Applied Ontology 1 (2005) 1–5
IOS Press

Introduction

Applied Ontology: Focusing on content

Nicola Guarino and Mark A. Musen
Editors-in Chief

In a world that is overflowing with journals and other outlets for scientific publication, the appearance of any new periodical requires some justification. There are already more journals than we can read and more conferences than we can attend. In the case of *Applied Ontology*, we believe that the creation of a new journal not only is completely justifiable, it is downright exciting.

For too long, workers in computer science have assumed that content comes for free. "Theory" in computer science has always meant the theory of processes and of computation. We measure the complexity of computer programs in terms of how long it takes machines to execute them, not in terms of how long it takes people to understand and to represent the data on which those programs might operate. We typically describe computer code in terms of algorithms that operate on formal parameters, often without pausing to discuss where the data that might satisfy those parameters come from.

This journal was founded on the premise that workers in computer science, informatics, and information science are overdue in paying as much attention to *content* as they do to algorithms.

There are journals that emphasize methods for representing knowledge and data. There are journals that emphasize processing of knowledge and data. *Applied Ontology* takes the position that, before we choose representation systems or algorithms, we first must create models of the world and clarify "what exists." Ontology is no longer perceived as an arcane branch of metaphysics, the province only of philosophers; the study of ontology now fits squarely into the study of modern computer science and informatics. Building ontologies is now an essential activity that underlies nearly everything we do in the development of computational systems.

The dramatic growth of information and communication technologies in recent years has led to a new focus on information content, and to an increasingly interdisciplinary approach to research and development activities in these fields. Linguists and philosophers now work hand-in-hand with traditional computer scientists to build complex information systems with explicit, examinable conceptual models of the environments in which they are intended to operate, of the organizations in which they will be used, and of the data and knowledge that they will process. In order to be understandable and reusable, these models need to combine the precision of formal semantics with the efficacy of cognitive transparency, as they incorporate increasingly sophisticated and heterogeneous modeling paradigms. The recent explosion of interest in ontologies is a fundamental component of this trend. Researchers in disciplines such as knowledge engineering, information-systems modeling, artificial intelligence, formal and computational linguistics, information retrieval, library science, and knowledge management have come to realize that a solid foundation for their research calls for serious work in ontology, understood as a general theory of the types of entities and relations that make up their respective domains of inquiry. The unification of these diverse disciplines by a shared desire to elucidate, to evaluate, to apply, and to

reuse models of the domains on which information systems operate is a remarkable outgrowth of this new emphasis on making information content a "first-class entity".

Applied Ontology is the first journal with a central and exclusive focus on ontological analysis and conceptual modeling as an interdisciplinary science. It is our goal that, as a new journal in a new area, it will establish a unique niche in the realm of scientific publication. Indeed, there is no current periodical that can embrace conceptual modeling issues both in artificial intelligence and in conventional software engineering, and that is equally at home publishing papers on the theoretical aspects of ontological analysis and pragmatic papers that delve deep into the nuances of modeling real content. Existing computer-science journals tend to define their scope rather narrowly, avoiding on the one hand general ontological questions (which are perceived as too "philosophical") and on the other hand questions of domain analysis (which are perceived as too "pedestrian"). We consequently are convinced that there is an expanding need for a scholarly publication that can address the broad issues of conceptual modeling head on, and that can draw on both theoretical and applied research that transcends the entire enterprise of building and managing information systems.

We anticipate that authors of theoretical contributions will show the relevance of their theory for applications, and that authors of more application-oriented papers will show the relevance of a well-founded theoretical perspective. Moreover, the journal will publish papers focusing on representation languages only where these languages address relevant content issues, whether at the level of practical application or of theoretical understanding. Similarly, we will publish descriptions of tools or implemented systems only where a contribution to the practice of ontological analysis or conceptual modeling is clearly established. The scope of the journal thus will range from the philosophical to the psychological, from the theoretical to the pragmatic. In sum, the journal aims to be a highly accessible, highly informative place to publish scholarly work related to all aspects of ontological analysis, engineering, and evaluation.

We believe that the time is right to begin such a journal. Ten years ago, academic workers in computer and information science spoke of ontology carefully and cautiously, almost embarrassed to utter the "o" word. Now, there are nearly 5 million Web pages that include the word "ontology". Most revealing, typing the word "ontology" into Google brings up several *sponsored links* that advertise jobs for people to work on projects such as the construction of catalogs for e-commerce, the implementation of natural-language systems, or the development of online ontologies. The advent of model-driven architectures in software engineering, of model-based approaches for information integration, and of terminological standards for the annotation of experimental data in the sciences has brought the notion of ontology to the center of attention in a range of disciplines. Ontologists are everywhere, earning a living, and working in every domain of human activity imaginable. At the same time, a burgeoning cadre of scientists now see ontology as a rapidly expanding area of scholarly activity in its own right, seeking to refine principles both for modeling the world and for using those models to drive a wide range of information systems.

