# Architecting a Spatial Group Memory with Plexus

**Mark Foltz** and **William Neveitt**
{mfoltz,jackal}@ai.mit.edu
MIT Artificial Intelligence Lab
545 Technology Square
Cambridge, MA, USA 02139

**Rebecca Xiong**
becca@graphics.lcs.mit.edu
MIT Laboratory for Computer Science
545 Technology Square
Cambridge, MA, USA 02139

**Abstract**

Supporting the collaborative construction of shared knowledge bases is an exciting design challenge. We present Plexus, which functions as a group memory by recording user contributions in a shared, visual information space.

Plexus uses the notion of a shared virtual blackboard to guide both user contributions and information seeking activity. Users position questions and answers in physical proximity to related items on the blackboard. Interaction with concurrent users is supported via real-time updates of changes to the blackboard and a synchronous chat capability.

Plexus is distinguished from other systems that provide a shared memory (such as AnswerGarden, TeamInfo, and gIBIS (Ackerman 1998; Berlin *et al.* 1993; Conklin & Begeman 1988)) by the pervasive use of the spatial metaphor to guide user interaction. By encouraging users to consider themselves as being at a particular "place," we believe the metaphor enables users to better situate their contributions, find the information they seek, and remain aware of others. The architecture of the space, designed around real-world navigation features like landmarks, paths, and regions (Lynch 1960), further assists navigation and guides contribution to the memory.

To illustrate Plexus we have created OfficeHours, an information space of students' questions and answers about an undergraduate course. OfficeHours provides students with the means to structure their questions, find related answers, and assess their understanding of course material in the context of others. Preliminary user reactions to OfficeHours have been positive.

**Keywords:** human-computer interaction, multimedia, knowledge engineering, group memory, knowledge navigation, knowledge management
**AAAI Tracking Number:** A404

## Introduction

We are interested in exploring the potential of presenting an important artifact of group work, a *group memory*, as a shared, visual information space. A group memory holds shared knowledge that might otherwise be lost because no mechanism was in place to record it.

In our system, Plexus, a group memory is a shared, visual information space in which participants can navigate to locate previously stored knowledge elements and can position new contributions in relation to those elements. The metaphor encourages users to consider themselves as being at a particular "place," and, we believe, enables them to better situate their contributions and to remain aware of others.

Plexus is distinguished from other systems that provide information spaces for group memory (for example, Answer Garden (Ackerman 1998)) by the pervasive use of the spatial metaphor, which we believe supports navigation, contribution, and social awareness.

We begin by characterizing a notion of a "group memory." Next, we list desirable properties of a shared, visual information space, including: *persistence*; *cognitive, visual,* and *temporal coherence*; and *social awareness.* We then argue that these properties are especially appropriate for group memories.

The spatial metaphor alone, however, is not enough, because it does not provide affordances that assist memorability and navigability of the space. In response, we have augmented the spatial metaphor with navigation features, such as landmarks, paths, and regions, which have been shown to improve the navigability and memorability of information spaces (Foltz 1998; Lynch 1960). These features help to solve the problem in group memories that "people can generally find things that are on their own messy desks and file systems, but not on other people's" (Berlin *et al.* 1993).

We illustrate Plexus through OfficeHours, a shared visual information space of answers to students' questions about an undergraduate course.

### What is a Group Memory?

People keep personal memories in computers, in hardcopy, and in their heads. In many circumstances, knowledge stored in one person's memory that is useful to other members of a group is shared opportunistically, through one-to-one or one-to-many communication. If the medium permits it, the recipient may archive a selected portion according to a personal system of organization. However, this model of knowledge sharing has several disadvantages:

- Group knowledge remains with the sender and receiver, and is not accessible to other members of the group.
- Group knowledge is poorly indexed. Interesting relationships between one person's archive and another's typically remain undiscovered.
- A new member of the group must ask many individuals to learn what the group knows.
- When a member of the group leaves, important unrecorded knowledge leaves with her.

