usually puts far more emphasis on proof techniques than on models -- and nowhere would he even dream of making your wild-eyed claims that it is sufficient to "define" the individuals in a model to be things in the world.

Another of Quine's papers that I strongly urge you to read is his response to Saul Kripke, which is reprinted in Theories and Things.

Quine's vocabulary and general style are remarkably free of jargon and claims about what is or is not an official part of "model theory" or any other kind of received gospel. I suggest it as a "model" for our proposed document.

John

From ???@??? Mon Oct 12 16:44:34 1998

Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id NAA02044;
Mon, 12 Oct 1998 13:39:49 -0500 (CDT)

X-Sender: phayes@mail.coginst.uwf.edu (Unverified)

Message-Id: <v04003a07b247ea2cbc3d@[143.88.7.118]>

In-Reply-To: <199810101558.LAA26722@west>

Mime-Version: 1.0

Date: Mon, 12 Oct 1998 13:36:59 -0500

To: sowa@west.poly.edu (John F. Sowa)

From: Pat Hayes <phayes@coginst.uwf.edu>

Subject: Re: Standard arithmetics

Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net

Content-Type: text/plain; charset="us-ascii"

Content-Length: 2922

Status:

>Pat wrote

>

>>..., I think that our ontological responsibility can be

>>just to *say* what the relation names are and what computations are

>>supposed to be attached to them.

>

>Saying "what the relation names are" sounds like what Cyc did by

>contributing the names (and partial ordering) of their top level

>without giving the associated axioms. But saying "what computations

>are supposed to be attached to them" requires some kind of specification,

>which sounds very much like giving axioms.

What I meant was that for functions like plus and times we could just say something like 'using any normal arithmetic', or 'in their usual senses', and allow either axioms or procedural attachment, or both. The suggestion was meant only to save us effort polishing old wheels and as a way to not get embroiled in unnecessary debates about which is the 'correct' axiomatization of arithmetic. But if there is a consensus that we must provide arithmetic axioms, I withdraw the suggestion.

>And Bill replied

>

>>Actually, Cyc (as a computational thing) doesn't need to have them.

>>In my work I use Cyc (as a specification) to create other computational

>>things (deductive databases) which try to approximate as much as

>>possible the semantics of Cyc as a first order theorem prover. So, the

>>lack of them in Cyc is bad news for any system trying to do what mine

>>does.

>

>This raises questions about the scope of our ontology efforts and

>about what should be provided to support our customers. In fact, we

>have never actually discussed the question of what kinds of customers

>or users we are trying to serve. Some of the controversies seem to

>arise from different opinions about what we should be doing for whom.

Yes, good point. It might be useful for our Organizer(s) to give us an idea about what they think our customer base is supposed to be. What are we up to here? Axiomatizing Truth? Sketching data models for databases? Helping standardize action reasoning for a robot? (All of the above?) Or...what?

>>.... So, Cyc does the

>>right thing computationally but the wrong thing as a specification.

>
>I think that we have to do both. There is no reason why we can't
>support procedural attachments while keeping the axioms that define
>the procedures. Java, for example, supports what they call "introspection",
>which allows the Java programs to be asked metalevel questions about
>their specifications. We should require any attached procedures to
>support something like the Java introspection facilities.
>

Is the idea that the axiom to which the procedure is attached *is* the
Java-type introspection description?

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Thu Oct 08 12:38:08 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA29470;
Thu, 8 Oct 1998 11:01:49 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a02b24290f2dac8@[143.88.7.118]>
In-Reply-To: <199810080250.WAA10870@west>
Mime-Version: 1.0
Date: Thu, 8 Oct 1998 10:58:54 -0500
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Standard arithmetics (was: Re: Tropes)
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2128
Status:

>>Why dont we simply state that certain well-known mathematical structure are
>>being *assumed* by our ontology, rather than providing ontologies for them?

>

>I agree. I think that we should list them and either copy the basic
>axioms (e.g. Peano's axioms for arithmetic) or cite the best source
>we can find. But we do have a responsibility for showing where they
>fit in relationship to other related categories.

>

Purely as a public-relations point, there is a danger here of claiming too much authority. For example, suppose we include Peano's axioms; many of our customers will complain that we have then somehow prevented them from using a calculator. There are two different reasons for providing axioms; to give a foundation, or to sketch a way to get inferences done. Peano is a good foundation, but its not a lot of practical use for doing one's sums. We dont want to seem to be suggesting that users of our ontology must actually *use* Peano's axioms in order to multiply 245 by 463.

That is why I suggested literally not providing any ontology for arithmetic, but just saying that we are assuming that the user has some resources for making inferences/performing calculations involving numbers, and that the power of the ontology is to some extent limited by the power of the arithmetic apparatus which they provide, whatever that is. We should of course give some guidance about what they might need to provide (for example, that a conventional calculator provides only ground truths of arithmetic, and gives no access to anything that requires induction) and maybe offer some pointers to alternative ways of formalising theories of numbers. But if come to the marketplace offering a 'standard', and include Peano (or any other) arithmetic, then a lot of people will take us to be saying that this must be the standard arithmetic.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Fri Oct 09 10:01:06 1998

Received: from smtp3.erols.com (smtp3.erols.com [207.172.3.236])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id WAA19547

for <phayes@coginst.uwf.edu>; Thu, 8 Oct 1998 22:39:35 -0500 (CDT)
Received: from [207.172.131.10] (207-172-131-10.s10.tnt4.col.erols.com
[207.172.131.10])
by smtp3.erols.com (8.8.8/8.8.5) with ESMTP id XAA10207;
Thu, 8 Oct 1998 23:35:32 -0400 (EDT)
Message-Id: <103130300b24295270df8@[207.172.129.128]>
In-Reply-To: <v04003a02b24290f2dac8@[143.88.7.118]>
References: <199810080250.WAA10870@west>
Mime-Version: 1.0
Date: Thu, 8 Oct 1998 13:02:05 -0400
To: Pat Hayes <phayes@coginst.uwf.edu>
From: Bill Andersen <chezewiz@pop.erols.com>
Subject: Re: Standard arithmetics (was: Re: Tropes)
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu,
"guarino@ladseb.pd.cnr.it chezewiz@erols.com Nancy Lawler"
<e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2310
Status:

>Purely as a public-relations point, there is a danger here of claiming too
>much authority. For example, suppose we include Peano's axioms; many of our
>customers will complain that we have then somehow prevented them from using
>a calculator. There are two different reasons for providing axioms; to give
>a foundation, or to sketch a way to get inferences done. Peano is a good
>foundation, but its not a lot of practical use for doing one's sums.

>That is why I suggested literally not providing any ontology for
>arithmetic, but just saying that we are assuming that the user has some
>resources for making inferences/performing calculations involving numbers,

There's a danger in elevating this bit about arithmetic to a
general principle. The same could be applied to any other subject
matter which could be axiomatized but has a nice alternative
implementation (like a calculator).

For example, let's apply the principle to whether we should include
axioms defining transitive closure in the ontology. Transitive closure
is definable in FOL with a least fixed-point operator or in SOL. I
sure as hell don't want to use those to compute closures - I'll use
Warshall's algorithm or do a search over some pointer structure.
So, by the principle, we won't include a notion of transitive

closure, or things with transitive definitions, in the ontology.

Even if one doesn't plan to use axioms to do actual computation, it's still a good idea to include them. Cyc does this (but see below). It has axiomatic definitions for some things which are computed by special purpose modules for efficiency.

Also, we as designers of the ontology, will not be able to predict to which uses it will be put. I personally ran into this problem with Cyc - my stuff which extracts chunks of Cyc and turns them into logic programs missed great big classes of axioms which the logic program needed but didn't have because Cyc provided a procedural implementation and omitted the axioms from the ontology. Doug Lenat never thought anyone would use Cyc in this way.

All of that said, I realize Pat is making a subtle point, there being a number of theories of arithmetic, but that doesn't mean we shouldn't include them. They should be appropriately encapsulated in some kind of context mechanism.

...bill

From ???@??? Fri Oct 09 10:01:06 1998
Received: from smtp3.erols.com (smtp3.erols.com [207.172.3.236])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id WAA19547
for <phayes@coginst.uwf.edu>; Thu, 8 Oct 1998 22:39:35 -0500 (CDT)
Received: from [207.172.131.10] (207-172-131-10.s10.tnt4.col.erols.com
[207.172.131.10])
by smtp3.erols.com (8.8.8/8.8.5) with ESMTP id XAA10207;
Thu, 8 Oct 1998 23:35:32 -0400 (EDT)
Message-Id: <103130300b24295270df8@[207.172.129.128]>
In-Reply-To: <v04003a02b24290f2dac8@[143.88.7.118]>
References: <199810080250.WAA10870@west>
Mime-Version: 1.0
Date: Thu, 8 Oct 1998 13:02:05 -0400
To: Pat Hayes <phayes@coginst.uwf.edu>
From: Bill Andersen <chezewiz@pop.erols.com>
Subject: Re: Standard arithmetics (was: Re: Tropes)
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu,

"guarino@ladseb.pd.cnr.it chezewiz@erols.com Nancy Lawler"
<e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2310
Status:

>Purely as a public-relations point, there is a danger here of claiming too
>much authority. For example, suppose we include Peano's axioms; many of our
>customers will complain that we have then somehow prevented them from using
>a calculator. There are two different reasons for providing axioms; to give
>a foundation, or to sketch a way to get inferences done. Peano is a good
>foundation, but its not a lot of practical use for doing one's sums.

>That is why I suggested literally not providing any ontology for
>arithmetic, but just saying that we are assuming that the user has some
>resources for making inferences/performing calculations involving numbers,

There's a danger in elevating this bit about arithmetic to a
general principle. The same could be applied to any other subject
matter which could be axiomatized but has a nice alternative
implementation (like a calculator).

For example, let's apply the principle to whether we should include
axioms defining transitive closure in the ontology. Transitive closure
is definable in FOL with a least fixed-point operator or in SOL. I
sure as hell don't want to use those to compute closures - I'll use
Warshall's algorithm or do a search over some pointer structure.
So, by the principle, we won't include a notion of transitive
closure, or things with transitive definitions, in the ontolgy.

Even if one doesn't plan to use axioms to do actual computation,
it's still a good idea to include them. Cyc does this (but see below).
It has axiomatic definitions for some things which are computed by
special purpose modules for efficiency.

Also, we as designers of the ontology, will not be able to predict
to which uses it will be put. I personally ran into this problem with
Cyc - my stuff which extracts chunks of Cyc and turns them into logic
programs missed great big classes of axioms which the logic program
needed but didn't have because Cyc provided a procedural
implementation and omitted the axioms from the ontology. Doug Lenat
never thought anyone would use Cyc in this way.

All of that said, I realize Pat is making a subtle point, there

being a number of theories of arithmetic, but that doesn't mean we shouldn't include them. They should be appropriately encapsulated in some kind of context mechanism.

...bill

From ???@??? Fri Oct 09 10:01:07 1998
Received: from smtp3.erols.com (smtp3.erols.com [207.172.3.236])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id WAA19548
for <phayes@coginst.uwf.edu>; Thu, 8 Oct 1998 22:39:35 -0500 (CDT)
Received: from [207.172.131.10] (207-172-131-10.s10.tnt4.col.erols.com
[207.172.131.10])
by smtp3.erols.com (8.8.8/8.8.5) with ESMTP id XAA10256;
Thu, 8 Oct 1998 23:35:37 -0400 (EDT)
Message-Id: <103130300b24334f59ec9@[207.172.129.128]>
Mime-Version: 1.0
Date: Thu, 8 Oct 1998 23:35:33 -0400
To: Pat Hayes <phayes@coginst.uwf.edu>
From: Bill Andersen <chezewiz@pop.erols.com>
Subject: Re: Standard arithmetics (was: Re: Tropes)
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu,
"guarino@ladseb.pd.cnr.it chezewiz@erols.com Nancy Lawler"
<e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2310
Status:

>Purely as a public-relations point, there is a danger here of claiming too
>much authority. For example, suppose we include Peano's axioms; many of our
>customers will complain that we have then somehow prevented them from using
>a calculator. There are two different reasons for providing axioms; to give
>a foundation, or to sketch a way to get inferences done. Peano is a good
>foundation, but its not a lot of practical use for doing one's sums.

>That is why I suggested literally not providing any ontology for
>arithmetic, but just saying that we are assuming that the user has some
>resources for making inferences/performing calculations involving numbers,

There's a danger in elevating this bit about arithmetic to a general principle. The same could be applied to any other subject matter which could be axiomatized but has a nice alternative implementation (like a calculator).

For example, let's apply the principle to whether we should include axioms defining transitive closure in the ontology. Transitive closure is definable in FOL with a least fixed-point operator or in SOL. I sure as hell don't want to use those to compute closures - I'll use Warshall's algorithm or do a search over some pointer structure. So, by the principle, we won't include a notion of transitive closure, or things with transitive definitions, in the ontology.

Even if one doesn't plan to use axioms to do actual computation, it's still a good idea to include them. Cyc does this (but see below). It has axiomatic definitions for some things which are computed by special purpose modules for efficiency.

Also, we as designers of the ontology, will not be able to predict to which uses it will be put. I personally ran into this problem with Cyc - my stuff which extracts chunks of Cyc and turns them into logic programs missed great big classes of axioms which the logic program needed but didn't have because Cyc provided a procedural implementation and omitted the axioms from the ontology. Doug Lenat never thought anyone would use Cyc in this way.

All of that said, I realize Pat is making a subtle point, there being a number of theories of arithmetic, but that doesn't mean we shouldn't include them. They should be appropriately encapsulated in some kind of context mechanism.

...bill

From ???@??? Fri Oct 09 15:28:46 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA15971;
Fri, 9 Oct 1998 11:55:52 -0500 (CDT)
<199810080250.WAA10870@west>
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a02b243d8cbec31@[143.88.7.118]>
In-Reply-To: <103130300b24295270df8@[207.172.129.128]>
References: <v04003a02b24290f2dac8@[143.88.7.118]>

<199810080250.WAA10870@west>

Mime-Version: 1.0

Date: Fri, 9 Oct 1998 11:52:58 -0500

To: Bill Andersen <chezewiz@pop.erols.com>

From: Pat Hayes <phayes@coginst.uwf.edu>

Subject: Re: Standard arithmetics (was: Re: Tropes)

Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net

Content-Type: text/plain; charset="us-ascii"

Content-Length: 4059

Status:

[Pat:]

>>Purely as a public-relations point, there is a danger here of claiming too
>>much authority. For example, suppose we include Peano's axioms; many of our
>>customers will complain that we have then somehow prevented them from using
>>a calculator. There are two different reasons for providing axioms; to give
>>a foundation, or to sketch a way to get inferences done. Peano is a good
>>foundation, but its not a lot of practical use for doing one's sums.

>

>>That is why I suggested literally not providing any ontology for
>>arithmetic, but just saying that we are assuming that the user has some
>>resources for making inferences/performing calculations involving numbers,

>

[Bill:]

> There's a danger in elevating this bit about arithmetic to a
>general principle. The same could be applied to any other subject
>matter which could be axiomatized but has a nice alternative
>implementation (like a calculator).

Good point. OK, by all means lets include Peano arithmetic if people feel that it would be useful, but lets be very careful to say that this isnt THE arithmetic. Many useful arithmetics are much weaker than Peano (eg have no general induction.)

....

>

> Even if one doesn't plan to use axioms to do actual computation,
>it's still a good idea to include them. Cyc does this (but see below).
>It has axiomatic definitions for some things which are computed by
>special purpose modules for efficiency.

If the axiomatic definitions are never used at all, I wonder why CYC needs

to have them? Is it so that the relation names can be in the axiomatic vocabulary in order to provide something to attach the computational method to (as JMC outlined yesterday) so that unification can supply the argument-delivery and result-recovery interface for the attached computations? If so, I think that our ontological responsibility can be just to *say* what the relation names are and what computations are supposed to be attached to them. (Eg 'times' is a polytyped variadic function on $(\text{integer} \rightarrow \text{integer} \cup \text{real} \rightarrow \text{real})$, with the usual meaning.) We don't have to provide syntactic mechanisms for procedural attachment and reflection, any more than we have to give inference rules.

Here's a test question for arithmetic. Should we expect that the fact that adding and multiplying (but not dividing) integers by integers always gives an integer, to be *provable* within our ontology? And if we do, should we expect that this proof be short and almost trivially easy, or should we expect that it might be quite subtle and involve the use of the induction schema? I'd prefer to answer yes, and easy; but I don't think it would be a problem if someone wanted to use an arithmetic in which this couldn't even be stated, let alone proved (if they just want to get some sums done.)

Another (more ontological) question: Do integers and reals have different criteria of individuation? For example, consider the integer 4 and the real number 4.0: are these the same thing or two different things (which happen to be intersubstitutable in arithmetic)?

...

>... theories of arithmetic... should be appropriately encapsulated in >some kind of context mechanism.

I think that JMC's reflexion idea is better. If different axiomatisations of the same topic are in different contexts, then we have yet another meaning of "context" to deal with.

Pat

PS. Hmm. I guess one could have a kind of contextual attachment (implemented by attaching a context? Or if the fact of being attached is expressible in the logic, by doing the attachment in the context.) Maybe this has already been done. (If not, it might make a neat thesis topic. Desirable lemma: if you can do it both ways they ought to be equivalent, or at least the different kinds of attaching ought to commute.)

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office

Pensacola, FL 32514
phayes@ai.uwf.edu
http://www.coginst.uwf.edu/~phayes

(850)474 3023 fax

From ???@??? Mon Oct 12 16:44:32 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA21179;
Mon, 12 Oct 1998 12:12:59 -0500 (CDT)
<103130300b24295270df8@[207.172.129.128]>
<v04003a02b24290f2dac8@[143.88.7.118]> <199810080250.WAA10870@west>
X-Sender: phayes@mail.coginst.uwf.edu (Unverified)
Message-Id: <v04003a04b247e5499622@[143.88.7.118]>
In-Reply-To: <103130304b244a12132f9@[207.172.129.160]>
References: <v04003a02b243d8cbec31@[143.88.7.118]>
<103130300b24295270df8@[207.172.129.128]>
<v04003a02b24290f2dac8@[143.88.7.118]> <199810080250.WAA10870@west>
Mime-Version: 1.0
Date: Mon, 12 Oct 1998 12:10:08 -0500
To: Bill Andersen <chezewiz@pop.erols.com>
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Standard arithmetics (was: Re: Tropes)
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 1504
Status:

Hi Bill

>For example, in Cyc there is a predicate called 'kleenePlusOf' which
>is a predicate defined on binary predicate symbols in Cyc. When one
>states in Cyc:

>

> kleenePlusOf(P,Q)

>

>it is supposed act as an axiom schema:

>

> $P(x,y) \rightarrow Q(x,y)$

> $P(x,y) \ \& \ Q(y,z) \rightarrow Q(x,z)$

>

>But, if you assert, say,
>
> kleenePlusOf(parent,ancestor),
>
>the axioms
>
> parent(x,y) -> ancestor(x,y)
> parent(x,y) & ancestor(y,z) -> ancestor(x,z)
>
>never show up as being explicitly asserted. This is because
>kleenePlusOf is handled using procedural attachment - in fact a
>very efficient set of graph search algorithms. So, Cyc does the
>right thing computationally but the wrong thing as a specification.

Point taken, and your example illustrates it nicely. I think it also illustrates that Cyc's claim to have an 'axiom schema' is just plain false: kleenePlusOf(P,Q) *isnt* an axiom schema! There are good computational reasons to not have axiom-schemas in the language (of a computational thing) itself, notably that unification then doesnt work.

Pat

PS. Can CYC discover this inconsistency by graph-searching:

```
{parent(a,b)
 (forall x)(not(ancestor(x,b)))
 kleenePlusOf(parent, ancestor) }
? (Maybe Fritz knows the answer?)
```

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Mon Oct 12 16:44:34 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id NAA04068;
Mon, 12 Oct 1998 13:44:53 -0500 (CDT)
X-Sender: phayes@mail.coginst.uwf.edu (Unverified)
Message-Id: <v04003a08b247fce2223b@[143.88.7.118]>
In-Reply-To: <199810100000.UAA23296@west>
Mime-Version: 1.0

Date: Mon, 12 Oct 1998 13:42:03 -0500
To: sowa@west.poly.edu (John F. Sowa)
From: Pat Hayes <phayes@coginst.uwf.edu>
Subject: Re: Standard arithmetics (was: Re: Tropes)
Cc: Piek.Vossen@let.uva.nl, cmenzel@tamu.edu, doug@csi.uottawa.ca,
fritz@cyc.com, p.m.simons@leeds.ac.uk, phayes@coginst.uwf.edu,
sowa@west.poly.edu, guarino@ladseb.pd.cnr.it,
Bill Andersen <chezewiz@erols.com>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
JMcCarthy <jmc@cs.stanford.edu>, skydog@pacbell.net
Content-Type: text/plain; charset="us-ascii"
Content-Length: 1082
Status:

>I agree with John that the ontology must allow procedural attachment.
>
>That doesn't require any new features to the logic, since you could
>just assume that the procedural attachment is equivalent to an infinite
>(possibly even uncountable) conjunction of axioms that assert the value
>of the function for every possible combination of input values.

That's not quite the same (unless one uses Prolog as the procedural language), since that set of ground sentences would also allow one to infer things about the arguments, given the value. But I tend to agree in spirit: we needn't be too concerned about the machinery of procedural attachment.

Pat

PS. I probably shouldn't let my chain get pulled, but it really is impossible to have uncountably many finite ground sentences over a countable vocabulary.

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Mon Nov 02 12:24:49 1998
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id RAA15274
for <phayes@coginst.uwf.edu>; Fri, 30 Oct 1998 17:07:47 -0600 (CST)
Received: from west (west.poly.edu [128.238.20.21])

by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id SAA20958;
Fri, 30 Oct 1998 18:03:35 -0500 (EST)
Received: by west (SMI-8.6/SMI-SVR4)
id RAA21731; Fri, 30 Oct 1998 17:57:54 -0500
Date: Fri, 30 Oct 1998 17:57:54 -0500
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199810302257.RAA21731@west>
To: phayes@coginst.uwf.edu
Subject: Splitters vs. Lumpers
Cc: Piek.Vossen@let.uva.nl, chezewiz@erols.com, cmenzel@tamu.edu,
doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu, fritz@cyc.com,
guarino@ladseb.pd.cnr.it, jmc@cs.stanford.edu, p.m.simons@leeds.ac.uk,
skydog@pacbell.net, sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 13323
Status:

Pat,

I did a search on AltaVista with the input "splitters near lumpers" and got 69,996 responses. I glanced at a few of them, and they seemed to cluster in biological taxonomy and anthropology. Following are two typical discussions: the first is an outline of somebody's lecture on intelligence, and the second is an email note about dinosaurs.

John

Theories of Intelligence

I. Introductory Discussion

- * A. What does it mean to be intelligent? Compare answers to list generated by experts (Snyderman and Rothman, 1987).
 - o 1. All said: abstract reasoning, problem solving, and capacity to acquire knowledge.
 - o 2. More than half said: memory, adaptation to one[s environment, mental speed, linguistic competence, mathematical competence, general knowledge, and creativity.
 - o 3. One-fourth said: sensory acuity, goal directedness, and achievement motivation.

- * B. What ability (or abilities) comprise(s) intelligence? Do answers suggest that intelligence is an unitary ability or a set of interrelated abilities.
- * C. What behaviors provide outward signs of intelligence? Compare to what is typically asked in IQ tests. Heavy focus on verbal skills? What about performance or mechanical aspects?
- * D. How can we measure intelligence? Do the supposed tests of intelligence measure intelligence as you have defined it?

II. What is intelligence?

- * A. The editors of the Journal of Educational Psychology (1921) asked psychologists what intelligence was. Every one has a different view of what intelligence was but two themes emerged:
 - o 1. intelligence is the capacity to learn from experience
 - o 2. intelligence is adaptation to one's environment
- * B. Part of adaptation is learning, so these themes really may be referring to the same process. Suggests that intelligence may be a general ability of learning.
- * C. Implicit definitions by lay people. Sternberg, et al (1981) asked people at train stations, college library, grocery store. As in this class people gave many different answers but the fell in three general categories:
 - o 1. verbal intelligence - good vocabulary, reads w/high comprehension, and is verbally fluent, converses easily on variety of subjects.
 - o 2. problem solving - able to apply knowledgeto problem at hand, makes good decisions, plans ahead, poses problems in an optimal way
 - o 3. practical intelligence - sizes up situations well, determines how to achieve goals, displays awareness of world around him/her, and displays interest in world at large.
- * D. How one views intelligence will influence the type of theory one suggests for understanding intelligence. Historically there have been two general views of intelligence: Lumpers and Splitters.

- o 1. Lumpers - feel intelligence is a general capacity for acquiring knowledge - sounds like the themes of the 1921 survey.
- o 2. Splitters - feel intelligence is composed of many separate abilities that operate more or less independently. Splitters vary in how fragmented they view intelligence, a few different abilities or a lot of different abilities - sounds like the "lay persons" definition.

III. Theories of intelligence - Lumpers: Galton and Spearman

- * A. Both can be described as lumpers. And both were interested in how to measure intelligence. That is why they are included in the psychometric section of your book.
- * B. In a nutshell, Galton and Spearman thought that intelligent behavior or intelligence was governed by some underlying ability. He suggested something like mental speed was the underlying basis of intelligent behavior and he suggested that this was biologically passed on from parents to children (research ID lots of famous people had famous fathers).
- * C. Galton measured a number of characteristics of people hoping to find relationships among these measures. For example, he measured sensory and motor skills and hoped to find that people who were high on one measure scored high on other measures. He was hoping to find positive correlations among these measures. More detail on correlations? Galton wasn't very successful. The correlations were very low.
- * D. Spearman took a similar approach although he measured different things. Used mental tests rather than physical characteristics. Felt that there was some underlying ability that governed performance on all these tests. Called this underlying ability general intelligence or g. Correlations among tests varied but were which Spearman interpreted as showing varying degrees of g on that particular test. He also concluded that there was not one specific test that was a good measure of g.

IV. Theories of intelligence - Splitters: Gall and Gardner

- * A. First splitter was named Gall. Founded the science of Phrenology. According to Gall, human skills differ from one another and their variations reflect differences in the size and shape of the brain. He proposed that since different areas in the brain subserved discrete functions. He also suggested that by carefully examining the bumps on

the skull, you should be able to determine the person's strengths and weaknesses and the idiosyncronies of his/her mental abilities. So, different abilities are governed by different areas of the brain, the bigger the bump, the more advanced the abilities, the more aptitude in that area. Obviously, Gall wasn't right since we don't go around feeling people's heads to find out how smart they are.

* B. Problems with Gall's approach.

- o 1. Size of the brain has no clear-cut correlation with intelligence. Ex: Walt Whitman had a very small brain and achieved great success. Einstein's brain isn't any larger than "normal".
- o 2. The size and configuration of the skull itself is an inexact measure of the important configurations of the cortex. The outside does not exactly mirror the inside.

