

AAAI-97 Workshop on AI and Knowledge Management

Bradley L. Whitehall

The Workshop on AI and Knowledge Management, held at the Fourteenth National Conference on Artificial Intelligence (AAAI-97), focused on knowledge structure, the interaction of distinct knowledge sources, and the relevance of knowledge to a particular problem. AI research has investigated many of these areas for some time. However, the demands of knowledge management systems place different constraints on the problems. The purpose of the workshop was to explore how AI can contribute to the emerging area of knowledge management.

Knowledge management is concerned with systematically and actively creating, collecting, managing, and leveraging the knowledge and information in an organization. This knowledge often is unstructured, scattered, inconsistent, and incomplete. However, if made available in the proper context, it is relevant to solving business problems. Knowledge management involves people, content, process, and technology. AI research obviously has much to contribute to the technology aspects of knowledge management, but it can also provide insight into the content that might be most beneficial.

As many questions were raised in this workshop as were answered, which is expected in a new area. Many questions centered on how much structure should be imposed on knowledge bases and what the appropriate kinds of structure are. Much of the knowledge in business is in the form of documents that contain little (formal) structure. Forcing

knowledge creators to use a formal structure improves the usability of the knowledge but constrains the creation process. Based on the realization that it is impractical to force such constraints on users, the matter

This article describes a one-day workshop entitled AI and Knowledge Management that was held at the Fourteenth National Conference on Artificial Intelligence. The workshop was successful in identifying areas where AI techniques can be used to help those working on knowledge management and identifying areas for future work in this area.

becomes what the most useful structure is to add. Some participants illustrated how small amounts of structure can provide important benefits in knowledge retrieval.

Other issues centered on evaluating the relevance of knowledge and

determining when to present information to a user. It is clear that people are more open to interruption on some tasks than others, but constant interruption with irrelevant information is not tolerated for long. A related issue is how the user should interact with a knowledge management system. How much pull versus push is appropriate?

An important concern focused on the value proposition for knowledge management systems. Knowledge management systems must provide value for the creators of the knowledge as well as someone else farther down the information food chain who might want to use this information. Designing systems that provide benefit at all stages of the knowledge creation and collection process is not trivial.

The workshop had a panel that provided interesting corporate perspectives on the knowledge management problem. Panelists included Amy Rice (Brightware), Phil Klahr (Inference), and Ramasamy Uthurusamy (General Motors).

The organizing committee for the workshop included Bradley Whitehall (Johnson Controls, Inc.), Michael Gruninger (University of Toronto), Ed Rogan (United Technologies Research Center), Jude Shavlik (University of Wisconsin at Madison), Brock Barkley (Johnson Controls), and Aron Dutta (Silicon Valley Internet Partners).

Bradley L. Whitehall is a principal research engineer in the Research and Development Group at Johnson Controls, Inc. Whitehall received his M.S. and Ph.D. from the University of Illinois at Urbana-Champaign in 1987 and 1990, respectively. His research interests include knowledge management, data mining, collaborative work environments, and the embedding of knowledge into products. His e-mail address is bradley.l.whitehall@ci.com.