Editor: Terry R. Payne
University of Southampton
trp@ecs.soton.ac.uk

Is a Semantic Web Agent a Knowledge-Savvy Agent?

Valentina Tamma, *University of Liverpool*
Terry R. Payne, *University of Southampton*

The issue of knowledge sharing has permeated the field of distributed AI and, in particular, its successor, multiagent systems. Through the years, many research and engineering efforts have tackled the problem of encoding and sharing knowledge without the need for a single, centralized knowledge base. The proposed solutions to this problem are based, more or less, on stringent assumptions, such as static, shared ontological models, or the existence of a common blackboard (or "Linda Space") environment where entities can share knowledge.

However, the uptake of the World Wide Web and the emergence of modern computing paradigms, such as distributed, open systems, have highlighted the importance of sharing distributed and heterogeneous knowledge on a larger scale—possibly on the scale of the Internet.

Interoperability

Research frameworks such as the Knowledge Sharing Effort (KSE)[1] previously focused on sharing heterogeneous knowledge between "knowledge-savvy" entities. While software agents weren't explicitly part of the initial KSE vocabulary, the problem of addressing interoperability at different levels of granularity (syntactic, semantic, and semiotic) soon became very relevant within the multiagent-system paradigm. (We use the term *agents* for those entities that possess and use knowledge of the environment and a given domain of expertise to make decisions.[2]) In the KSE, interoperability was based on

- the definition of a common model that abstracted the shared knowledge (ontology);
- a formalism for representing such knowledge (Knowledge Interchange Format, or KIF); and
- a transport-independent mechanism for querying and communicating this knowledge (Knowledge Query Manipulation Language, or KQML[3]).

Thus, ontologies provided a means of achieving semantic interoperability between different applications committed to different knowledge bases, and KIF provided a way to overcome syntactic differences in the way the applications represented knowledge.

The Semantic Web manifesto, which originally appeared in *Scientific American* in 2001, took the same model of interoperability a step further.[4] This vision was enticing and seemingly easy to follow:

- represent knowledge formally and explicitly by means of ontologies,
- use them to annotate data sources with semantic markup, and
- use software agents to process the resulting markup and reason over this knowledge.

Thus, agents can infer new facts, link scattered but conceptually related data, and provide higher-level algorithms necessary for managing complex systems.

However, while the vision appeared promising, realizing this dream has been highly challenging. Some analyses have tried to understand why the vision has only partially been fulfilled.[5] Others have tried to refine and reduce in scope the original vision, to better align the challenges and achievements to the current solutions for global information and knowledge management.[6] Many have questioned whether large-scale, agent-based mediation was an important factor in the original Semantic Web manifesto, and some conclude that "agents can only flourish when standards are well established," suggesting that the provision of Web standards for expressing shared meaning was more than sufficient, having progressed steadily over the past five years.[6] However, although such standards (including those for Web services) have facilitated the emergence and adoption of service-oriented methodologies, these advances have yet to address how we can share and collaboratively exploit distributed, heterogeneous knowledge in a scalable way.

Autonomy

The very characteristics that define the Semantic Web—

1541-1672/08/$25.00 © 2008 IEEE
Published by the IEEE Computer Society

that is, dynamic, distributed, incomplete, and uncertain knowledge—suggest the need for autonomy in distributed software systems. Autonomy emerges owing to the scale and heterogeneity of the knowledge and service environment. Different stakeholders with varying aims and objectives will own the entities in an open system (such as e-commerce or information services). In addition, the stakeholders might evolve these objectives and change their services' offerings over time. Although services providing software systems might not initially be autonomous, expecting users to have direct control over their behavior is unrealistic as the number and diversity of such systems increase. As ontologies evolve, it's also unrealistic to expect a business entity to suspend transactions with other partners while a human knowledge engineer revises some model of a dynamic domain. So, behaviors that support decision making and adaptation will ultimately have to occur autonomously.

Likewise, the variety and number of disparate stakeholders will result in a corresponding diversity in the way the software entities differ and in the way they each model some common domain. Jerome Euzenat and Pavel Shvaiko have already identified the extent to which knowledge engineers can model heterogeneous knowledge in open knowledge systems.[7] This type of heterogeneity can range from terminological to semiotic and can be represented, for example, using different syntaxes and different representation formalisms.