* C. Modern-day splitter - Howard Gardner. Developed Multiple Intelligences Theory from his work with brain damaged people. He is a neuropsychologist and dealt with a lot of people who suffered brain damage. Was intrigued by the fact that some abilities are lost and others are not. His theory has 7 components:

- o 1. Linguistic: verbal ability in general, especially the ability to understand subtle shades of meaning.
- o 2. Musical: singing, playing an instrument, composing music,
- o 3. Logical-mathematical: reasoning
- o 4. Spatial: the ability to perceive and draw spatial relationships
- o 5. Bodily-kinesthetic: ability to control one[s] muscle movements in graceful and skillful ways, as in dance, athletics, and tool use
- o 6. Interpersonal; social skills, including the ability to understand and respond appropriately to others[nonverbal messages
- o 7. Intrapersonal: the ability to understand oneself, including one[s] own emotions and wishes, and to use that understanding effectively in guiding behavior.

* D. #1,3, 4 are part of standard intelligence tests. Others suggest that Gardner should call #2, 5, 6, 7 talents rather than intelligences. He disagrees.

* E. Gardner feels that people excel in some areas and not in others and it shouldn't be expected that all people will be good in all areas. As

a result, education should be tailored toward their strengths.

- * F. Criticisms: the theory was based on brain damaged people -- individual cases and hasn't been tested on general people.

V. Theories of Intelligence: Processing Approach - Sternberg

- * A. The above theories describe the components of intelligence. There is another approach to understanding intelligence -- understanding the processing aspects of intellectual skills. So instead of saying a person is using verbal reasoning skills, the focus is on how does a person use his/her verbal reasoning skills or solve a problem that involves verbal reasoning. These two approaches are complementary.
- * B. Answers the how question by suggesting there are three different types of information processing components. Metacomponents, knowledge acquisition components, and performance components. Detailed below.
- * C. Metacomponents. Higher order control processes that define the problem, decide whether or not it is worth solving, select the lower components needed to solve it, control the order in which those components are activated, monitor the progress toward solution, and decide when the problem is solved.
- * D. Knowledge acquisition components. Involved in learning new information and storing it in memory.
- * E. Performance components. Used in execution of various strategies for task performance.
- * F. Example: Solving an anagram. DERYHETI.
 - o 1. Metacomponent: decide on a tentative strategy for unscrambling the letters
 - o 2. Activation of performance components: write down possible combinations of letters, place vowels and consonants together. Feedback of success goes to PC and MC.
 - o 3. Knowledge acquisition components: as strategies are tried person is learning what works and what doesn't work in solving this type of problem. Storage of this new information takes place.

RaptorRKC@aol.com
Sun, 29 Oct 1995 01:41:02 -0500

- * Next message: RaptorRKC@aol.com: "Cryptozoology?"
- * Previous message: RaptorRKC@aol.com: "Please translate this into English?"
- * Next in thread: Dinogeorge@aol.com: "Re: Lumpers and Splitters"

Dinogeorge wrote:

>Not very many, actually. The trend now seems to be turning to synonyming
>genera and species (lumping) rather than naming or renaming new genera
>(splitting). This is a statistical trend, by the way, stemming perhaps
>from new concepts of what a genus or a species might be--not a conscious
>effort on the part of paleontologists, as in, "Well, we've done
>enough splitting; it's time to lump."

I am way sick of specimens being split apart, renamed, placed in different genera and families, etc. Lumping can sometimes be the way to go. Many lumpings are reasonable. Splitters tend to forget, or WANT to forget, that SOME species can be so closely related that they belong in the same genus. And we damn well know that more than one species of dino belonged in each genus, and I am sure we have found some of them. Splitters seem to want to put each specimen in its own genus -- how lame! We should lump when necessary, and split when necessary. When anatomical details are minor, distinct species are closely related, and geological separation (by locale and time period) are minor, sometimes it is reasonable to lump. Sometimes it may take a little study to figure out which name is the best to lump specimens under.

I am a lumper at heart, but I do disagree with many lumpings. Such as many of GSP's, which include Daspletosaurus = Tyrannosaurus, Deinonychus = Velociraptor, etc. Some of them are reasonable. I favor the lumping of all species (remember A CLUTTER OF DUCKBILLS?) of Corythosaurus, Lambeosaurus, and Hypacrosaurus into Hypacrosaurus -- it quite a reasonable lumping. I think that until more fossils are discovered, Ultrasauros should be placed in Brachiosaurus. Struthiomimus and Ornithomimus, I believe, should be united. Ditto Gryposaurus, Hadrosaurus, and Kritosaurus. And, for some reason, IMHO Orodromeus looks suspiciously similar to Hyspilophodon...

Some groups of dinosaurs confuse me. For example, the Styracosaurus, Eucentrosaurus, and Monoclonians. Their skulls are all startlingly similar -- it appears to me that the horn and spike ornamentation is what keeps them separated. I would like to be more informed on these dinosaurs and their status. (Info would be appreciated -- but please give me the info directly,

don't point me towards refs!!!)

Some splittings are startlingly inappropriate. Why is Lufengosaurus, an Asian prosauropod, kept separate from Plateosaurus? Why is a species that had been originally referred to Brachiosaurus (*B. brancai*) now in its own genus, Giraffatitan? Why is the original Coelophysis specimen kept separate from the Ghost Ranch "rioarribasaurus"? Why are all of those Asian tyrannosaurs in different genera? Why, suddenly, are some species of Aublysodon placed into a new genus, Sinraptor? There are many more appropriate examples.

Why why why???

Why must the conflicts between lumpers and splitters keep the phytodinosaur family tree in such a jumble? Oh, what a tangled web they weave!!

Raptor RKC (Rachel Clark)

From ???@??? Mon Nov 02 12:24:59 1998
Received: from vapor.stanford.edu (vapor.Stanford.EDU [171.64.71.11])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id SAA20318
for <phayes@coginst.uwf.edu>; Fri, 30 Oct 1998 18:04:36 -0600 (CST)
Received: (from jmc@localhost)
by vapor.stanford.edu (8.8.8/8.8.8) id QAA02737;
Fri, 30 Oct 1998 16:00:08 -0800 (PST)
Date: Fri, 30 Oct 1998 16:00:08 -0800 (PST)
Message-Id: <199810310000.QAA02737@vapor.stanford.edu>
From: John McCarthy <jmc@Steam.Stanford.EDU>
To: sowa@west.poly.edu
CC: phayes@coginst.uwf.edu, Piek.Vossen@let.uva.nl, chezewiz@erols.com,
cmenzel@tamu.edu, doug@csi.uottawa.ca, e6nl001@coe.coppin.umd.edu,
fritz@cyc.com, guarino@ladseb.pd.cnr.it, p.m.simons@leeds.ac.uk,
skydog@pacbell.net, sowa@west.poly.edu
In-reply-to: <199810302257.RAA21731@west> (sowa@west.poly.edu)
Subject: Re: Splitters vs. Lumpers
Reply-to: jmc@cs.Stanford.EDU
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 337
Status:

Just to confuse matters a bit, I am a lumper w/r human intelligence and a splitter w/r AI. I accept the experimental evidence that human intellectual abilities are strongly correlated. On the other hand, there are a large number of intellectual mechanisms, and we understand

only a few well enough to put them into computer programs.

From ???@??? Wed Nov 25 13:33:14 1998

Received: from pdadr1.pd.cnr.it (pdadr1.pd.cnr.it [150.178.1.2])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id KAA28967

for <phayes@coginst.uwf.edu>; Wed, 25 Nov 1998 10:54:15 -0600 (CST)

Received: from [150.178.2.93] (duoguarino.ladseb.pd.cnr.it [150.178.2.93])

by mail.pd.cnr.it (PMDF V5.2-27 #30074)

with ESMTP id <01J4LL3E6EK699EIDO@mail.pd.cnr.it> for phayes@coginst.uwf.edu;

Wed, 25 Nov 1998 17:49:39 MET

From: Nicola.Guarino@ladseb.pd.cnr.it

Date: Wed, 25 Nov 1998 17:52:33 +0100

Subject: Re: Results of the Villa Bosch workshop

In-reply-to: <v04003a06b280a6c7c668@[143.88.7.118]>

X-Sender: guarino@mail.ladseb.pd.cnr.it

To: Pat Hayes <phayes@coginst.uwf.edu>

Cc: "Reuter, Andreas" <Andreas.Reuter@eml.org>, skydog@pacbell.net,

"Tschira, Klaus" <Klaus.Tschira@kts.villa-bosch.de>

Message-id: <v03102804b281e033f653@[150.178.2.93]>

MIME-version: 1.0

References: <v03102808b27f498d3dbd@[150.178.2.93]>

<E17818F052FCD111ADFF00609793D170082895@tigershark.villa-bosch.de>

Content-Transfer-Encoding: 8bit

X-MIME-Autoconverted: from quoted-printable to 8bit by nuts.coginst.uwf.edu id

KAA28967

Content-Type: text/plain; charset="us-ascii"

<E17818F052FCD111ADFF00609793D170082895@tigershark.villa-bosch.de>

Content-Length: 2403

Status:

At 12:26 PM -0600 11/24/98, Pat Hayes wrote:

>Hey, Nicola!! Surely things aren't as bad as you paint them here. Its true

>that there has been a lack of gung-ho whip-crackin' leadership initiative,

>but then we appointed the two nice guys to be the leaders, so what did you

>expect?

>

>I think that work has progressed. I seem to be in constant communication

>with the people Im supposed to be writing articles with (Chris, Peter and

>John.) So I dont think you should say that we have missed an opportunity.

>We are behind schedule, but then I never took that schedule very seriously.

>We all work busily for other employers and have active professional lives,

>so getting us to work together is like herding cats.

Well, Pat, I am happy to see that you are less disappointed than me, but your message does not answer the many organizational questions I raised, especially for what concerns the relationships with all the workshop participants.

Moreover, I understand you are aware of a schedule and of a task assignment to the various people of the theory group, but I have definitely no idea of that. I remember having been asked by Bob on the phone to contribute on a number of issues including dependence, identity criteria, and so on, but I never got a clarification about how such a work should proceed. I notified with two separate messages (10/12/98 and 11/2/98) that I had not been able to read the table sent around by Bob on 10/7. Moreover, I explicitly asked about the deadline and the editing process, but nobody answered. You (thanks!) have been the only one who answered my last message to the theory group entitled "has this mailing list any sense?"

Finally, you are right in pointing out that we all have active professional lives, but exactly because of this, this post-workshop activity (which was not intended to be an extra-professional divertimento) could have been organized in a bit more professional way...

All the best,

-- Nicola

Nicola Guarino
National Research Council phone: +39 049 8295751
LADSEB-CNR fax: +39 049 8295763
Corso Stati Uniti, 4 email: Nicola.Guarino@ladseb.pd.cnr.it
I-35127 Padova (old address guarino@ladseb.pd.cnr.it still valid)
Italy

Warning: you are now required to dial "0" before the area code!

<http://www.ladseb.pd.cnr.it/infor/ontology/ontology.html>

From ???@??? Wed Dec 02 13:26:59 1998
Received: from pdadr1.pd.cnr.it (pdadr1.pd.cnr.it [150.178.1.2])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA02875
for <phayes@coginst.uwf.edu>; Wed, 2 Dec 1998 11:46:28 -0600 (CST)
Received: from [150.178.2.93] (duoguarino.ladseb.pd.cnr.it [150.178.2.93])
by mail.pd.cnr.it (PMDF V5.2-27 #30074)
with ESMTP id <01J4VEXXOWRE99EQBP@mail.pd.cnr.it> for
phayes@coginst.uwf.edu;
Wed, 2 Dec 1998 18:41:25 MET
From: Nicola.Guarino@ladseb.pd.cnr.it
Date: Wed, 02 Dec 1998 18:44:33 +0100

Subject: Logically vs. linguistically motivated ontologies
In-reply-to: <Pine.GSO.3.96.981125174024.4158A-100000@cosmos>
X-Sender: guarino@mail.ladseb.pd.cnr.it
To: Josiah Lee Auspitz <lee@textwise.com>, Pat Hayes <phayes@coginst.uwf.edu>
Cc: "Reuter, Andreas" <Andreas.Reuter@eml.org>,
robert grayson spillers <skydog@pacbell.net>, EHovy <hovy@isi.edu>,
geo@thought.princeton.edu, john sowa <sowa@west.poly.edu>,
polanyi@pal.xerox.com, Piek Vossen <Piek.Vossen@let.uva.nl>,
peters@csl.stanford.edu, fellbaum@thought.princeton.edu,
cmenzel@tamu.edu, Fritz Lehmann <fritz@cyc.com>,
TSUJII Junichi <tsujii@is.s.u-tokyo.ac.jp>,
Peter Simons <p.m.simons@leeds.ac.uk>, Martin.v.d.Berg@let.uva.nl,
doug@csi.uottawa.ca, jmc@cs.stanford.edu, lreeker@nsf.gov,
wahlster@cs.uni-sb.de, Giovanni.VARILE@LUX.DG13.cec.be,
Antonio Sanfilippo <antonio@anite-systems.lu>,
Graeme Hirst <gh@cs.toronto.edu>, Bill Andersen <chezewiz@erols.com>,
nancy lawler <E6NL001@coe.coppin.umd.edu>,
"Tschira, Klaus" <Klaus.Tschira@kts.villa-bosch.de>,
sophia.ananiadou@i-u.de,
"Mack, =?UNKNOWN?Q?B=E4rbel?=" <Baerbel.Mack@eml.villa-bosch.de>

Message-id: <v03102802b2897eefe659@[150.178.2.93]>

MIME-version: 1.0

References: <v04003a05b280a0fa6950@[143.88.7.118]>

Content-Transfer-Encoding: 8bit

X-MIME-Autoconverted: from quoted-printable to 8bit by nuts.coginst.uwf.edu id
LAA02875

Content-Type: text/plain; charset="us-ascii"

Content-Length: 5059

Status:

Dear Lee,

At 9:19 AM -0500 11/26/98, you wrote:

>... the point with which you conclude goes to the heart of my
>concern: whether and how we can bridge the disparities between a logically
>and a linguistically motivated upper-level ontology.

I don't think there is a logically *motivated* ontology. I agree with John when he says that logic and philosophy are more "criteria of soundness" than "motivations" for ontologies. In the logico-philosophical community, the motivations for ontological work are of two kinds: the desire to build a theory of the structure of reality *per se*, or otherwise to build a theory of the structure of reality as perceived by human beings and referred to by natural language (in other words, the structure of common sense). A lot of ontological work in the tradition of analytic philosophy (so-called "analytic ontology") roughly belongs to the latter category (see for instance Strawson, Wiggins, Davidson). In this case, linguistic tests are often used to accept or refuse a certain argument.

It is true, however, that in many cases the "ontologies" developed for linguistic purposes (like WordNet, Pangloss, or Mikrokosmos) do not share any of the motivations above. The reason of this, in my opinion, is mainly due to a difference of domain: on one hand you have "things" in your domain, on the other hand you have words. Many "semantic structures" used by linguists describe relations among words rather than relations among world entities. So the hypernym relation in WordNet (a lexical relation) does not necessarily coincide with a subsumption relation between classes of entities (an ontological relation).

In fact, lexical and ontological relations are often intermixed in current "upper level semantic resources". One of the advantages of a *clean* upper level ontology should be the clarification of the boundary between ontology and language. Surely, this clarification would benefit from a theory of the links between language and ontology: such a theory, in my modest understanding of Peirce, seems to be exactly within the scope of semiotics. As such, however, it is separate from ontology in the proper sense (although you can picture it as an "ontology of signs").

Notice that, once this separation between language and ontology is established, some linguists may conclude that they can do without ontology; at the Heidelberg workshop we have seen however that many others really need it, and are convinced of the high utility of a purely ontological, language-independent upper level.

Let us now come back to your observation concerning the "three-pier" architecture:

>The group structure of our meeting encouraged a three-tiered approach to
>the problem: 1. a philosophers' and logicians' upper level, 2. a
>linguists' upper level (from EuroWordNet or something like it), 3.
>applications.
>
>Now, where level 3 applications involve free text or natural language, as
>opposed to the more structured databases for which the impressive Ontek
>ontology was designed, we shall likely find that though we can link level
>one to level two and level two to level three, we cannot go
>from one to three in certain NLP tests. This is the problem in
>non-transitivity-- or perhaps it is better seen as blocked inheritance--
>that interests me.

To me, it is clear that we shall never "inherit" from level one to level three *in certain NLP tests*. This is because, in general, NLP tests depend *both* on language *and* on ontology. So ontology alone can give only a limited contribution to level 3 (as one may expect). Yet, such a contribution can play a crucial role when trying to integrate different level-3 applications, or to establish the appropriate "grounding relations" between the application and the real world.

In other words, only the "ontological component" (if taken separate from the linguistic component) will inherit from level one to level three: NLP tests require more than ontology.

>Ed Hovy led a small sub-group considering 1-3
>inheritance by posing a number of cases and mapping them onto the emerging
>lattice structure. I tried a few further tests on my own, using examples
>from a tool developed by the TextWise lab in Syracuse. It seemed to me
>that at some point after our group had the courage of a communique and an
>agenda for prospective funders it would be worth discussing the reasons
>for any inheritance glitches.

Collecting some examples of these glitches may be quite useful to better understand the relationship between language and ontology. It may be nice to have a list of the "tests" you considered.

-- Nicola

Nicola Guarino
National Research Council phone: +39 049 8295751
LADSEB-CNR fax: +39 049 8295763
Corso Stati Uniti, 4 email: Nicola.Guarino@ladseb.pd.cnr.it
I-35127 Padova (old address guarino@ladseb.pd.cnr.it still valid)
Italy

Warning: you are now required to dial "0" before the area code!

<http://www.ladseb.pd.cnr.it/infor/ontology/ontology.html>

From ???@??? Mon Dec 07 10:41:21 1998
Received: from ns1.textwise.com (ns1.mnis.net [208.7.177.50])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id IAA24902
for <phayes@coginst.uwf.edu>; Sun, 6 Dec 1998 08:04:11 -0600 (CST)
Received: from cosmos.textwise.com (cosmos [208.17.42.2])
by ns1.textwise.com (8.9.0/8.9.0) with SMTP id IAA03363;
Sun, 6 Dec 1998 08:54:10 -0500 (EST)
Received: from localhost by cosmos.textwise.com (SMI-8.6/SMI-SVR4)
id IAA10745; Sun, 6 Dec 1998 08:58:16 -0500
Date: Sun, 6 Dec 1998 08:58:15 -0500 (EST)
From: Josiah Lee Auspitz <lee@textwise.com>
X-Sender: lee@cosmos

Reply-To: Josiah Lee Auspitz <lee@textwise.com>
To: Nicola.Guarino@ladseb.pd.cnr.it
cc: Pat Hayes <phayes@coginst.uwf.edu>,
"Reuter, Andreas" <Andreas.Reuter@eml.org>,
robert grayson spillers <skydog@pacbell.net>, EHovy <hovy@isi.edu>,
geo@thought.princeton.edu, john sowa <sowa@west.poly.edu>,
polanyi@pal.xerox.com, Piek Vossen <Piek.Vossen@let.uva.nl>,
peters@csl.stanford.edu, fellbaum@thought.princeton.edu,
cmenzel@tamu.edu, Fritz Lehmann <fritz@cyc.com>,
TSUJII Junichi <tsujii@is.s.u-tokyo.ac.jp>,
Peter Simons <p.m.simons@leeds.ac.uk>, Martin.v.d.Berg@let.uva.nl,
doug@csi.uottawa.ca, jmc@cs.stanford.edu, lreeker@nsf.gov,
wahlster@cs.uni-sb.de, Giovanni.VARILE@LUX.DG13.cec.be,
Antonio Sanfilippo <antonio@anite-systems.lu>,
Graeme Hirst <gh@cs.toronto.edu>, Bill Andersen <chezewiz@erols.com>,
nancy lawler <E6NL001@coe.coppin.umd.edu>,
"Tschira, Klaus" <Klaus.Tschira@kts.villa-bosch.de>,
sophia.ananiadou@i-u.de,
"Mack, =?UNKNOWN?Q?B=E4rbel?=" <Baerbel.Mack@eml.villa-bosch.de>
Subject: Re: Logically vs. linguistically motivated ontologies
In-Reply-To: <v03102802b2897eefe659@[150.178.2.93]>
Message-ID: <Pine.GSO.3.96.981205075721.1450C-100000@cosmos>
MIME-Version: 1.0
Content-Type: TEXT/PLAIN; charset=US-ASCII
Content-Length: 10450
Status:

Nicola,

Now that Bob Spillers has kindly brought us up to date on the two tracks of post-Heidelberg activity-- theoretical and monetary-- I am even more inclined to follow my original instinct to await a resolution of some of the theoretical issues within the theory group before introducing either the semiotic issue at the theoretical level or the "glitches" at the applications level. This is not because I am in any sense holding back a criticism-- I flagged it in the pre-Heidelberg e-mails, but rather, because I believe it to be irrelevant to the task at hand.

On the theoretical level, the task at hand is to produce a philosophically informed, domain-independent upper level set of type labels for computational use that reflect a consensus among a group that is largely but not unanimously inclined toward an analytic (or really, nominalistic in the old, scholastic sense) perspective. To produce such a set of type labels would be a significant achievement with an excellence of its own. As we come to apply it, we shall find that it is more useful in some areas

than in others, and there will then be a greater reason to consider why this is so. At that time, the anticipation of some underlying theoretical issues in (only partially developed) correspondence between Charles Peirce and Victoria Welby to which I have referred may command some attention. My interest theoretically is in exploring this semiotic ground. It forces us to examine more self-consciously what happens when we attempt to use an ontology in the philosophical sense to inform a set of variously tagged, annotated and axiomatized type labels computationally. We are not yet at that stage.

At the level of applications, I have been engaged by a company that has developed and licensed tools for dealing with both free text and structured databases. So the promise of an artifact that would enable us to integrate such heterogeneous materials is of more than academic interest to me. Here, too, the "glitches" that interest me are of a systemic kind. That is, they may require other adaptations that work around rather than through the categories being developed.

Interlinear comments follow, with initials added to the >> markers to clarify who is responsible for previous rounds of emails that are quoted.

Josiah Lee Auspitz
lee@textwise.com
17 Chapel Street
Somerville, MA 02144
617-628-6228
fax -9441

On Wed, 2 Dec 1998 Nicola.Guarino@ladseb.pd.cnr.it wrote:

> Dear Lee, >> At 9:19 AM -0500 11/26/98, you wrote: >>... the point with which you conclude goes to the heart of my >>concern: whether and how we can bridge the disparities between a logically >>and a linguistically motivated upper-level ontology.

NG:

>> I don't think there is a logically *motivated* ontology. I agree with John when he says that logic and philosophy are more "criteria of soundness" than "motivations" for ontologies. In the logico-philosophical community, the motivations for ontological work are of two kinds: the desire to build a theory of the structure of reality *per se*, or otherwise to build a theory of the structure of reality as perceived by human beings and referred to by natural language (in other words, the

structure of common sense). A lot of ontological work in the tradition of analytic philosophy (so-called "analytic ontology") roughly belongs to the latter category (see for instance Strawson, Wiggins, Davidson). In this case, linguistic tests are often used to accept or refuse a certain argument. >

agreed. "Philosophically informed" would have been a more accurate description, though "logically motivated" would still apply to inference engines.

NG:

> It is true, however, that in many cases the "ontologies" developed for linguistic purposes (like WordNet, Pangloss, or Mikrokosmos) do not share any of the motivations above. The reason of this, in my opinion, is mainly due to a difference of domain: on one hand you have "things" in your domain, on the other hand you have words. Many "semantic structures" used by linguists describe relations among words rather than relations among world entities. So the hypernym relation in WordNet (a lexical relation) does not necessarily coincide with a subsumption relation between classes of entities (an ontological relation).

What you say seems valid in the very loose sense of "domain", but in the stricter sense the differentiating factors are purpose, sign modality and data structure rather than whether or not we are dealing with "things". Words qua tokens are also things.

NG:

>> In fact, lexical and ontological relations are often intermixed in current "upper level semantic resources". One of the advantages of a *clean* upper level ontology should be the clarification of the boundary between ontology and language. Surely, this clarification would benefit from a theory of the links between language and ontology: such a theory, in my modest understanding of Peirce, seems to be exactly within the scope of semiotics. As such, however, it is separate from ontology in the proper sense (although you can picture it as an "ontology of signs").

Yes, this is why it would be better to postpone discussion until we have a "clean" ontology in hand.

The term you put quotes-- an "ontology of signs"-- could be misleading if it suggests that a sign type must exist to be useful. A typology of hypothetical sign-modalities-- a "possible signs semiotics" in the lingo of those who talk about "possible worlds semantics"-- is really what is called for.

>> Notice that, once this separation between language and ontology is established, some linguists may conclude that they can do without ontology; at the Heidelberg workshop we have seen however that many others really need it, and are convinced of the high utility of a purely ontological, language-independent upper level.

This is precisely the point that requires examination, since the committee organization of our session, by eliminating early a systematic discussion of applications, hid from view the sign-shift that occurred looking "downward" toward NL applications and upward to language-independent categories. By a language-independent ontology some linguists may mean what you are developing, others may mean an interlingual or "concept ontology", others still may mean generalized syntactic terms and tropes-- and in point of fact, the upper levels developed for machine translation will properly combine all of these.

>> Let us now come back to your observation concerning the "three-pier" architecture:

By substituting the word "architecture" for "approach" in my e-mail you misconstrue me as making an observation about the final computational artifact we produce rather than about the committee organization of our meetings.

JLA:

>

>>The group structure of our meeting encouraged a three-tiered approach to
>>the problem: 1. a philosophers' and logicians' upper level, 2. a
>>linguists' upper level (from EuroWordNet or something like it), 3.
>>applications.

>>

>>Now, where level 3 applications involve free text or natural language, as
>>opposed to the more structured databases for which the impressive Ontek
>>ontology was designed, we shall likely find that though we can link level
>>one to level two and level two to level three, we cannot go
>>from one to three in certain NLP tests. This is the problem in
>>non-transitivity-- or perhaps it is better seen as blocked inheritance--
>>that interests me.

NG:

>> To me, it is clear that we shall never "inherit" from level one to level three *in certain NLP tests*. This is because, in general, NLP tests depend *both* on language *and* on ontology. So ontology alone can give only a limited contribution to level 3 (as one may expect). Yet, such a contribution can play a crucial role when trying to integrate different

level-3 applications, or to establish the appropriate "grounding relations" between the application and the real world.

This is well put, leaving aside the underlying things-words disjunction on which it rests.

NG:

>> In other words, only the "ontological component" (if taken separate from the linguistic component) will inherit from level one to level three: NLP tests require more than ontology.

This is precisely the point that seems to me to require more discussion once a clean ontology is in hand. What you call the ontological component will not be present unless the data is pre-tagged for it, and then we must consider Peirce's argument on whether certain pre-taggings are blind alleys not worth the effort.

JLA:

>>> Ed Hovy led a small sub-group considering 1-3 >> inheritance by posing a number of cases and mapping them onto the emerging >> lattice structure. I tried a few further tests on my own, using examples >> from a tool developed by the TextWise lab in Syracuse. It seemed to me >> that at some point after our group had the courage of a communique and an >> agenda for prospective funders it would be worth discussing the reasons >> for any inheritance glitches.

NG:

>> Collecting some examples of these glitches may be quite useful to better understand the relationship between language and ontology. It may be nice to have a list of the "tests" you considered.

