Available online at www.sciencedirect.com
ScienceDirect
Int. J. Human-Computer Studies 65 (2007) 563–568
International Journal of Human-Computer Studies
www.elsevier.com/locate/ijhcs

Editorial

Knowledge representation with ontologies: Present challenges—Future possibilities

Christopher Brewster[a,*], Kieron O'Hara[b]

[a]Department of Computer Science, University of Sheffield, S1 4DP, UK
[b]School of Electronics and Computer Science, University of Southampton, SO17 1BJ, UK

Available online 18 April 2007

Abstract

Ontologies have become the knowledge representation medium of choice in recent years for a range of computer science specialities including the Semantic Web, Agents, and Bio-informatics. There has been a great deal of research and development in this area combined with hype and reaction. This special issue is concerned with the limitations of ontologies and how these can be addressed, together with a consideration of how we can circumvent or go beyond these constraints. The introduction places the discussion in context and presents the papers included in this issue.

Keywords: Ontologies; Knowledge representation; Semantic web; Artificial intelligence; Bio-informatics; Agents

1. Knowledge representation with ontologies

There is a long tradition in computer science and artificial intelligence which equates knowledge with facts. This view draws on a philosophical and scientific tradition going back at least as far as Aristotle and certainly passing through the renaissance and enlightenment efforts to systematise knowledge ranging from Giordano Bruno's Art of Memory to Linaeus's classification schema for plants and animals. The early efforts at machine translation were famously criticised for their lack of 'world knowledge.' From this (in part) were born efforts to 'represent knowledge' where knowledge was taken to mean knowledge of the world. Since knowledge equated with facts, knowledge representation was largely seen as the task of managing collections of facts about the world.

It is within this tradition that we need to place the contemporary explosion of interest in ontologies, and to which this special issue is a contribution. Our intention has been to provide an opportunity to go beyond the hype, to address core issues in the technology, above all to identify the challenges and limitations and thereby identify ways and means to go beyond the current state of play.

Ontologies have traditionally reflected a number of assumptions about the nature of knowledge and we will identify two here. First, they reflect a fundamentally monolithic view of knowledge. Starting with Bacon and continuing on through Locke up to most modern day scientists, the prevalent view has been that knowledge is a single edifice to which new bricks are continuously being added. This has been challenged by a number of writers. Quine held that knowledge was like a 'field of force' which impinged on experience only along the edges, while Thomas Kuhn considered that there were long periods where one scientific 'paradigm' or 'disciplinary matrix' dominated with brief 'revolutions' which shattered the prevailing mind set. Under this view, a scientist has necessarily to be conservative in order to work within a paradigm and yet innovation might trigger a revolution (Kuhn, 1977). It follows from this that we identify the ontology with the current theory or hypothesis and thus Kuhn might expect ontologies at 'revolutionary' moments to be re-written. The reality, however, is that rarely are there revolutions in ontologies and no one has ever 'proven' an ontology to be false.

*Corresponding author. Tel.: +44 1142221967; fax: +44 1142221810.
E-mail address: C.Brewster@dcs.shef.ac.uk (C. Brewster).

1071-5819/$ - see front matter
doi:10.1016/j.ijhcs.2007.04.003

A second assumption is that 'concepts' are the key building blocks, and we manipulate concepts with words. In effect, all ontologies, in spite of pretensions to the contrary, use human language to 'represent' the world. Some small functional element may be off-loaded to some form of logic but the major interpretive load—the semantics—is always carried by language. It was to eliminate the vagueness and ambiguity of language that John Wilkins in the 17th century created his 'Real Character' assigning each concept a numerical value derived from its position in the taxonomic tree. It was with much the same intent that, 400 years later, Lenat conceived his monumental enterprise to capture all of human common sense knowledge in Cyc (Lenat and Guha, 1990). The question remains how much, if at all, can items in the world or experiences which do not lend themselves readily to verbal expression be modelled?

It is an open question as to how far modern relatively lightweight philosophies of ontologies adopt these assumptions. For instance, the Semantic Web anticipates the use of ontologies to facilitate data sharing. That requires agreement or translation between the relevant parts of the ontologies underlying the data formats, but no agreement wider than this and certainly no claim to global consistency or completeness.