Decentralized software agents can benefit from exploiting knowledge represented through the Semantic Web paradigm. In a previous discussion of the similarity between autonomous agents and Web services,[8] Terry Payne argued that autonomy can emerge only through the intelligent management and evolution of knowledge (including ontological and epistemological knowledge) of the environment and its knowledge resources. Facilitating the integration of such knowledge with mechanisms that can intelligently exploit it to support collaborative behavior will be crucial in achieving efficient and scalable knowledge manipulation (including reasoning-based query answering) in large-scale, open environments.[9] In particular, agent-based solutions are well suited to cope with the highly dynamic, uncertain, and large-scale nature of these environments. In this article, we therefore discuss some of the challenges that truly autonomous and rational agents face to provide semantic-middleware mechanisms that can scale to the size of open systems and can autonomously and dynamically handle heterogeneity and uncertainty.

Emerging challenges

The challenges faced by *knowledge-savvy agents* (agents that make intensive use of knowledge in open environments) are primarily rooted in the dynamism and heterogeneity of open environments and mainly affect the exchange and evolution of knowledge. These problems can't simply be addressed using traditional AI or knowledge-engineering methods, because many traditional assumptions are no longer valid in this type of environment. For example, the assumption that agents possess complete knowledge of the environment at design time is no longer valid, due to the environment's dynamism and scale. Likewise, assumptions about the global use of a single, consistent ontological model for some domain by all stakeholders are unrealistic owing to the diversity of tasks and stakeholders in that domain.

> The openness and decentralization introduced by multiagent systems give rise to new knowledge management models, about which we can't make a priori assumptions.

Semantic Web research promises more than mere management of ontologies and data through the definition of machine-understandable languages. The openness and decentralization introduced by multiagent systems and service-oriented architectures give rise to new knowledge management models, for which we can't make a priori assumptions about the type of interaction an agent or a service might be engaged in, or about the message protocols and vocabulary used. This opens up communication models where the first step effectively involves establishing whether a transaction between two services or agents is feasible, and whose ontologies are either the same or can be reconciled in part or entirely. The following challenges therefore represent several issues related to the exchange and evolution of knowledge in open environments. We will need to resolve them if agents are to make extensive and successful use of knowledge. While the list of issues is far from exhaustive, it arguably represents some of the more pertinent challenges shared by agents and Semantic Web resources alike.

Discovering resources

Providing mechanisms for discovering a system's agents is a crucial challenge. Discovery of new agents (or of the services they provide) does not only affect methods of interaction. As a consequence of heterogeneity, knowledge about the rules of agent engagement (for example, how to query each other) and the mechanisms for requesting agents' services also require further discovery of the permissible protocols and formalisms known to the agents.

In addition, not all the agents in the system might be able or willing to comply with a request. This can be due to limited resources (for example, in grid systems), overprovisioning (for example, because of high popularity or an imbalance in supply and demand), or simply performing a task that might result in negative utility (that is, the cost of performing that specific task is greater than any utility gained in achieving the task). For these reasons, the behavior of agents in the system can only be considered at a local level, and knowledge about the system in its entirety should be considered only when requested.

Determining ontology identity

Establishing a potential collaborator's presence requires efficient mechanisms of discovery that enable agents or services to become aware of other agents and services in the same environment, and to determine whether some form of communication might be possible. This raises an important issue. Determining whether two ontologies are the same ultimately implies assigning some notion of identity to URIs (universal resource identifiers) and providing identity conditions that let us establish whether two ontologies (that might have been expressed differently and stored in different locations) are effectively the same. This resembles the referential-integrity problem—determining whether the person John Doe referred to by one agent is the same as the employee John Doe known by another.

Dynamic semantic interoperability

Semantic interoperability is grounded in ontology reconciliation: finding relationships between entities belonging to different ontologies. This reconciliation usually relies on the existence of correspondences (or mappings) between different ontologies (ontology alignment), which agents can use to interpret or translate messages exchanged with collaborator agents. The Semantic Web community has been actively pursuing mechanisms to support ontology alignment and matching at design time,[7] but there's still no single, unified framework for this.