For Ed's tests, with Graeme Hirst and me as sounding boards, we might ask him, once a clean ontology is in hand, to revive the notes he was discouraged from presenting to the general group, on the grounds that they were premature. They conform to your observation that an "ontological component" will be attributable but not usually salient for linguistic purposes.

My informal tests were with a set of about 20 quasi-logical relations (from a TextWise tool called KNOW-IT) that can be extracted from free text with rules using the most obvious linguistic markers: for example, commas marking off an appositional phrase signal the is-a relationship, an apostrophe + s in English marks the possessive or has-a relationship, a certain limited number of prepositions mark various time relationships. I found that when these are tested against the emerging upper level the outcomes are sometimes very useful, sometimes perverse, sometimes-- like

most of the Hovy tests-- bland. Explaining these outcomes is actually a project that might drive us into interesting and largely unexplored theoretical territory, not to mention the applications to heterogeneous databases.

Lee

From ???@??? Mon Aug 17 11:23:01 1998
Received: from mail-gw3.pacbell.net (mail-gw3.pacbell.net [206.13.28.55])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id AAA16440
for <phayes@coginst.uwf.edu>; Mon, 17 Aug 1998 00:44:11 -0500 (CDT)
Received: from pacbell.net (ppp-206-170-7-85.rdcy01.pacbell.net [206.170.7.85]) by
mail-gw3.pacbell.net (8.8.8/8.7.1+antispam) with ESMTP id WAA21466; Sun, 16 Aug
1998 22:38:47 -0700 (PDT)
Message-ID: <35D7C1EC.63FEFAB0@pacbell.net>
Date: Sun, 16 Aug 1998 22:38:52 -0700
From: Robert Spillers <skydog@pacbell.net>
X-Mailer: Mozilla 4.04 [en] (Win95; I)
MIME-Version: 1.0
To: John Sowa <sowa@west.poly.edu>
CC: Bill Andersen <chezewiz@erols.com>,
Chris Menzel <cmenzel@turing.stanford.edu>,
Doug Skuce <doug@site.uottawa.ca>, Fritz Lehmann <fritz@cyc.com>,
John McCarthy <jmc@cs.stanford.edu>,
Nancy Lawler <E6NL001@coe.coppin.umd.edu>,
Nicola Guarino <guarino@ladseb.pd.cnr.it>,
Pat Hayes <phayes@coginst.uwf.edu>,
Peter Simons <p.m.simons@leeds.ac.uk>,
Adam Farquhar <axf@KSL.Stanford.EDU>,
Robert Spillers <skydog@pacbell.net>
Content-Type: multipart/alternative; boundary="-----
-0B2BA1810DC489779FF25EC1"
Content-Length: 4029
Status:

<x-html><!x-stuff-for-pete base="" src="" id="0"><HTML>

<P>John Sowa writes:

To answer Adam's question, we can use the following criterion for

classifying entities according to prototypes:

<P>Given a similarity measure $m(x,y)$, every category c can be assigned

a prototype or typical instance p .

Then an entity x can be classified

by the following recursive procedure:

<P> 1. Suppose that x has been classified as an instance of some category c ,

 which has subcategories s_1, s_2, \dots, s_N .

<P> 2. Measure that similarity $m(x, p_i)$ of x to each prototype p_i for
the

 subcategory s_i .

<P> 3. Classify x as an instance of that subtype s_i for which $m(x, p_i)$

 indicates the greatest similarity.

<P>For any entity x , this procedure is invoked with x compared to the immediate

subtypes of the universal type T . After x has been classified
as an instance

of any category c , the procedure is invoked recursively to classify
 x further

by one of the subcategories of c . The procedure stops when x
is classified

as an instance of some category whose only proper subtype is the absurd

type at the bottom of the lattice.

<P>Is this measure of similarity, $m(x, y)$, a consistent one?
Should one attempt to classify a cow in the same way one would classify
a contract or a lamp? If an entity can be classified in more than
one way - a dog is a pet and a canine - how does one know when it is properly
and/or exhaustively classified?

<P>Does a measure of similarity imply a measure of distance? Since
ontologies can be of arbitrary size and uneven in their granularity,
by distance I mean something more than simply the shortest path length.
Are these measures related? Would two entities similarity or distance
be different if the size or granularity of the ontology were different?

<P>Bob</HTML>

</x-html>

From ???@??? Tue Dec 08 17:46:13 1998

Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id RAA10023;

Tue, 8 Dec 1998 17:44:45 -0600 (CST)

X-Sender: phayes@mail.coginst.uwf.edu

Message-Id: <v04003a07b29367e8604d@[143.88.7.118]>

Mime-Version: 1.0

Date: Tue, 8 Dec 1998 17:43:11 -0600

To: phayes@ai.uwf.edu

From: Pat Hayes <phayes@coginst.uwf.edu>

Subject:

removal

Cc: Andreas.Reuter@eml.org, Baerbel.Mack@eml.villa-bosch.de,
E6NL001@coe.coppin.umd.edu, Giovanni.VARILE@LUX.DG13.cec.be,
Klaus.Tschira@kts.villa-bosch.de, Martin.v.d.Berg@let.uva.nl,
Piek.Vossen@let.uva.nl, antonio@anite-systems.lu, chezewiz@erols.com,
cmenzel@tamu.edu, fellbaum@thought.princeton.edu, fritz@cyc.com,
geo@thought.princeton.edu, gh@cs.toronto.edu, guarino@ladseb.pd.cnr.it,
hovy@isi.edu, jmc@cs.stanford.edu, lreeker@nsf.gov, phayes@ai.uwf.edu,
p.m.simons@leeds.ac.uk, peters@csl.stanford.edu,
polanyi@pal.xerox.com, skydog@pacbell.net, Sophia.Ananiadou@eml.org,
sowa@west.poly.edu, tsujii@is.s.u-tokyo.ac.jp, wahlster@cs.uni-sb.de

Content-Type: text/plain; charset="us-ascii"

Content-Length: 374

Status:

Please remove Doug Skuce (doug@csi.uottawa.ca) from the CC list of future messages.

Thanks.

Pat

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Dec 09 11:09:43 1998

Received: from tigershark.villa-bosch.de (tigershark.villa-bosch.de [195.185.79.67])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id CAA29591

for <phayes@coginst.uwf.edu>; Wed, 9 Dec 1998 02:13:22 -0600 (CST)

Received: by TIGERSHARK with Internet Mail Service (5.0.1457.3)

id <X3TKNB3Q>; Wed, 9 Dec 1998 09:12:58 +0100

Message-ID: <E17818F052FCD11ADFF00609793D170134CD4@TIGERSHARK>

From: "Reuter, Andreas" <Andreas.Reuter@eml.org>

To: Pat Hayes <phayes@coginst.uwf.edu>

Cc: "Ananiadou, Sophia" <Sophia.Ananiadou@eml.org>,
Iris Flechtner

<Iris.Flehtner@villa-bosch.de>

Subject: RE: web sites

Date: Wed, 9 Dec 1998 09:12:52 +0100

X-Priority: 3

MIME-Version: 1.0

X-Mailer: Internet Mail Service (5.0.1457.3)

Content-Type: text/plain

Content-Length: 1457

Status:

Pat,

yes, we are still keeping the web site that we had installed for the workshop. It has not been touched since, but if you want to use for the ongoing discussions - great. The URL is <http://www.ontology.villa-bosch.de/>. Right now it is password protected, and we should probably keep it this way until the results are finalized. The user name is "ontology", and the password is "villabosch". We will send you the templates for the pages, so you can edit them. Whenever there is something new, just mail them back to us, and we will put them on the server. That would certainly be the easiest way of doing it.

Best wishes

Andreas

> -----Original Message-----

> From: Pat Hayes [SMTP:phayes@coginst.uwf.edu]

> Sent: Monday, December 07, 1998 11:29 PM

> To: Andreas.Reuter@eml.villa-bosch.de

> Subject: web sites

>

> Andreas, greetings. Tell me, do you already have a website set up for
> use

> by the participants in the 'ontology' workshop? If so, we could

> perhaps use

> that instead of making a new one (?). In any case, could you send me

> the

> current URL for that site, just so that I can keep all the records

> straight?

>

> Many thanks

>

> Pat Hayes

>

> -----

> IHMC, University of West Florida (850)434 8903 home

> 11000 University Parkway (850)474 2091 office

> Pensacola, FL 32514 (850)474 3023 fax

> phayes@ai.uwf.edu
> http://www.coginst.uwf.edu/~phayes

>

>

From ???@??? Mon Dec 14 17:15:46 1998
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id RAA00546;
Mon, 14 Dec 1998 17:03:19 -0600 (CST)

X-Sender: phayes@mail.coginst.uwf.edu

Message-Id: <v04003a07b29b429edef2@[143.88.7.118]>

Mime-Version: 1.0

Date: Mon, 14 Dec 1998 17:01:53 -0600

To: Andreas.Reuter@eml.org, Baerbel.Mack@eml.villa-bosch.de,
E6NL001@coe.coppin.umd.edu, Giovanni.VARILE@LUX.DG13.cec.be,
Klaus.Tschira@kts.villa-bosch.de, Martin.v.d.Berg@let.uva.nl,
Piek.Vossen@let.uva.nl, antonio@anite-systems.lu, chezewiz@erols.com,
cmenzel@tamu.edu, fellbaum@thought.princeton.edu, fritz@cyc.com,
geo@thought.princeton.edu, gh@cs.toronto.edu, guarino@ladseb.pd.cnr.it,
hovy@isi.edu, jmc@cs.stanford.edu, lreeker@nsf.gov, phayes@ai.uwf.edu,
p.m.simons@leeds.ac.uk, peters@csl.stanford.edu,
polanyi@pal.xerox.com, skydog@pacbell.net, Sophia.Ananiadou@eml.org,
sowa@west.poly.edu, tsujii@is.s.u-tokyo.ac.jp, wahlster@cs.uni-sb.de

From: Pat Hayes <phayes@coginst.uwf.edu>

Subject: timing

Content-Type: text/plain; charset="us-ascii"

Content-Length: 1126

Status:

Ahem.

Ladies and Gentlemen

A few days ago I said it would take a few days to construct a webpage for our communal use. Now, I want to say that I was NOT LYING when I said that, nor have I ever asked anyone else to lie. However, the meaning of 'few' is very flexible, as any man approaching the late middle years knows in his heart only too well. So while it was only a few days ago I said that, I have to tell you that it may still be a few more days (in the same sense) before this actually gets done.

I know that there is something WRONG with this, and I apologise for it unreservedly. I have now made my peace with God, and so if you don't like it then you can just [undecipherable].

Pat Hayes

PS. Of course, when I said 'day', what I meant was 'working day'. Happy holiday.

PPS. We will probably use the villa-bosch website. Thanks, Andrea.

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Tue Jan 05 11:01:17 1999
Received: from tigershark.villa-bosch.de (tigershark.villa-bosch.de [195.185.79.67])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA29128
for <phayes@ai.uwf.edu>; Wed, 16 Dec 1998 12:47:52 -0600 (CST)
Received: by tigershark.villa-bosch.de with Internet Mail Service (5.0.1457.3)
id <ZA3B4MXN>; Wed, 16 Dec 1998 19:47:14 +0100
Message-ID:
<E17818F052FCD111ADFF00609793D17013AE4B@tigershark.villa-bosch.de>
From: "Reuter, Andreas" <Andreas.Reuter@eml.org>
To: E6NL001@coe.coppin.umd.edu, Giovanni.VARILE@LUX.DG13.cec.be,
Martin.v.d.Berg@let.uva.nl, Piek.Vossen@let.uva.nl,
"Ananiadou, Sophia"
<Sophia.Ananiadou@eml.org>,
antonio@anite-systems.lu, chezewiz@erols.com, cmenzel@tamu.edu,
fellbaum@thought.princeton.edu, fritz@cyc.com,
geo@thought.princeton.edu, gh@cs.toronto.edu, guarino@ladseb.pd.cnr.it,
hovy@isi.edu, jmc@cs.stanford.edu, lreeker@nsf.gov,
p.m.simons@leeds.ac.uk, peters@csl.stanford.edu, phayes@ai.uwf.edu,
polanyi@pal.xerox.com, skydog@pacbell.net, sowa@west.poly.edu,
tsujii@is.s.u-tokyo.ac.jp, wahlster@cs.uni-sb.de
Cc: "Tschira, Klaus" <Klaus.Tschira@kts.villa-bosch.de>
Subject: Re: Season's Greetings
Date: Wed, 16 Dec 1998 19:47:13 +0100
X-Priority: 3
MIME-Version: 1.0
X-Mailer: Internet Mail Service (5.0.1457.3)
Content-Type: text/plain
Content-Length: 282
Status:

Dear workshop participants,
I don't want to discuss whether christmas is physical object or an

event, whether merry can be defined in logical terms - I just want to wish you a Merry Christmas from Villa Bosch, and the all the best for your plans in the New Year.

Sincerely

Andreas

From ???@??? Tue Jan 05 11:01:44 1999

Received: from tigershark.villa-bosch.de (tigershark.villa-bosch.de [195.185.79.67])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id HAA16825
for <phayes@coginst.uwf.edu>; Fri, 18 Dec 1998 07:25:18 -0600 (CST)

Received: from picassotrigger.villa-bosch.de by tigershark.villa-bosch.de with SMTP
(Microsoft Exchange Internet Mail Service Version 5.0.1457.7)
id Z1AFLQ9Q; Fri, 18 Dec 1998 14:25:51 +0100

Message-ID: <367A58D5.1515CC4B@eml.org>

Date: Fri, 18 Dec 1998 14:29:57 +0100

From: Iris Flechtner <iris.flechtner@eml.org>

X-Mailer: Mozilla 4.05 [de] (WinNT; I)

MIME-Version: 1.0

To: Pat Hayes <phayes@coginst.uwf.edu>

Subject: ontology-web-sites

Content-Type: multipart/mixed; boundary="-----E7C74C507063DF8F50A68011"

Content-Length: 1311

Status:

Hello Mr. Hayes,

I am the webmaster at the eml.

Andreas Reuter told me to send you all the ontology-web-sites so you could edit them. I don't think it is necessary. You just have to save the pages you want to edit and then send the edited one to me so I can put them on the server, resp. replace the old ones with the new ones. Nevertheless I can send you the sites if you prefer that. Just mail me how you want to handle it.

Greetings,

Iris Flechtner

Content-Type: text/x-vcard; charset=us-ascii; name="vcard.vcf"

Content-Transfer-Encoding: 7bit

Content-Description: Visitenkarte f,r Iris Flechtner

Content-Disposition: attachment; filename="vcard.vcf"

Attachment converted: lonestar:vcard.vcf 1 (TEXT/ttxt) (00007EF6)

From ???@??? Sun Feb 14 09:47:26 1999

Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id VAA03706

for <phayes@ai.uwf.edu>; Wed, 10 Feb 1999 21:41:22 -0600 (CST)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id WAA11063;
Wed, 10 Feb 1999 22:36:24 -0500 (EST)
Received: by west (SMI-8.6/SMI-SVR4)
id WAA17577; Wed, 10 Feb 1999 22:27:09 -0500
Date: Wed, 10 Feb 1999 22:27:09 -0500
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199902110327.WAA17577@west>
To: fritz@cyc.com
Subject: Ontology Book
Cc: Andreas.Reuter@eml.org, Baerbel.Mack@eml.villa-bosch.de,
E6NL001@coe.coppin.umd.edu, Giovanni.VARILE@LUX.DG13.cec.be,
Klaus.Tschira@kts.villa-bosch.de, Martin.v.d.Berg@let.uva.nl,
Sophia.Ananiadou@eml.org, chezewiz@erols.com, cmenzel@tamu.edu,
fellbaum@thought.princeton.edu, geo@thought.princeton.edu,
gh@cs.toronto.edu, guarino@ladseb.pd.cnr.it, hovy@isi.edu,
jmc@cs.stanford.edu, lreeker@nsf.gov, phayes@ai.uwf.edu,
skydog@pacbell.net, sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 1941
Status:

Fritz,

I agree with you that something should be done to get the show on the road, and the best way to get something done is to establish a firm deadline with some mechanism to enforce the deadline. But we also need a plan for making sure that what gets done will be good. I agree with you that a small good book is much better than a large mediocre book.

Following is a suggestion:

1. Decide on a fixed time for an editing meeting -- a 5-day meeting of anyone and everyone from last June's group who is willing to work for 5 days on putting together all the material that has been contributed; decide what will be accepted, rejected, and revised; and make the revisions during the meeting so that a completed book is produced by Friday afternoon, with nothing left to do but the final formatting and copyediting.
2. Before the EM date, all submissions from all authors must be submitted in machine-readable form. Ideally, everyone should have a chance to read every submission long before the EM date, but we can be sure of two inevitable facts: some people will send in their submissions at

the very last minute, and many of the participants won't get around to reading all the submissions before they arrive at the EM.

3. Before the EM, we can correspond by email to do as much organizing and allocation of chapters, authors, etc., as possible. But we can also be sure that there will be last minute changes, etc., that have to be made during the EM.
4. As a proposed date for the meeting, I would suggest some time this summer, say in August -- exact date and location to be determined by email discussion.

This is a first-pass suggestion, and the details are open to negotiation. But I think that two things are important: first, a firm date for an editing meeting; and the beginning of the EM as the deadline for all submissions to the book.

John Sowa

From ???@??? Tue Feb 23 10:58:07 1999

Received: from mail-gw.pacbell.net (mail-gw.pacbell.net [206.13.28.25])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id RAA21677
for <phayes@ai.uwf.edu>; Wed, 17 Feb 1999 17:23:16 -0600 (CST)

Received: from pacbell.net (ppp-206-171-33-86.rdcy01.pacbell.net [206.171.33.86]) by
mail-gw.pacbell.net (8.8.8/8.7.1+antispam) with ESMTP id PAA15094; Wed, 17 Feb
1999 15:16:46 -0800 (PST)

Message-ID: <36CB4DC5.F801DB0A@pacbell.net>

Date: Wed, 17 Feb 1999 15:16:21 -0800

From: Robert Spillers <skydog@pacbell.net>

X-Mailer: Mozilla 4.5 [en] (Win95; I)

X-Accept-Language: en

MIME-Version: 1.0

To: "John F. Sowa" <sowa@west.poly.edu>

CC: fritz@cyc.com, Andreas.Reuter@eml.org, Baerbel.Mack@eml.villa-bosch.de,
E6NL001@coe.coppin.umd.edu, Giovanni.VARILE@LUX.DG13.cec.be,
Klaus.Tschira@kts.villa-bosch.de, Martin.v.d.Berg@let.uva.nl,
Sophia.Ananiadou@eml.org, chezewiz@erols.com, cmenzel@tamu.edu,
fellbaum@thought.princeton.edu, geo@thought.princeton.edu,
gh@cs.toronto.edu, guarino@ladseb.pd.cnr.it, hovy@isi.edu,
jmc@cs.stanford.edu, lreeker@nsf.gov, phayes@ai.uwf.edu

Subject: Re: Ontology Book

References: <199902110327.WAA17577@west>

Content-Transfer-Encoding: 7bit

Content-Type: text/plain; charset=us-ascii

Content-Length: 2730

Status:

John,

I agree that a meeting for the purpose of reviewing and final editing of a proposed draft is appropriate and useful. However, the meeting should not be scheduled until the entire draft(s) has been circulated and commented on by all of the attendees (via e-mail). The meeting should deal with the refinement of a text and not an open discussion of the underlying ideas (this should occur prior to the meeting via e-mail). If this seems a good idea I will attempt to arrange such a meeting - probably at Stanford. We could hold it in conjunction with a meeting of the ANSI Ad Hoc Group where the text could be reported to the committee and the standards process begun.

Comments?

Bob

"John F. Sowa" wrote:

> Fritz,

>

> I agree with you that something should be done to get the show on the road,
> and the best way to get something done is to establish a firm deadline
> with some mechanism to enforce the deadline. But we also need a plan for
> making sure that what gets done will be good. I agree with you that a small
> good book is much better than a large mediocre book.

>

> Following is a suggestion:

>

> 1. Decide on a fixed time for an editing meeting -- a 5-day meeting of
> anyone and everyone from last June's group who is willing to work for
> 5 days on putting together all the material that has been contributed;
> decide what will be accepted, rejected, and revised; and make the
> revisions during the meeting so that a completed book is produced
> by Friday afternoon, with nothing left to do but the final formatting
> and copyediting.

>

> 2. Before the EM date, all submissions from all authors must be submitted
> in machine-readable form. Ideally, everyone should have a chance to
> read every submission long before the EM date, but we can be sure of
> two inevitable facts: some people will send in their submissions at
> the very last minute, and many of the participants won't get around to
> reading all the submissions before they arrive at the EM.

>

> 3. Before the EM, we can correspond by email to do as much organizing and
> allocation of chapters, authors, etc., as possible. But we can also be

> sure that there will be last minute changes, etc., that have to be made
> during the EM.
>
> 4. As a proposed date for the meeting, I would suggest some time this summer,
> say in August -- exact date and location to be determined by email
> discussion.
>
> This is a first-pass suggestion, and the details are open to negotiation.
> But I think that two things are important: first, a firm date for an
> editing meeting; and the beginning of the EM as the deadline for all
> submissions to the book.
>
> John Sowa

From ???@??? Tue Feb 23 10:58:10 1999
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTTP id GAA17285
for <phayes@ai.uwf.edu>; Thu, 18 Feb 1999 06:22:17 -0600 (CST)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id HAA21225;
Thu, 18 Feb 1999 07:17:14 -0500 (EST)
Received: by west (SMI-8.6/SMI-SVR4)
id HAA00335; Thu, 18 Feb 1999 07:07:54 -0500
Date: Thu, 18 Feb 1999 07:07:54 -0500
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199902181207.HAA00335@west>
To: skydog@pacbell.net
Subject: Re: Ontology Book
Cc: Andreas.Reuter@eml.org, Baerbel.Mack@eml.villa-bosch.de,
E6NL001@coe.coppin.umd.edu, Giovanni.VARILE@LUX.DG13.cec.be,
Klaus.Tschira@kts.villa-bosch.de, Martin.v.d.Berg@let.uva.nl,
Sophia.Ananiadou@eml.org, chezewiz@erols.com, cmenzel@tamu.edu,
fellbaum@thought.princeton.edu, fritz@cyc.com,
geo@thought.princeton.edu, gh@cs.toronto.edu, guarino@ladseb.pd.cnr.it,
hovy@isi.edu, jmc@cs.stanford.edu, lreeker@nsf.gov, phayes@ai.uwf.edu,
sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 22024
Status:

Bob,

I recently came across the following discussion by Guha, who worked on Cyc as the associate director for many years before leaving to go to Apple and now to Netscape, where he has been working on RDF (Resource Description

Framework). About halfway through the note, he makes the following comment:

Building Cyc turned out to be much more difficult than ever imagined. It turns out that a lot of fundamental research needs to be done before we can go about actually building something like Cyc. My guess would be that in 10 or 15 years the time will be right to try again.

I don't agree that we have to wait 10 or 15 years, but Guha's opinion is one that we can't ignore. I believe that more of the fundamental research has been done than Guha has considered. But before we can present a unified approach to a reference ontology, we have to take stock of what has been done, what remains to be done, and what makes us think that we can accomplish something that Guha abandoned after working on it for seven and a half years.

I believe that we can make a good case for saying that now is the time to work towards building a reference ontology. But our email discussions indicate that we don't yet have a unified vision of what it would mean or how to proceed on building one. Before we write a book that presents the reference ontology, we have to tell the world (and ourselves first of all) where we are now, where we want to go, how we propose to get there, and last, but not least, why.

John

INNOVATORS OF THE NET: RAMANATHAN V. GUHA AND RDF

By Marc Andreessen, Chief Technology Officer
Netscape Communications Corporation

January 8, 1999 - Before Johannes Gutenberg used movable type to create mass-produced books, people read large, heavy, handwritten manuscripts that were often chained to desks. But even Gutenberg's mass-produced books were oversized and difficult to carry around. It was a subsequent printer, Aldus Manutius, who came up with the idea of making books pocket-sized so they would be portable. Similarly, the web has been a revolutionary new form of communication, but it's the Resource Description Framework (RDF) and Extensible Markup Language (XML) that will help make it accessible from anywhere, in a wide range of formats. You'll no longer have to be chained to your PC to view web content.

Browsers were revolutionary because they allowed web content to be accessed by everyone, regardless of the computing platform they were using. And now, because RDF and XML allow content to be separated from

its presentation, I think they're the standards-based technologies that will really take the web to the next level, enabling many different kinds of devices to become web clients and enabling new ways of organizing and navigating information. With his pioneering work on RDF, Ramanathan V. Guha, known as Guha, just may be the Aldus of the web.

Currently a principal engineer at Netscape, Guha's work on RDF really began in 1995, when he wrote the Meta Content Framework (MCF) while at Apple. MCF is like a table of contents for a web site. Also at Apple, Guha wrote Project X, a browser plug-in that Apple turned into a web navigation system called HotSauce. Project X used MCF to let web users fly through 3D maps of web sites.

Guha came to Netscape in 1997, and after meeting consultant Tim Bray, who was working on XML, he decided to turn MCF into an XML application. The result, the Resource Description Framework (RDF), is currently under development as a World Wide Web Consortium (W3C) standard. RDF is designed to provide an infrastructure to support metadata across a variety of web-based activities. RDF can be used for sitemaps, content ratings systems, search engine data collection systems, digital library collections, and distributed authoring systems. RDF changes the way people interact with the web.

At Netscape, Guha has focused his efforts on creating the architecture for the RDF-enabled Smart Browsing features in Netscape Communicator 4.5 and on building Aurora, an RDF reader planned for the next generation of Netscape browsers. Aurora will allow users to organize diverse types of content in a personalized way, instead of based on the application used to access the content.

Guha has been characterized by some people as having more ideas than the rest of the Netscape Client Product Division combined. He's a programming maniac and is often found in a hacking frenzy, trying to produce a prototype of his most recent idea. He spends a lot of time spreading his ideas around and working with various industry groups. Underneath his obvious intelligence and the energy field that surrounds him, Guha is a great guy who's fun to know.

He received a bachelor of science degree in mechanical engineering from the Indian Institute of Technology in Madras, India, in 1986; a master of science degree in mechanical engineering from the University of California at Berkeley in 1987; and a Ph.D. in computer science from Stanford in 1991. With Douglas B. Lenant, he coauthored a book, *Building Large Knowledge-Based Systems*, which was published by Addison-Wesley in 1989. In addition to that book, he's written numerous papers and technical reports, and he's applied for five patents. In

April 1998, Guha received one of the first Web Innovator Awards from cnet's Builder.com for his contributions to RDF technology.

I recently met with Guha to talk about RDF.

Marc Andreessen: You began your career by doing research in artificial intelligence (AI). Tell me how you got started with that.

Ramanathan V. Guha: While I was at UC Berkeley working on my master's degree in mechanical engineering, I took one course in AI. I liked it and didn't have a summer job, so I sent a resume to a research consortium in Austin, Texas, called Microelectronics and Computer Technology Corporation (MCC). They were working on the Cyc project, which was an attempt to build a commonsense knowledge base for AI. They called me to say no, they didn't need me, but they were so sweet. They said, "Well, we really don't have room. We already have all of our students." I said, "OK, no problem. I'll work at night." They said, "But we don't have an office for you," and I said, "That's okay, I'll work wherever." So they hired me. After three weeks, they said, "You get your own office, and we want you to stay." I ended up staying there for seven-and-a-half years.