A simplistic view would say that an ontology is a model of the world which can be used to reason about it. However, there is a whole range of functions, assumptions and aspirations encoded in a given type or instance of an ontology. Following Davis et al. (1993), we can consider five functions of a knowledge representation, re-interpreted in the light of ontologies.

An ontology is a surrogate: Intelligent entities go through processes of reasoning about the world, often in order to plan actions. The reasoning involves internal representations but the objects reasoned about are usually external, outside in the world. Consequently, the ontology or knowledge representation is a surrogate standing for the objects and relations outside in the world. The 'fidelity' of the representation depends on what the ontology captures from the real thing and what it omits. Perfect fidelity is impossible.

One of the major claims for ontologies is that they will facilitate the interchange of knowledge between (for example) agents, or the reuse in different systems. However, if each agent or system has an imperfect model of its universe, knowledge interchange or sharing may increase or compound errors which were not visible in the initial use of an ontology.

An ontology is a set of ontological commitments: This may appear tautological but the choice of ontology is also a "decision about how and what to see in the world" (Davis et al., 1993, p. 19). This is both unavoidable because representations are imperfect and useful because it allows the representation to focus on that which the representation's author considers relevant or interesting. They see these choices as allowing us to cope with the overwhelming complexity and detail of the world. Consequently, it is the content of the representation, i.e. the set of concepts chosen and their inter-relation which provides a particular perspective on the world. The choice of notation (logic, LISP, or OWL) is unimportant.

It is interesting that, with respect to ontologies, an immense amount of effort has been expended in developing and defining ontology representation languages, and in contrast relatively little effort has been made to analyse what ontological commitments particular ontologies make. The only exception to this has been Guarino's critique of structures such as WordNet for not conforming to a logician's world view in terms of consistency and logical rigour (Guarino, 1998; Gangemi et al., 2001).

An ontology or knowledge representation is a fragmentary theory of intelligent reasoning: The way a knowledge representation is conceived reflects a particular insight or understanding of how people reason.[1] The selection of any of the currently available representation technologies (such as logic, frames, knowledge bases, or connectionism) commits one to fundamental views on the nature of intelligent reasoning and consequently very different goals and definitions of success. For example, Minsky himself in his original paper on frames noted that his work was a "partial theory of thinking" (Minsky, 1975), and later noted that "I want AI researchers to appreciate that there is no one 'best' way to represent knowledge. Each kind of problem requires appropriate types of thinking and reasoning—and appropriate kinds of representations" (Minsky and Laske, 1992).

The OWL language which has been developed by the W3C consortium as a standard language for describing ontologies for the Semantic Web comes in three flavours.

The existence of these three flavours reflects differing traditions that have merged in the current effort to construct a standard, and Horrocks et al. (2003) provide a fascinating insight into the various conceptual and technical pressures on OWL's developers. The provision of standards for ontology expression is a new slant on the idea of KR as theory; a standard, as opposed to an imposed KR formalism, can be ignored, adapted or adopted as appropriate. The knowledge engineer has, in effect, a managed choice of formalism. The extent to which such management imposes a theory is another open question.

An ontology is a medium for efficient computation: In the final analysis an ontology must allow for computational processing, and consequently issues of computational efficiency will inevitably arise. For example, using taxonomic hierarchies both "suggests taxonomic reasoning and facilitates its execution" (Davis et al., 1993, p. 27). Clearly, the development of the different flavours of OWL are a recognition of this fact, but in seeking sufficient speed

[1] The very use of a *representation* is a commitment to symbolic AI or what Haugeland calls "Good Old Fashioned AI" (Haugeland, 1985, p. 112ff.).

severe restrictions on the reasoning capacity of the representation have had to be made. More generally, it could be noted that since all ontologies depend on a propositional view of knowledge in order to begin to be computationally tractable, already a very restricted view of what it is possible to represent has arisen. The fact that *OWL Full* is not guaranteed to be 'decidable' unfortunately does not guarantee it to be sufficiently powerful to represent the whole gamut of what we can consider to be knowledge.