In addition, dynamic discovery of new agents at runtime necessitates the subsequent discovery and selection of correspondences to facilitate interpretation and utilization of exchanged messages. To have more realistic mechanisms of knowledge management, where knowledge becomes part of an agent's or a service's assets only when it's beneficial to the agent or service, we face two obstacles: the fact that correspondence selection mechanisms depend on the task for which ontologies must be matched, and the need to define and take into account some notion of utility (as the agent community defines it). A small body of work has tried to address this problem by assuming that agents can exploit mapping repositories when negotiating for a mutually acceptable set of ontology alignments.[10,11]

Dynamic evolution of agent ontologies

Assumptions that agents are equipped with a priori knowledge of the environment (including knowledge of other agents and the ontologies they use) no longer hold when we consider the Web's inherent dynamic and distributed nature. Therefore, to relax this assumption, we need novel mechanisms for discovering the other agents' presence (as discussed earlier). A large body of work has to date been published on discovering new agents on the basis of their capabilities. However, an agent can engage in conversation with these newly discovered agents in an open environment only if they can reconcile their vocabularies first. In this type of scenario, the idea of agents committing to a single, monolithic, shared ontology is highly impractical. Analogously, a large body of work has looked at the problem of ontology evolution from the engineering perspective,[12] but very few efforts have looked at how ontologies dynamically evolve through their use. The notion of semantics dynamically emerging through a large number of local, peer-wise interactions has attracted recent attention;[13] semantic interoperability results when interacting agents converge on

> Recently the Semantic Web community has focused on defining a standard notion of uncertainty; a first step in this endeavor is an ontology of uncertainty.

a set of mutually accepted statements. Although some researchers propose that the notion of emergent semantics can facilitate the emergence of a consensual ontology that enables information integration among (human and software) agents, this notion is particularly relevant for multiagent systems in distributed and possibly large-scale environments. To realize this emergence, recent approaches have started to look at estimating the cost of changing an ontology before this occurs.[14,15]

Ontologies for describing dialogues and protocols

Agents typically need to model and exchange knowledge about interaction protocols, obligations, and norms within agent communities, as well as the constraints regulating these interactions. Although a few research efforts have proposed a number of representational formalisms to describe dialogues or negotiation frameworks, in most cases these mechanisms are hard coded within the participating agents. For example, the operational characteristics and protocols assumed in auctions (including the type of auction, rules, social norms, and obligations) are hard coded within the participants, making such marketplaces closed to other agents. These constraints are overly restrictive when agents are in open environments such as the Web. These environments are characterized by flexible, dynamic scenarios, where agents search the Internet for suitable marketplaces to sell or buy goods, and the interaction rules that govern the way agents communicate with the marketplace can change within an interaction or between interactions. Open environments require that agents are free to engage in interactions; the only prerequisite is that the agents must share some background knowledge and commit to some common rules of encounter. Assuming a fixed and immutable interaction mechanism is no longer viable, these environments need representations and models that support the sharing of interaction mechanisms.

Representing and reasoning with uncertain information

Agents interacting in open environments must deal with uncertain information coming from a multiplicity of sources. Uncertainty is intrinsic in the World Wide Web (and thus in the Semantic Web): resources publish statements independently of whether other (similar) statements exist, and contradictory statements coexist on the Web. In this type of scenario, agents must be able to resolve contradictions if they need to integrate conflicting viewpoints during an interaction. In some cases—for example, in service provision—service requesters can use the level of trust and the reputation of an information source to assign a stronger degree of preference to the statements they publish. In addition to uncertainty about a knowledge source's reliability and the consequences inferred by distributed statements, agents only possess partial knowledge about the environment they operate in. So, agents must be able to reason with what they know about other agents and, in particular, what they know that other agents know and what they can assume is known. Recently the Semantic Web community has focused on defining a standard notion of uncertainty (see www.w3.org/2005/Incubator/urw3/XGR-urw3-

20080331 for more details); a first step in this endeavor is an ontology of uncertainty that captures the different causes and effects of uncertainty.