AI is one of the grand challenges today, and I was in my twenties then and easily influenced, so I wholeheartedly flung myself into that effort. Cyc, a ten year project at MCC, was fascinating. Here's the basic idea: There's a lot of stuff that a five-year-old knows that a computer doesn't know. How would we go about teaching a computer that corpus of knowledge? The body of knowledge is pretty much the same across all people. If you ask someone "If I drop something, what will happen to it? Will it fall?" they will say, "Yes." Or if you ask them "What color is the sky on a clear day?" they'll say "Blue." There is a core consensus. It doesn't matter if the person is a computer scientist or a doctor or a farmer or whatever, there's a substrate of consensus knowledge or common sense that we all share. Interestingly, that's exactly the kind of stuff that computers don't know. So if you want computers to go from being the stupid kinds of things they are today to being natively interesting, it's not specialized expertise, such as drawing vector graphics that they need. Rather, they need to know the zillions of little things that enable us to function as intelligent beings.

The Cyc project was about hunkering down and building a machine-understandable corpus of all this knowledge. Building Cyc turned out to be much more difficult than ever imagined. It turns out that a lot of fundamental research needs to be done before we can go about actually building something like Cyc. My guess would be that in

10 or 15 years the time will be right to try again. Such a beast will be needed if computers are ever going to be able to understand human languages, such as English, and do 1700 of the other things that we absolutely take for granted that people can do.

I think that the core issue of whether a computer is intelligent is not going to be an interesting question. The question of what does it mean to be alive was a primary question in biology for many hundreds of years. We still don't know what it means to be alive. Is a virus alive? No...well, maybe. It sure can cause havoc, but it can't reproduce by itself. That question turned out to be largely irrelevant, because there was no real answer. It was the wrong question; nobody asks that question anymore. I think the same thing is going to happen to the concept of intelligence in computers.

But you went back to school after you started working on the Cyc project?

I went back and finished my master's thesis, then decided if I was going to stay in that field, I might as well get a Ph.D. So I went to Stanford and got my Ph.D. in computer science. For a while, I was living this crazy life: teaching a class about the Cyc project at Stanford and the University of Texas at Austin, managing 25 people in Austin and Palo Alto, and writing my thesis - all at the same time.

In 1995, after my work on the Cyc project, I tried to do a startup. I had an idea for a heterogeneous database integration engine that I called Babelfish. The idea was to find a way to describe the semantics of the schemas of relational databases so that a program could transparently query a large heterogeneous set of databases. I wrote the program, but could not sell it. That kind of product is an extremely high-end-enterprise kind of thing - difficult for a lone kid in Texas to sell. I also realized that I was more interested in building the product than in building a business. After I built the product prototype, I decided to return to research.

When did your work on RDF begin to take shape?

Right around the time I was finishing up with Babelfish, Alan Kay, who was then an Apple fellow, convinced me to go to Apple, which was where I developed the Meta Content Framework (MCF). MCF was a way to represent metadata structures - information about information - to bridge the gaps in information flow created by various heterogeneously structured software products. The goal of MCF was not unlike the goal of Babelfish. But research at Apple imploded when Jobs came back to Apple in 1997; within a month or so, Apple just got out of all research.

There were two places that were clearly the places to go if I wanted to be able to reach millions of people. One was Microsoft, the other was Netscape. They were the two big platforms: the operating system and the browser. I chose Netscape.

So in February 1997 I came to Netscape, where I met Tim Bray, coauthor of the W3C's XML 1.0 specification and, at the time, a consultant for Netscape. Tim and I decided to adapt MCF using XML. This ultimately resulted in the W3C spec for the Resource Description Framework (RDF), which is based significantly on MCF.

One of the interesting things about the Internet is that the stakes are so high. Whenever a market is growing exponentially like this, the stakes are so high that people try crazy things and things get accelerated. I realized that there was more innovation going on in companies like Netscape that were trying to get products out and change people's lives. They didn't think of it as research, by any stretch of the imagination. It just had to be done. As a research guy I would have said, "Yeah, that's a very interesting problem. Let's work on it for the next six months and write a couple of papers about it." At Netscape we say, "Oh, this seems like a hard problem. Let's see if we can solve it by dinner and ship it next week." It's a different attitude. More often than not, when you have to solve it by dinner, you don't solve it by dinner, but you do have an 80 percent solution in three days. It's more exciting, and you have more impact on people's lives. You actually make more progress this way.

What's the fundamental concept behind RDF?

If you look at different structures that organize information on your computer, you have a file system, you have mailboxes, you have online directories, you have browser bookmarks and history - you have what I call a Balkanization of information across all these structures. Wouldn't it be great if you were able to organize all this information based on your own terms, instead of based on the application you use to access the information? The RDF reader codenamed Aurora is planned for the next generation of Netscape browsers. You can use Aurora to create a folder called "Cooking," and the items in that folder can be your local files related to cooking, mail messages related to cooking, bookmarks related to cooking, and parts of an online directory related to cooking. It's completely a content-oriented organization. It's organization on your own terms.

How can RDF change the way people experience the web?

The web, as we know it today, is totally and completely in its infancy.

When television first came out, people went on TV and read from a piece of paper. They took what they did on radio, except they sat in front of a television camera and did it. When the telephone came out, people thought it was a replacement for town criers. They would put a telephone in the middle of town, and somebody would call in and make announcements. When Gutenberg used the printing press, all he did was take those big books - the books in those days were big, heavy, handwritten things that were attached to desks with chains - and enable them to be copied a little bit faster. It took Aldus Manutius, many decades later, to realize that the printing press could enable a new kind of communication. Aldus was the one who changed the size of the book to something that people could carry around; he based the size of the book on the size of the horse saddle. He also invented the concept of random access to the book. Before Aldus, we didn't have page numbers, chapters, or indexes. So it takes a while before people realize they can do other things with a new medium. The way we're using the web today is a lot like print media. You pick up a magazine and you go to a page, then you go to another page and another page. People are excited because they can jump around random pages and jump from magazine to magazine. Ho hum. It's still the same fundamental structure, which is that you have one person talking to N thousand people.

But that's beginning to change. In Netscape Communicator 4.5, we introduced Smart Browsing. The way you browse the web with Smart Browsing is different. Say you go to www.nfl.com. If I see that you're looking at the National Football League site, I may know a few other places you'd be interested in. Is it in the NFL's interest to point you to those places? No, because you'll leave their site. But your friends would definitely tell you about other, related sites. With Smart Browsing, your browser provides access to a service that functions like a friend's recommendations, giving you alternate suggestions like, "Based on what you're doing, these are other places you might want to go." But it's not just other sites that you might want to look at. Imagine, as you're about to submit your credit card number to buy something on the Net, having Dun & Bradstreet come along and give you information about the company, such as its Dun & Bradstreet rating, its location, and how many employees it has. Or say you go to a politician's Web page and don't know if what the politician is saying is true or false. ABC News might have done a story about the politician stating that the information is false, but you, the person visiting the politician's home page, would have no way of knowing that. Wouldn't it be nice if you had one or more third parties acting on your behalf, saying, "Since you're looking at this, here's a contradictory view that you might be interested in"?

We're working with Alexa, a service that harnesses the experiences of

the Internet community and provides useful information about the sites being viewed and suggests related sites. We're also working to create our own databases - from directories, through the Open Directory, and in a whole bunch of other ways. It's a big stew. If a user types something into the browser that is not an URL, such as a keyword, the browser will try to figure out what information the user is really looking for by querying these Smart Browsing providers. This gets users in touch more quickly with the information, products, and services that Netscape Netcenter serves as a hub for.

Once you have the basic infrastructure for doing something like that, you can do all kinds of things. The fundamental difference to the user is that it's no longer a bilateral relationship between you and the content provider. It's a multilateral relationship between you, the content provider, and one or more of these third-party providers who are watching what's happening and are participating actively. This does raise the issue of privacy. The presumption is that the people who are watching what you're doing can be trusted. You're not going to give that ability to just anybody. You'll want to set it as an option. Ultimately, what's required is a better system for guarantees of privacy of information. We'll need to have an organization that serves as an independent auditor of information dissemination.

The basic difference between the way we're consuming information today and the kinds of features RDF allows is that with RDF you're making use of the fact that everybody is together on a network. It's not a one-to-many communications model, it's a many-to-many model, unencumbered by the physical limitations of geography. Everybody around the world can be connected together. You can talk to this person, this other person, then get something from here, and piece it all together. It's so radically different from a print journal experience.

What about the server side of Smart Browsing?

We've put all of the client-side hooks for Smart Browsing features in Communicator 4.5. Now we've got to build up the service on the server side. It just makes sense. Instead of baking all the functionality into the client, it is far better for it to come down from the server. This means that a lot of the smarts on how the client behaves come from the network. The advantage is that it can be adapted easily. You don't have to go through a lot to upgrade it. Another advantage is that servers can be a lot smarter. A client is isolated, but a server, serving a lot of people, can pull together information and become smart very fast. After all, if the purpose of the application is to enable interconnectivity, why not have aspects of the program itself come off the Net? So John Giannandrea [another Netscape principal engineer] and

I am working with the Netscape Netcenter team to build the basic infrastructure for Smart Browsing integration on Netcenter's servers, and we'll go from there.

Why would you want to put such a large part of the intelligence of the browser on the server side? Information is on the server and compute cycles are on the client. Information is far more valuable than compute cycles. It's far easier to update the data and software running on a couple of servers than it is to update 70 million clients. Additionally, an Internet service provider or a corporate customer can create its own Smart Browsing databases, modifying or adding to the ones Netcenter provides, or it can defer to the Netcenter databases.

One of the great things about Netscape Navigator is that it allowed content providers to communicate with their users irrespective of which platform they were on. This platform heterogeneity is going to substantially increase over the next couple of years. With Palm Pilots, Windows CE devices, and cell phones all becoming first-class Internet clients, we are going to see a new breed of Internet content. RDF and XML are going to become very important players in this new world. These formats allow a complete separation of the content from the presentation. When you are delivering a piece of data that could be rendered on anything from a 40x60-pixel cell phone to a 1600x1200-pixel PC monitor, it is really important to deliver the data in a fashion that doesn't make any assumptions about how it is going to be presented.

I think it was Bertrand Russell who said that physical restrictions, such as geography, have constrained large sets of people. The whole structure of society is based around these limitations. Most communications that we have, in terms of mass communications, are bilateral. We gave up on multilateral communications because of our physical limitations. The Net changes all these rules. The limits just go away, and there are some fundamental changes that are going to take place because of that. So, I'm rallying against this tyranny of place and time and holding on to the thought that, somewhere - it could be a mile away, it could be a hundred thousand miles away - someone has the information that you need. With the work we are doing at Netscape, our goal is that the information can be shared.

Marc Andreessen is cofounder and chief technology officer at Netscape Communications. Andreessen developed the idea for the NCSA Mosaic browser for the Internet while he was an undergraduate student at the University of Illinois and a staff member at the university's National Center for Supercomputing Applications in Champaign, Illinois. He created the friendly, easy-to-use navigational tool for the Internet with a team of students and staff at NCSA in early 1993.

From ???@??? Tue Feb 23 10:58:18 1999
Received: from mail-gw2.pacbell.net (mail-gw2.pacbell.net [206.13.28.53])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id TAA19046
for <phayes@ai.uwf.edu>; Thu, 18 Feb 1999 19:44:35 -0600 (CST)
Received: from pacbell.net (ppp-206-170-6-33.rdcy01.pacbell.net [206.170.6.33]) by
mail-gw2.pacbell.net (8.8.8/8.7.1+antispam) with ESMTP id RAA05548; Thu, 18 Feb
1999 17:38:26 -0800 (PST)
Message-ID: <36CCC07F.F052C52C@pacbell.net>
Date: Thu, 18 Feb 1999 17:38:08 -0800
From: Robert Spillers <skydog@pacbell.net>
X-Mailer: Mozilla 4.5 [en] (Win95; I)
X-Accept-Language: en
MIME-Version: 1.0
To: "John F. Sowa" <sowa@west.poly.edu>
CC: Andreas.Reuter@eml.org, Baerbel.Mack@eml.villa-bosch.de,
E6NL001@coe.coppin.umd.edu, Giovanni.VARILE@LUX.DG13.cec.be,
Klaus.Tschira@kts.villa-bosch.de, Martin.v.d.Berg@let.uva.nl,
Sophia.Ananiadou@eml.org, chezewiz@erols.com, cmenzel@tamu.edu,
fellbaum@thought.princeton.edu, fritz@cyc.com,
geo@thought.princeton.edu, gh@cs.toronto.edu, guarino@ladseb.pd.cnr.it,
hovy@isi.edu, jmc@cs.stanford.edu, lreeker@nsf.gov, phayes@ai.uwf.edu
Subject: Re: Ontology Book
References: <199902181207.HAA00335@west>
Content-Transfer-Encoding: 7bit
Content-Type: text/plain; charset=us-ascii
Content-Length: 3350
Status:

John,

I think we must have very different ideas about the content of the book. My view is the book is about the theoretical (logical, philosophical, mathematical) basis for the construction of an ontology. In particular, what I have in mind is a cogent (and hopefully concise) discussion of the definition and usage of ontological factors that the authors agree should be used to create an upper level Reference Ontology - not identify all the nodes of such an ontology. Further, I hope, the recommended structure of the Reference Ontology will be a factored combinatorial graph that results in a lattice.

What I mean by factors are the topics assigned to authors in the theory group such as independent/dependent, universal/particular, individual/plurality, abstract/concrete, continuant/occurrent, etc. Once there is a published text of their definitions and usage, it will still require a substantial additional effort to use them to create an upper level Reference Ontology. However, I believe the work done to create the current Reference Ontology (work done primarily by Ed Hovy and Fritz Lehmann) will not be wasted, but it will need to

be restructured into a new form.

In the attachment to your note, Guha refers to the CYC representation of "real world" knowledge. My understanding is that, in fact, it is a very large ontology consisting of hundreds of thousands of assertions of the type mentioned by Guha. A project of this size is far beyond anything I have contemplated. However, it is also my understanding that CYC will seriously consider converting this ontology into the form we recommend if they (Doug, Fritz & company) agree that the form makes sense and makes a business difference to CYC.

Bob

"John F. Sowa" wrote:

> Bob,
>
> I recently came across the following discussion by Guha, who worked on
> Cyc as the associate director for many years before leaving to go to Apple
> and now to Netscape, where he has been working on RDF (Resource Description
> Framework). About halfway through the note, he makes the following comment:
>
> Building Cyc turned out to be much more difficult than ever imagined.
> It turns out that a lot of fundamental research needs to be done before
> we can go about actually building something like Cyc. My guess would be
> that in 10 or 15 years the time will be right to try again.
>
> I don't agree that we have to wait 10 or 15 years, but Guha's opinion
> is one that we can't ignore. I believe that more of the fundamental
> research has been done than Guha has considered. But before we can
> present a unified approach to a reference ontology, we have to take
> stock of what has been done, what remains to be done, and what makes
> us think that we can accomplish something that Guha abandoned after
> working on it for seven and a half years.
>
> I believe that we can make a good case for saying that now is the time
> to work towards building a reference ontology. But our email discussions
> indicate that we don't yet have a unified vision of what it would mean
> or how to proceed on building one. Before we write a book that presents
> the reference ontology, we have to tell the world (and ourselves first
> of all) where we are now, where we want to go, how we propose to get there,
> and last, but not least, why.
>
> John

(Attachment deleted)

From ???@??? Tue Feb 23 10:58:19 1999
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id VAA02093
for <phayes@ai.uwf.edu>; Thu, 18 Feb 1999 21:41:01 -0600 (CST)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id WAA22357;
Thu, 18 Feb 1999 22:35:56 -0500 (EST)
Received: by west (SMI-8.6/SMI-SVR4)
id WAA01478; Thu, 18 Feb 1999 22:26:47 -0500
Date: Thu, 18 Feb 1999 22:26:47 -0500
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199902190326.WAA01478@west>
To: skydog@pacbell.net
Subject: Re: Ontology Book
Cc: Andreas.Reuter@eml.org, Baerbel.Mack@eml.villa-bosch.de,
E6NL001@coe.coppin.umd.edu, Giovanni.VARILE@LUX.DG13.cec.be,
Klaus.Tschira@kts.villa-bosch.de, Martin.v.d.Berg@let.uva.nl,
Sophia.Ananiadou@eml.org, chezewiz@erols.com, cmenzel@tamu.edu,
fellbaum@thought.princeton.edu, fritz@cyc.com,
geo@thought.princeton.edu, gh@cs.toronto.edu, guarino@ladseb.pd.cnr.it,
hovy@isi.edu, jmc@cs.stanford.edu, lreeker@nsf.gov, phayes@ai.uwf.edu
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 1562
Status:

Bob,

Your summary is pretty close to what we've been talking about for the past couple of years. I would be happy to do something along those lines. But I think that we should also include a rationale that explains why we are doing this project and how it differs from what Guha (and quite a few other people ranging from Martin Kay to George Lakoff) don't believe is possible right now.

There are still some areas of disagreement that come up among some of us in email discussions. For example, Pat thinks that an ontology suitable for NL would be quite different from one suitable for KB. But I believe that it would be possible to have a common base ontology that could be used for both purposes. And Ed has raised questions about whether the combinatorial method can support the superstructure for WordNet. These are some fundamental issues that we have to resolve among ourselves before we can propose them to the whole world.

There are some related points that should be ironed out and presented in an introductory section that explains how this ontology can be used to support more specialized domain ontologies. And there should be lots of examples that show people how to make those connections.

Basically, I agree that the definition of the factors and the method for combining them to form the more detailed categories should be the centerpiece for this book. But I think that we also need some introductory and explanatory material that tells people how this relates to what other groups are doing, such as Express, UMLS, RDF, etc.

John

From ???@??? Tue Mar 02 09:37:43 1999
Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id XAA05561
for <phayes@ai.uwf.edu>; Mon, 1 Mar 1999 23:35:14 -0600 (CST)
Received: from scratchy (scratchy [207.207.8.118])
by catbert.cyc.com (8.8.8/8.8.8) with SMTP id XAA18207;
Mon, 1 Mar 1999 23:28:55 -0600 (CST)
Message-Id: <3.0.32.19990301232911.00f52868@catbert.cyc.com>
X-Sender: fritz@catbert.cyc.com
X-Mailer: Windows Eudora Pro Version 3.0 (32)
Date: Mon, 01 Mar 1999 23:29:13 -0600
To: biederman@tu-darmstadt.de, wille@tu-darmstadt.de, phayes@ai.uwf.edu
From: Fritz Lehmann <fritz@cyc.com>
Subject: Tropes and triadic contexts (2nd try)
Cc: fritz@cyc.com
Mime-Version: 1.0
Content-Type: text/plain; charset="us-ascii"
Content-Length: 1590
Status:

Dear Klaus, Rudolf and Pat,

I'd like to direct your attention to the Stanford Internet Encyclopaedia article on "Tropes". It is at:

<http://plato.stanford.edu/entries/tropes/>

A trope appears to me to be exactly an "X" in the 2-dimensional relational cross-table of the "contexts" used in Formal Concept Analysis, relating individual and property. I think the tropes community in philosophy would benefit from knowing about Formal Concept Analysis. All the issues discussed in the article cited above will be very familiar to anyone who knows Formal Concept Analysis.

The statement in the article about possible worlds, wherein a trope relates property, individual and world, is very suggestive of the "triadic contexts". The problems with the algebra of triadic contexts, triadic Galois connections, and trilattices are: A. nobody understands them except Klaus and maybe Rudolf (not me anymore!), and B. nobody cares, or sees why they are important. The tropes community in philosophy might recognize the importance.

One person to contact might be John Bacon at

<http://www-personal.usyd.edu.au/~jbacon/homepage.html>

Pat, I mention this to you so you'll be well-armed to deal with tropism when you're hit with it. I think I had had some difficulty describing tropes clearly in an earlier email message after our meeting in Heidelberg.

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

=====
=====

From ???@??? Tue Mar 02 11:10:18 1999

Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id KAA22628
for <phayes@ai.uwf.edu>; Tue, 2 Mar 1999 10:48:31 -0600 (CST)

Received: from scratchy (scratchy [207.207.8.118])
by catbert.cyc.com (8.8.8/8.8.8) with SMTP id KAA08099;
Tue, 2 Mar 1999 10:42:36 -0600 (CST)

Message-Id: <3.0.32.19990302104257.00f42050@catbert.cyc.com>

X-Sender: fritz@catbert.cyc.com

X-Mailer: Windows Eudora Pro Version 3.0 (32)

Date: Tue, 02 Mar 1999 10:42:57 -0600

To: Pat Hayes <phayes@ai.uwf.edu>

From: Fritz Lehmann <fritz@cyc.com>

Subject: Formal Concept Analysis

Cc: Fritz Lehmann <fritz@cyc.com>

Mime-Version: 1.0

Content-Type: text/plain; charset="us-ascii"

Content-Length: 2145

Status:

At 09:46 AM 3/2/99 -0600, Pat Hayes wrote:

>Fritz, thanks for the info and pointers. However I confess to never having
>heard of Formal Concept Analysis, either, so do you have any pointers to

>*that*?
>Pat

Dear Pat,

You might try

<http://www.math.tu-dresden.de/~ganter/EFCAlinks.html>

FCA is primarily the work of mathematicians at Darmstadt, led by Rudolf Wille. The original article was "Restructuring Lattice Theory", around 1980 in the first of the NATO-series books on "Ordered Sets". The standard introduction in English is Wille's article in my collection on "Semantic Networks in Artificial Intelligence", which was a reprint of the special issue of "Computers and Mathematics with Applications" journal, vol. 23, no.s 2-9, 1992. Since then they have expanded far beyond the Galois connection between bundles of properties/attributes and sets of individuals. A newer book by Ganter & Wille is supposed to be translated into English and published by Springer.

Triadic concept analysis was invented (I'm proud to say) but not really understood (I'm ashamed to say) by me -- and developed mathematically by Rudolf Wille and Klaus Biedermann. It's pretty fancy mathematics, about the hierarchies (actually quasiorders in the triadic case, and "trilattices") induced by a triadic relation among object, property and "context" (possible world, source, whatever) which creates a "triadic Galois connection".

Work on triadic FCA and trilattices has been reported in mathematics journals, and (barely) more accessibly in the proceedings of recent Conceptual Structures conferences. The FCA people filled a gap in the work of the CG people, because FCA tells you what the type lattice is going to be, based on the features of things, instead of letting you just make it up as an arbitrary poset (as in Cyc).

I'm going to send that message a third time -- I keep getting their email wrong.

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

=====
=====

From ???@??? Mon Mar 08 10:52:23 1999
Received: from pdadr1.pd.cnr.it (pdadr1.pd.cnr.it [150.178.1.2])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTTP id GAA23461
for <phayes@ai.uwf.edu>; Sat, 6 Mar 1999 06:43:49 -0600 (CST)
Received: from [150.178.99.15] (adr15.ppp.pd.cnr.it [150.178.99.15])
by mail.pd.cnr.it (PMDF V5.2-27 #23415)
with ESMTTP id <01J8IFQSY0VC000HQM@mail.pd.cnr.it> for phayes@ai.uwf.edu; Sat,
6 Mar 1999 13:38:31 MET
From: Nicola.Guarino@ladseb.pd.cnr.it
Date: Sat, 06 Mar 1999 13:41:13 +0100
Subject: On the meaning of "Conceptualization"
In-reply-to: <199903061131.GAA28377@west>
X-Sender: guarino@mail.ladseb.pd.cnr.it
To: sowa@west.poly.edu (John F. Sowa), fritz@cyc.com
Cc: Andreas.Reuter@eml.org, Baerbel.Mack@eml.villa-bosch.de,
E6NL001@coe.coppin.umd.edu, Giovanni.VARILE@LUX.DG13.cec.be,
Klaus.Tschira@kts.villa-bosch.de, Martin.v.d.Berg@let.uva.nl,
Piek.Vossen@let.uva.nl, Sophia.Ananiadou@eml.org,
andersen@knowledgebus.com, antonio@anite-systems.lu, cmenzel@tam.u.edu,
fellbaum@thought.princeton.edu, geo@thought.princeton.edu,
gh@cs.toronto.edu, hovy@isi.edu, jmc@cs.stanford.edu, lreeker@nsf.gov,
p.m.simons@leeds.ac.uk, peters@csl.stanford.edu, phayes@ai.uwf.edu,
polanyi@pal.xerox.com, skydog@pacbell.net, sowa@west.poly.edu,
tsujii@is.s.u-tokyo.ac.jp, wahlster@cs.uni-sb.de
Message-id: <v03102802b306cf21e0f7@[150.178.99.15]>
MIME-version: 1.0
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2368
Status:

Fritz:

>That's not an issue I care about much, but I'd like to lobby in general
>against reliance on the word "conceptualization". Some people think that
>the statement "An ontology is a specification of a conceptualization." is
>>true and instructive. I don't. First, it's false (since the same ontology
>can have many different "specifications" in Italian, English, Loglan, LOOM,
>CGs, CycL, MELD, LINCOS, KIF, XAQ, NUDE, Buginese, etc.), and second,
>nobody understands what you've said anyway because now you have to define
>"conceptualization" --- which generally turns out to mean: an ontology.
>Lotsa' help.

John:

>
>Re "conceptualization": I agree that that word has been bandied about
>without any clear definition. I don't use it myself for that very reason.
>

I have EXTENSIVELY discussed the meaning of the word "conceptualization" and its relationship with "ontology" in various papers (all accessible from our web site). The first paper was:

Guarino, N. and Giaretta, P. 1995. Ontologies and Knowledge Bases: Towards a Terminological Clarification. In N. Mars (ed.) Towards Very Large Knowledge Bases: Knowledge Building and Knowledge Sharing 1995. IOS Press, Amsterdam: 25-32.

The most recent one is my introduction to the FOIS'98 conference:

Guarino, N. 1998. Formal Ontology in Information Systems. In N. Guarino (ed.) Formal Ontology in Information Systems. Proceedings of FOIS'98, Trento, Italy, 6-8 June 1998. IOS Press, Amsterdam: 3-15.

In these papers I argue that Gruber's definition of "ontology" can only be accepted if the term "Conceptualization" is understood in an intensional way (that is, differently from the definition of "conceptualization" given in the Nilsson and Genesereth textbook on AI, and cited by Gruber). A formal definition of conceptualization, ontology and ontological commitment is given in the FOIS paper. Since in these years I got many positive feedbacks to these papers and NO observations, I believe that the issue "conceptualization vs. ontology" should be considered as settled...