An ontology is a medium of human expression: All forms of knowledge representation including ontologies are both mediums of expression for human beings *and* ways for us to communicate with machines in order to tell them about the world. That knowledge representation is a form of human expression is something frequently forgotten in the field.

As new applications of ontologies have appeared (from the Semantic Web to knowledge sharing), it has become clearer that whereas some types of knowledge are eminently suitable to representation by ontologies—taxonomic information most obviously—others might not be. Furthermore, others have argued that ontologies as artefacts are unsuited to the rough-and-tumble of real-world applications once they get beyond a certain level of complexity; in other words, lightweight ontologies are OK, but there is a trade-off between expressivity and usability. As ontologies migrate from the lab to the Web scale, there will, so this argument goes, be inevitable problems. Related to this pragmatic argument is another, that ontologies are hard to maintain, and reify a particular point of view of the domain knowledge. In a dynamic, complex world, they will always be one step behind the current state of domain knowledge. Many point to the emergent structures that arise from the practice of tagging content in Web 2.0 applications, and argue that such so-called 'folksonomies' are preferable, in that (a) they are created at zero cost—they arise as a by-product of an activity that people are prepared to do for free, and that (b) because of that they produce conceptualisations that are strongly connected with the actual use of material, and which therefore reflect the use of such resources automatically.

In reply, defenders of ontologies argue that these perceptions are simply incorrect. The effort required to craft an ontology is often overestimated, and indeed will become more manageable, and the trade-offs less challenging, as larger communities use them more. And ontologies cannot be replaced by folksonomies, which are not intended to organise knowledge and facilitate its sharing in the same way; indeed, one of the aims of ontologies is to cut across the ambiguities and redundancies that are inevitable with folksonomies. For more discussion of ontologies and folksonomies, and Web 2.0 and the Semantic Web, see Berners-Lee et al. (2006).

Our aim in this issue was to delve deeper into these arguments, and to try to assemble a set of papers that would set out a robust debate and enable us to see whether the same questions are being raised across domains and applications. Have those who are sceptical about ontologies identified serious problems? Are there work-arounds? And if we can specify limits to the applicability of ontologies, surely that makes them more, not less, usable (since we would then have guidelines for their application). Is the fuss about ontologies wholly, partly, barely justified—or not at all?

2. Structure of this issue

These are deep questions, and the eight papers selected for this issue approach them from a variety of angles. The first two papers explore the problems of using ontologies in real-world situations, and attempt to use those experiences to draw lessons for the disciplines of ontology development and management. The next three suggest approaches to inform ontology development with input from other disciplines, modelling methods or types of information, to attempt to bridge perceived gaps between ontologies and different ways of modelling. The next two papers attempt to extend the areas of application of ontologies to new fields of endeavour. The final paper follows the strategy of going beyond ontologies, looking at the epistemological characteristics of knowledge discovery and developing a formalism that can capture that.

2.1. *Ontologies in their natural habitat*

The first strategy for investigating the limits of ontologies is to study the use of ontologies in real-world situations, to discover their pragmatic limits, and to engineer them for real-world application. Bo Hu and colleagues reflect on their experiences in building an ontology for the specific context of providing knowledge services for breast cancer screening assessment, a complex domain in which there is a large amount of fast-altering knowledge, much procedural knowledge (involved in identifying potential cancers from medical instrumentation), and a large number of heterogeneous points of view with associated variation in nomenclature (such as radiologists, histologists, surgeons, etc). Not only that, the ontology had to be readable by both machines and humans.

As they describe their approach to the ontology-development task, they argue against an attempt to model the breast cancer domain generically, and describe their design decisions and trade-offs in the context of the particular application and resource limitations. Compromise is essential, they argue. In particular, they make a plea for ontology engineering methodologies and criteria for success in application contexts.

Hu *et al.* make a methodological discussion. Robert Stevens and colleagues make similar points by a different route, setting out their experiences with building a bioinformatics ontology of protein phosphatases. Rather than focusing on the fitness of the ontology for use in the real-world context, Stevens *et al.* look at the ability of the Web ontology language OWL, which included description logic amongst its various influences, to represent what the bioinformaticians need. They discover that much of the domain is representable perfectly well using OWL-DL, and in many cases where it is not, simple extensions would suffice. But they also discuss several types of knowledge and representational apparatus that would lead to undecidability. Some complex property restrictions would have this effect. And certain *desiderata* such as fuzziness stemming from the biological notion of similarity, defaults, exception handling and prototypes have very little reasoning support available and it is an open question as to how incorporating them will affect OWL-DL's decidability.