We invite you to become a part of this adventure. By reading and contributing to the pages of this journal, you will help to advance our field and to share in the dissemination of knowledge about conceptual modeling and domain analysis. We find it remarkable that an activity that traces its origins to the work of philosophers who lived more than two millennia ago has become central to the development of modern information technology. We find it exciting to be able to articulate broadly applicable principles for ontological analysis and to see how to apply them in new domains. We believe it is essential to look at the details of how modeling may have been done in particular domains and in particular situations in order to extract those generalizable principles. At the core of *Applied Ontology* is a desire to understand the nature of reality and how people construe their world. There is nothing more fundamental to human thought or to translating our thoughts into computational artifacts. With links to a vast philosophical

tradition, to nuances in human language and cognition, and to emerging methods for building the information systems that will affect the future of our society, the study of ontology has never been more important. We and our editorial board are enormously excited about this new journal, and about what the future of our discipline holds in store.

The contents of this issue

The papers we have assembled for this inaugural issue are intended to show the journal's character, and the overall attitude we shall adopt. They are organized into two sections: short editorials coming from the journal's editorial board, and regular papers that have undergone rigorous peer review. We shall briefly comment all of them, taking this as a first opportunity to engage in some debate.

The section on editorials is indeed intended to stimulate discussion, presenting different points of view coming from the large and heterogeneous group of distinguished researchers who have agreed to serve in our editorial board. Further contributions are expected in the coming issues; we encourage all our readers to submit short position statements for potential publication in the journal.

The first editorial note comes from Doug Lenat, perhaps the first computer scientist to realize the importance of assembling a large, comprehensive ontology, and who (successfully) invested his own money in such an enterprise. Lenat describes the "slippery slope" from Ontology to Knowledge Representation, encouraging articles that discuss inferencing issues that bear directly on ontology, such as those concerning the ontology of contexts. He then presents some common pitfalls of ontological engineering, making some provocative points that will probably raise some debate. For instance, he raises the old puzzle concerning the ontological difference (if any) between a table and the wood it is made of by assuming that the same entity of the real world can be taken as an instance of Table or WoodStuff, depending on the relevant context. The OntoClean methodology (which is extended in this issue in the paper by Welty and Andersen) would solve the conundrum by clarifying that Table and WoodStuff have incompatible unity conditions, corresponding, therefore, to disjoint classes.

The paper by Mike Uschold focuses on the need to establish well-designed metrics, methodologies, and experiments for evaluating the practical benefits of ontologies. He sees our journal as a way to foster the development of a pipeline that disseminates academic research results into industry and government. Boeing was probably the first company (after Cycorp) to hire ontologists, and Mike knows very well what it means to put ontologies to work. A serious issue he raises concerns the practical limits of axiomatization, and the need to rely on natural language – despite the problems of ambiguity and vagueness.

Brandon Bennett addresses straightforwardly these worries concerning vagueness and ambiguity. Coming from the tradition of formal logic and AI representations, Brandon confesses his puzzlement with the position of many information-system designers, who appear to be so negative about the possibility of achieving precision and universality in ontologies. He then offers us a way to understand the reasons of this puzzlement, by analyzing in detail the different ways in which we describe and classify objects, and the different underlying factors that make up the phenomenon of conceptual vagueness. Bennett's conclusion is that the apparatus of classical logic, within which ontologies traditionally have been defined, cannot by itself account for the meanings of conceptual terms in natural language. Axiomatic ontologies rather should focus on specific modes of classification (corresponding to single aspects of complex term senses), while an auxiliary superstructure should represent the conflation of such different modes in the intuitive sense of natural-language terms.

The short note by Massimo Poesio focuses on the link between ontologies and computational lexicons. He makes an interesting reconstruction of the last thirty years or so of research in natural-language processing and knowledge representation, showing the emergence of two classic schools of thought: those defending formal semantics and rigorous philosophical analysis for the purpose of building ontologies, on one hand, and those focusing on concept formation and machine learning techniques for the purpose of building computational lexica, on the other hand. As we may expect, given his presence in this journal, he argues convincingly for taking advantage of both approaches in practical NLP applications, by exploiting (and further investigating) the subtle links between language and ontology.

The joint paper by Nathalie Aussenac-Gilles and Dagobert Sörgel, both members of the *Applied Ontology* editorial board, focuses again on the relationship between ontologies and natural language, from the point of view of a different community: that of terminology engineering. The authors discuss in detail the role of domain- and task – specific ontologies, arguing that, in many cases, such ontologies can be more useful than general ontologies, especially if developed through document content analysis by using ontology-learning techniques. The authors then present their point of view on the way to extract an ontology from texts, including Web documents, followed by a compact and yet comprehensive review of natural language processing tools for ontology design and population.