-- Nicola

Nicola Guarino
National Research Council phone: +39 049 8295751
LADSEB-CNR fax: +39 049 8295763
Corso Stati Uniti, 4 email: Nicola.Guarino@ladseb.pd.cnr.it
I-35127 Padova
Italy
<http://www.ladseb.pd.cnr.it/infor/ontology/ontology.html>

From ???@??? Mon Mar 08 10:52:26 1999
Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id LAA22419
for <phayes@ai.uwf.edu>; Sat, 6 Mar 1999 11:31:43 -0600 (CST)
Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id MAA07337;

Sat, 6 Mar 1999 12:26:22 -0500 (EST)
Received: by west (SMI-8.6/SMI-SVR4)
id MAA29697; Sat, 6 Mar 1999 12:16:36 -0500
Date: Sat, 6 Mar 1999 12:16:36 -0500
From: sowa@west.poly.edu (John F. Sowa)
Message-Id: <199903061716.MAA29697@west>
To: Nicola.Guarino@ladseb.pd.cnr.it, fritz@cyc.com
Subject: Re: On the meaning of "Conceptualization"
Cc: Andreas.Reuter@eml.org, Baerbel.Mack@eml.villa-bosch.de,
E6NL001@coe.coppin.umd.edu, Giovanni.VARILE@LUX.DG13.cec.be,
Klaus.Tschira@kts.villa-bosch.de, Martin.v.d.Berg@let.uva.nl,
Piek.Vossen@let.uva.nl, Sophia.Ananiadou@eml.org,
andersen@knowledgebus.com, antonio@anite-systems.lu, cmenzel@tamu.edu,
fellbaum@thought.princeton.edu, geo@thought.princeton.edu,
gh@cs.toronto.edu, hovy@isi.edu, jmc@cs.stanford.edu, lreeker@nsf.gov,
p.m.simons@leeds.ac.uk, peters@csl.stanford.edu, phayes@ai.uwf.edu,
polanyi@pal.xerox.com, skydog@pacbell.net, tsujii@is.s.u-tokyo.ac.jp,
wahlster@cs.uni-sb.de
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 678
Status:

Bill, Nicola, et al.,

These email exchanges make it abundantly clear that we are not using words in the same way among ourselves. Even for terms that some of us have considered "settled".

Recommendation: If there are any terms that any of us feel are necessary for use in this book in more than an offhand mention, then I suggest that the person who wants to have us all use the term write a one-paragraph definition (which may be quoted or adapted from some published work by himself, herself, or anybody else). If after some suitable discussion and/or revision, we agree to use the word in that sense, then we can declare it "settled" for the purpose of this book.

John

From ???@??? Mon Mar 08 16:51:43 1999
Received: from [143.88.7.118] (eels.coginst.uwf.edu [143.88.7.118])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id QAA09248;
Mon, 8 Mar 1999 16:33:12 -0600 (CST)
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04003a07b309fcb7699d@[143.88.7.118]>
In-Reply-To: <199903061716.MAA29697@west>
Mime-Version: 1.0

Date: Mon, 8 Mar 1999 16:33:21 -0600

To: sowa@west.poly.edu (John F. Sowa), Nicola.Guarino@ladseb.pd.cnr.it

From: Pat Hayes <phayes@ai.uwf.edu>

Subject: Re: On the meaning of "Conceptualization"

Cc: Andreas.Reuter@eml.org, Baerbel.Mack@eml.villa-bosch.de,
E6NL001@coe.coppin.umd.edu, Giovanni.VARILE@LUX.DG13.cec.be,
Klaus.Tschira@kts.villa-bosch.de, Martin.v.d.Berg@let.uva.nl,
Piek.Vossen@let.uva.nl, Sophia.Ananiadou@eml.org,
andersen@knowledgebus.com, antonio@anite-systems.lu, cmenzel@tamu.edu,
fellbaum@thought.princeton.edu, geo@thought.princeton.edu,
gh@cs.toronto.edu, hovy@isi.edu, jmc@cs.stanford.edu, lreeker@nsf.gov,
p.m.simons@leeds.ac.uk, peters@csl.stanford.edu, phayes@ai.uwf.edu,
polanyi@pal.xerox.com, skydog@pacbell.net, tsujii@is.s.u-tokyo.ac.jp,
wahlster@cs.uni-sb.de

Content-Type: text/plain; charset="us-ascii"

Content-Length: 1947

Status:

>Bill, Nicola, et al.,

>

>These email exchanges make it abundantly clear that we are not using words

>in the same way among ourselves. Even for terms that some of us have

>considered "settled".

>

>Recommendation: If there are any terms that any of us feel are necessary

>for use in this book in more than an offhand mention, then I suggest that

>the person who wants to have us all use the term write a one-paragraph

>definition (which may be quoted or adapted from some published work

>by himself, herself, or anybody else). If after some suitable discussion

>and/or revision, we agree to use the word in that sense, then we can

>declare it "settled" for the purpose of this book.

>

An excellent suggestion!

Nicola, in the above spirit, could you send me a couple of paragraphs on
'ontology' and 'conceptualization'?

A thought which may be helpful (if it isn't, forget it): on this particular
issue we may have some extra confusion which can be ultimately laid at Tom
Gruber's door. Until recently, "ontology" referred to something independent
of, and conceptually prior to, a particular formalization, while
"conceptualization" (the noun) has often been taken in database terminology
to refer to what a philosopher would call a formal vocabulary. Post-Gruber,
"ontology" has been used to refer to a particular axiomatization, and
"conceptualization" has, as it were, retreated into the territory formerly

occupied by "ontology", so that the contrast between them has reversed itself along the axis (conceptual/syntactically particular).

Pat

PS. I'm beginning to feel like Hamlet. (The time is out of joint. O cursed spite/ that ever I was born to set it right.)

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

From ???@??? Tue Mar 09 11:10:32 1999

Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id SAA19879
for <phayes@ai.uwf.edu>; Mon, 8 Mar 1999 18:08:34 -0600 (CST)

Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id TAA12732;
Mon, 8 Mar 1999 19:03:13 -0500 (EST)

Received: by west (SMI-8.6/SMI-SVR4)
id SAA12511; Mon, 8 Mar 1999 18:53:01 -0500

Date: Mon, 8 Mar 1999 18:53:01 -0500

From: sowa@west.poly.edu (John F. Sowa)

Message-Id: <199903082353.SAA12511@west>

To: Nicola.Guarino@ladseb.pd.cnr.it, phayes@ai.uwf.edu

Subject: Glossary

Cc: Andreas.Reuter@eml.org, Baerbel.Mack@eml.villa-bosch.de,
E6NL001@coe.coppin.umd.edu, Giovanni.VARILE@LUX.DG13.cec.be,
Klaus.Tschira@kts.villa-bosch.de, Martin.v.d.Berg@let.uva.nl,
Piek.Vossen@let.uva.nl, Sophia.Ananiadou@eml.org,
andersen@knowledgebus.com, antonio@anite-systems.lu, cmenzel@tamu.edu,
fellbaum@thought.princeton.edu, geo@thought.princeton.edu,
gh@cs.toronto.edu, hovy@isi.edu, jmc@cs.stanford.edu, lreeker@nsf.gov,
p.m.simons@leeds.ac.uk, peters@csl.stanford.edu,
polanyi@pal.xerox.com, skydog@pacbell.net, sowa@west.poly.edu,
tsujii@is.s.u-tokyo.ac.jp, wahlster@cs.uni-sb.de

Mime-Version: 1.0

Content-Type: text/plain; charset=us-ascii

Content-Length: 14406

Status:

Pat,

Following are some terms that I have included in Appendix B of my forthcoming book. This is an updated version of the list that I circulated to the onto-std mailing list in Dec. 1997. I am hereby contributing them to the gang for anyone who wants to adopt, adapt, consider, reject, or shoot at them.

John

ALIGNMENT. A mapping of concepts and relations between two ontologies A and B that preserves the partial ordering by subtypes in both A and B. If an alignment maps a concept or relation x in ontology A to a concept or relation y in ontology B, then x and y are said to be *_equivalent_*. The mapping may be partial: there could be many concepts in A or B that have no equivalents in the other ontology. Before two ontologies A and B can be aligned, it may be necessary to introduce new subtypes or supertypes of concepts or relations in either A or B in order to provide suitable targets for alignment.

No other changes to the axioms, definitions, proofs, or computations in either A or B are made during the process of alignment. Alignment does not depend on the choice of names in either ontology. For example, an alignment of a Japanese ontology to an English ontology might map the Japanese concept Go to the English concept Five. Meanwhile, the English concept for the verb *_go_* would not have any association with the Japanese concept Go.

AXIOMATIZED ONTOLOGY. A terminological ontology whose categories are distinguished by axioms and definitions stated in logic or in some computer-oriented language that could be automatically translated to logic. There is no restriction on the complexity of the logic that may be used to state the axioms and definitions. The distinction between terminological and axiomatized ontologies is one of degree rather than kind. Axiomatized ontologies tend to be smaller than terminological ontologies, but their axioms and definitions can support more complex inferences and computations. Examples of axiomatized ontologies include formal theories in science and mathematics, the collections of rules and frames in an expert system, and specifications of conceptual schemas in languages like SQL.

DIFFERENTIAE. The properties that distinguish a subtype from other

types that have a common supertype. The term comes from Aristotle's method of defining new types by stating the _genus_ or supertype and stating the properties that distinguish the new type from its supertype. Aristotle's method of definition has become the de facto standard for natural language dictionaries, and it is also widely used for AI knowledge bases and object-oriented programming languages.

HIERARCHY. A partial ordering of entities according to some relation.

A _type hierarchy_ is a partial ordering of concept types by the type-subtype relation. In lexicography, the type-subtype relation is sometimes called the _hypernym-hyponym_ relation. A _meronymy_ is a partial ordering of concept types by the part-whole relation. Classification systems sometimes use a _broader-narrower hierarchy_, which mixes the type and part hierarchies: a type A is considered narrower than B if A is subtype of B or any instance of A is a part of some instance of B. For example, Cat and Tail are both narrower than Animal, since Cat is a subtype of Animal and a tail is a part of an animal. A broader-narrower hierarchy may be useful for information retrieval, but the two kinds of relations should be distinguished in a knowledge base because they have different implications.

IDENTITY CONDITIONS. The conditions that determine whether two different appearances of an object represent the same individual. Formally, if *c* is a subtype of Continuant, the identity conditions for *c* can be represented by a dyadic predicate Id_{c} . Two instances *x* and *y* of type *c*, which may appear at different times and places, are considered to be the same individual if the predicate $Id_{c}(x,y)$ is true. As an example, a predicate Id_{Human} , which determines the identity conditions for the type HumanBeing, might be defined by facial appearance, fingerprints, DNA, or some combination of all those features. At the atomic level, the laws of quantum mechanics make it difficult or impossible to define precise identity conditions for entities like electrons and photons. If a reliable identity predicate Id_{t} cannot be defined for some type *t*, then *t* would be considered a subtype of Occurrent rather than Continuant.

INTEGRATION. The process of finding commonalities between two different ontologies A and B and deriving a new ontology C that facilitates interoperability between computer systems that are based on the A and B ontologies. The new ontology C may replace A or B, or it may be used only as an intermediary

between a system based on A and a system based on B. Depending on the amount of change necessary to derive C from A and B, different levels of integration can be distinguished: alignment, partial compatibility, and unification. Alignment is the weakest form of integration: it requires minimal change, but it can only support limited kinds of interoperability. It is useful for classification and information retrieval, but it does not support deep inferences and computations. Partial compatibility requires more changes in order to support more extensive interoperability, even though there may be some concepts or relations in one system or the other that could create obstacles to full interoperability. Unification or total compatibility may require extensive changes or major reorganizations of A and B, but it can result in the most complete interoperability: everything that can be done with one can be done in an exactly equivalent way with the other.

KNOWLEDGE BASE. An informal term for a collection of information that includes an ontology as one component. Besides an ontology, a knowledge base may contain information specified in a declarative language such as logic or expert-system rules, but it may also include unstructured or unformalized information expressed in natural language or procedural code.

LEXICON. A knowledge base about some subset of words in the vocabulary of a natural language. One component of a lexicon is a terminological ontology whose concept types represent the word senses in the lexicon. The lexicon may also contain additional information about the syntax, spelling, pronunciation, and usage of the words. Besides conventional dictionaries, lexicons include large collections of words and word senses, such as WordNet from Princeton University and EDR from the Japan Electronic Dictionary Research Institute, Ltd. Other examples include classification schemes, such as the Library of Congress subject headings or the Medical Subject Headers (MeSH).

MIXED ONTOLOGY. An ontology in which some subtypes are distinguished by axioms and definitions, but other subtypes are distinguished by prototypes. The top levels of a mixed ontology would normally be distinguished by formal definitions, but some of the lower branches, such as plants, animals, and common household objects might be distinguished by prototypes.

ONTOLOGY. The study of the categories of things that exist or may

exist in some domain. The product of such a study, called an ontology, is a catalog of the types of things that are assumed to exist in a domain of interest D from the perspective of a person who uses a language L for the purpose of talking about D. The types in the ontology represent the predicates, word senses, or concept and relation types of the language L when used to discuss topics in the domain D. An uninterpreted logic, such as predicate calculus, conceptual graphs, or KIF, is ontologically neutral. It imposes no constraints on the subject matter or the way the subject may be characterized. By itself, logic says nothing about anything, but the combination of logic with an ontology provides a language that can express relationships about the entities in the domain of interest.

PARTIAL COMPATIBILITY. An alignment of two ontologies A and B that supports equivalent inferences and computations on all equivalent concepts and relations. If A and B are partially compatible, then any inference or computation that can be expressed in one ontology using only the aligned concepts and relations can be translated to an equivalent inference or computation in the other ontology.

PROTOTYPE-BASED ONTOLOGY. A terminological ontology whose categories are distinguished by typical instances or prototypes rather than axioms and definitions in logic. For every category c in a prototype-based ontology, there must be a prototype p and a measure of semantic distance $d(x,y,c)$, which computes the dissimilarity between two entities x and y when they are considered instances of c. Given such a measure, an entity x can be classified by the following recursive procedure:

- o Suppose that x has already been classified as an instance of some category c, which has subcategories s_{1} , ..., s_{n} .
- o For each subcategory s_{i} with prototype p_{i} , measure the semantic distance $d(x, p_{i}, c)$.
- o If $d(x, p_{i}, c)$ has a unique minimum value for some subcategory s_{i} , then classify x as an instance of s_{i} , and call the procedure recursively to determine whether x can be further classified by some subcategory of s_{i} .

- o If c has no subcategories or if $d(x, p_{i'}, c)$ has no unique minimum for any $s_{i'}$, then the classification procedure stops with x as an instance of c , since no finer classification is possible with the given selection of prototypes.

As an example, a black cat and an orange cat would be considered very similar as instances of the category Animal, since their common catlike properties would be the most significant for distinguishing them from other kinds of animals. But in the category Cat, they would share their catlike properties with all the other kinds of cats, and the difference in color would be more significant. In the category BlackEntity, color would be the most relevant property, and the black cat would be closer to a crow or a lump of coal than to the orange cat. Since prototype-based ontologies depend on examples, it is often convenient to derive the semantic distance measure by a procedure that learns from examples, such as statistics, cluster analysis, or neural networks.

QUINE'S CRITERION. A test for determining the implicit ontology that underlies any language, natural or artificial. The philosopher Willard van Orman Quine proposed a criterion that has become famous: "To be is to be the value of a quantified variable." That criterion makes no assumptions about what actually exists in the world. Its purpose is to determine the implicit assumptions made by the people who use some language to talk about the world. As stated, Quine's criterion applies directly to languages like predicate calculus that have explicit variables and quantifiers. But Quine extended the criterion to languages of any form, including natural languages, in which the quantifiers and variables are not stated as explicitly as they are in predicate calculus. For English, Quine's criterion means that the implicit ontological categories are the concept types expressed by the basic content words in the language: nouns, verbs, adjectives, and adverbs.

TERMINOLOGICAL ONTOLOGY. An ontology whose categories need not be fully specified by axioms and definitions that determine the necessary and sufficient conditions for their use. WordNet is an example of a terminological ontology, whose categories are partially specified by relations such as subtype-supertype or part-whole, which determine the relative positions of the concepts with respect to one another, but which do not

completely define them. Although a terminological ontology may be expressed in logic, the versions of logic required are usually simpler, less expressive, and more easily computable than full first-order predicate calculus.

UNIFICATION. A one-to-one alignment of all concepts and relations in two ontologies A and B. After A and B have been unified, any inference or computation that can be expressed in one can be mapped to an equivalent inference or computation in the other.

From ???@??? Wed Apr 28 09:36:03 1999
Received: from catbert.cyc.com (catbert.cyc.com [207.207.8.5])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id AAA14195
for <phayes@coginst.uwf.edu>; Wed, 28 Apr 1999 00:09:16 -0500 (CDT)
Received: from scratchy (scratchy [207.207.8.118])
by catbert.cyc.com (8.8.8/8.8.8) with SMTP id XAA07520;
Tue, 27 Apr 1999 23:58:24 -0500 (CDT)
Message-Id: <3.0.32.19990427235852.00b007b8@catbert.cyc.com>
X-Sender: fritz@catbert.cyc.com
X-Mailer: Windows Eudora Pro Version 3.0 (32)
Date: Tue, 27 Apr 1999 23:58:54 -0500
To: andersen@knowledgebus.com
From: Fritz Lehmann <fritz@cyc.com>
Subject: Re: Putting out our shingle...
Cc: Bob Spillers <skydog@pacbell.net>,
Bruce Bargmeyer <bargmeyer.bruce@epa.gov>,
Henry Heffernan <henry_heffernan@nih.gov>,
Jim Hendler <jhendler@darpa.mil>,
Mark Westling <westling@perceptech.com>,
Mike Wiener <mike@servtech.net>,
Nancy Lawler <e6nl001@coe.coppin.umd.edu>,
Woojin Paik <woojin@textwise.com>, John Sowa <sowa@west.poly.edu>,
Pat Hayes <phayes@coginst.uwf.edu>, Sean Luke <seanl@cs.umd.edu>,
Fritz Lehmann <fritz@cyc.com>, Brian Peterson <bjpeter@erols.com>,
Josh Engel <engel@erols.com>, Paul Brinkley <gamma@clark.net>
Mime-Version: 1.0
Content-Type: text/plain; charset="us-ascii"
Content-Length: 2891
Status:

At 09:40 PM 4/27/99 -0400, Bill Andersen wrote:

> Folks,
> We just started up the Knowledge Bus, Inc. website. I'm
> sending this to you (a friendly crowd) of technical people
> to:
> a) let you know we're officially in business,

- > b) let you know what we're in business to do, and
- > c) ask you to send comments on the content. It's
- > rough right now but we had to put it up because
- > we got a plug from Sun on their Java website.
- > Looking forward to hearing from you. Thanks!
- > ...bill

Dear Bill,

Nice website! On first glance I had one minor suggestion: in Brian's How and Why page he contrasts syntactic and semantic integration. He emphasizes the wrong words though -- I'd emphasize the "syntactic" and the "semantic" rather than the "interoperability" and "integration". Also, syntax and semantics are rather technical terms, and a lot of DB people have never really figured out what they mean -- so you might want to explain them more plainly, with examples.

A key issue for semantic DB integration is that of composite descriptions of data elements, coded values, etc. I'm not sure what Knowledge Bus does with these, but from my recent experiences at IEEE Metadata and in DB work, I think it very important that DB people not think that they will be able to anticipate every combination of concepts that will be needed for DB applications. This has been the problem of those quixotic DoD Data/Enterprise models/standards, and in medical thesauri. You need to be able to compose a description of, say, a column-meaning, as a logical structure made up of agreed-upon ontological "primitives". (Not ultimate semantic primitives, just easily-agreed-upon basic concepts.) Lambda descriptions, Kappa definitions and Cyc's NATs are among the ways to do this.

My standard example these days is "Hire Date"; initially it may seem fine to just have #HireDate as a relation or slot in the ontology, but that's no good. People can be re-hired, and more than once. Do you mean "the date we last hired the employee" or "the date we first hired the employee"? The initial response of "global data model" types is to say Oh, we'll have both. But the problem is always compounded further, because you cannot anticipate (in your data dictionary or your ontology) every combination of concepts that will be needed to define things in new applications. You need to be able to define a data element with the composite logical description equivalent to the English "the date on which we last hired the employee", etc.

The same applies to XML tag repositories. Lists of "all the tags you'll ever need" won't cut it.

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

=====
=====

From ???@??? Fri May 07 11:03:16 1999

Received: from tigershark.villa-bosch.de (l47.eml.org [195.185.79.56] (may be forged))
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id GAA10519

for <phayes@coginst.uwf.edu>; Fri, 7 May 1999 06:36:30 -0500 (CDT)

Received: by tigershark.villa-bosch.de with Internet Mail Service (5.5.2448.0)
id <JSL0D0X7>; Fri, 7 May 1999 13:33:44 +0200

Message-ID:

<11EDDAB052CBD211AE2800609793D1700465B4@tigershark.villa-bosch.de>

From: "Reuter, Andreas" <Andreas.Reuter@eml.villa-bosch.de>

To: E6NL001@coe.coppin.umd.edu, Giovanni.VARILE@LUX.DG13.cec.be,
Martin.v.d.Berg@let.uva.nl, chezewiz@erols.com, cmenzel@tamu.edu,
fritz@cyc.com, "skydog@pacbell.net" <skydog@pacbell.net>,
"hovy@isi.edu" <hovy@isi.edu>,
"geo@clarity.princeton.edu"
<geo@thought.princeton.edu>,
"sowa@west.poly.edu" <sowa@west.poly.edu>,
"polanyi@pal.xerox.com" <polanyi@pal.xerox.com>,
"guarino@ladseb.pd.cnr.it" <guarino@ladseb.pd.cnr.it>,
"Piek.Vossen@let.uva.nl" <Piek.Vossen@let.uva.nl>,
"peters@csl.stanford.edu" <peters@csl.stanford.edu>,
"fellbaum@clarity.princeton.edu" <fellbaum@thought.princeton.edu>,
"cmenzel@tamu.edu" <cmenzel@tamu.edu>,
"Fritz Lehmann" <fritz@cyc.com>,
"TSUJII Junichi" <tsujii@is.s.u-tokyo.ac.jp>,
"Pat Hayes"
<phayes@coginst.uwf.edu>,
"Peter Simons" <p.m.simons@leeds.ac.uk>,
"Martin.v.d.Berg@let.uva.nl" <Martin.v.d.Berg@let.uva.nl>,
"doug@csi.uottawa.ca" <doug@csi.uottawa.ca>,
"jmc@cs.stanford.edu"
<jmc@cs.stanford.edu>,
"lreeker@nsf.gov" <lreeker@nsf.gov>,
"wahlster@cs.uni-sb.de" <wahlster@cs.uni-sb.de>,
"Giovanni.VARILE@LUX.DG13.cec.be" <Giovanni.VARILE@LUX.DG13.cec.be>,
"Antonio Sanfilippo" <antonio@anite-systems.lu>,
"Graeme Hirst"
<gh@cs.toronto.edu>,
"Josiah Lee Auspitz" <lee@textwise.com>,
"Bill Andersen" <chezewiz@erols.com>,
"E6NL001@coe.coppin.umd.edu"
<E6NL001@coe.coppin.umd.edu>

Cc: "Reuter, Andreas" <Andreas.Reuter@eml.villa-bosch.de>,

Baerbel Mack
<Baerbel.Mack@eml.villa-bosch.de>,
"Tschira, Klaus"
<Klaus.Tschira@kts.villa-bosch.de>

Subject: Re: FYI
Date: Fri, 7 May 1999 13:33:43 +0200
MIME-Version: 1.0
X-Mailer: Internet Mail Service (5.5.2448.0)
Content-Type: text/plain
Content-Length: 593
Status:

Dear all,
regarding the follow-up activities of our workshop last year and future plans in the area of ontologies I have to inform you that Dr. Sophia Ananiadou no longer works for EML. She asked for termination of her contract last week, and we complied with this request. This is a serious setback for our research activities in that field, but nevertheless we are determined to implement our plans according to the original design. We will try to find a successor for the vacant position as soon as possible and will keep you informed about any progress we make.

Best wishes

Andreas Reuter

From ???@??? Mon Sep 27 11:45:34 1999

Received: from tigershark.villa-bosch.de (147.eml.org [195.185.79.56] (may be forged))
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id HAA26536
for <phayes@ai.uwf.edu>; Sun, 26 Sep 1999 07:34:34 -0500 (CDT)

Received: by tigershark.villa-bosch.de with Internet Mail Service (5.5.2448.0)
id <S1A5BXTX>; Sun, 26 Sep 1999 14:40:30 +0200

Message-ID:

<11EDDAB052CBD211AE2800609793D1700A6F1B@tigershark.villa-bosch.de>

From: "Reuter, Andreas" <Andreas.Reuter@eml.villa-bosch.de>

To: Pat Hayes <phayes@ai.uwf.edu>

Cc: "Reuter, Andreas" <Andreas.Reuter@eml.villa-bosch.de>,
"Porzel, Robert" <robert.porzel@eml.villa-bosch.de>

Subject: RE: transcript?

Date: Sun, 26 Sep 1999 14:40:30 +0200

MIME-Version: 1.0

X-Mailer: Internet Mail Service (5.5.2448.0)

Content-Type: text/plain

Content-Length: 1539

Status:

Dear Pat,
thanks for your message. I have forwarded it to one of our researchers, Robert Porzel, who is working on the (major) task of condensing the bulk of

recorded material into something digestible. He will contact you and explain what is available that might be useful for you. A written transcript of all that was said during the ten days does not exist.

Best wishes
Andreas

> -----Original Message-----

> From: Pat Hayes [SMTP:phayes@ai.uwf.edu]
> Sent: Friday, September 24, 1999 6:36 PM
> To: Reuter, Andreas
> Subject: transcript?

>

> Andreas, greetings.

>

> I hope all is well with you.

>

> Bob Spillers tells me that there is a transcript of the recordings of
> the ontology meeting. If such a thing exists (even if only in a
> sketchy or un-edited form) then I would greatly appreciate having
> access to a copy, especially of what was said at the 'language'
> subgroup. Is it accessible electronically? I will of course keep it
> confidential if you wish me to, but it would be of considerable help
> in putting together a coherent summary, if only by serving to remind
> me of the many points which should be checked against other opinions.

>

> I enjoyed reading the EML report recently.

>

> Best wishes

>

> Pat Hayes

>

>

> -----

> IHMC, University of West Florida (850)434 8903 home
> 11000 University Parkway (850)474 2091 office
> Pensacola, FL 32514 (850)474 3023 fax
> phayes@ai.uwf.edu

> <http://www.coginst.uwf.edu/~phayes>

From ???@??? Wed Oct 06 11:24:43 1999

Received: from tigershark.villa-bosch.de (147.eml.org [195.185.79.56] (may be forged))
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id EAA00725
for <phayes@ai.uwf.edu>; Wed, 6 Oct 1999 04:55:53 -0500 (CDT)

Received: from eml.villa-bosch.de (195.185.79.53 [195.185.79.53]) by
tigershark.villa-bosch.de with SMTP (Microsoft Exchange Internet Mail Service Version
5.5.2448.0)

id S1A5B7GW; Wed, 6 Oct 1999 12:02:41 +0200

Message-ID: <37FB1ADC.CEC62BF@eml.villa-bosch.de>
Date: Wed, 06 Oct 1999 11:48:12 +0200
From: Robert Porzel <Robert.Porzel@eml.villa-bosch.de>
Reply-To: Robert.Porzel@eml.villa-bosch.de
Organization: EML
X-Mailer: Mozilla 4.6 [en] (WinNT; I)
X-Accept-Language: en
MIME-Version: 1.0
To: Pat Hayes <phayes@ai.uwf.edu>
Subject: Ontology Videos
Content-Transfer-Encoding: 7bit
Content-Type: text/plain; charset=us-ascii
Content-Length: 682
Status:

Dear Mr. Hayes,

Andreas Reuter informed me of Your interests. I am currently going through the video material of the ontology workshop, with the intend to condense the material into smaller presentations thereof. Please let me know if there is anything I can do to help You. Specifically in which parts you would be interested.

