

A semantic future for AI

Rudi Studer

Institute AIFB, University of Karlsruhe
FZI Research Center for Information Technologies
Karlsruhe, Germany
`{studer}@aifb.uni-karlsruhe.de`

Anupriya Ankolekar, Pascal Hitzler and York Sure
Institute AIFB, University of Karlsruhe
Karlsruhe, Germany
`{ankolekar, hitzler, sure}@aifb.uni-karlsruhe.de`

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 making use of heterogeneous services. The amount of information to be processed by each individual is constantly growing, making it increasingly difficult to control, channel, share and make constructive use of it. To mitigate this, computing needs to become much more human-centred, e.g. by presenting personalised information to users and by respecting personal preferences in controlling multiple devices or invoking various services. Appropriate representation of the semantics of the information and functionality of devices and services will be critical to such personalised computing. Symbolic artificial intelligence (AI) techniques provide the method of choice for the required semantic representation and reasoning capabilities. The challenge for symbolic AI is to be able 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. Below, 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, discuss experiences and share 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 new members of a community. For example, assume Maria, a budding photographer, is looking to buy a digital camera. Searching the Web for the best digital camera for amateurs, she has just found a community where a lengthy discussion has been taking place about digital cameras for different kinds of users. Now imagine the underlying Web could become an intelligent medium and store of collaboration, automatically processing each contribution from the community, organising it meaningfully within the current set of contributions, such that the community

can see the current state of the contributions. In this case, in addition to the textual discussion, Maria would see an automatically-generated table with the digital cameras mentioned so far. Each entry in the table would indicate some of the key advantages and disadvantages of the camera as well as indicating which type of user, it is best suited for. This would allow Maria to quickly compare the different cameras, as well as judge how comprehensive the discussion has been. In addition, each of the above information would be present as a link to the actual discussion that generated that information, making it easy for Maria to verify for herself, the actual point being made. This kind of view may also be customised for particular people, roles and tasks within the community. Maria, for example, as a newcomer to the community, might require more historical information on people's contributions than a regular in the community, who already knows the people in the community and their background and 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 organisation of information can also raise the quality of the contributions.

The realisation of this scenario requires explicit representation of the meaning of the content, structure and interactions of a community. The underlying Web infrastructure must be able to extract semantics from the multimedia documents that comprise the community interactions. In addition, the Web must be able to reason over a large knowledge of user contributions, to aggregate and merge knowledge, potentially stemming from multiple communities. A key challenge will be in dealing with information that is globally inconsistent.

Intelligent Personalised Environments

Personal devices, such as mobile digital assistants (MDA), increasingly contain information about a person's schedules and preferences. As physical environments are increasingly augmented with embedded computing power and RFID tags, personal devices can make use of ambient services to provide unprecedented personalisation of the user's space and support for the user's activities. For example, assume Stephan wants to organise a party. Stephan's MDA has Stephan's schedule and contacts the MDAs of his friends, negotiating with them to schedule an evening convenient for all. Further, Stephan has told his MDA that a party typically involves food and music. It 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 takes charge of selecting the music and queries the guests' MDAs for their music preferences, taking them into account when selecting the next track to play through the music system service. Noting that fewer than expected guests have arrived, Stephan's MDA identifies a couple of friends who might be interested and checks with Stephan before sending them an impromptu invitation.

To realise this scenario, Stephan's MDA needs to know about Stephan's schedules, food preferences and typical behaviour as well as keeping track of his current context, for example, at the party. In addition, the environment of his home must provide ambient services that can be invoked, such as the refrigerator and music system services. Furthermore, the services must be organised in a service-oriented architecture that enables self-organisation of ambient services to support the user's goals and activities. This also requires appropriate knowledge

representation of activities, the context of a person, user preference policies and rules that determine actions to be taken within the context. The MDA must be able to learn Stephan's preferences from his behaviour and use these to automatically adapt to the changing context. Much of this knowledge will be fuzzy or probabilistic in nature and corresponding reasoning capabilities are required.

Market Coalitions

Semantic descriptions of desired products and services will enable better matching of buyers and sellers as well as enabling dynamic coalitions of customers. Imagine, for example, in the previous scenario, that Stephan's refrigerator wishes to purchase snacks from a neighbourhood online supermarket. Similarly, there may be several other customers looking to purchase various kinds of groceries in the vicinity. If the requests of each of these customers is semantically described with respect to an ontology of grocery items, Stephan's refrigerator can identify those customers looking for similar items and propose a coalition to procure purchase or delivery discounts from the supermarket. Thus, since Stephan's refrigerator knows that crisps are a kind of snacks, it can join a dynamic coalition of customers purchasing crisps.

Such dynamic coalitions will become increasingly common in the near future and, as 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 the fields of economics and law, in developing techniques for dynamic coalition building for multiple products and auctions for buyers and sellers with partially matching requests and offers.

The representation of complex information and the reasoning over the *semantics* of such information is key to realising the scenarios described above, in particular to enable generic, flexible, self-organising solutions as opposed to hard-wired, customised systems. However, several technical challenges must be addressed before knowledge representation can be applied successfully in these contexts. To begin with, *robust and reliable knowledge acquisition* is required to extract knowledge semi-automatically from the myriad information sources and from the content of users' interactions within an environment. Furthermore, knowledge acquisition must be able to learn from users' behaviour, and to represent the knowledge obtained in an appropriate form for reasoning. Another challenge is to manage the *expressivity-scalability tradeoff* 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 is uncertain, e.g. 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, allowing intelligent agents to discover patterns across people, roles and tasks.

Initial steps to address these challenges are being taken within several European projects, such as NeOn¹, Nepomuk² in the context of collaborative work

¹Lifecycle Support for Networked Ontologies: <http://www.neon-project.org>

²Networked Environment for Personalized, Ontology-based Management of Unified Knowl-

and SmartWeb³ for access to an intelligent Web via mobile devices. However, it will require a persistent, concerted effort on the part of the artificial intelligence community to address these challenges in their full complexity.

edge - The Social Semantic Desktop: <http://www.nepomuk.semanticdesktop.org>
³<http://www.smartweb-project.org/>