The Knowledge Engineering Review, Vol. 17:1, 1–5. © 2002, Cambridge University Press
DOI: 10.1017/S0269888902000310 Printed in the United Kingdom

Introduction to the special issue on ontologies in agent systems

STEPHEN CRANEFIELD,[1] STEVEN WILLMOTT[2] and TIM FININ[3]

[1] *Department of Information Science, University of Otago, Dunedin, New Zealand*
[2] *Laboratoire d'Intelligence Artificielle, Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland*
[3] *Department of Computer Science and Electrical Engineering, University of Maryland Baltimore County, USA*

1 Introduction

It is now more than ten years since researchers in the US Knowledge Sharing Effort envisaged a future where complex systems could be built by combining knowledge and services from multiple knowledge bases and the first agent communication language, KQML, was proposed (Neches *et al*., 1991). This model of communication, based on speech acts, a declarative message content representation language and the use of explicit ontologies defining the domains of discourse (Genesereth & Ketchpel, 1994), has become widely recognised as having great benefits for the integration of disparate and distributed information sources to form an open, extensible and loosely coupled system. In particular, this idea has become a key tenet in the multi-agent systems research community.

However, to date most multi-agent systems have been of a small scale and confined to individual research laboratories. Despite a recognition of the central role of explicitly modelled ontologies, there has been little need for agent developers to define and share explicit ontologies. Although some research groups (notably those projects using the Ontolingua ontology server[1]) have made ontologies available to the wider community, for many other projects an implicit agreement on the terms that can be used within a message has no doubt been sufficient due to an emphasis on other research goals.

Work towards standards for agent communication[2,3,4], and the availability of open-source software supporting these standards[5,6], has meant that multi-agent systems have become easier to develop and more interoperable than in previous years and the amount of research and development in this area is increasing dramatically. Research efforts such as CoABS[7] and Agentcities[8] are working to link large numbers of heterogeneous agents. This is bringing a renewed focus within multi-agent research community on the development of explicitly modelled ontologies and tools and the infrastructure to support their use.

At the same time, industry uptake of architectures for distributed computing – such as CORBA,[9] J2EE,[10] NET[11] and Web Services[12] – has stimulated the development of business standards for

[1] http://www-ksl-svc.stanford.edu:5915/doc/ontology-server-projects.html.
[2] http://logic.stanford.edu/kif/dpans.html.
[3] http://grouper.ieee.org/groups/suo/KIF/.
[4] http://www.fipa.org/.
[5] http://www.cs.umbc.edu/kqml/software/.
[6] http://www.fipa.org/resources/livesystems.html.
[7] http://coabs.globalinfotek.com/.
[8] http://www.agentcities.org/.
[9] http://www.omg.org/.
[10] http://java.sun.com/j2ee/.
[11] http://www.microsoft.com/NET/.
[12] http://www.w3.org/2002/ws/.

https://doi.org/10.1017/S0269888902000310 Published online by Cambridge University Press

modelling application domains and services to facilitate information flow between distributed and heterogeneous system components. Technologies such as UDDI,[13] WSDL,[14] ebXML,[15] BizTalk[16] and the OMG Model Driven Architecture[17] address many of the same problems as multi-agent systems, and their emergence signals that the development of conceptual domain models (whether they be called ontologies, schemas or data models) is now an important aspect of business computing. Furthermore, these industry initiatives provide both an opportunity and a challenge to the agent research community to demonstrate the additional benefits that their notions of ontology and knowledge-level communication bring and the potential business applications of their technology.

Another recent development has been increased activity towards the creation of the Semantic Web[18]. Formally defined ontologies are a key component of the Semantic Web, which seeks to enhance the potential for automated search, information integration and the location and composition of services on the Web by including machine-readable information in Web resources. This work also addresses some of the key concerns of multi-agent system architectures.

Against this background, research on the use of ontologies in agent systems has gained rapidly in importance over the past two to three years, corresponding to a convergence between research on explicitly modelled ontologies and the systems that are expected to use them. However, researchers looking for an appropriate forum to discuss their work in this area have had to choose between conferences and workshops with ontology-related themes and those addressing aspects of agent-based systems. To provide a more focused discussion for researchers working in the intersection of these areas the Workshop on Ontologies in Agent Systems (OAS)[19] (Cranefield *et al*., 2001) was held as part of the Fifth International Conference on Autonomous Agents in Montreal, Canada. This workshop addressed issues such as the comparison of different approaches to the representation in ontologies for agent systems, the practical considerations of designing applications using these techniques and the infrastructural support required for their effective use.