In response, communities build shared artifacts that serve the purpose of a *group memory*,[1] in which personal knowledge relevant to the needs of the community is archived and indexed for use by others. Key needs for the users of a group memory are *cognitive coherence*, i.e. sensible organization of its contents; *navigability*, the ability to find and browse knowledge relevant to the users' needs; *guidance for contribution*, indicating when and where to contribute to the memory; and *social awareness* of the memory's other users.

One familiar real world group memory is a library. It serves a community's need for a shared repository of factual knowledge and literature, is designed to be relatively easy to navigate, and is organized to give the collection a measure of cognitive coherence, placing related items in close physical adjacency. As a physical space visited by many people at once, a library also promotes social awareness among its patrons. And, well-designed libraries contributed to regularly by their users offer guidance on how to do so.

We would like our virtual group memories to provide similar functionality, and we believe a shared, visual information space can do this effectively.

## Shared Visual Information Spaces

To understand how such a space can do this it is useful to consider the properties desired in particular categories of information spaces (see Figure 1).

**Information Spaces.** The basic properties of any information space are *persistence* and *cognitive coherence*. Information placed into the space should remain until explicitly removed, and the user should be able to organize the space sensibly. These facilities are provided by most conventional information technologies (for example, a hierarchical file system viewed as nested folders).

**Visual Information Spaces.** A visual information space adds the property of *visual coherence* and *navigability*. In visual information spaces, the user can view the contents of the space in a coherent graphical (2-D or 3-D) representation. Information-seeking behavior then becomes literal navigation through a continuous space. Visual coherence requires that the graphical organization of the visualization reflect a meaningful organization of the knowledge contained in the space. Together, these properties help the user find relevant information and learn the structure of the space.

[1]Our focus is *group* memory, in which participants are unified by a shared intellectual interest or enterprise, as distinguished from *organizational* memory, in which participants are members of a formal organization.

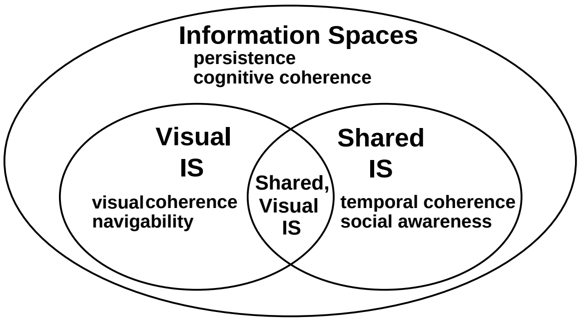


Figure 1: Properties of information spaces.

**Shared, Visual Information Spaces.** In a shared information space, multiple users participate on a regular basis. One desirable property for shared spaces is *social awareness* – an awareness of who else uses the space and what they have contributed. Another desirable property is *temporal coherence*, the provision of a coherent view of the space, so that one user's changes to the space are immediately reflected in other users' views of the space. A shared, visual space should enable visualizing both the information and the activities of other users.

Users' concerns about privacy and trust are additionally important to address in shared spaces, but are not the main focus of this work. Instead, we focus on enabling users to easily archive knowledge of value to the group in a shared space.

## Architecting a Spatial Group Memory

Previously, we identified three key needs of the users of a group memory: navigability, guidance for contribution, and social awareness. We believe the spatial metaphor is significantly better at meeting these needs for group memory than conventional hypertext approaches.

A group memory is ineffective if users cannot navigate to locate others' contributions. The spatial metaphor assists navigability by giving the user a sense of place within the information space, helping him stay oriented with respect to the information within it. The spatial representation also permits overviews, providing the user an understanding of the scope and organization of the group memory. In a space organized using a spatial metaphor, navigation for information-seeking becomes wayfinding through an analogue of physical space, as opposed to navigation through a folder hierarchy or a hypertext.

