# Semantic Networks for Knowledge Representation in an IE

**Stephen Peters** SLP@AI.MIT.EDU
MIT Artificial Intelligence Laboratory, Cambridge, MA, 02139, USA

## 1. Introduction

For many years now, research in intelligent spaces has grown, exploring different ways that a room can react to one or more users and their actions. As usage of these intelligent environments (IEs) grows, however, they will by necessity collect ever-increasing amounts of data about their users, in order to adapt to the user's desires. Information will be collected on the users' interests, who they communicate with, their location, web pages they visit, and numerous other details that we may not even notice. All this information needs to be collected and organized, so that the IE can make quick, correct assumptions about what the user would like to do next.

At the Intelligent Room project (Hanssens et al., 2002), we are beginning to define one such knowledge representation (KR), using semantic networks as the basis for the representation. This creates inherent advantages, both in ease of adding and changing information as well as inference generation.

## 2. Knowledge Representations and Intelligent Environments

Davis et al. (Davis et al., 1993) examined the roles of KRs within intelligent systems, and divided them into several categories, some of which have implications for the role of a KR in an IE:

**They are computational surrogates for the real world.** For an IE, this means that we need to have representations for real-world entities, such as users, projectors, cameras, appliances, as well as more abstract notions like groups of people, roles, actions, etc. This requires that the KR quickly change to accurately track the current state of the real-world entities.

**They impose a filter on the world view.** For an IE, we require many levels of information, since we may want to track people not only as high-level objects associated with a name, email address, and favorite pet, but also as blobs of color moving throughout a camera's visual field, and probably at many levels of abstraction in between. The KR needs to be able to link between these levels as inferences are generated.

**They are a medium for human expression.** A KR is a way of representing our view of the world as well as the computers. Therefore, it needs to be simple to use and turn into representations that humans can parse. Even though IEs can incorporate a great deal of information, editing it needs to be done in as localized a manner as possible, so that updating information requires changes to only a few small, well-defined places.

## 3. Using a Semantic Network in the Intelligent Room

We have begun deploying a KR based on semantic networks (Quillian, 1968) within the Intelligent Room project. Our current implementation works on top of a SQL database storing the semantic networks information, but could be extended to use object-oriented databases or even a Resource Description Framework (RDF) back-end if desired. Metaglue agents (Coen et al., 1999) are able to create, query, and update the semantic network information as the user requires.

While developing this system, we have identified several areas in which using a semantic network-based representation is an appropriate and valuable piece of infrastructure for intelligent environments such as this one, most notably in the areas of user information and meeting management.

### 3.1 User Knowledge

One of the key pieces of knowledge for any intelligent space is that of the users and the individual spaces they work with. At a simplistic level, this can simply be a set of objects comprising spaces, how they are encapsulated within each other, and the user's current location. This gives access to straightforward queries like "where is Steve located" and allows for simple resource management dependent on the task and space involved (Gajos, 2001).

However, in order to make a system that truly acts as an "intelligent assistant", you need to include far more information about people, their relationships, their interests, and

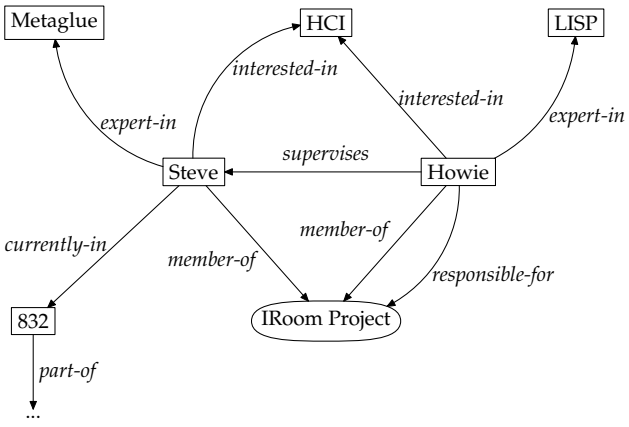


*Figure 1.* Knowledge base for users, interests, expertise, hierarchies, and groups

more. Using the semantic net, the system can respond to queries and requests such as:

- "Send this information to everyone in the group."
- "Send this information to group members interested in HCI."
- "Do I know anyone who is an expert in writing LISP code?"
- "Who do I know who can introduce me to a LISP expert?"

We can also use this knowledge as part of the IE access model, providing for more interesting, situation-dependent control of components. For example, even if all users can control the room's devices most of the time, it may be that special access is set up for meetings – for instance, allowing only the current meeting facilitator to grant access to the room's devices for presenting information.

### 3.2 Meeting Management

We are now using the semantic network structures to capture meetings as they occur, linking together the main meeting topics along with their contributors and attendees. Typical information that gets captured during a meeting includes agenda topics, action items, supporting and dissenting arguments, and documents such as presentations or web references. People are linked in as meeting attendees, document authors, and as issue-raisers. When meetings take place in instrumented environments, they can be linked to a video or audio capture of the meeting in progress.

Using the philosophy that meetings are not the primary piece of information, but merely a framework for examining and disseminating information, the discussion topics within the meeting management software can be linked together. This makes it possible to review proposals as they travel through a long-term set of meetings, and to ask the system questions regarding previous meetings discussing the current topic.

The latest incarnation of our MeetingManager application uses these concepts a great deal, allowing us to capture much of the aforementioned information from a meeting as it occurs.

## 4. Conclusions and Future Work

As we have found, the semantic network buys us a great deal of flexibility in terms of the data we capture and interpret. Because of the nature of the semantic networks, adding and changing information is localized and clear-cut, and inference generation can be done extremely efficiently.

We are exploring different ways to enrich the interactions we can provide to the user. As one example, we would like to introduce a network activation system into the network's searching capabilities. This would allow a user (or an agent operating on the user's behalf) to specify a segment of network that it is looking for, and get notified whenever a matching piece of network is created.

Further work will also likely include integrating the knowledge framework with the Haystack project (Adar et al., 1999).

## References

Adar, E., Karger, D., & Stein, L. (1999). Haystack: Per-user information environments. *Proceedings of the 1999 Conference on Information and Knowledge Management, CIKM*.

Coen, M., Phillips, B., Warshawsky, N., Weisman, L., Peters, S., & Finin, P. (1999). Meeting the computational needs of intelligent environments: The Metaglue system. In P. Nixon, G. Lacey and S. Dobson (Eds.), *Managing interactions in smart environments: 1st international workshop on managing interactions in smart environments (manse '99)*, 201–212. London: Springer-Verlag.

Davis, R., Shrobe, H. E., & Szolovits, P. (1993). What is a knowledge representation? *AI Magazine, 14*, 17–33.

Gajos, K. (2001). Rascal - a resource manager for multi agent systems in smart spaces. *Proceedings of CEEMAS 2001*.

Hanssens, N., Kulkarni, A., Tuchinda, R., & Horton, T. (2002). Building agent-based intelligent workspaces. *ABA Conference Proceedings*. To appear.

Quillian, M. (1968). Semantic memory. In M. Minsky (Ed.), *Semantic information processing*, 227–270. Cambridge, MA: MIT Press.