Stevens *et al.* observe that OWL is currently oriented towards modelling biological facts as they are observed, rather than as biologists think the world is (a similar point to that made by Hu *et al.*). They also point out that it is not possible *a priori* to determine in advance what expressivity biology will demand, but are optimistic that further exploration of this question—rooted in actual experience of building complex ontologies in real domains—will enhance the disciplines of both bioinformatics and ontology development. Furthermore, they address some of Hu *et al.*'s methodological questions with the use of ontology *design patterns*, modelled on design patterns from OOP.

2.2. Augmenting ontologies

The second idea is that of providing new sources of input to the ontology development and application process. Stephen Cranefield and Jin Pan suggest an augmentation potentially leading to a large extension of ontologies' user base, by the Object Management Group's (OMG) model driven architecture (MDA). MDA is an industry standard for writing platform-independent specifications of applications or systems in enterprise contexts, abstracting away from implementation and focusing on functionality and behaviour, and can be used to manage enterprise architectures and the transfer of information, and is supported by automated services and tools; it is a key industrial modelling standard. Cranefield and Pan's work describe the process of making the MDA toolset available for ontology modelling, via a tool for converting ontologies with a metamodel defined using MOF (meta-object facility—an OMG technology) to representations of the same metamodel in RDF. This links ontology and enterprise modelling by providing a transformation mechanism between traditional ontology development formalisms and industrial-strength MDA models, which are of course language-independent. Bringing ontologies into the MDA arena will facilitate interoperability with complex industrial enterprises.

Walid Saba suggests methods for modelling commonsense knowledge and feeding those methods into an ontology. This goal, long associated with the CYC project (Lenat and Guha, 1990), is pursued by exploring the cues available in natural language use. Ordinary language use, Saba claims, can provide evidential support for the development of ontologies for significant fragments of discourse at least, which are well typed and avoid multiple inheritance, using Frege's principle of compositionality, which states that our understanding of a particular sentence is based on our knowledge of the senses of the words in that sentence together with an understanding of their method of combination in the sentence. He explores a number of common linguistic phenomena using this method. Perhaps most intriguing in the context of this special issue is his sketch of the processes behind metaphor based on Lakoff's theory Lakoff (1986), with mappings between isomorphic spaces (the space of the domain to be described and the space of discourse to be exploited for metaphorical terms).

Meanwhile, Henrik Eriksson notes that while knowledge to be shared between machines can be stored in an ontology, knowledge for sharing between humans tends to be stored in documents (Hu *et al.* highlighted the difficulties in using ontologies for both these sharing modes). Eriksson's response is to develop the notion of *semantic documents*, which are aimed at bridging the gap between knowledge stored in documents and that stored in ontologies. Semantic documents combine the two, allowing the knowledge stored in them to be accessed via different routes, and indeed allowing joint ontology/documentation development. The implementation of the approach is based on an Acrobat plug-in to Protégé, so has the advantage of exploiting and combining well-known technologies to annotate PDF documents with OWL statements, for example.

2.3. Extending ontologies

Stevens *et al.* identify a number of types of knowledge and inference common in bioinformatics that either can not be represented, or represented only with difficulty, in description logic-based ontology languages. The next two papers explore the limitations and potential of ontologies by extending their areas of application, learning how to apply them outside their comfort zone of taxonomic and classificatory knowledge. There are many areas where it is hoped that ontologies could be applicable or useful, often areas where technologies that lean heavily on ontologies (such as the Semantic Web) are meeting obstacles. Such areas include, for example, alternative ways of presenting information such as narratives, different representational modalities such as multimedia, or problematic contexts such as enterprise modelling. The extension of ontologies' application spaces also varies with the task that an ontology is supposed to perform.

