A Semantic Future for AI

Rudi Studer, *Institute AIFB, University of Karlsruhe; FZI Research Center for Information Technologies; and Ontoprise GmBH*
Anupriya Ankolekar, Pascal Hitzler, and York Sure, *Institute AIFB, University of Karlsruhe*

In our modern information society, people need to manage ever-increasing numbers of personal devices and conduct more of their work and activities online, often using heterogeneous services. The amount of information each individual has to process is constantly growing, making this information increasingly difficult to control, channel, share, and constructively use. To mitigate this, computing must become much more human centered—for example, by presenting personalized information to users and by respecting personal preferences when controlling multiple devices or invoking various services. Appropriate representation of the information's semantics and of the functionality of devices and services will be critical to such personalized computing.

Symbolic artificial intelligence techniques provide the method of choice for the required semantic-representation and reasoning capabilities. The challenge for symbolic AI is to support large-scale, distributed, dynamic knowledge bases enabling highly adaptive and evolving systems. AI must also look to specific application contexts and develop real-world solutions for problems in those domains. Here, we present some examples of such application contexts.

Collaboration in communities

As online communities proliferate within intranets and the Internet, people—bound by common identities and interests—share experiences and knowledge through mailing lists, discussion forums, and wikis, and even collaborate to construct artifacts, such as documents and software. However, the large number of dynamic interactions can be overwhelming, particularly for a community's new members.

For example, assume Maria, a budding photographer, wants to buy a digital camera. Searching the Web for the best digital camera for amateurs, she has just found a community that's been holding an ongoing discussion about digital cameras for different kinds of users. Now imagine that the underlying Web has become an intelligent medium and repository for collaboration. The Web automatically processes each contribution from the community, organizing it meaningfully in the current set of contributions, such that the community can see the current state of contributions. In this case, in addition to the textual discussion, Maria can see an automatically generated table of the digital cameras mentioned so far. Entries in the table indicate each camera's key advantages and disadvantages as well as which type of user it's best suited for. This lets Maria quickly compare the cameras and judge how comprehensive the discussion has been. In addition, each piece of information serves as a link to the discussion that generated it, letting Maria easily verify the point being made.

This kind of view can also be customized for particular people, roles, and tasks in the community. Maria, for example, as a newcomer, might require more historical information on people's contributions than would a regular in the community, who already knows the community members, their background, and their opinions. By making it easier for community members to judge the scope and quality of the contributions as a whole before they contribute themselves, the automatic organization of information can also raise the contributions' quality.

The realization of this scenario requires explicit representation of the meaning of a community's content, structure, and interactions. The underlying Web infrastructure must be able to extract semantics from the multimedia documents constituting the community interactions. In addition, the Web must be able to reason over a large knowledge base of user contributions, to aggregate and merge knowledge, some of which could stem from multiple communities. A key challenge will be to deal with information that's globally inconsistent.

Intelligent personalized environments

Personal devices, such as *mobile digital assistants*, increasingly contain information about a person's schedules and preferences. As physical environments become increasingly augmented with embedded computing power and RFID tags, personal devices can use ambient services to provide unprecedented personalization of the user's space and to support the user's activities.

1541-1672/06/$20.00 © 2006 IEEE
Published by the IEEE Computer Society

For example, assume Stephan wants to organize a party. Stephan's MDA has his schedule, and it contacts his friends' MDAs, negotiating with them to schedule an evening convenient for all. Furthermore, Stephan has told his MDA that a party typically involves food and music. The MDA invokes the refrigerator service, checking whether the refrigerator is well stocked with the food Stephan typically orders for parties. The fridge notes that some essential items are missing and orders them from the online supermarket. During the party, Stephan's MDA selects the music and queries the guests' MDAs for their preferences, taking them into account when selecting the next track to play through the music system service. Noting that fewer guests have arrived than expected, Stephan's MDA identifies a couple of friends who might be interested in joining the party and checks with Stephan before sending them an impromptu invitation.