The challenges we've identified are inexorably linked to the characteristic of open environments. They require the agent, the Semantic Web, and the knowledge community to rethink some of the approaches so far assumed. Projects such as LarkC are investigating reasoning mechanisms for large-scale and uncertain environments, where human intervention is neither feasible nor practical in most cases.[9] In these types of situations, the agents' ability to acquire new capabilities and therefore to achieve new tasks (or answer new queries, in the case of knowledge-based systems) must be offset by the cost of the change in terms of employing scarce resources,[9] and with partial knowledge of the environment (typical of distributed, open systems[8]). When the change concerns agent ontologies, agents must be able to deliberate rationally regarding the usefulness of that change in order to cope with the complexity of the environment and the modifications required. The notion of bounded rationality will become prominent because it assumes that the agent's decision-making process is optimized to work with partial knowledge and limited computational resources.

The research challenges we've illustrated will benefit from a closer collaboration between the multiagent-systems and Semantic Web communities. The evolution of standards and data integration alone can't be the answer. ■

References

1. R. Neches et al., "Enabling Technology for Knowledge Sharing," *AI Magazine*, vol. 12, no. 3, 1991, pp. 36–56.
2. M. Wooldridge and N. Jennings, "Intelligent Agents: Theory and Practice," *Knowledge Engineering Rev.*, vol. 10, no. 2, 1995, pp. 115–152.
3. T. Finin et al., *Specification of the KQML Agent Communication Language*, DARPA Knowledge Sharing Initiative External Interfaces Working Group, 1993.
4. T. Berners-Lee, J. Hendler, and O. Lassila, "The Semantic Web," *Scientific American*, May 2001, pp. 35–43.
5. J. Hendler, "Introduction to the Special Issue: AI, Agents, and the Web," *IEEE Intelligent Systems*, vol. 21, no. 1, 2006, p. 11.
6. N. Shadbolt, T. Berners-Lee, and W. Hall, "The Semantic Web Revisited," *IEEE Intelligent Systems*, vol. 21, no. 3, 2006, pp. 96–101.
7. J. Euzenat and P. Shvaiko, *Ontology Matching*, Springer, 2007.
8. T.R. Payne, "Web Services from an Agent Perspective," *IEEE Intelligent Systems*, vol. 23, no. 2, 2008, pp. 12–14.
9. D. Fensel and F. van Harmelen, "Unifying Reasoning and Search to Web Scale," *IEEE Internet Computing*, vol. 11, no. 2, 2007, pp. 96, 94–95.
10. L. Laera et al., "Reaching Agreement over Ontology Alignments," *Proc. 5th Int'l Semantic Web Conf.* (ISWC 06), Springer, 2006, pp. 371–384.
11. C. Trojahn, P. Quaresma, and R. Vieira, "An Extended Value-Based Argumentation Framework for Ontology Mapping with Confidence Degrees," *Argumentation in Multi-Agent Systems: 4th Int'l Workshop*, LNAI 4946, Springer, 2007, pp. 69–81.
12. N.F. Noy and M.C.A. Klein, "Ontology Evolution: Not the Same as Schema Evolution," *Knowledge Information Systems*, vol. 6, no. 4, 2004, pp. 428–440.
13. K. Aberer et al., "Emergent Semantics Principles and Issues," *Proc. 9th Int'l Conf. Database Systems for Advanced Applications* (DASFAA 04), LNCS 2973, Springer, 2004, pp. 25–38.
14. I. Palmisano et al., "Dynamic Ontology Evolution in Open Environments," tech. report ULCS-08-012, Univ. of Liverpool, 2008.
15. M. Sensoy and P. Yolum, "A Cooperation-Based Approach for Evolution of Service Ontologies," *Proc. 7th Int'l Conf. Autonomous Agents and Multiagent Systems* (AAMAS 08), Int'l Foundation for Multiagent Systems, 2008, pp. 837–844.

Valentina Tamma is a lecturer in the Department of Computer Science at the University of Liverpool. Contact her at v.tamma@liverpool.ac.uk.

Terry R. Payne is a lecturer in the University of Southampton's School of Electronics and Computer Science. Contact him at trp@ecs.soton.ac.uk.