This special issue contains two such attempts to extend ontologies' applications, in different tasks, representation and knowledge acquisition. K.L. Clark and F.G. McCabe use a formal ontology as a framework for the belief store of an agent. The background beliefs of the agent are modelling in the TBox, which is relatively unproblematic, but the dynamic beliefs that the agent holds about its environment are put in the ABox, where reasoning is harder and has less support. The use of ontologies in the agents world is spreading and new research questions are being raised. Clark and McCabe draw attention to the importance of consistency maintenance, for example; in their case the question is how to revise ontology facts to maintain consistency with the axioms within an agent. In a system of distributed agents with constant querying, consistency maintenance (and the various interpretations of what that might mean) is likely to be an important challenge.

An alternative direction in which to extend the application of ontologies is in the type of domain knowledge they can be used to acquire. Sangun Park and Jae Kyu Lee use an ontology-based approach to extract rules from Web pages. Although XRML provides some means of automating rule generation, rule acquisition is still a bottleneck for rule-based systems. Park and Lee have developed an ontology of structures for rule components and rule identification which they apply to Web pages to extract rules from them, in a bottom-up manner mapping from the identified variables and objects. The method works in particular with Web pages that have a certain conventional structure in common (Park and Lee use bookselling sites, such as amazon.com and barnesandnoble.com, for their examples and evaluation).

2.4. Ontotherapy

Finally, one can accept the thrust of arguments that ontologies are fundamentally inadequate for representing the knowledge that system builders need. That means either accepting limits to the extent of the knowledge applications that one wishes to use, or exploiting alternative methods of knowledge representation. The latter approach is used by William Pike and Mark Gahegan, whose aim is to represent scientific knowledge in such a way as to include within the representation as much information as possible relating to the history and justification of the knowledge, in other words, information that helps contextualise the knowledge by showing how it is situated. If human knowledge is irretrievably situated, then the straightforward logical structure of an ontology, which records the knowledge at the end-point of what might have been a complex and contentious process of discovery, will miss this information about provenance. The result is a system that aims to get the best of all worlds, the richness of representation of bottom-up methods (from the cooperative knowledge construction community) and the richer semantics of top-down methods.

Not only does Pike and Gahegan's system store descriptions of abstract ideas or concepts, but related to those are profiles of the people who have either created or applied resources, files in which there is information about or relevant to a concept, records of tools used to analyse information and to construct concepts, places where research has taken place, and the workflows relevant to the concept. This is not merely a complex set of metadata for an ontology, partly because all this data can be stored independently of the concepts in the domain: 'concept' is at the same level in the system as 'tools', 'places' or 'tasks', for instance, and partly because although all the resources are described in OWL-DL to facilitate interoperability with other Web representations, the OWL statements in the system are not full ontologies but rather whatever pieces of information relevant to the user's current work.

Pike and Gahegan's approach does not—cannot—meet the full force of the arguments from a total sceptic about ontologies. But by allowing aspects of context and provenance of knowledge to be represented as full-fledged pieces of knowledge themselves, rather than annotations to ontologies, they allow the contextualisation of ontologies, as opposed to the inclusion of context within ontologies. This may suggest a way forward to those who are sceptical about perceived overclaiming by the ontology community, but who still see no reason to reject ontologies altogether.

2.5. Conclusions

Ontologies are extremely important tools for the organisation and contextualisation of knowledge, particularly in well-bounded contexts, such as scientific research, or within individual organisations. However, perceptions of ontologies are not uniformly positive, with many thinking they demand too much work, and others believing they are too rigid to be of use in dynamic, complex real-world contexts. The appearance on the scene of Web 2.0 and its folksonomies has also charged the debate still further. Of course, a special issue like this one cannot decide the arguments. But the papers here help to clarify them, showing where ontologies are difficult to use, where there are opportunities to use them in new contexts, and where they can be stretched. Only by collaboratively exploring the space and presenting results will the best engineering solutions be discovered, and we hope that this special issue has made a contribution to that.

3. Final thoughts and Acknowledgements

The editors have been concerned with ontologies for some years, as a result of their work as researchers on the UK Engineering and Physical Sciences Research Council's Interdisciplinary Research Collaboration Advanced Knowledge Technologies (AKT, Grant no. GR/N15764/01.)[2] AKT, which has focused on providing tools and methods for managing knowledge across its whole lifecycle from acquisition to maintenance, has naturally focused on the Semantic Web as an important technology.