Best wishes,

Robert Porzel

--

=====
Robert Porzel
European Media Laboratory - EML
Villa Bosch, Schloss-Wolfsbrunnenweg 33
D - 69118 Heidelberg, Germany
Tel: +49 - 6221 - 533 203
Fax: +49 - 6221 - 533 298
E-Mail: robert.porzel@eml.villa-bosch
WWW: <http://www.eml.villa-bosch.de>
=====

From ???@??? Thu Oct 07 11:40:15 1999
Received: from tigershark.villa-bosch.de (147.eml.org [195.185.79.56] (may be forged))
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id DAA20953
for <phayes@ai.uwf.edu>; Thu, 7 Oct 1999 03:27:03 -0500 (CDT)

Received: from eml.villa-bosch.de (195.185.79.53 [195.185.79.53]) by tigershark.villa-bosch.de with SMTP (Microsoft Exchange Internet Mail Service Version 5.5.2448.0)

id S1A5B7V1; Thu, 7 Oct 1999 10:31:47 +0200
Message-ID: <37FC5715.D0FAD030@eml.villa-bosch.de>
Date: Thu, 07 Oct 1999 10:17:25 +0200
From: Robert Porzel <Robert.Porzel@eml.villa-bosch.de>
Reply-To: Robert.Porzel@eml.villa-bosch.de
Organization: EML
X-Mailer: Mozilla 4.6 [en] (WinNT; I)
X-Accept-Language: en
MIME-Version: 1.0
To: pat hayes <phayes@ai.uwf.edu>
Subject: Re: Ontology Videos
References: <37FB1ADC.CEC62BF@eml.villa-bosch.de>
<v04210100b421286d5d4f@[143.88.7.173]>
Content-Transfer-Encoding: 7bit
Content-Type: text/plain; charset=us-ascii
Content-Length: 2345
Status:

Hello, next to the sound on the video tapes additional sound was recorded on DAT tape for obtaining a better quality, if you let me know, which sessions you are interested in, I can check whether those sessions have been recorded separately and send the audio material to you on DAT or MC, otherwise I could also take the sound from the video and put that on tape, if the quality was to suffice.

Greetings Robert Porzel

pat hayes wrote:

>>Dear Mr. Hayes,
>>
>>Andreas Reuter informed me of Your interests. I am currently going
>>through the video material of the ontology workshop, with the intend to
>>condense the material into smaller presentations thereof. Please let me
>>know if there is anything I can do to help You. Specifically in which
>>parts you would be interested.
>>
>>Best wishes,
>>
>>Robert Porzel
>
> Hi, thanks for getting in touch. What I was most interested in was a

> transcript of what was said at the meetings, rather than the video
> (though I am sure that would be very entertaining.) There are two
> reasons: one to refresh my own memory of what was said during the
> part of the meeting I was actually present at, but more seriously to
> find out what was said at the other session (the 'language' group
> which met in the larger of the two meeting-rooms). So far neither
> subgroup has written a summary; I am still trying to do this for the
> 'ontology' subgroup, but based more on subsequent email than on the
> meeting itself, and I have no way to discover the details of what was
> being discussed at the other subgroup meeting.

>
> Are there audio tapes of what was said at the meetings, as well as
> the video record?

>
> Pat Hayes

>
> -----
> IHMC, University of West Florida (850)434 8903 home
> 11000 University Parkway (850)474 2091 office
> Pensacola, FL 32514 (850)474 3023 fax
> phayes@ai.uwf.edu
> <http://www.coginst.uwf.edu/~phayes>

--

=====
Robert Porzel
European Media Laboratory - EML
Villa Bosch, Schloss-Wolfsbrunnenweg 33
D - 69118 Heidelberg, Germany
Tel: +49 - 6221 - 533 203
Fax: +49 - 6221 - 533 298
E-Mail: robert.porzelt@eml.villa-bosch
WWW: <http://www.eml.villa-bosch.de>
=====

From ???@??? Fri Oct 08 13:30:38 1999
Received: from [143.88.7.173] (apple.coginst.uwf.edu [143.88.7.173])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id MAA00827;
Fri, 8 Oct 1999 12:37:50 -0500 (CDT)
Mime-Version: 1.0
X-Sender: phayes@mail.coginst.uwf.edu
Message-Id: <v04210101b423d7548497@[143.88.7.173]>
In-Reply-To: <199910080024.UAA24521@west>
References: <199910080024.UAA24521@west>

Date: Fri, 8 Oct 1999 12:33:20 -0500
To: sowa@west.poly.edu (John F. Sowa)
From: pat hayes <phayes@ai.uwf.edu>
Subject: Re: Uppel Level Reference Ontology
Cc: E6NL001@coe.coppin.umd.edu, Tony.Sarris@unisys.com,
andersen@knowledgebus.com, cmenzel@philebus.tamu.edu, doug@cyc.com,
fritz@cyc.com, geo@thought.princeton.edu, jmc@cs.stanford.edu,
larry.reeker@nist.gov, lee@textwise.com, peters@csl.stanford.edu,
phayes@ai.uwf.edu, piek.vossen@let.uva.nl, polanyi@pal.xerox.com,
skydog@pacbell.net, vdberg@pal.xerox.com
Content-Type: text/plain; charset="us-ascii" ; format="flowed"
Content-Length: 359
Status:

<x-flowed>

For the record, and remarkably enough, I agree with everything John says.

Pat Hayes

IHMC, University of West Florida (850)434 8903 home
11000 University Parkway (850)474 2091 office
Pensacola, FL 32514 (850)474 3023 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

</x-flowed>

From ???@??? Fri Oct 08 12:33:20 1999
To: sowa@west.poly.edu (John F. Sowa)
From: pat hayes <phayes@ai.uwf.edu>
Subject: Re: Uppel Level Reference Ontology
Cc: E6NL001@coe.coppin.umd.edu, Tony.Sarris@unisys.com,
andersen@knowledgebus.com, cmenzel@philebus.tamu.edu, doug@cyc.com,
fritz@cyc.com, geo@thought.princeton.edu, jmc@cs.stanford.edu,
larry.reeker@nist.gov, lee@textwise.com, peters@csl.stanford.edu,
phayes@ai.uwf.edu, piek.vossen@let.uva.nl, polanyi@pal.xerox.com,
skydog@pacbell.net, vdberg@pal.xerox.com
Bcc:
X-Attachments:
Message-Id: <v04210101b423d7548497@[143.88.7.173]>
In-Reply-To: <199910080024.UAA24521@west>
References: <199910080024.UAA24521@west>

For the record, and remarkably enough, I agree with everything John says.

Pat Hayes

From ???@??? Fri Oct 08 15:01:40 1999
Received: from catbert.cyc.com (root@catbert.cyc.com [207.207.8.5])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id OAA17829
for <phayes@ai.uwf.edu>; Fri, 8 Oct 1999 14:01:50 -0500 (CDT)
Received: from scratchy (scratchy [207.207.8.118])
by catbert.cyc.com (8.8.8/8.8.8) with SMTP id NAA11901;
Fri, 8 Oct 1999 13:53:19 -0500 (CDT)
Message-Id: <3.0.32.19991008135247.03c4ef38@catbert.cyc.com>
X-Sender: fritz@catbert.cyc.com
X-Mailer: Windows Eudora Pro Version 3.0 (32)
Date: Fri, 08 Oct 1999 13:52:49 -0500
To: sowa@west.poly.edu (John F. Sowa)
From: Fritz Lehmann <fritz@cyc.com>
Subject: Re: Uppel Level Reference Ontology
Cc: E6NL001@coe.coppin.umd.edu, Tony.Sarris@unisys.com,
andersen@knowledgebus.com, cmenzel@philebus.tamu.edu, doug@cyc.com,
fritz@cyc.com, geo@thought.princeton.edu,
andreas.reuter@eml.villa-bosch.de, jmc@cs.stanford.edu,
larry.reeker@nist.gov, lee@textwise.com, peters@csl.stanford.edu,
phayes@ai.uwf.edu, piek.vossen@let.uva.nl, polanyi@pal.xerox.com,
skydog@pacbell.net, vdberg@pal.xerox.com, sowa@west.poly.edu
Mime-Version: 1.0
Content-Type: text/plain; charset="us-ascii"
Content-Length: 7042
Status:

At 08:24 PM 10/7/99 -0400, John F. Sowa wrote:

>Bob et al.,

>

>Following are my comments on your questions. To clarify my responses,

>I would like to summarize my general position about where we are and

>where I think we should be going:

>...

>> b. Can consensus be achieved in the working group?

> That depends on what we are trying to do. Right now, I don't think

> that there exists a consensus, but we might be able to achieve one

> after some preliminary work. The Heidelberg week, for example, did not

> lead to a clear consensus -- as evidenced by the fact that we still do

> not have a written statement of what, if anything, was accomplished.

Dear John et al.,

I disagree with this last point. I consider that the Heidelberg theory group, at least, in which you and I participated, came very close to consensus. The remaining issues, though interesting, were side issues (like whether isomorphism=identity for abstract things, whether much can be

said about contexts, whether there can be genuine abstract occurrents, etc.) some of which may be resolvable in email

>

>> c. Will this standard ever be approved by NCITS, ANSI, ISO?

>

> Example: The first draft of the COBOL specification was written
> in November 1959 and accepted by the COBOL working group in 1960.
> A few tentative implementations were tried, and the spec's were
> rewritten in 1961. None of the early implementations of COBOL were
> widely used until after the first ANSI standard came out in 1965.

>

> Note that six-year time lag. Such a time lag (perhaps greater) will
> exist before a widely accepted upper-level ontology is produced and
> approved.

>

>> d. How will non-U.S. and non-English stakeholders feel about such a
>> standard?

>

> That depends on what kind of standard is proposed.

>

>> e. Are there any valid reasons why a single upper level ontology
>> should NOT become a world standard?

>

> That depends on what is meant by "a single upper level ontology".
> I don't believe that any current proposal, certainly not any one
> that any of us has produced or seen so far comes close. I would not
> approve Cyc as it currently stands, I would not approve WordNet or
> any current variation of it, I would not approve the attempted merger
> of Cyc with other ontologies, and I would not even approve my own
> ontology, as published in the KR book, as it now stands. I believe
> that all those efforts are important things to consider, but designing
> an upper level that SHOULD be approved as a world standard is still
> a research effort.

>

>> f. Could this effort drag on for years with little progress?

>

> That depends on how you define progress. I believe that the Heidelberg
> week made a lot of progress for some of us. In my case, it led to a
> much improved upper level (compare the ontology in my recent book with

- > the version that was distributed in draft form before Heidelberg).
- > But we still have a lot more research to do.
- >
- > In the planning session for Heidelberg (Jan. 1998 at CSLI), I said that
- > we should aim to publish a series of volumes about the ontology and how
- > we were progressing. If we had agreed to publish a collection of
- > individually authored papers after the Heidelberg week, I believe that
- > we would have had a very useful contribution. But it was (and still is)
- > premature to claim that there exists a SINGLE unified consensus.
- >
- >> g. Any other risks or problems we should know about?
- >
- > The issues in (f) are the major problem. We don't have a consensus
- > now, and we won't get it unless we have a major research project.
- >
- >> 2. What is the status of the merged upper level ontology? Was it
- >> completed? How many concepts? Is anyone using it?
- >
- > I was disappointed that we never had an opportunity to take a serious
- > look at the upper levels of the merged ontology. The version that was
- > produced in 1998 had all of the serious flaws of the Cyc upper levels,
- > which no one (other than the Cyclers) fully understands or agrees with.
- >
- >> 3. What is the success of other large ontologies? Why would this one
- >> be any better?
- >
- > There is no upper level ontology that has been seriously studied,
- > analyzed, and accepted by anyone other than the originator(s). That
- > includes both large ones and small ones.
- >
- >> 4. Would this upper level ontology and architecture help in other areas
- >> in need of semantic interoperability? Such as:
- >
- > Such an upper level is ABSOLUTELY ESSENTIAL for the success of all of
- > the following.
- >
- >> a. XML metadata tags?
- >
- > Right now, the XML work is in utter chaos. I have seen some of the
- > work, and I have talked to some of the people doing the work, and
- > they are in far worse shape and know (on average) much, much less
- > about what they are doing and where they are going than the people
- > who met in Heidelberg last year. I am sure that I would make similar
- > comments about groups (b) through (d) if I talked with them.
- >
- > And by the way, I distinguish XML, which as a simplified version of SGML

- > is fine, from the choice of metadata tags that are being developed for
- > RDF and a plethora of other proposals, which are in general chaos.
- >
- >> 5. How would we deal with legacy data, metadata, schemes, etc.
- >
- > Without good guidelines for an upper level, there is no chance
- > of any success at dealing with legacy data coming from any of the
- > groups (a) through (d) mentioned in question 4. With such guidelines,
- > there is some hope.
- >
- >> 6. How is the Ontolingua server coming? Any major success stories?
- >
- > A server is not an ontology. A language is not an ontology. They
- > might be useful if you have an ontology, but by themselves, they
- > won't grow into anything other than a collection of unrelated modules.
- >
- >> 7. If we achieve a standard arch and upper level ontology, but we still
- have
- >> a proliferation of middle level, lower level and undocumented ontologies,
- >> will we be any better off?
- >
- > Yes, definitely, because you would then have a possibility of
- > coordinating and aligning those lower levels. The upper level won't
- > guarantee success, but without one, failure is guaranteed.
- >
- >> 8. Why didn't DARPA fund this? This seems like such a good project, but
- why
- >> aren't others pursuing funding?
- >
- > It is still a research effort, and the short-term efforts for specific
- > projects are always a more pressing priority for the funding agencies.
- >
- >> 9. Besides the people that would be funded to do this work, who else knows
- >> enough about this project to comment on its feasibility?
- >
- > The efforts for this project must be published openly at every stage,
- > with every knowledgeable party able to and encouraged to comment. That
- > is the ISO and ANSI policy on standards.
- >
- >

Yours truly, Fritz Lehmann

Fritz Lehmann, Cycorp, 3721 Executive Center Dr., Austin, TX 78731 USA
email: fritz@cyc.com telephone: (512) 342-4013 fax: (512) 342-4040

=====
=====
From ???@??? Tue Oct 12 12:26:43 1999

Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id SAA08658
for <phayes@ai.uwf.edu>; Fri, 8 Oct 1999 18:21:45 -0500 (CDT)

Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.8.8/8.8.8) with SMTP id TAA10538;
Fri, 8 Oct 1999 19:14:41 -0400 (EDT)

Received: by west (SMI-8.6/SMI-SVR4)
id TAA00416; Fri, 8 Oct 1999 19:14:11 -0400

Date: Fri, 8 Oct 1999 19:14:11 -0400

From: sowa@west.poly.edu (John F. Sowa)

Message-Id: <199910082314.TAA00416@west>

To: fritz@cyc.com

Subject: Re: Uppel Level Reference Ontology

Cc: E6NL001@coe.coppin.umd.edu, Tony.Sarris@unisys.com,
andersen@knowledgebus.com, andreas.reuter@eml.villa-bosch.de,
cmenzel@philebus.tamu.edu, doug@cyc.com, geo@thought.princeton.edu,
jmc@cs.stanford.edu, larry.reeker@nist.gov, lee@textwise.com,
peters@csl.stanford.edu, phayes@ai.uwf.edu, piek.vossen@let.uva.nl,
polanyi@pal.xerox.com, skydog@pacbell.net, sowa@west.poly.edu,
vdberg@pal.xerox.com

Mime-Version: 1.0

Content-Type: text/plain; charset=us-ascii

Content-Length: 1179

Status:

Dear Pat, Fritz, et al.,

Pat, I'm amazed, but pleased that we agree. We usually have enough
in common that we can find lots of arcane details to fight about.
While it lasts, we should savor this bit of agreement.

Fritz, I agree with your point in general:

>I disagree with this last point. I consider that the Heidelberg theory
>group, at least, in which you and I participated, came very close to
>consensus. The remaining issues, though interesting, were side issues
>(like whether isomorphism=identity for abstract things, whether much can be
>said about contexts, whether there can be genuine abstract occurrents,
>etc.) some of which may be resolvable in email.

I think that we did come to a general agreement on a lot of important
issues, but unfortunately, we didn't have enough time to work out the
details, to communicate them properly to the rest of the group, to find

out what the others were thinking about, and to apply them to a detailed analysis of some of the important systems we were considering.

So I'm optimistic that something useful can be done, but more time and effort will have to be devoted to it before we have a publishable consensus.

John

From ???@??? Tue Oct 26 10:06:10 1999
Received: from emlsapc02.villa-bosch.de (l44.eml.org [195.185.79.53] (may be forged))
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id EAA00819
for <phayes@ai.uwf.edu>; Tue, 26 Oct 1999 04:31:36 -0500 (CDT)
Received: from eml.villa-bosch.de (andersen@localhost [127.0.0.1])
by emlsapc02.villa-bosch.de (8.9.3/8.9.3) with ESMTP id LAA03033
for <phayes@ai.uwf.edu>; Tue, 26 Oct 1999 11:26:01 +0200
Sender: andersen@emlsapc02.villa-bosch.de
Message-ID: <381573A9.B8ED46C9@eml.villa-bosch.de>
Date: Tue, 26 Oct 1999 11:26:01 +0200
From: Bill Andersen <bill.andersen@eml.villa-bosch.de>
Organization: EML
X-Mailer: Mozilla 4.61 [en] (X11; I; Linux 2.2.10 i686)
X-Accept-Language: en
MIME-Version: 1.0
To: phayes@ai.uwf.edu
Subject: Greetings from Germany
Content-Transfer-Encoding: 7bit
Content-Type: text/plain; charset=us-ascii
Content-Length: 1150
Status:

Hey, Pat...

We're here at EML working with them on their biochemical ontology and have already run up against a nice ontological problem. The question is:

What is the relation between a chemical reaction construed as combining a substance A and a substance B to form some other substance C and a chemical reaction construed as taking, say, a molecule of A and an atom of B and forming a molecule C? In the former, heat, pressure, and other mass properties are relevant where in the latter, only chemical properties such as electron valence, polarity, etc are important.

All of this ties into, of course, a general theory about the nature of substances and their being viewed as continuous in one context and as collections of discrete particles in another.

Do you know of any good work on axiomatized ontologies of substances that have been done? We're working on our own but I don't like reinventing the wheel.

Any advice you could give would be appreciated.

Cheers... bill

P.S. I'm going to talk with the EML folks about cutting loose some consulting \$ (or DM as the case may be) for you to work on such questions....

From ???@??? Tue Oct 26 14:07:19 1999

Received: from [205.160.76.86] (betelgeuse.coginst.uwf.edu [205.160.76.86])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id KAA08459;
Tue, 26 Oct 1999 10:39:29 -0500 (CDT)

Mime-Version: 1.0

X-Sender: phayes@mail.coginst.uwf.edu

Message-Id: <v04210100b43b75309810@[205.160.76.86]>

In-Reply-To: <381573A9.B8ED46C9@eml.villa-bosch.de>

References: <381573A9.B8ED46C9@eml.villa-bosch.de>

Date: Tue, 26 Oct 1999 10:35:29 -0500

To: Bill Andersen <bill.andersen@eml.villa-bosch.de>

From: pat hayes <phayes@ai.uwf.edu>

Subject: Re: Greetings from Germany

Cc: phayes@ai.uwf.edu

Content-Type: text/plain; charset="us-ascii" ; format="flowed"

Content-Length: 4866

Status:

<x-flowed>> Hey, Pat...

>

> We're here at EML working with them on their biochemical
>ontology and have already run up against a nice ontological
>problem. The question is:

>

> What is the relation between a chemical reaction construed
>as combining a substance A and a substance B to form some other
>substance C and a chemical reaction construed as taking, say,
>a molecule of A and an atom of B and forming a molecule C? In
>the former, heat, pressure, and other mass properties are
>relevant where in the latter, only chemical properties such as

>electron valence, polarity, etc are important.

>

> All of this ties into, of course, a general theory about
>the nature of substances and their being viewed as continuous
>in one context and as collections of discrete particles in
>another.

Yep, you just walked right into Heraclitus' river, among other things.

> Do you know of any good work on axiomatized ontologies of
>substances that have been done? We're working on our own but
>I don't like reinventing the wheel.

Not that would be adequate for you. All the 'substance' stuff that I know of basically thinks of a piece of stuff as being a spatial region with a kind-of-stuff property associated with it, so even mixtures are a problem, let alone reactions.

Wait...Forbus and his students did something some time ago which connected molecular activity with larger-scale stuff. They were thinking of temperature changes and fluid flow rather than reactions, but the molecular/mass distinction was there. As I recall, their approach was basically to "associate" the molecular theory with the large-scale volume, rather than try to have a theory of *how* large-scale reduced to small-scale. So they would make inferences like: these molecular goings-on are taking place in this macroscopic volume; molecular goings-on of this sort entail macroscopic properties of that sort in the same place; so macro events of that sort are happening in this volume. Ask Forbus for details (it was about a decade ago, so my recollection is slightly dim now) but I don't think you will get much more than this.

BTW, how complicated do you expect this to eventually get? For example if you look at what is going on inside a blast furnace there are about five or six different reactions taking place in various parts of the furnace, some producing gas which is moving upward and there reacting with something else, and others increasing the density of liquid which is then moving down and getting hotter, and so on. Do you want to be able to get into this kind of stuff?

> Any advice you could give would be appreciated.

I'd suggest starting with the above kind of approach, where one relies on pretty simple 'equivalence' laws directly relating micro-events happening inside a space to macro-changes in the large-scale stuffs which occupy those spaces. (Blast furnaces are

more complicated because one has to consider rates of flow.) Then move to stating those connections in terms of (ie so they follow from) a kind of molecular theory (so one could conclude things like that the molecules of the macro stuff *constituted* the macrostuff, and maybe get some simple thermodynamics into the picture, ie heat = molecular motion etc. . That would be a fail-soft strategy since the inferences will still be workable even if you can't do the full molecular theory, and you can always just *assert* things like conservation of mass, if it turns out to be too hard to derive them from molecular theory.

HOwever i would suggest taking a look at my old ontology of liquids, right at the end, where it discusses the piece-of-liquid/liquid-object distinction. I think that you will need to keep that distinction (or something like it) straight at the macro level to avoid getting the ontology confused. So for example be very careful of saying something like, 'this lake is made of water molecules'. At a given moment, the lake consists of a liquid object and that LO is made up of water molecules, but the very next moment the (same) lake is a different collection of water molecules. When reactions are taking place it gets even more complicated, eg a flame isn't even a thing at the molecular level, its a process.

>P.S. I'm going to talk with the EML folks about cutting loose
>some consulting \$ (or DM as the case may be) for you to work
>on such questions....

Whatever. Keep me in touch, \$ or no \$, as this stuff is more fun than anything else I'm doing here :-)

Pat

OUR SNAILMAIL ADDRESS HAS CHANGED !

IHMC (850)434 8903 home
40 South Alcaniz St. (850)202 4416 office
Pensacola, FL 32501 (850)202 4440 fax
phayes@ai.uwf.edu
<http://www.coginst.uwf.edu/~phayes>

</x-flowed>

From ???@??? Wed Oct 27 12:03:21 1999

Received: from emilsapc02.villa-bosch.de (144.eml.org [195.185.79.53] (may be forged))
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id DAA23683

for <phayes@ai.uwf.edu>; Wed, 27 Oct 1999 03:12:30 -0500 (CDT)
Received: from eml.villa-bosch.de (andersen@localhost [127.0.0.1])
by emlsapc02.villa-bosch.de (8.9.3/8.9.3) with ESMTP id LAA03621
for <phayes@ai.uwf.edu>; Wed, 27 Oct 1999 11:16:04 +0200
Sender: andersen@emlsapc02.villa-bosch.de
Message-ID: <3816C2D4.3E02E0E2@eml.villa-bosch.de>
Date: Wed, 27 Oct 1999 11:16:04 +0200
From: Bill Andersen <bill.andersen@eml.villa-bosch.de>
Organization: EML
X-Mailer: Mozilla 4.61 [en] (X11; I; Linux 2.2.10 i686)
X-Accept-Language: en
MIME-Version: 1.0
To: pat hayes <phayes@ai.uwf.edu>
Subject: Re: Greetings from Germany
References: <381573A9.B8ED46C9@eml.villa-bosch.de>
<v04210100b43b75309810@[205.160.76.86]>
Content-Transfer-Encoding: 7bit
Content-Type: text/plain; charset=us-ascii
Content-Length: 3716
Status:

pat hayes wrote:

> BTW, how complicated do you expect this to eventually get? For
> example if you look at what is going on inside a blast furnace there
> are about five or six different reactions taking place in various
> parts of the furnace, some producing gas which is moving upward and
> there reacting with something else, and others increasing the density
> of liquid which is then moving down and getting hotter, and so on. Do
> you want to be able to get into this kind of stuff?

Their immediate goals are pretty basic - storage and visualisation of biochemical pathway information. But they say they also want to support digestion of natural language, simulation of biochemical reactions, and hypothesizing of potential new pathways. A tall order and pretty far beyond what anyone else has tried. They originlly viewed the ontology as an aid to the NL system but have since realized that it needs to be made pervasive. It's a great project.

>> Any advice you could give would be appreciated.
>
> I'd suggest starting with the above kind of approach, where one
> relies on pretty simple 'equivalence' laws directly relating
> micro-events happening inside a space to macro-changes in the
> large-scale stuffs which occupy those spaces.

What they have mostly in the biochemical world is data on the macro-level reactions. What I had in mind was using the micro-level to support eventual simulation but in the near term as a source of constraints on the macro-level (e.g., if a molecule of A and a molecule of B *can't* react, then no amount of A-stuff and B-stuff will react...)

> However I would suggest taking a look at my old ontology of liquids,

I don't have it and can't find it in any form that doesn't include buying a big book from MKP. So I guess I will have to wait a while. Is there an online copy floating around (no pun intended) anywhere???

> right at the end, where it discusses the
> piece-of-liquid/liquid-object distinction. I think that you will need
> to keep that distinction (or something like it) straight at the macro
> level to avoid getting the ontology confused. So for example be very
> careful of saying something like, 'this lake is made of water
> molecules'. At a given moment, the lake consists of a liquid object
> and that LO is made up of water molecules, but the very next moment
> the (same) lake is a different collection of water molecules.

So, you're saying the lake has a separate identity than simply the set of molecules that make it up at any given time. I was wondering what you mean then by 'consist' - is it that there is a time-varying mereological relation between the 'lake' and any one of its molecules?? Is there some liquid whole intervening between the 'lake' and the molecules?

> When
> reactions are taking place it gets even more complicated, eg a flame
> isn't even a thing at the molecular level, it's a process.

That's the way we viewed reactions. We have this now:

reactions among molecules are events
reactions among substances have the above as sub-events
metabolisms have the above as sub-events

Funny thing is that this leads to the view that all an organism is is a metabolism and that all the organism is is the metabolism.

