The Knowledge Engineering Review, Vol. 13:1, 1998, 1–3
Printed in the United Kingdom

Putting ontologies to use

MIKE USCHOLD[1] and AUSTIN TATE

Artificial Intelligence Applications Institute, 80 South Bridge, Edinburgh, EM1 1HN, Scotland

Interest in the nature, development and use of ontologies is becoming increasingly widespread. Since the early nineties, numerous workshops have been held. Representatives from historically separate disciplines concerned with philosophical issues, knowledge acquisition and representation, planning, process management, database schema integration, natural language processing and enterprise modelling, came together to identify a common core of issues of interest. There was highly varied and inconsistent usage of a wide variety of terms, most notably, ''ontology'', rendering cross-discipline communication difficult. However, progress was made toward understanding the commonality among the disciplines. Subsequent workshops addressed various aspects of the field, including theoretical issues, methodologies for building ontologies, as well as specific applications in government and industry.

The field is maturing, as evidenced by the increasing quality and quantity of published results, as well as the upcoming *International Conference on Formal Ontology in Information Systems* in June 1998, the first conference specifically devoted to this topic.

Among workers in the field, there is now broad agreement as to what people are talking about when discussing ontologies. This is not to say that everyone agrees on a definition, but rather that the various definitions that do exist are understood, and relatively little heated discussion takes place arguing about what *the* definition of ontology is or ought to be.

A wide variety of ontologies have been reported in the literature, an important development in the past five years. Web-based tools for creating, editing, browsing and using ontologies are available. A number of representation formalisms exist, and there are some preliminary methodologies for constructing and evaluating ontologies. Tools for performing some syntactic translation between representation languages are also available, for inter-operability. However, as our understanding of the issues increases, more and more difficult unsolved problems emerge, ensuring that this will be an active research area for some time.

Ontologies have been used within commercial applications already. However, details are rarely disseminated. This is due to commercial sensitivity, or the fact that applications are often deeply embedded, so it is difficult to demonstrate clearly the important benefits of an ontology in a practical context. It is still to a large extent a matter of faith that they will deliver the goods.

There is much faith, however, as both the research and applications communities are hard at work addressing these issues. The importance of ontologies has been recognised in fields and industries as diverse as semi-conductor manufacturing, aircraft design, enterprise process management, heterogeneous database integration, knowledge engineering and planning, to name a few.

Ironically, despite the fact that many ontologies are specifically built to be *re-used*, there are few reported cases of this. There is a temptation for people to build their own, even if others exist in that domain. This may be due to the fact that the *purposes* for which ontologies are built varies a great deal, as well as the fact that there is little uniformity in how ontologies are *presented* to the community.

[1]Now at Boeing Applied Research and Technology, P.O. Box 3707 M/S 7L–40 Seattle, WA, USA 98124.

We still have very little understanding of how the specific purpose for an ontology affects decisions about its form and content, how one goes about building it, what is the right level of abstraction, and the extent to which it may be re-used.

The diversity of presentation styles for reporting ontologies also limits reuse. The degree of formality varies, as well as the range of languages used when an ontology is formalised. Also, and very importantly, relatively few are published in the literature with detailed descriptions of the structure and content, along with the specific technical details illustrated with examples of how they were used. While it is true that *direct* reuse of terms and definitions of a pre-existing ontology is rare, there is now considerable reuse of the *ideas* underpinning the terms and definitions. This theme will be illustrated in this special issue.

In the June 1996 issue of *The Knowledge Engineering Review*, a comprehensive introduction to the field of ontologies was presented. It included a working definition of the term "ontology", examples of how ontologies are intended to be used, and examined some of the existing projects and ongoing research issues. However, because there have been very few ontologies whose contents have been fully described and presented in the literature, it is difficult to really grasp the concept of where and how ontologies are going to have impact in the forseeable future.

The papers in this special issue are intended to help rectify this and to achieve the following objectives:

1 To report and summarise *de facto* standard usage of terminology that is emerging in the area of ontologies in particular, and in knowledge level modelling more generally.
2 To present in full the contents of two substantial ontologies including some discussion of the process and experiences of constructing them.
3 To describe and illustrate with specific examples the intended and actual uses of these ontologies.
4 To illustrate cases where important ideas from existing ontologies were reused as the basis for subsequent development of more substantial ontologies.

The first paper, "Knowledge level modelling: Concepts and terminology", addresses the first objective. In it, the authors attempt to clarify some of the confusion that has arisen due to inconsistent usage of terms in the area. First, a knowledge level model is defined, and two main types are distinguished: ontologies and problem solving models. For each, detailed definitions are given, followed by concise definitions of a variety of important related terms. Importantly, this is a descriptive, rather than prescriptive exercise. The intention is *not* to dictate a standard that should be conformed to, but rather, to present a snapshot summarising how terms are currently used. The work reported was conducted as part of the ESPRIT EuroKnowledge project which is the European equivalent of the US Knowledge Sharing Effort.

The remaining papers address the last three objectives. All describe significant items of work in creating ontologies intended for practical use. All build upon earlier work.

In "The Enterprise Ontology", the authors provide a comprehensive description of the development and use of the Enterprise Ontology, a collection of terms and definitions relevant to business enterprises. After a brief introduction to the context of the work, a glossary of terms with natural language definitions is given. Existing ontologies for enterprise modelling and planning heavily influenced the choice of terms and definitions. This is followed by a discussion of how the the informal definitions were converted into a formal language. The experiences of using the Enterprise Ontology are described. These are compared with the intended uses, and successes and failures are discussed. The paper concludes with an overall evaluation and summary of what was learned. The work described was part of a large collaborative project involving industry and research groups in the UK. Practical applications for the industrial partners resulted from the work.

Next, "The PIF Process Interchange Format and Framework" presents the full contents of an ontology for representing process information. The intended purpose is to capture in an implementation-language independent way, the key concepts that are required to represent processes. PIF was created by an international group concerned with process information interchange between diverse tools and systems. PIF also reused ideas from previous ontologies,

including the Enterprise Ontology, the ontologies for processes, time, and resources from the Toronto Virtual Enterprise (TOVE) project, the Process Handbook project at the MIT Center for Coordination Science, and others.

Reuse is ongoing; the PIF team is currently collaborating with a National Institute of Standards (NIST) effort to develop a Process Specification Language (PSL). There is already agreement that core elements of these ontologies will be explicitly shared. This is elaborated in more detail in the final paper, which helps meet the fourth objective of this special issue. In ''Roots of SPAR – Shared Planning and Activity Representation'' new work is described, in which existing ontologies (including Enterprise and PIF) are being combined to create a practical Shared Planning and Activity Representation (SPAR) that can be used across technical and applications communities. This is indicative of the benefits that can be obtained by using well founded work on ontologies as the basis of representations intended to meet specific needs.

The work reported in this special issue of *The Knowledge Engineering Review* should serve as a guide to the current level of work in the field and to show the early stages of work that aims to put ontologies to practical use.