[2]http://www.aktors.org.

Our concern to understand how ontologies should and should not be used initially provoked us into editing an edition of the 'Trends and Controversies' section of the journal *IEEE Intelligent Systems* (Brewster et al., 2004) under the title of ''Knowledge representation with ontologies: the Present and future''. A very interesting and lively debate ensued, for which we thank the participants, who included Simon Buckingham Shum, Jeremy Ellman, Enrico Franconi, Steve Fuller, Mark A. Musen and Yorick Wilks.

Editing the special issue was harder work than we had bargained for. We got many more papers than we expected, and were able to accept very few. There is one paper which we regret could not be included. In the course of revising his provocative and stimulating paper, 'Ontology, Society and Ontotheology', Professor Joseph Goguen of the University of California at San Diego, died, in July 2006, only a few days after attending a Festschrift to honour his 65th birthday. Prof. Goguen was one of the most-cited authors in computer science, and an expert in the field of ontologies and database integration, as well as algebraic specification of programs (he coincidentally supervised the MSc thesis of one of the special issue editors at the University of Oxford in 1989, on the OBJ3 language), but his range of interests was incredibly wide, taking in sociology, music and philosophy as well as many other sub-disciplines of computing. An earlier version of his paper was presented at the 2004 International Conference on Formal Ontology in Information Systems, in Torino, Italy, and—for a highly sceptical view of ontologies, informed by wide reading and long experience in the field—we heartily recommend this paper to readers of this issue. We extend our sympathies to his family and colleagues.

We would like to thank all those who submitted papers, which were almost without exception stimulating and interesting. Finally, many thanks to our authors, to whom we now hand over, not without a small sigh of relief.

References

Berners-Lee, T., Hall, W., Hendler, J.A., O'Hara, K., Shadbolt, N., Weitzner, D.J., 2006. A Framework for Web Science. Foundations and Trends in Web Science 1 (1), 1–130.

Brewster, C., O'Hara, K., Fuller, S., Wilks, Y., Franconi, E., Musen, M.A., Ellman, J., Shum, S.B., 2004. Knowledge representation with ontologies: the present and future. IEEE Intelligent Systems 19 (1), 72–81 URL: ⟨http://csdl.computer.org/comp/mags/ex/2004/01/x1072abs.htm⟩.

Davis, R., Shrobe, H., Szolovits, P., 1993. What is a knowledge representation. AI Magazine 14 (1), 17–33.

Gangemi, A., Guarino, N., Oltramari, A., 2001. Conceptual analysis of lexical taxonomies the case of WordNet top-level. In: Proceedings of the International Conference on Formal Ontology in Information Systems. ACM Press, New York, pp. 285–296.

Guarino, N., 1998. Some ontological principles for designing upper level lexical resources. In: Proceedings of the First International Conference on Language Resources and Evaluation, LREC 98. URL: ⟨http://www.loa-cnr.it/Papers/LREC98.pdf⟩.

Haugeland, J., 1985. Artificial Intelligence The Very Idea. 1985. MIT Press, Cambridge, MA.

Horrocks, I., Patel-Schneider, P.F., van Harmelen, F., 2003. From SHIQ and RDF to OWL the making of a Web ontology language. Journal of Web Semantics 1 (1). URL: ⟨http://www.websemanticsjournal.org/volume1/issue1/Horrocksetal2003/⟩.

Kuhn, T.S., 1977. The Essential Tension. Chicago University Press, Chicago.

Lakoff, G., 1986. Women, Fire, and Dangerous Things. The University of Chicago Press, Chicago.

Lenat, D.B., Guha, R.V., 1990. Building Large Knowledge-based Systems. Addison-Wesley, Reading, MA.

Minsky, M., 1975. A framework for representing knowledge. In: Winston, P. (Ed.), The Psychology of Computer Vision. McGraw-Hill, New York, pp. 211–277 URL: ⟨ftp://publications.ai.mit.edu/ai-publications/pdf/AIM-306.pdf⟩.

Minsky, M., Laske, O., 1992. A conversation with Marvin Minsky. AI Magazine 13 (3), 31–45.