To realize this scenario, Stephan's MDA must know about Stephan's schedules, food preferences, and typical behavior as well as keep track of his current context—for example, at the party. In addition, his home's environment must provide ambient services that can be invoked, such as the refrigerator and music system services. Furthermore, these services must be organized within a service-oriented architecture that enables self-organization of ambient services to support the user's goals and activities. This also requires appropriate knowledge representation of activities, a person's context, user preference policies, and rules that determine actions to take within the context. The MDA must be able to learn Stephan's preferences from his behavior and use them to automatically adapt to the changing context. Much of this knowledge will be fuzzy or probabilistic, so the MDA needs corresponding reasoning capabilities.

Market coalitions

Semantic descriptions of desired products and services will enable better matching of buyers and sellers and enable dynamic coalitions of customers. Imagine that Stephan's refrigerator wishes to purchase snacks from an online neighborhood supermarket. Several other customers in the vicinity might similarly be looking to purchase various kinds of groceries. If each of these customers' requests is semantically described with respect to a grocery item ontology, Stephan's refrigerator can identify those customers and propose a coalition to procure purchase or delivery discounts from the supermarket. So, because Stephan's refrigerator knows that crisps are a kind of snack, it can join a dynamic coalition of customers purchasing crisps.

Such dynamic coalitions will become increasingly common in the near future and, as we discussed, will require a semantic representation of the products as well as a framework for comparing product requests and offers. In addition, there are interdisciplinary technical challenges, involving economics and law, in developing techniques both for dynamic coalition building for multiple products and for auctions involving buyers and sellers with partially matching requests and offers.

The representation of complex information and the reasoning over that information's semantics are key to realizing the scenarios we've just described—in particular, to enable generic, flexible, self-organizing solutions as opposed to hard-wired, customized systems. However, we need to address several technical challenges before we can successfully apply knowledge representation in these contexts.

To begin with, to extract knowledge semiautomatically from the myriad information sources and from the content of users' interactions in an environment, we need robust, reliable knowledge acquisition. Furthermore, knowledge acquisition must be able to learn from users' behavior and represent the obtained knowledge in an appropriate form for reasoning. Another challenge is to manage the expressivity-scalability trade-off of reasoning over declarative knowledge, enabling reasoning over large-scale distributed knowledge bases for suitably expressive knowledge representations. Automated knowledge acquisition will typically yield knowledge that's uncertain—for example, fuzzy or probabilistic. Such knowledge must be represented and reasoned with in an adequate and scalable way. As knowledge from distributed knowledge bases is aggregated, a deeper semantics can emerge, letting intelligent agents discover patterns across people, roles, and tasks.

Several European and German projects are taking initial steps to address these challenges—for example,

- NeOn (Lifecycle Support for Networked Ontologies, www.neon-project.org), which explores the merging of distributed, evolving, and context-specific knowledge bases;
- Nepomuk (Networked Environment for Personalized, Ontology-Based Management of Unified Knowledge—The Social Semantic Desktop, nepomuk.semanticdesktop.org), in the context of collaborative work; and
- SmartWeb (www.smartweb-project.org), for access to an intelligent Web via mobile devices.

However, addressing these challenges in their full complexity will require a persistent, concerted effort from the AI community. ■

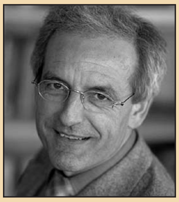

Rudi Studer is the head of the Knowledge Management research group at the Institute AIFB, University of Karlsruhe. Contact him at studer@aifb.uni-karlsruhe.de.

Anupriya Ankolekar is an assistant professor and a project leader at the Institute AIFB, University of Karlsruhe. Contact her at ankolekar@aifb.uni-karlsruhe.de.

Pascal Hitzler is an assistant professor and a project leader at the Institute AIFB, University of Karlsruhe. Contact him at hitzler@aifb.uni-karlsruhe.de.

York Sure is an assistant professor and a project leader at the Institute AIFB, University of Karlsruhe. Contact him at sure@aifb.uni-karlsruhe.de.