>>P.S. I'm going to talk with the EML folks about cutting loose
>>some consulting \$ (or DM as the case may be) for you to work

> >on such questions....

>

> Whatever. Keep me in touch, \$ or no \$, as this stuff is more fun than

> anything else I'm doing here :-)

Talked to the project manager about this. She seems more willing to cut funds for a visit than for consulting. Interested in coming to Germany for a few days sometime to tear up our ontology??

...bill

From ???@??? Thu Oct 28 12:11:01 1999

Received: from emlsapc02.villa-bosch.de (l44.eml.org [195.185.79.53] (may be forged))
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id IAA10520
for <phayes@ai.uwf.edu>; Thu, 28 Oct 1999 08:31:20 -0500 (CDT)

Received: from eml.villa-bosch.de (andersen@localhost [127.0.0.1])
by emlsapc02.villa-bosch.de (8.9.3/8.9.3) with ESMTP id PAA04289
for <phayes@ai.uwf.edu>; Thu, 28 Oct 1999 15:24:35 +0200

Sender: andersen@emlsapc02.villa-bosch.de

Message-ID: <38184E93.BD4EE512@eml.villa-bosch.de>

Date: Thu, 28 Oct 1999 15:24:35 +0200

From: Bill Andersen <bill.andersen@eml.villa-bosch.de>

Organization: EML

X-Mailer: Mozilla 4.61 [en] (X11; I; Linux 2.2.10 i686)

X-Accept-Language: en

MIME-Version: 1.0

To: pat hayes <phayes@ai.uwf.edu>

Subject: Re: Greetings from Germany

References: <381573A9.B8ED46C9@eml.villa-bosch.de>

<v04210100b43b75309810@[205.160.76.86]>

<3816C2D4.3E02E0E2@eml.villa-bosch.de>

<v04210102b43d0ff58d1a@[205.160.76.86]>

Content-Transfer-Encoding: 7bit

Content-Type: text/plain; charset=us-ascii

Content-Length: 5903

Status:

pat hayes wrote:

>> What they have mostly in the biochemical world is data on the

>>macro-level reactions. What I had in mind was using the micro-

>>level to support eventual simulation but in the near term as a

>>source of constraints on the macro-level (e.g., if a molecule

>>of A and a molecule of B *can't* react, then no amount of A-stuff

>>and B-stuff will react...)


>

> OK, but why do you need to go to the molecular level to express this

> kind of constraint? Why not just say that A stuff and B stuff won't
> react, and leave it at that?

Partly an ontological gestapoism on my part and partly because they will eventually need to access molecular-level representations in order to construct simulations at multiple levels.

We need the molecular/atomic/particle view in order to talk about things like valency (a property of molecules and atoms, not of substances), polarity (molecules), and charge (all three). Also, the way stoichiometric equations are expressed pressures one toward a micro-view, at least wrt some properties. Example:



You can't read this as "2 moles of A and one of B give you one of C and one of D" because if you dump the A and B in a tub and set up the right conditions for the reaction, the stoichiometric equation only tells you what the ideal situation would be, not what you would actually observe. But any individual molecular level reaction which actually does happen in the soup of A and B will actually look as the equation describes.

> What do you want to say about the
> situation when A molecules will react with B molecules only vary
> rarely, but more often under pressure and temperature increase? (The
> Villa Bosch museum is a lovely testament to that particular kind of
> insight!)
> Also it occurs to me that there are going to be reactions that can
> take place only when other conditions hold, eg when there is ionic
> water, or catalysts, or a high enough temperature or in contact with
> certain kinds of surface, etc.. (as you can tell, my chemical
> intuitions tend to be INorganic), and these won't be expressible as
> constraints arising from molecular reactions.

Right. In this case it's best to talk about the A-Stuff and B-Stuff view... Obviously, pressure, temperature, etc are only properties at the substance level...

> If you want to have POL's as the 'primary' kind of liquid entity,
> then you could maybe think of LO's as liquid processes (occurents
> rather than continuants). However, that has the peculiar consequence
> that the most 'object-like' liquid things in the macroscopic world -
> cups of coffee, rivers, lakes, even oceans - aren't objects at all,
> and the 'real' objects are almost impossible to capture for more than
> an instant. This has a Ripleyesque believe-it-or-not feel about it to

> me, but it may be more in line with what you need.

You've convinced me. I think we will go with the LO's as primary but use the POL (at least its parts - molecules) to handle certain constraints. The vast majority of the information in the database will look like:

"A reacts with B in the presence of C ..."

which I'm now taking to be type-level statements about what particular blobs (instances) of (LO's) A, B, and C would do if you dumped them in a test tube.

>>> When

>>> reactions are taking place it gets even more complicated, eg a flame

>>> isn't even a thing at the molecular level, its a process.

>>

>> That's the way we viewed reactions. We have this now:

>>

>> reactions among molecules are events

>> reactions among substances have the above as sub-events

>> metabolisms have the above as sub-events

>

> Hmmm, I'm suspicious of saying that molecular events are subevents of

> macro-events. I see that this is correct, in a sense, but there are

> macro properties that aren't mereosums of micro-events in the way

> that this suggests. Temperature, for example, is an *average* of

> molecular motions, not a kind of assembly of them.

>

> At the very least you ought to allow that macro-events have

> 'components' that aren't molecular events as well.

Yup. We anticipated that. Things like reaction rate are really stochastic properties of some reacting LO's under certain ambient conditions.

>> Funny thing is that this leads to the view that all an organism

>> is is a metabolism and that all the organism is is the metabolism.

>

> Yes, good illustration of the problems I was thinking about. After

> all, surely a lot of what's important about how a cell works is the

> physical shapes of the parts involved, and things like osmotic

> pressure on membranes and so forth. Or, look at it another way: even

> if someone were frozen solid and all his molecules stopped reacting,

> he wouldn't *vanish*; so if his body has to be made up of events, then

> there must be something else there as well as the reactions. Just
> being there, filling up space, is doing *something*.

Great point. I had ignored that..

>> Talked to the project manager about this. She seems more willing
>>to cut funds for a visit than for consulting. Interested in coming
>>to Germany for a few days sometime to tear up our ontology??
>
> Ah, yet another trip to the old Schloss, what a prospect. Well, I
> have to say, not in the immediate future; I have travel booked
> extending into next year. Will you still be there early in the new
> millenium? In the meantime, lets keep in touch on the electronic
> autobahn.

No problem. We're supposed to be working on this until April 15th.
After that we don't know. I think they're testing us to see if this
ontology stuff is of any use to a practical project. Maybe next
year, when you want to cool off in Germany you could come over to do
a talk and review what the project has produced. Meanwhile, we'll
keep sending you stuff as long as you want us to keep sending it, ok?

...bill

From ???@??? Tue Nov 09 14:41:47 1999

Received: from mta1.snfc21.pbi.net (mta1.snfc21.pbi.net [206.13.28.122])

by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id AAA07671

for <phayes@ai.uwf.edu>; Tue, 9 Nov 1999 00:49:29 -0600 (CST)

Received: from pacbell.net ([206.170.6.192])

by mta1.snfc21.pbi.net (Sun Internet Mail Server sims.3.5.1999.09.16.21.57.p8)

with ESMTP id <0FKX00K1I3YDTG@mta1.snfc21.pbi.net> for phayes@ai.uwf.edu;

Mon,

8 Nov 1999 22:42:22 -0800 (PST)

Date: Mon, 08 Nov 1999 22:39:27 -0800

From: Robert Spillers <skydog@pacbell.net>

Subject: Ontology Workshop Nov. 18th and 19th

To: Andreas Reuter <andreas.reuter@EML.villa-bosch.de>,

Bob Engelmores <rse@ksl.stanford.edu>, Frank Olken <olken@lbl.gov>,

John McCarthy <jmc@cs.stanford.edu>, Lee Auspitz <lee@textwise.com>,

Livia Polanyi <polanyi@pal.xerox.com>,

Nancy Lawler <E6NL001@coe.coppin.umd.edu>,

Richard Fikes <fikes@ksl.stanford.edu>, Pat Hayes <phayes@ai.uwf.edu>,

Christiane Fellbaum <fellbaum@clarity.Princeton.EDU>,

Chris Menzel <cmenzel@tamu.edu>,

Bill Andersen <andersen@knowledgebus.com>,

Larry Reeker <larry.reeker@nist.gov>,

Tony Sarris <Tony.Sarris@unisis.com>,

Martin van den Berg <vdberg@pal.xerox.com>, "Y.T." <ytchien@nsf.gov>,
Susan Gauch <sgauch@ittc.ukans.edu>,
Deborah McGuinness <dln@ksl.stanford.edu>,
Shelia McIlraith <sam@ksl.stanford.edu>,
Larry Fagan <fagan@SMI.stanford.EDU>,
Mark Musen <Musen@SMI.stanford.EDU>
Cc: Rick Morris <MORRISR1@LEAV-EMH1.ARMY.MIL>,
James Hendler <jhendler@darpa.mil>,
Jim Schoening <schoenin@mail1.monmouth.army.mil>,
Robert Mayes <rmayes@hcfa.gov>
Message-id: <3827C19E.D4453C95@pacbell.net>
MIME-version: 1.0
X-Mailer: Mozilla 4.6 [en] (Win98; I)
X-Accept-Language: en
Content-Type: MULTIPART/MIXED;
BOUNDARY="Boundary_(ID_vMUMCyjuKClyUU3v/BJxWA)"
Content-Length: 38717
Status:

Following up on our phone discussions, here is some information about the ontology meeting. An agenda is attached.

On Nov. 18, 1999 several people from the U.S. Army, Center for Army Lessons Learned (CALL), will attend an invitation only meeting at Stanford University at the Center for Integrated Systems Extension - room CISX 338. Also attending will be representatives of several other government agencies with an interest in ontology. At this meeting we hope the Army will agree to consider financing an upper level Reference Ontology (perhaps in conjunction with other government agencies). If this occurs, we believe start up funds will be available quickly. These funds will be to develop a plan to create a Reference Ontology (RO). If the Army likes the resulting proposal they will ask congress for an increase in their appropriation in March/April and, if approved, major funds will be available about this time next year. This particular group has an excellent track record on this type of funding approach.

Jim Hendler, DARPA Program Manager, will attend and present the new DAML BAA on the 18th. This is a \$30 million program to create an ontological markup language for agents with emphasis on its use over the Internet ^ a sort of ontological extension of XML. Jim would like to coordinate his effort with work done on the Reference Ontology.

There will be a separate but related meeting on Nov. 19th in room X-275

at the Stanford Medical School Office Building to discuss the use of ontology in medicine. We wish to begin building the case with the appropriate government agencies - Health Care Finance Administration (HCFA), the CPR Workgroup, NIH, NSF, NIST, etc. - that work should begin on an ontology for medicine / healthcare. Much terminology work has been done in medicine ^ some of it with an ontological basis ^ and prior investments can be leveraged. A particular interest of HCFA is the auditing of the quality of care given patients by HMOs, hospitals, etc.

I view these meetings as a series of discussions. There are only two presentations scheduled ^ one by the Rick Morris on CALL and the other by Jim Hendler on his new BAA. If charts facilitate the discussion that is fine, but I prefer avoid the atmosphere of a series of lectures.

I look forward to seeing you.

Bob

Attachment converted: Poppy:AGENDA Nov 18.doc (WDBN/MSWD) (0000A8D2)
From ???@??? Tue Nov 09 14:41:49 1999
Received: from mta3.snfc21.pbi.net (mta3.snfc21.pbi.net [206.13.28.141])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTTP id KAA24719
for <phayes@ai.uwf.edu>; Tue, 9 Nov 1999 10:31:11 -0600 (CST)
Received: from pacbell.net ([206.170.6.92])
by mta3.snfc21.pbi.net (Sun Internet Mail Server sims.3.5.1999.09.16.21.57.p8)
with ESMTTP id <0FKX001RBUKWO@mta3.snfc21.pbi.net> for phayes@ai.uwf.edu;
Tue,
9 Nov 1999 08:22:03 -0800 (PST)
Date: Tue, 09 Nov 1999 08:19:09 -0800
From: Robert Spillers <skydog@pacbell.net>
Subject: Nov. 18th Agenda (text)
To: Andreas Reuter <andreas.reuter@EML.villa-bosch.de>,
Bob Engelmores <rse@ksl.stanford.edu>, Frank Olken <olken@lbl.gov>,
John McCarthy <jmc@cs.stanford.edu>, John Sowa <sowa@west.poly.edu>,
Lee Auspitz <lee@textwise.com>, Livia Polanyi <polanyi@pal.xerox.com>,
Nancy Lawler <E6NL001@coe.coppin.umd.edu>,
Richard Fikes <fikes@ksl.stanford.edu>, Pat Hayes <phayes@ai.uwf.edu>,
Christiane Fellbaum <fellbaum@clarity.Princeton.EDU>,
Chris Menzel <cmenzel@tamu.edu>,
Bill Andersen <andersen@knowledgebus.com>,
Larry Reeker <larry.reeker@nist.gov>,
Tony Sarris <Tony.Sarris@unisys.com>,
Martin van den Berg <vdberg@pal.xerox.com>, "Y.T." <yтчien@nsf.gov>,
Susan Gauch <sgauch@ittc.ukans.edu>,
Deborah McGuinness <dml@ksl.stanford.edu>,

Shelia McIlraith <sam@ksl.stanford.edu>,
Larry Fagan <fagan@SMI.stanford.EDU>,
Mark Musen <Musen@SMI.stanford.EDU>
Cc: Rick Morris <MORRISR1@LEAV-EMH1.ARMY.MIL>,
James Hendler <jhendler@darpa.mil>,
Jim Schoening <schoenin@mail1.monmouth.army.mil>,
Robert Mayes <rmayes@hcfa.gov>

Message-id: <3828497C.75510862@pacbell.net>

MIME-version: 1.0

X-Mailer: Mozilla 4.6 [en] (Win98; I)

Content-transfer-encoding: 7bit

X-Accept-Language: en

Content-Type: text/plain; charset=us-ascii

Content-Length: 1674

Status:

For those who can't read the msword.doc attachment here is the agenda.

Bob

Reference Ontology Workshop Meeting
Stanford University
Center for Integrated Systems Extension
Room CISX 338

AGENDA

November 18, 1999

1:00 Open / Welcome /
Introductions Bob Spillers

1:15 Center for Army Lessons Learned (CALL)
Rick Morris
Mission / Objectives / Issues

1:45 Discussion of How Ontology is Relevant to CALL Open
Discussion

2:30 Break

2:45 Reference

Ontology
Discussion

Open

3:45 Darpa Agent Markup Language (DAML) - BAA Jim
Hendler

4:15 Other Ontology Related
Programs Open Discussion

4:45 How to
Proceed
Bob Spillers, et al

5:30 Close

Visitor Information about Stanford

Here are the URLs for visitor information at Stanford.

<http://www.stanford.edu/home/visitors/index.html>

This visitors page has links to directions, information on area hotels, parking and maps, etc.

<http://www.stanford.edu/home/visitors/maps.html>

This page has links to maps of the area, the campus, parking and a searchable map of the campus

http://www.stanford.edu/home/map/search_map.html

This page has a searchable map that will zoom in to the building location. There is a scrollable menu for academic and administrative buildings. Choose the Center for Integrated Systems CIS CISX. There is a parking structure near by.

From ???@??? Fri Nov 12 10:02:47 1999

Received: from mta4.snfc21.pbi.net (mta4.snfc21.pbi.net [206.13.28.142])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id QAA17148
for <phayes@ai.uwf.edu>; Thu, 11 Nov 1999 16:31:30 -0600 (CST)

Received: from pacbell.net ([206.170.6.195])

by mta4.snfc21.pbi.net (Sun Internet Mail Server sims.3.5.1999.09.16.21.57.p8)
with ESMTP id <0FL2002LK0RI28@mta4.snfc21.pbi.net> for phayes@ai.uwf.edu; Thu,
11 Nov 1999 14:21:38 -0800 (PST)

Date: Thu, 11 Nov 1999 14:18:35 -0800

From: Robert Spillers <skydog@pacbell.net>

Subject: [Fwd: FW: Standard Reference Ontology]

To: Andreas Reuter <andreas.reuter@EML.villa-bosch.de>,

Bob Engelmores <rse@ksl.stanford.edu>, Frank Olken <olken@lbl.gov>,

John McCarthy <jmc@cs.stanford.edu>, John Sowa <sowa@west.poly.edu>,

Lee Auspitz <lee@textwise.com>, Livia Polanyi <polanyi@pal.xerox.com>,

Nancy Lawler <E6NL001@coe.coppin.umd.edu>,

Richard Fikes <fikes@ksl.stanford.edu>, Pat Hayes <phayes@ai.uwf.edu>,

Christiane Fellbaum <fellbaum@clarity.Princeton.EDU>,

Chris Menzel <cmenzel@tamu.edu>,

Bill Andersen <andersen@knowledgebus.com>,

Larry Reeker <larry.reeker@nist.gov>,

Tony Sarris <Tony.Sarris@unisys.com>,

Martin van den Berg <vdberg@pal.xerox.com>, "Y.T." <yтчien@nsf.gov>,

Susan Gauch <sgauch@ittc.ukans.edu>,

Deborah McGuinness <dml@ksl.stanford.edu>,

Shelia McIlraith <sam@ksl.stanford.edu>,

Larry Fagan <fagan@SMI.stanford.EDU>,

Mark Musen <Musen@SMI.stanford.EDU>

Message-id: <382B40BB.9F561AA0@pacbell.net>

MIME-version: 1.0

X-Mailer: Mozilla 4.6 [en] (Win98; I)

X-Accept-Language: en

Content-Type: multipart/mixed; boundary="-----637B0B1FE1393B1110B9FF1D"

Content-Length: 105852

Status:

I was asked to testify before the committee appointed by the Secretary of HHS to draft recommendations on how HHS should handle electronic patient records and what recommendations the Secretary should make to Congress. I was unable to attend the hearing but submitted the attached paper instead.

I sent a copy to Rick Morris with CALL. He thought it might be a good starting point for our discussions on Nov. 18.

fyi - Jeff Blair is the chair of the HHS committee.

Bob

Jeff Blair wrote:

> Hello everyone,
> Bob Spillers was unable to present the attached ontological principles to us at our last CPR Workgroup Hearings. However he has submitted a well written description of these principles below. These principles are extremely relevant to the development of a system of medical terminologies which can enable comparable patient medical record information. These principles have been explicitly or implicitly recognized by many of those who have testified to us including Jim Cimino, Keith Campbell, Kent Spackman, Tim McNamara, and Judy Ozbolt. I would like to thank Bob Spillers for sharing this article to us.

>
> Kind regards,
> Jeff Blair

>
> -----Original Message-----
> From: Robert Spillers [SMTP:skydog@pacbell.net]
> Sent: Friday, October 29, 1999 9:59 PM
> To: Jeff Blair; Michael Fitzmaurice; Robert Mayes
> Cc: Henry Heffernan
> Subject: Standard Reference Ontology

>
> Jeff / Michael / Bob,
> I am sorry that I was unable to attend the hearing earlier this month.
> I have attached a document that discusses how ontology might be
> helpful. I hope you find it useful.
>
> Bob

Attachment converted: Poppy:Reference Ontology.doc (WDBN/MSWD) (0000A94A)

From ???@??? Fri Nov 12 10:02:48 1999

Received: from hobbes.poly.edu (hobbes.poly.edu [128.238.1.20])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id SAA09922
for <phayes@ai.uwf.edu>; Thu, 11 Nov 1999 18:14:01 -0600 (CST)

Received: from west (west.poly.edu [128.238.20.21])
by hobbes.poly.edu (8.9.0/8.9.0) with SMTP id TAA01502;
Thu, 11 Nov 1999 19:06:26 -0500 (EST)

Received: by west (SMI-8.6/SMI-SVR4)
id TAA08076; Thu, 11 Nov 1999 19:06:19 -0500

Date: Thu, 11 Nov 1999 19:06:19 -0500

From: sowa@west.poly.edu (John F. Sowa)

Message-Id: <199911120006.TAA08076@west>

To: E6NL001@coe.coppin.umd.edu, Musen@SMI.stanford.EDU,

Tony.Sarris@unisys.com,

andersen@knowledgebus.com, andreas.reuter@EML.villa-bosch.de,
cmenzel@tamu.edu, dlm@ksl.stanford.edu, fagan@SMI.stanford.EDU,
fellbaum@thought.princeton.edu, fikes@ksl.stanford.edu,
jmc@cs.stanford.edu, larry.reeker@nist.gov, lee@textwise.com,

olken@lbl.gov, phayes@ai.uwf.edu, polanyi@pal.xerox.com,
rse@ksl.stanford.edu, sam@ksl.stanford.edu, sgauch@ittc.ukans.edu,
skydog@pacbell.net, vdberg@pal.xerox.com, ytchien@nsf.gov
Subject: Re: [Fwd: FW: Standard Reference Ontology]
Mime-Version: 1.0
Content-Type: text/plain; charset=us-ascii
Content-Length: 165
Status:

Bob,

Is it possible to send files in plain text? Those .doc files just come across as gibberish. In general, it's a good idea to avoid proprietary formats.

John

From ???@??? Fri Nov 12 10:02:49 1999

Received: from mta1.snfc21.pbi.net (mta1.snfc21.pbi.net [206.13.28.122])
by nuts.coginst.uwf.edu (8.8.6/8.8.6) with ESMTP id UAA02822
for <phayes@ai.uwf.edu>; Thu, 11 Nov 1999 20:25:18 -0600 (CST)

Received: from pacbell.net ([206.171.33.103])
by mta1.snfc21.pbi.net (Sun Internet Mail Server sims.3.5.1999.09.16.21.57.p8)
with ESMTP id <0FL200GU5BCCA@mta1.snfc21.pbi.net> for phayes@ai.uwf.edu;

Thu,

11 Nov 1999 18:09:56 -0800 (PST)

Date: Thu, 11 Nov 1999 18:07:07 -0800

From: Robert Spillers <skydog@pacbell.net>

Subject: Re: [Fwd: FW: Standard Reference Ontology]

Cc: "John F. Sowa" <sowa@west.poly.edu>, E6NL001@coe.coppin.umd.edu,
Musen@SMI.stanford.EDU, Tony.Sarris@unisys.com,
andersen@knowledgebus.com, andreas.reuter@EML.villa-bosch.de,
cmenzel@tamu.edu, dlm@ksl.stanford.edu, fagan@SMI.stanford.EDU,
fellbaum@thought.princeton.edu, fikes@ksl.stanford.edu,
jmc@cs.stanford.edu, larry.reeker@nist.gov, lee@textwise.com,
olken@lbl.gov, phayes@ai.uwf.edu, polanyi@pal.xerox.com,
rse@ksl.stanford.edu, sam@ksl.stanford.edu, sgauch@ittc.ukans.edu,
vdberg@pal.xerox.com, ytchien@nsf.gov

Message-id: <382B7649.1B024BE6@pacbell.net>

MIME-version: 1.0

X-Mailer: Mozilla 4.6 [en] (Win98; I)

Content-transfer-encoding: 8BIT

X-Accept-Language: en

References: <199911120006.TAA08076@west>

Content-Type: text/plain

Content-Length: 15326

Status:

"John F. Sowa" wrote:

> Bob,

>

> Is it possible to send files in plain text? Those .doc files just come across
> as gibberish. In general, it's a good idea to avoid proprietary formats.

>

> John

Standard Reference Ontology

Robert Spillers

One of the problems of representing knowledge in a complex domain is that so much knowledge is available. It has been developed over a long period of time for highly specific purposes and only later was thought given to how it might be integrated. Creation of a set of terms and the rules for their expression naturally focused on the specifics of the issues they were developed to address. An unintended consequence of this process is the inability to apply all of the knowledge that is currently available to any particular problem. This is true even when the knowledge is highly relevant. It is also likely that acquisition of the knowledge was costly and laborious. The organizations that funded and created it almost certainly intended it to be used for larger purposes.

A solution to this problem is to integrate existing knowledge and provide a foundation for the organization of new knowledge. Ontology makes this integration possible at both a syntactic and semantic level. All of the knowledge an organization possesses can be used in any of their applications. This is true even if the knowledge is incomplete. The hierarchical organization of ontology, its ability to represent knowledge in multiple dimensions, its foundations in logic and set theory and its combinatorial structure, allow inference engines to reason about the domain and draw conclusions that, in many cases, create new knowledge.

Ontology's roots go back to Aristotle and classical philosophy. However its, modern implementation as a philosophical, logical, mathematical tool in computer science is a cutting edge technology whose significance is just beginning to be recognized. It is roughly equivalent to the early stages of development of the Internet. Most of the pieces exist (in varying degrees of quality) but have not been systematically implemented.

Ontology

Because ontology allows the integration of many forms of knowledge at a semantic level, the results can be examined in many dimensions both for an individual instance and comparatively. An analysis could review financial costs, resource distributions, outcomes, the application of ethical standards or any other dimension where the data are gathered.

This does require the domain to be modeled and an ontology created. It will also require participation by subject matter experts (SME), experts in the existing terminology structures and experts in ontology. It enables all forms of available knowledge to be used in integrated applications that provide semantic understanding and the ability to be transparently used in inference engines.

In order to allow different domains to share knowledge, a standard upper level Reference Ontology (RO) must be created. The Reference Ontology (RO) contains the most general categorization of concepts used by all domains. Ontology is not a replacement for terminology or for the various forms of meta data. (Angelo Rossi-Mori has an excellent discussion of many of these issues in the section on third generation formal systems in his 1997 paper) Terminology can be mapped to concepts and the extent of equivalence (term to term, term to code, etc.) can be determined. This is a robust information model that provides the structure to enable all forms of available knowledge to be used in integrated applications with semantic understanding. It also is a road map for the restructuring of terms, elimination of unneeded or confusing terms, the creation of new terms, and (at least in controlled vocabularies) the permitted range of usage. At a minimum, ontology provides a model centric mapping that is 1 to n rather than n to n.

Architecture

The architecture for ontology should include at least three levels.

1. An upper level Reference Ontology (RO) - a mathematically rigorous and comprehensive theory of ontology construction, structure and interpretation. It is where the most general concepts reside. These are the overarching concepts that are incorporated in all other (lower level) ontologies. A Reference Ontology (RO) provides the inter-model fulcrum that allows systems using different sets of terminology to share fundamental information. Most of the concepts at this level are philosophical such as abstract/concrete, universal/particular, continuant/occurrent, independent/dependent, identity criteria, context, granularity, meretopology, etc. It is also the place where most of the lexical concepts of commonly used English words one finds in a well-written magazine will be connected to the more general philosophical concepts.

Each term in the RO should be accompanied by several fields of information including taxonomic relations, an English description, relations linking synonyms and antonyms etc., axioms, and (partial) definitions. Each statement in these (partial) definitions should be expressed in an unambiguous notation that is easy to transform into any of the AI and Database notations in current use, in so far as they (the AI / DB notations) have the required expressive power.

An application programmers, interface (API) must be created that will support both standard functionality and new functionality required by future applications. Information access and update must be provided not only for first-order knowledge, but also for higher-order, modal, non-monotonic, probabilistic and fuzzy inferences. The API should be able to support efficient scheduling for distributed querying, indirect queries launched through a superordinate ontology (the RO), or through subordinate (but conformant) domain ontologies. Both the syntax and semantics of the API should rest on fully developed, mathematically rigorous foundations.

