The Knowledge Engineering Review, Vol. 10:1, 1995, 67–68

Comparison of implemented ontologies

Report on the workshop held during the *European Conference on Artificial Intelligence,* August 9 1994, Amsterdam

NICOLAAS J. I. MARS

Knowledge-based Systems Group, University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands

A number of groups developing knowledge-based systems have found (or at least posited) that the design and representation of a limitative set of concepts and relations, a so-called ontology, can contribute to sharing and reusing knowledge bases. However, very few descriptions of implemented ontologies have appeared in the literature. No comparison of competing proposals is available, let alone an empirical determination of the benefits of using an ontology. There is no accepted method for designing and building such ontologies, nor is it clear how ontologies can best be evaluated.

During the European Conference on Artificial Intelligence in 1992 in Vienna, a workshop was held where the emphasis was on the potential benefits to be obtained from using an ontology in allowing multiple use and reuse of knowledge bases (Mars, 1992). At that time, very few research groups could report on actual experience in designing, implementing and using ontologies; a notable exception being the work done in the Electronic Dictionary Research Project in Tokyo, Japan. The aim of the 1994 Amsterdam workshop, on which we report here, was to gather those researchers who have actually designed and implemented ontologies for the purpose of improved reuse and sharability of knowledge bases.

Requiring submission of a machine-readable implementation of an ontology as a condition for participation in the workshop was felt to be necessary but risky by the organizers (Tom Gruber of Stanford University, Palo Alto, USA; Otto Kuehn of DFKI, Kaiserslautern, Germany; and Nicolaas Mars, University of Twente, Enschede, The Netherlands). It was thought to be necessary to prevent presentations of planned projects, and risky because it was unclear whether enough contributions would be obtained. A surprisingly large number of submissions was in fact obtained. Thirteen ontologies were accepted for discussion at the workshop and publication in the workshop notes (Mars, 1994). They cover a wide range of application domains, and an ever wider variety of opinions on what ontologies are, how they should be developed, and what they are good for.

Requiring submission of an implemented ontology was not intended to lead to a discussion of the contents of the ontologies themselves. In most cases, too much domain expertise is needed to understand the intricacies of the ontologies. Rather, as was stated in the acceptance letter to the authors, two questions were to be the focus of discussion during the workshop: What is the use of an ontology? and How should one build an ontology? Luckily, most authors treated these two issues explicitly in their presentations.

To allow contrasting approaches to be brought out, the thirteen contributions have been grouped according to their application domain, or the intended use of the ontology. Two contributions, by Bateman et al. and Horacek, were aimed at probably the oldest area where the use of implemented ontologies in AI appeared: natural-language analysis. However, where one group aims to cover "the world", the other analyses a small fraction of the world in considerable detail.

One pair of papers was characterized by the application domain of materials science: Bachman and Stevens treated composite materials, and Van der Vet et al. industrial ceramics. Here, the pairing proved fruitful in identifying differences caused by the differences in the domains from differences caused by the intended use of the ontologies. In particular, the strong bond between composition and fabrication in composite materials was found to be much looser in ceramics, with considerable influence on the resulting structure of the ontologies.

Another strongly related pair of ontologies were Gruber and Olsen's ontology of a large part of engineering mathematics, and Mars' ontology of a small part thereof, namely measurement and measurement units. These ontologies are closely linked to (and in some cases actually used in) ontologies aimed at modelling physical artifacts, especially in engineering design, by Alberts and Wognum, by Borst et al., and by Kuehn.

Four contributions formed the unavoidable category "others", with ontologies of medicine by Falasconi and Stefanelli; of expertise for problem solving by Mizoguchi and Vanwelkenhuysen; of enterprise models by Gruninger and Fox; and of system development methodologies by Van der Touw.

A lively debate was held after the presentation on the intended use and the development strategy used. The scale of the meeting, with 25 participants, contributed to this liveliness. The discussion ranged from highly philosophical issues (the difference between term and concept being a recurring issue) to highly pragmatic ones (How much domain knowledge should the designer of an ontology possess? Consensus answer: a lot, but not so much as to loose his/her ability to abstract.)

Consensus was reached on the first issue raised by the organizers: what is the use of an ontology? Three positive effects of developing an ontology were identified. The first is the effect already mentioned during the Vienna 1992 workshop: to ease multiple use and reuse of knowledge bases. In the long run, this may be the economically most important effect, as reuse of knowledge bases will allow the amortisation of the high costs of building knowledge bases over several applications. Also, this may allow the emergence of knowledge bases as a commodity. However, as was mentioned in passing during the workshop, considerable legal problems will have to be solved to make this possibility a reality. Also, it should be said that, like in Vienna, the positive effect on reuse and sharing of knowledge bases was not shown empirically yet.

A second effect is helping structuring a knowledge base, even when this knowledge base is not (*a priori*) intended for multiple use. Especially in the case of very large knowledge bases, as needed in future complex applications in engineering and science, this structuring will become crucial. Finally, a third effect, certainly not shared by all participants, could be that developing a well-thought out ontology may contribute to insight in the domain modelled. It was interesting to note that proponents of this vision came from those domains generally thought to be well-organized: engineering and mathematics. The most extreme form of this position was represented by researchers from the University of Twente, who defended that a ontology should be publishable (and indeed published) in scientific journals in the domain concerned.

Concerning the second issue of the workshop "How should one develop an ontology?" the consensus was that it is much too early to speak of a method (often called methodology) for ontology development. Currently, the best recommendation is "just do it". Only after much more experience in developing ontologies has been gained, and after several more workshops for comparing experiences, will it be possible to formulate recommendations for the development process. There was, however, agreement that an ontology should be designed for a particular class of applications. Very few participants believed in a universal ontology.

References

Mars, NJI (Ed.), 1992. "Workshop knowledge sharing and reuse: ways and means". In: *Workshop notes; 10th European Conference on Artificial Intelligence,* 3–7 August, Vienna, Austria.

Mars, NJI (Ed.), 1994. "Workshop comparison of implemented ontologies". In: *Workshop notes; 11th European Conference on Artificial Intelligence*, 8–12 August, Amsterdam, The Netherlands.