The motivation for this special issue came out of the discussions at the 2001 workshop and the clearly growing relevance of the research area. Many of the recurring themes at the workshop touched upon domains and problems that are quite familiar to the knowledge representation community but with a particular focus: enabling meaningful interactions between systems in open and possibly highly dynamic environments. Some of the issues raised include the following:

- How much of the semantics of a domain needs to be explicitly encoded and how much can be implicit in an agent implementation?
- How crucial and how practical is inference using ontological information and what trade-offs should be made between expressiveness and tractability?
- What is the relationship between ontologies and agent communication and content languages?
- How do we achieve ‘ontological critical mass’ and how can we generate enough ontologies and encourage their use?
- Can we expect ontologies to be consistent; if not, how can agents deal with the inconsistency?
- Ontologies will change over time (particularly in a dynamic agent world). How will agents deal with this change?
- There is a need for industrial use cases to facilitate technology transfer. These should focus on realistic agent applications and not to try to solve the general information integration problem.

Although these issues are far from being resolved at present, there was general agreement on the following three conclusions about the current state of the art and immediate needs for the future:

[13] http://www.uddi.org/.
[14] http://www.w3.org/TR/wsdl.
[15] http://www.ebxml.org/.
[16] http://www.biztalk.org/.
[17] http://www.omg.org/mda/.
[18] http://www.w3.org/2001/sw/.
[19] http://autonomousagents.org/2001/oas/.

https://doi.org/10.1017/S0269888902000310 Published online by Cambridge University Press

- There is convergence on the use (and perhaps adaptation) of existing representation languages such as DAML + OIL[20] and UML.[21] Researchers recognise that to encourage the sharing of ontologies, shared formalisms are vitally important.
- Shared ontologies are needed to build realistic and useful agent systems. We need to find ways to publish, post and host these more widely to begin the sharing process. Initially it is expected there will emerge islands of shared terminology and representation, but these may later become inter-linked.
- Reusable services and tools are emerging for the use of ontologies in an agent context. These include editors, agent platform support, translation and hosting services. Such initiatives should become more visible and widely used to encourage reuse and to make it easier for agent researchers to make use of ontologies in their systems.

2 This volume

In this special issue we have combined a sample of research in the area of ontologies and agents with a snapshot of current opinion on the tools, applications, opportunities and challenges in the field. We begin by including three articles that are extended versions of papers presented at the OAS workshop. They represent current research in three very different areas of the field:

- Bailin and Truszkowski address the problem of ontology negotiation between agents that do not share the same ontology or perhaps even the same conceptualisation of a domain. They present an overview of different methods for resolving ontology mismatches and then, focusing on inter-agent dialogues for information search tasks, describe a protocol allowing agents to identify and resolve ontological mismatches.
- Cranefield and Purvis use the Unified Modeling Language (UML) to investigate the relationship between ontologies and content languages used in agent communication. In particular, they examine a current trend for using object structures in the form of ontology-specific expressions within messages, and suggest that this can be understood as the use of an ontology-specific content language that can be automatically generated from the ontology. The article discusses some requirements this places on ontology representation languages and presents a UML profile that includes these features.
- Tamma and Bench-Capon propose an extension of a frame-based knowledge model to allow the representation of additional information about the properties describing a concept. This information includes a characterisation of the way each attribute is expected to change over time, a qualitative description of the likelihood of each property of a frame to be inherited by specialisations of that frame, and an indication of which properties are considered to be prototypical and exceptional. It is argued that this information will assist with the task of integrating diverse ontologies, which is crucial in multi-agent systems.

In addition to these three extended workshop papers this volume includes a series of invited short contributions from leading researchers. These articles provide visions or updates on the state of the art in particular focused areas of the domain. The first two articles consider particular models or technologies:

- While many people associate ontologies with traditional knowledge representation languages based on frames or description logic, or with the Semantic Web languages RDF and DAML + OIL, there are other formalisms that could be used to express ontological knowledge for the Web and other applications. Kogut *et al*. outline the advantages of applying the Object Management Group's widely used Unified Modeling Language (UML) to ontology design and sharing and discuss current work in this area.

[20] `http://www.daml.org/2001/03/reference.html`.
[21] `http://www.omg.org/uml/`.

https://doi.org/10.1017/S0269888902000310 Published online by Cambridge University Press

- Pease and Niles present a view of work to date on the IEEE Standard Upper Ontology. This project is an ongoing effort to design a standard high-level model of generic concepts that will be relevant to a broad range of domains. In particular, the article gives an overview of one ontology under discussion by the IEEE SUO Working Group: the Suggested Upper Merged Ontology.