2. Industry level domain ontologies (e.g. medical, finance, insurance, etc.) ^ domain ontologies that are constructed to conform to the theory and standards of the RO and its, API. They are comprehensive in their representation of an industry domain (medical) or sub-domain (internal medicine, cardiology, lab tests, etc.) In contrast to the upper level RO, construction of a domain ontology requires significant input from subject matter experts (SMEs ^ e.g. a cardiologist, a lab technician, etc.) plus experts in the domain terminology and experts in ontology.

All domain (or sub-domain) ontologies that are designed and constructed using this architecture will be able to align, integrate and inter-operate with both the upper level RO and any other conforming ontology. Concepts ^ particularly complex concepts like time, money, etc.- can be imported from any other conforming ontology and used without modification. If modification is desired, one can begin with the imported concept and change it to suit one,s needs. If these changes are made in conformance with the design standards of the RO, it can be used with any conforming ontology.

3. Application level ontologies conform to the RO and domain ontologies, but have been further extended to meet the needs of specific application requirements. Application developers will be able to leverage the very considerable investments made in the construction of higher level ontologies and concentrate on a much more narrow area where they (hopefully) have unique expertise. Without the RO and industry ontologies, it is extremely unlikely that an organization or company ^ particularly a small one- could muster either the resources or the expertise to compete.

Integration of Structured and Unstructured Data

Ontology offers a general structure of knowledge that is universal ^ it is a conceptualization that places the resulting concepts in the proper conceptual structure. This structure may include many things, but it is primarily these concepts and their relations. Although the context (i.e. a natural language such as English, a relational database, various file formats, coded data, etc.) in which the concepts are used is very important, their relations are independent of the context. This makes possible applications that can determine semantics within the context. It also allows the creation of applications that can truly integrate knowledge from many different types and sources while accurately maintaining their semantics.

An application taking advantage of ontology would know the semantic relationship among concepts used in everyday English, the schema of relational databases, professional terminology, coded data, etc. for every concept in the ontology. Applications that require this level of integration would now be possible. However, it is still true that the utility of the application will be dependent on the quality of its own design, software engineering, etc.

An application would be able to integrate information from:

- o Physician,s notes regarding the examination of a patient that used everyday English terms, professional terms or standardized terms from an accepted list of terminology.
- o Clinical data (in various codes and formats) from laboratory tests
- o The language, terms and coding used to describe a diagnosis ^ perhaps on an insurance form
- o Treatment prescribed that may include terms for medication (including its timing and use), diets, medical procedures, etc.
- o Information regarding the non-medical assistance offered patients such as completion of insurance forms, transportation, referral to organizations for psychological or financial support (e.g. Cancer societies, Red Cross, AA, state or local government assistance programs), long term or home care resources.
- o Financial /insurance information regarding all of the above

Applications that integrate at this level will require both an upper level Reference Ontology and domain ontology for each domain that participates in the application. Modeling a domain and creating a domain ontology will require subject matter experts (SMEs), experts in the existing terminology structures, and experts in ontology.

Natural Language Processing (NLP)

All natural languages permit considerable ambiguity in both their syntax and

semantics. Much of the meaning of words is determined from context ^ identical character strings can have radically different meanings that can only be determined by context. Even in formal languages there are slight differences, for example, in the way mathematicians, logicians and computer scientists use terms relating to subjects common to their fields (e.g. set theory). Many NLP tools have been created over the years to assist people and computer programs to understand some part of these issues. These tools include dictionaries, lexicons, thesauri, part of speech taggers, parsers, etc. They have been created both academically and commercially using widely disparate theories and technologies. Although no single tool (in fact no combination of tools) completely solves these problems, significant progress has been made in recent years.

Many of the applications that use NLP are, in effect, some form of information retrieval. There has been a strong push to make these tools effective in internet applications, particularly internet search engines ^ with varying degrees of success. Ontology is being used in leading edge Knowledge Management (KM) applications that attempt a more sophisticated form of retrieval ^ conceptual retrieval. Conceptual retrieval returns documents that are about the concept and not merely documents that mention the concept ^ or its, synonyms and morphological extensions.

Conceptual retrieval marries ontology with advanced statistical methods that identify concepts within documents. These statistical methods vary in complexity and results from rather simple vector analysis to advanced implementations of Hidden Markov Models (HMM). All implementations require „truth training%“. Truth training is the identification (by humans) of exemplars - documents that are about a specific concept (i.e. truth) ^ normally 10+ examples for each concept. The system examines the documents identified as truth (for a concept) then builds a statistical model of a true document. This is sometimes called topic spotting.

If one trains a topic spotter on the concepts that are the nodes of the domain ontology, then the system retrieves documents that are about the concept ^ not just mentions the concept. It is not uncommon for this type of a retrieval engine to return documents that are truly about a concept but never actually mentions the concept or its, synonyms etc. In large collections, there may still be too many documents retrieved. Because ontology can be n dimensional, many other factors may be expressed in the ontology that may be used as filters to return the most useful document(s).

Because statistical engines use tokens (instead of words) as the basis of their analysis, everyday English, professional terminology, controlled vocabularies or alphanumeric codes receive similar analysis. Documents containing concepts expressed in any of these forms may be modeled and retrieved. This is a particularly useful combination of statistical and symbolic processing.

Using this technology, an application could analyze documents ^ e.g. an insurance

form or some element of a patient record ^ that contained formatted spaces for everyday English, lab data, professional terminology and insurance codes, etc. The application could perform the same analysis on an unstructured document ^ the formatting actually helps the analysis. A document (or a collection) can be reviewed in many dimensions both for individual instances and comparatively across geographic regions, hospitals, various service organizations, etc. The analysis could review financial costs, resource distributions, outcomes, the application of ethical standards or any dimension where the data are gathered.

Ontology addresses the significant issues of

- o Integration of structured and unstructured data
- o Mapping where there is no parity or exact semantic equivalence among data elements and terms due to definitions, granularity, usage context, etc. Ontology highlights deficiencies and provides a roadmap for corrections.
- o An inter-model fulcrum to serve as a robust information model that enables the understanding of semantic relationships
- o The ability to use semantic understanding of concepts, relations, terms and structured data in inference engines.
- o The use of inference engines to analyze in n dimensions (costs, quality, ethical standards, etc.) for any area where data are gathered

Ontology has been used in a wide variety of classified applications at DoD. The aircraft industry has a number of applications involving both design and manufacturing. Ontology is becoming an important enabling technology for knowledge management particularly in customer care applications in the financial services and high tech manufacturing industries. It will become the competitive edge for many types of retrieval engines and the key technology for enterprise application integration.

Ontology lowers the barriers to market entry, especially for small companies. More open markets lead to lower costs for all consumers.

From ???@??? Wed Jan 17 10:43:45 2001
Return-Path: <graham.horn@aihw.gov.au>
Received: from mace.isecure.com.au (203.32.86.20) by mail.coginst.uwf.edu with ESMTP (Eudora Internet Mail Server 3.0.2) for <phayes@ai.uwf.edu>;
Wed, 17 Jan 2001 02:42:35 -0600
Received: from shield.isecure.com.au (shield.isecure.com.au [203.32.86.50])
by mace.isecure.com.au (8.10.2/8.10.2) with ESMTP id f0H8flj16082;
Wed, 17 Jan 2001 19:41:18 +1100
Received: (from uucp@localhost)
by shield.isecure.com.au (8.10.0.Beta12/8.10.0.Beta12) id f0H8fHs23230;
Wed, 17 Jan 2001 19:41:18 +1100 (EST)

Received: from nodnsquery(10.11.3.10) by shield.isecure.com.au via smap (V5.0)
id xma023215; Wed, 17 Jan 01 19:41:11 +1100

Received: from aihwgw.isecure.com.au (localhost [127.0.0.1])
by lance.isecure.com.au (8.10.2/8.10.2) with ESMTP id f0H8f6K23353;
Wed, 17 Jan 2001 19:41:06 +1100

Received: from master.AIHW ([203.2.120.232])
by aihwgw.isecure.com.au (8.10.2/8.10.2) with ESMTP id f0H8f2p18484;
Wed, 17 Jan 2001 19:41:02 +1100 (EST)

Received: by master.aihw.gov.au with Internet Mail Service (5.5.2650.21)
id <CB8SNT4Y>; Wed, 17 Jan 2001 19:37:30 +1100

Message-ID:

<EC943C5D0768D311AF8A00805F85B83EF82F64@master.aihw.gov.au>

From: "Horn, Graham" <graham.horn@aihw.gov.au>

To: "Frank Farance" <frank@farance.com>,
"Schoening CECOM DCSC4I James (E-mail)"
<James.Schoening@mail1.monmouth.army.mil>

Cc: "sofia@gia.ist.utl.pt" <sofia@gia.ist.utl.pt>,
"RASCHR@LEAVENWORTH.ARMY.MIL"

<RASCHR@LEAVENWORTH.ARMY.MIL>,
"Donald@prim.is" <Donald@prim.is>,
"Chris_Partridge@csi.com"
<Chris_Partridge@csi.com>,
"Peter.Eirich@jhuapl.edu"
<Peter.Eirich@jhuapl.edu>,
"apease@teknowledge.com"
<apease@teknowledge.com>,
"Gillman_D@bls.gov" <Gillman_D@bls.gov>,
"mfu@redwood.rt.cs.boeing.com" <mfu@redwood.rt.cs.boeing.com>,
"frank@farance.com" <frank@farance.com>,
"reKent@ontologos.org"
<reKent@ontologos.org>,
"cassidy@micra.com" <cassidy@micra.com>,
"Nicola.Guarino@ladseb.pd.cnr.it" <Nicola.Guarino@ladseb.pd.cnr.it>,
"ksundaram@vsnl.com" <ksundaram@vsnl.com>,
"martin_gladwell@UK.IBM.COM" <martin_gladwell@UK.IBM.COM>,
"dmckay@VERTICALNET.COM" <dmckay@VERTICALNET.COM>,
"lobrst@VERTICALNET.COM" <lobrst@VERTICALNET.COM>,
"btowle@knowledgeplanet.com" <btowle@knowledgeplanet.com>,
"kinshuk@mailandnews.com" <kinshuk@mailandnews.com>,
"finin@cs.umbc.edu" <finin@cs.umbc.edu>,
"skydog@pacbell.net"
<skydog@pacbell.net>,
"RogerB@vp.net" <RogerB@vp.net>,
"sgefelix@techst02.technion.ac.il" <sgefelix@techst02.technion.ac.il>,
"fikes@KSL.Stanford.EDU" <fikes@KSL.Stanford.EDU>,
"jerinic@uns.ns.ac.yu" <jerinic@uns.ns.ac.yu>,

"sambrosz@ipipan.waw.pl" <sambrosz@ipipan.waw.pl>,
"andersen@knowledgebus.com" <andersen@knowledgebus.com>,
"MannD@BATTELLE.ORG" <MannD@BATTELLE.ORG>,
"gangemi@saussure.irmkant.rm.cnr.it"
 <gangemi@saussure.irmkant.rm.cnr.it>,
"raross@rosscomp.com"
 <raross@rosscomp.com>,
"Steven.van.t.Veld@AIM.NL"
 <Steven.van.t.Veld@AIM.NL>,
"fischer@DARMSTADT.GMD.DE"
 <fischer@DARMSTADT.GMD.DE>,
"ocorcho@DELICIAS.DIA.FI.UPM.ES"
 <ocorcho@DELICIAS.DIA.FI.UPM.ES>,
"asun@FI.UPM.ES" <asun@FI.UPM.ES>,
"masuoka@FLAB.FUJITSU.CO.JP" <masuoka@FLAB.FUJITSU.CO.JP>,
"superb@FLAB.FUJITSU.CO.JP" <superb@FLAB.FUJITSU.CO.JP>,
"weltyc@HUB.CS.VASSAR.EDU" <weltyc@HUB.CS.VASSAR.EDU>,
"tmk@LSC.CO.UK" <tmk@LSC.CO.UK>,
"slussier@SIMS.BERKELEY.EDU"
 <slussier@SIMS.BERKELEY.EDU>,
"judy@STAFF.CS.USYD.EDU.AU"
 <judy@STAFF.CS.USYD.EDU.AU>,
"alozano@UNEX.ES" <alozano@UNEX.ES>,
"cschlenoff@VERTICALNET.COM" <cschlenoff@VERTICALNET.COM>,
"Thomas_A._Kalil@oa.eop.gov" <Thomas_A._Kalil@oa.eop.gov>,
"Matthew.R.West@is.shell.com" <Matthew.R.West@is.shell.com>,
"gamma@clark.net" <gamma@clark.net>,
"eric.jones@alphatech.com"
 <eric.jones@alphatech.com>,
"Dawidowicz, Edward CECOM RDEC C2D"
 <Edward.Dawidowicz@mail1.monmouth.army.mil>,
"epeterson@verticalnet.com" <epeterson@VERTICALNET.COM>,
"hans@isi.edu" <hans@isi.edu>,
"sluce@teknowledge.com"
 <sluce@teknowledge.com>,
"jtonkel@earthlink.net"
 <jtonkel@earthlink.net>,
"wmurray@teknowledge.com"
 <wmurray@teknowledge.com>,
"lan_J_Dickinson@hplb.hpl.hp.com"
 <lan_J_Dickinson@hplb.hpl.hp.com>,
"cmenzel@philebus.tamu.edu"
 <cmenzel@philebus.tamu.edu>,
"sowa@west.poly.edu" <sowa@west.poly.edu>,
"t.by@dcs.shef.ac.uk" <t.by@dcs.shef.ac.uk>,
"elst@dfki.uni-kl.de"

<elst@dfki.uni-kl.de>,
"James.Fulton@PSS.Boeing.com"
<James.Fulton@PSS.Boeing.com>,
"John.Thompson@PSS.Boeing.com"
<John.Thompson@PSS.Boeing.com>,
"mhonarc@grouper.ieee.org"
<mhonarc@grouper.ieee.org>,
"cmenzel@KBSI.com" <cmenzel@KBSI.com>,
"rschulz@teknowledge.com" <rschulz@teknowledge.com>,
"rrschulz@cris.com" <rrschulz@cris.com>, "mp@dfki.de" <mp@dfki.de>,
"lee@textwise.com" <lee@textwise.com>,
"naso@sirma.bg"
<naso@sirma.bg>,
"jfowler@pdsolutions.co.uk"
<jfowler@pdsolutions.co.uk>,
"mcdavid@us.ibm.com"
<mcdavid@us.ibm.com>,
"miz@ei.sanken.osaka-u.ac.jp"
<miz@ei.sanken.osaka-u.ac.jp>,
"fritz@cyc.com" <fritz@cyc.com>,
"mudcat@eil.utoronto.ca" <mudcat@eil.utoronto.ca>,
"auszok@kul.lublin.pl" <auszok@kul.lublin.pl>,
"kozaki@ei.sanken.osaka-u.ac.jp" <kozaki@ei.sanken.osaka-u.ac.jp>,
"a8626146@unet.univie.ac.at" <a8626146@unet.univie.ac.at>,
"thep@ei.sanken.osaka-u.ac.jp" <thep@ei.sanken.osaka-u.ac.jp>,
"jhalcomb@teknowledge.com" <jhalcomb@teknowledge.com>,
"dlm@KSL.Stanford.EDU" <dlm@KSL.Stanford.EDU>,
"PCassidy@verticalnet.com" <PCassidy@VERTICALNET.COM>,
"WhkIV@aol.com" <WhkIV@aol.com>,
"nivezic@nist.gov"
<nivezic@nist.gov>,
"pcassidy@vertical.net" <pcassidy@vertical.net>,
"cpeylo@ridcully.cl-ki.uni-osnabrueck.de"
<cpeylo@ridcully.cl-ki.uni-osnabrueck.de>,
"Carpenter_J@bls.gov"
<Carpenter_J@bls.gov>,
"David.Bruce-Steer@standards.com.au"
<David.Bruce-Steer@standards.com.au>,
"Horn, Graham"
<graham.horn@aihw.gov.au>,
"andreas.eberhart@i-u.de"
<andreas.eberhart@i-u.de>,
"kathy@sinic.kiev.ua" <kathy@sinic.kiev.ua>,
"kivs@bgcict.acad.bg" <kivs@bgcict.acad.bg>,
"lkagal1@cs.umbc.edu"
<lkagal1@cs.umbc.edu>,

"Chris.Angus@btinternet.com"
<Chris.Angus@btinternet.com>,
"WBurkett@pdit.com" <WBurkett@pdit.com>,
"JDickert@DTIC.MIL" <JDickert@DTIC.MIL>,
"heretic@ihug.co.nz"
<heretic@ihug.co.nz>,
"sar@mitre.org" <sar@mitre.org>,
"jli@teknowledge.com" <jli@teknowledge.com>,
"henry_heffernan@nih.gov" <henry_heffernan@nih.gov>,
"fmanola@mitre.org" <fmanola@mitre.org>,
"dick@canlinks.com"
<dick@canlinks.com>,
"sowa@bestweb.net" <sowa@bestweb.net>, "fnc@mitre.org" <fnc@mitre.org>,
"jdpark@etri.re.kr"
<jdpark@etri.re.kr>,
"thierry.bouron@rd.francetelecom.fr"
<thierry.bouron@rd.francetelecom.fr>,
"werner@landc.be"
<werner@landc.be>,
"P.Kocura@lboro.ac.uk" <P.Kocura@lboro.ac.uk>,
"sbj2@columbia.edu" <sbj2@columbia.edu>,
"fgm@fla.fujitsu.com"
<fgm@fla.fujitsu.com>,
"phil.jackson@computer.org"
<phil.jackson@computer.org>,
"vinay.chaudhri@sri.com"
<vinay.chaudhri@sri.com>,
"venkata.rebbapragada@mail1.monmouth.army.mil"
<venkata.rebbapragada@mail1.monmouth.army.mil>,
"DHeinze@alifemedical.com" <DHeinze@alifemedical.com>,
"zarri@ivry.cnrs.fr" <zarri@ivry.cnrs.fr>,
"baonysh@afterlife.ncsc.mil" <baonysh@afterlife.ncsc.mil>,
"John.Velman@HSC.com" <John.Velman@HSC.com>,
"stborgo@indiana.edu"
<stborgo@indiana.edu>,
"david@itfocus.demon.co.uk"
<david@itfocus.demon.co.uk>,
"mattj@newsblip.com" <mattj@newsblip.com>,
"lindner@igd.fhg.de" <lindner@igd.fhg.de>,
"freeman@pobox.com"
<freeman@pobox.com>, "leo@mmk.ru" <leo@mmk.ru>,
"ikeda@ei.sanken.osaka-u.ac.jp" <ikeda@ei.sanken.osaka-u.ac.jp>,
"xjwang@bupt.edu.cn" <xjwang@bupt.edu.cn>,
"mary.nunn@digitaldoor.net" <mary.nunn@digitaldoor.net>,
"naka@nttx.co.jp" <naka@nttx.co.jp>,
"tinwisle@bloomington.in.us"

<tinwisle@bloomington.in.us>,
"Rebbapragada, Lakshmi V CECOM RDEC C2D"
<Lakshmi.Rebbapragada@mail1.monmouth.army.mil>,
"bordogna@tin.it"
<bordogna@tin.it>,
"V.A.M.Tamma@csc.liv.ac.uk"
<V.A.M.Tamma@csc.liv.ac.uk>,
"iniles@teknowledge.com"
<iniles@teknowledge.com>,
"juanal@ns.ucb.edu.bo"
<juanal@ns.ucb.edu.bo>,
"whitten@lynx.eaze.net"
<whitten@lynx.eaze.net>,
"hauber.lists@aon.at" <hauber.lists@aon.at>,
"mxm@y12.doe.gov" <mxm@y12.doe.gov>,
"nc-engelre@netcologne.de"
<nc-engelre@netcologne.de>,
"Ronald_Krijgsheld@nl.compuware.com"
<Ronald_Krijgsheld@nl.compuware.com>,
"eneko@si.ehu.es"
<eneko@si.ehu.es>,
"MJaeger@lansing.cc.mi.us"
<MJaeger@lansing.cc.mi.us>,
"pm@uma.at" <pm@uma.at>,
"Thomas.A.Early@trw.com" <Thomas.A.Early@trw.com>,
"ms@ms.lt"
<ms@ms.lt>,
"martin_king@uk.ibm.com" <martin_king@UK.IBM.COM>,
"jsachs@cesdis40.gsfc.nasa.gov" <jsachs@cesdis40.gsfc.nasa.gov>,
"ronheld@xs4all.nl" <ronheld@xs4all.nl>,
"JerzyP@Bot.net"
<JerzyP@Bot.net>,
"jackpark@verticalnet.com"
<jackpark@VERTICALNET.COM>,
"Brian.Robinson@trw.com"
<Brian.Robinson@trw.com>,
"isanz@vub.ac.be" <isanz@vub.ac.be>,
"shams@pnu.ac.ir" <shams@pnu.ac.ir>,
"josef@ix.netcom.com"
<josef@ix.netcom.com>,
"dmiles@teknowledge.com"
<dmiles@teknowledge.com>,
"iciki@people.flab.fujitsu.co.jp"
<iciki@people.flab.fujitsu.co.jp>,
"pipherj@ncr.disa.mil"
<pipherj@ncr.disa.mil>,

"piph@erols.com" <piph@erols.com>,
"jim@spatial.maine.edu" <jim@spatial.maine.edu>,
"fuchs@ifi.unizh.ch"
 <fuchs@ifi.unizh.ch>,
"fred@comprecorp.com" <fred@comprecorp.com>,
"cargan@delrio.com" <cargan@delrio.com>,
"deanmel@microsoft.com"
 <deanmel@microsoft.com>,
"mask@ktl.mii.lt" <mask@ktl.mii.lt>,
"phayes@ai.uwf.edu" <phayes@ai.uwf.edu>,
"alexr@pcnet.ro"
 <alexr@pcnet.ro>,
"EJohnson@LGA-INC.COM" <EJohnson@LGA-INC.COM>,
"Robert.Meersman@vub.ac.be" <Robert.Meersman@vub.ac.be>

Subject: RE: Re: Revised definition of Base Document

Date: Wed, 17 Jan 2001 19:37:26 +1100

Return-Receipt-To: "Horn, Graham" <graham.horn@aihw.gov.au>

MIME-Version: 1.0

X-Mailer: Internet Mail Service (5.5.2650.21)

Content-Type: text/plain

Frank,

. I strongly support your approach. It is constructive, and aimed at achieving something worthwhile, yet does not require taking a non-standard (no pun intended) approach.

. If the overall vote is and/or remains NO, then I can see the following options:

1. await someone developing another ontology to overcome the shortcomings resulting in this one's rejection;
2. centre on one of the existing ontologies already published, with the objective of making it better apply to our purposes (eg. possibly make it more general in scope); or
3. abandon the whole project.

. Option #1 necessarily requires considerably more groundwork, in view of the reasons given for rejecting this one. Would anyone like to volunteer?

. The only other option of these with merit is #2. This was proposed some time ago (I supported it), but was roundly rejected.

. Of course, anyone supporting option #3 is free to leave the group. I hope we are not being sabotaged by guerillas or participants who have their noses out of joint over some issue or other.

Graham Horn
National Data Standards Unit
Australian Institute of Health and Welfare

=====
Phone: 02.6244.1094
Fax: 02.6244.1199
E-mail: Graham.Horn@aihw.gov.au <mailto:graham.horn@aihw.gov.au>

-----Original Message-----

From: Frank Farance [mailto:frank@farance.com]
Sent: Wednesday, January 17, 2001 3:02 PM
To: standard-upper-ontology@ieee.org
Subject: SUO: Re: Revised definition of Base Document

At 21:32 2001-01-16 -0500, Schoening, James R CECOM DCSC4I wrote:

> SUO Voters,

>

> The definition posted for Base Document appears to be the cause of
> some confusion. Since IEEE does not provide a definition, we are free to
> define it ourselves.

>

> How does this look?

In ISO, NCITS, IEEE, IETF, and elsewhere (Open Group, CommerceNet), the term "base document" is used consistently ... this definition is consistent with our original definition of a "base document". Here's some URLs (search for "base document" once the page is opened):

CommerceNet (ePay)

<http://lists.commerce.net/archives/ansi-epay/199904/msg00007.html>

Network News Transport Protocol (NNTP)

<http://www.academ.com/academ/nntp/ietf/1367.html>

Annotated C++ Reference Manual, ANSI Base Document

<http://www.softpro.com/0-201-51459-1.html>

VRML (Virtual Reality Modeling Language)

<http://www.vrml.org/Specifications/VRML97/part1/concepts.html>

COBOL base documents

<http://www.merant.com/Standards/x3j4m198.htm>

LTSC Documents
<http://ltsc.ieee.org>

- > Base Document: Document a standards working group has agreed to work on.
- > Viewed at having potential for successful completion. Has no official
- > weight (not a standard). Might not be in standards wording. Might contain
- > only small fraction of planned content and scope. Might be multiple base
- > documents. Most base documents represent substantial work created outside
- > committee and introduced into the WG. Changes to base document require
- > approval of group.

The above definition describes a Working Draft (a base document is usually Working Draft 1). I don't recommend that we use the above definition for Base Document.

Once a Base Document (in its original definition) is revised, it is no longer a Base Document, Revised Base Document, etc. -- it is now Working Draft 2.

A Base Document, as stated below, carries no official weight. Just like a Working Draft: it carries no official weight. The only implication for Working Drafts (which include Base Documents) is that the Working Group has agreed to work on improving the document towards developing a standard. Changes to Working Drafts (which include Base Documents) are under committee control, i.e., by committee vote.

In many cases, the final standard is substantially different than the Base Document: the original document might only be 15-20% of the final standard, the document may be completely reorganized.

In many other standards activities, when a group is starting and many people aren't familiar with the standards process, they get a bit worried that things are moving too quickly. In virtually all of these committees, everyone got caught up on the term "base document" because they thought (like participants in this Working Group) that it carried some official weight (it doesn't).

Later on, as more standards were developed in these other standards activities, virtually everyone felt comfortable with "base documents" because they recognized that (1) the work must begin at some point, and (2) the standards process allowed everyone to get all their comments incorporated ==> not much to worry about on "base documents" (that is, assuming one is working constructively and positively toward developing a standard).

One reason for designating a working paper as a Base Document is that it is identified as a starting point, i.e., one does not need to provide Rationale (informative standards wording) for words prior to a Base Document. Right now we don't need to worry too much about Rationale, but later on in the standards development we might consider publishing a Rationale as an Annex or a separate document. Since we have a fine E-mail reflector to record every comment we make, we don't really need to concern ourselves now ... but a year or two from now, we will need to understand (roughly) where the starting point(s) began.

I'm suggesting we keep the original definition:

> The prior definition was:

>

> base document: Starting point for standards words (it's clear the document
> *can* be transformed into standards words). Has no official weight (not a
> standard). Might not be in standards wording. Might not be complete with
> respect to the scope and/or final standard. Might be multiple base
> documents. Automatically allocated agenda time since this is the main
> product of the WG. Most base documents represent substantial work created
> outside the committee and introduced into the WG.

-FF

Frank Farance, Farance Inc. T: +1 212 486 4700 F: +1 212 759 1605
mailto:frank@farance.com http://farance.com
Standards, products, services for the Global Information Infrastructure