Antony Galton addresses a crucial issue, which is a the heart of any serious effort of ontological analysis: Although we are used to modeling the world in terms of 'things" or "objects", many real-world phenomena present themselves as more or less unbroken continua, with no clear-cut boundaries. This is particularly the case in the geographical domain, where object-based and field-based modeling approaches need to coexist. The author discusses the ontological aspects underlying the distinction between objects and fields, and concludes that the time is ripe for the development of a formal ontological framework concerning the relationship between them. In his, view, such formal work could be generalized to what he calls *multi-aspect* phenomena, such as in the case of a procession, which could be seen as an object, an event, or a process.

The first original article that we publish in this issue is by Maureen Donnelly, winner of the best paper prize for young researchers at the 2004 Conference on Formal Ontology in Information System. Donnelly presents here a revised and extended version of her conference paper on Basic Place Theory, which represents a first attempt at developing a comprehensive formal theory for reasoning on relative places and their changing relations to both other places and to material objects. The paper argues for the practical role of relative space, as contrasted with absolute space, not only in commonsense reasoning, but also in biology, engineering, meteorology, and other disciplines. Two foundational tasks are addressed: one is to provide a mechanism for dividing places into separate "location-complexes" (as in the human body), whose spatial relationships among their internal places are relatively stable in time; another task is to introduce time-dependent relations that may hold among places in different location-complexes.

The second original article is by Roman and colleagues, who describe the Web Services Modeling Ontology (WSMO). This paper presents an ontology that captures the distinctions necessary to allow an intelligent agent to discover, invoke, and chain together a sequence of Web-based services that will enable the agent to achieve a particular goal. The paper is a clear example of the importance of ontologies in the emerging Semantic Web, and of the kind of content knowledge that needs to be modeled in order for a computer-based agent to operate on the Web to automate different tasks. The promise of a future in which computer-based agents can "surf" the Web and can interact with Web sites in a manner similar to that in which people do continues to be a stimulus for much of the work that takes place on problems of ontology. Roman and colleagues offer a comprehensive solution to the problem of seeking out and invoking sequences of Web services. Their paper challenges the Semantic Web community to assess the

utility of the ontology by building WSMO-compliant tools and resources to study the effectiveness of the WSMO in actual applications.

Welty and Andersen build on the OntoClean methodology first proposed by Guarino and Welty. By providing a set of meta-properties for characterizing the ontological nature of an ontology's properties, OntoClean allows ontology builders not only to describe the semantics of their ontologies with more precision, but also to check for potential inconsistencies in the taxonomic structure. One of the meta-properties in OntoClean is known as *rigidity*. Welty and Andersen note that, since the time that OntoClean was originally described, a number of authors have noted problems in applying the Rigidity meta-property. The authors review and analyze what it means for a property to be Rigid, and they clarify how Rigidity can come in several flavors. Welty and Andsersen propose what they describe as a "formal ontology of rigid properties." The paper is important because it provides a cogent example of how work in applied ontology – even work in highly formal aspects of ontology description – evolves in response to challenge from the scientific community. It also is important because the refined meta-ontology proposed by the authors may have significant practical benefit for ontology developers with a more practical emphasis in their work who seek ways to provide clearer semantics in their ontologies and to check for potential inconsistenies.

The paper by Purao and Storey comes from the information-systems community, which traditionally has been rather separated from the community of scholars who perform research in logic-based ontologies (as is also typical of the AI and Semantic Web communities). Purao and Storey focus on a very relevant problem, understanding the semantics of relationships in order to compare (and possibly to integrate) different database schemes. They propose a comprehensive ontology of binary relationships based on the three fundamental categories inspired by Bunge's ontology, "Status, Change of Status, and Interaction", further articulated into more specialized primitives. They report in detail the results of an evaluation experiment designed to measure how well different researchers were able to classify a reference set of relationships according to the proposed ontology. Interestingly enough, different researchers could agree on the same ontology-based classification in most cases, and, in cases of disagreement, they were able to argue for their own interpretation and also to understand the alternative suggested by the investigator. This experiment shows very nicely how ontology can help people with different points of view to understand one another, even if they do not agree on everything.

A new forum for scientific discussion

As the editorials and research papers of this inaugural issue demonstrate so convincingly, the field of applied ontology is wide and draws on many heterogeneous disciplines. It is our goal to make accessible to a broad audience the range of activities characterized by the journal's initial offerings, by promoting exciting, readable articles that present all aspects of the field, chronicling new scientific advances and inspiring us with as yet unmet challenges.

We are very excited to be part of this new publication channel. It certainly is time for a scientific journal that can concentrate on the challenges of modeling, managing, communicating, and deploying ontologies. Our mission is nothing less than to focus on methods for developing and disseminating the content that is not only central to modern information systems, but also central to modern thought.