At the core of Core Ontologies

Preface to the Proceedings of the Workshop on "Core Ontologies in Ontology Engineering"

Research on ontological engineering has covered a fairly wide spectrum of types of ontologies. On the one extreme we find rich formal systems called foundational ontologies whose aim is to characterize explicitly and uniquely a viewpoint on the reality in the large. On the other side of the spectrum there are terminological ontologies (like glossaries and taxonomies) whose aim is to find a shared classification of terms used in some domain.

In assessing ontologies, it is quite easy to state the advantages and drawbacks of terminological systems. However, in considering more complex systems like foundational ontologies, it seems that their complexity makes hard to establish clear and shared evaluation judgments. Perhaps, a way to overcome this problem is to tackle some (perhaps the most) important objects that are based on these ontologies. That is, one can try to assess foundational ontologies through the study of the core ontologies based on them.

Core ontologies are ontologies that focus on a domain application without being restricted to specific applications. Our interest goes to the core ontologies that are built in agreement with foundational ontologies or that are based on general principles and well-founded methodologies. But, looking at the literature, we find a major problem from the start: just a few core ontologies have been actually developed in this way. This lack of examples might be related to the fact that foundational ontologies have been developed only recently. Nonetheless, we believe it is also a sign of a major gap between theoretical (here read 'foundational') and applicative ('taxonomy' driven) research in ontology. In the first case researchers look at principled distinctions and general notions only, in the latter case they concentrate too often on answering immediate and specific needs with ad hoc classifications. This being the situation, core ontologies find no space on either side.

Going back to the spectrum of ontologies, one should notice that content is at the center of the work on one side (foundational ontologies) while tools and languages take the spotlight in the other case (terminological ontologies). As a matter of fact, languages and tools enable the specification of ontologies, help in verifying their correctness, and endorse their suitability for information exchange. However, in our view, the choice for adopting an ontology should be based on its content and the possibility to share it unambiguously.

Content itself can be seen from different perspectives. For some, it is bound to context and, as a consequent, any effort to define widely reusable ontologies is mostly useless: agreement within communities of practice is the most one can obtain. Others hold that widely reusable ontologies can be built by looking at linguistic usage, standardization initiatives, or domain-independent principles isolated in specialized disciplines like philosophy and cognitive science.

We believe the study of core ontologies can help us even in dealing with these viewpoints. In an application domain, ontological research is supposed to provide a representation of the elements and relations specific to it, i.e., a core ontology for that domain. In doing this, it does not matter which position one takes in the ontology spectrum or which notion of content he embraces. At the end of the day, an ontological system has to be provides and it has to catch (formally or informally, uniquely or loosely) the central notions and relations of the domain at stake. These are the ontologies that will make interoperability and knowledge exchange concrete in applications; a good grasp of their potentialities (from either standpoint) would allow us to start delivering what the research community has been promising for a decade or so: reliable knowledge sharing, manipulation, and retrieving.

The usefulness of core ontologies is recognized even beyond the two major motivations mentioned above. But things move slowly in this area and we believe that there is a need of initiatives (like dedicated conferences and workshops) to bring wider attention to these objects and their study. Research on this topic would shred some light on other issues as well like the lack of cross fertilization among different areas where ontology is applied and the lack of agreed methodologies in building ontologies. Not to forget, it would overcome the (to us surprising) shortage of examples of deployment of core ontologies in real applications.

But what precisely are core ontologies asked to do? Here the answer may vary and we provide a partial list only:

a) They register an agreement on the types of entities (and their relationships) needed in a community of practice.
b) They support dynamic negotiation of the intended meaning across a distributed community.
c) They align (integrate, merge) several sources of metadata and terminologies.
d) They provide the backbone for multiple applications or services.
e) They furnish a template for specifying the content in some domain.

These considerations give some important motivations for the workshop on *Core Ontologies in Ontology Engineering* held at the Whittlebury Hall, Northamptonshire (UK) on October 8, 2004 in conjunction with the conference EKAW

2004. The (sub)title of the workshop makes clear that our interest goes to the broad picture:

> *(Un)Successful cases and best practices for ontology engineering: reusing well-founded ontologies for domain content specification.*

The workshop features a collection of six papers which span different perspectives. The atmosphere at the meeting was friendly and open, plenty of interesting questions brought out the commonalities across different areas of application. At the end, we have been asked to collect the the presentation slides as well. Our thanks to the authors for promptly providing this material as well.

During the discussion some themes have emerged as central issues for the development of core ontologies. Among these, we find

a) distinction and relationship between core ontology, domain ontology, and foundational ontology (are all these necessary? where are the borderlines?);
b) the nature of constraints, which leads to the distinction between ontological and contextual conditions;
c) levels of ontological commitment;
d) the embedding of basic distinctions (like three vs four dimensional entities);
e) the advantages and drawbacks in modeling properties as classes of classes and as values in quality spaces;
f) modeling distinctions (e.g. disease as process, object, quality, and perhaps situation).
g) the representation of situations (reality, conceptualization and knowledge levels);
h) the problem of integrating heterogeneous data models;
i) the advantages and drawbacks in mixing epistemology and ontology at the level of core ontologies;
j) what the evolutionary approach can offer to ontology;
k) how to include psychological evidence into ontology.

Some notions arise in a different perspective while developing core ontologies, others need to be clarified for this area to progress. Here are some examples:

a) functional object;
b) complex entities vs composite classes;
c) granularity;
d) discrete and mass terms;
e) location, parthood and containment;
f) duality (systematic polysemy).

We take the opportunity to thank all the people that submitted papers to the workshop, the organization of the EKAW 2004 conference for hosting our meeting, and the people that were part of the workshop program committee, namely, Hans Akkermans (VUA, Amsterdam), John Bateman (U. of Hamburg),

Brandon Bennett (U. of Leeds) ,Joost Breuker (UVA, Amsterdam), Oscar Corcho (UPM Madrid and ISOCO, Madrid), Michael Gruninger (U. of Maryland and NIST), Enrico Motta (KMI, Open U., Milton Keynes), Natasha Noy (Stanford U.), Sofia Pinto (Technical U. of Lisbon), Alan Rector (U. of Manchester), Guus Schreiber (VUA, Amsterdam), Steffen Staab (U. of Karlsruhe), Laure Vieu (IRIT, Toulouse and ISTC-CNR, Trento), Krisztof Wecel (U. of Poznan), Chris Welty (Watson RC, IBM).

Stefano Borgo and Aldo Gangemi
LOA - Laboratory for Applied Ontology
Institute of Cognitive Sciences and Technology (CNR), Italy

October 2004