Three authors address possible applications of ontologies and agents, discussing opportunities and challenges in these areas:

- Mayfield draws the analogy between current efforts to define semantics for application domains through ontologies and the use of thesauri in natural language applications. In his article he also identifies three key challenges in applying ontology based techniques to text retrieval: automatic generation of markup, ontology alignment and dealing with unstructured and semi-structured queries.
- Research and industry efforts in the life sciences deal with very large numbers of diverse data sources and processing techniques. Organisations in this field are amongst the earliest adopters of ontology techniques and some are beginning to see the potential of applying it in combination with agent technology. In his article, McEntire provides an insight into ongoing work in this field.
- Labrou's article discusses the application of agent and ontology technologies to enterprise systems. In particular, the article identifies the area of semantic integration (alignment of systems based on diverse domain models) as one of the key areas of opportunity for agent and ontology technology.

The final two articles consider the nature of semantics and ontology in the context of automated systems. Each author outlines elements of a research agenda and emphasises some of the challenges ahead:

- Uschold's article discusses the meaning of the term "semantics" in the context of the Semantic Web and highlights how the semantics in a system, an interaction or an environment can be encoded in different ways with varying degrees of formality. He also argues that progress along the so-called "semantic continuum" is likely to be gradual and must be motivated by application pull as well as technology push.
- In the final contribution to the collection, van Harmelen outlines some of the assumptions often made in traditional knowledge representation systems that need to be revisited in the new world of the Semantic Web.

3 Conclusions

We hope that this collection provides a useful primer for what is likely to be an active research area for some time to come. For future developments in the area we invite the reader to consult the proceedings of two workshops occurring this year: a second workshop on Ontologies in Agent Systems[22] to be held in July at the 1st International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS, 2002), Bologna, Italy, and the Workshop on Ontologies for Multi-Agent Systems[23] to be held in October at the 13th International Conference on Knowledge Engineering and Knowledge Management (EKAW, 2002), Siguenza, Spain.

The authors gratefully acknowledge the Ontologies in Agent Systems workshop programme committee and additional reviewers (listed below) as well as the invited speakers (Sheila McIlraith, Michael Uschold, Natalya Noy and Dominic Greenwood), authors of submitted papers, and other participants for making the workshop a success and thereby making this special issue possible. We also thank the Autonomous Agents 2001 Workshop Chair, Adele Howe, for her help with the workshop and this issue.

[22] http://autonomousagents.org/2002/oas/.
[23] http://www.csc.liv.ac.uk/~valli/OMAS-CFP.html.

https://doi.org/10.1017/S0269888902000310 Published online by Cambridge University Press

OAS Workshop Organising Committee
Stephen Cranefield, University of Otago, New Zealand
Tim Finin, University of Maryland, Baltimore County, USA
Steven Willmott, Ecole Polytechnique Fédérale de Lausanne, Switzerland

OAS Workshop Programme Committee
Federico Bergenti, Università Degli Studi di Parma, Italy
Monique Calisti, Ecole Polytechnique Fédérale de Lausanne, Switzerland
Patricia Charlton, Motorola Research, France
Ulises Cortes, Universitat Politècnica de Catalunya, Spain
Stefan Decker, Stanford University, USA
Stefan Haustein, Universität Dortmund, Germany
Pat Hayes, University of Western Florida, USA
James Hendler, DARPA/University of Maryland at College Park, USA
Noriaki Izumi, Shizuoka University, Japan
Matthias Klusch, Deutsche Forschungszentrum für Künstliche Intelligenz (DFKI), Germany
Yannis Labrou, Powermarket.com, USA
Frank McCabe, Fujitsu Laboratories of America, USA
Ryusuke Masuoka, Fujitsu Laboratories of America, USA
Martin Purvis, University of Otago, New Zealand
Frank van Harmelen, Vrije Universiteit Amsterdam, The Netherlands

Additional Reviewers
Michel Klein, Vrije Universiteit Amsterdam, The Netherlands
Jeen Broekstra, Vrije Universiteit Amsterdam, The Netherlands
Luis Botelno, Associação para o Desenvolvimento das Telecomunicações e Técnicas de Informática, Portugal

References

Cranefield, S, Finin, T and Willmott, S (eds), 2001, *Proceedings of the Workshop on Ontologies on Agent Systems, 5th International Conference on Autonomous Agents*, Volume 52 of *CEUR Workshop Proceedings*; http://CEUR-WS.org/Vol-52.

Genesereth, MR and Ketchpel, SP, 1994, "Software agents" *Communications of the ACM* **37**(7) 48–53.

Neches, R, Fikes, R, Finin, T, Gruber, T, Patil, R, Senator, T and Swartout, W, 1991, "Enabling technology for knowledge sharing" *AI Magazine* **12**(3) 36–56.

https://doi.org/10.1017/S0269888902000310 Published online by Cambridge University Press

https://doi.org/10.1017/S0269888902000310 Published online by Cambridge University Press