Simply creating an empty space to place contributions is not sufficient for navigability, however. In real world spaces, we use features in the environment to help us navigate. In Plexus, we provide virtual analogues to navigation features like landmarks, paths, and regions. Studies have shown that people tend to use these features to describe their mental images of real-world urban spaces (Lynch 1960). Their virtual analogues, if used to reflect useful relationships in the underlying knowledge, can enhance the memorability and navigability of information spaces (Chalmers, Ingram, & Franger 1996; Foltz 1998).

The organizing principle of the information space – the connections between the architecture of the space and the task and domain it represents – helps to build a cognitive map of the space in the memory's participants (Lokuge, Gilbert, & Richards 1996). This cognitive map, associating important concepts in the group memory with locations in the information space, can help to alleviate the messy desk problem in two ways.

First, by using the organizing principle of the space as shared model of the content of the group memory, users can reliably navigate through the space to find knowledge relevant to their needs. For example, in OfficeHours, to see what Questions relevant to the AI concept of "search" have been posted, a user could first locate the visually distinct Search landmark, and then browse the adjacent area of the space.

Second, the organizing principle of the space guides the placement of new contributions. Users can situate their contributions by literally placing them adjacent to similar or relevant contributions, using the space's navigation features to guide them.[2] Whether a user will place his contributions in a way that makes sense to others is a major question facing shared information spaces (Berlin *et al.* 1993). We hope that by providing a navigation framework of landmarks, paths, and regions users will be encouraged to place items in a way that helps other users navigate to find them.

The spatial metaphor can also enhance social awareness of other users. In non-visual group memories, users are aware of others indirectly by reading their contributions after they have been posted. However, the use of a shared visual space permits other ways to obtain social awareness. Visualization juxtaposes the contributions of multiple users, enabling a user to easily see her contributions in relation to others' contributions and the architecture of the space.

Other facilities in Plexus provide awareness among simultaneous users. Although Plexus is not envisioned as a tool for real-time collaboration (like TeamRooms (Roseman & Greenberg 1996)), we have provided means for lightweight, synchronous interaction. In Plexus, users' actions are immediately reflected in others' views, enabling users to directly observe others' activities. In Plexus, a real-time chat facility and list the users currently logged into the system promotes casual, synchronous interaction among simultaneous users.

Taken together, we believe the spatial metaphor offers several advantages over non-spatial approaches to group memories, by engaging users' visual and spatial abilities in the representation of the information space.

## The Plexus System

To evaluate the suitability of shared, visual information spaces for group memory, we have implemented Plexus, a shared two-dimensional blackboard that can be tailored for particular group memory applications. Users can post short text notes, which can be joined with directed links indicating a relationship between notes. A screen shot of Plexus is shown in Figure 2.

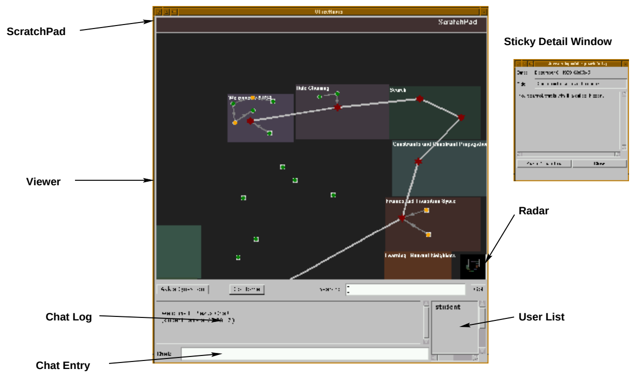


Figure 2: A screen shot of Plexus.

A Plexus information space consists of (1) titled text notes (stickies), (2) links between stickies, and (3) navigation features added to the space.

Every sticky is situated at an $(x, y)$ coordinate in a 2-D space of (conceptually) infinite extent (Figure 2). The user views this space through a 2-D multiscale interface (Bederson 1994) that presents an overview of the stickies and links in the space. A radar in the lower-right-hand corner of the space shows the relationship between the user's current view and the extents of the entire information space.

The user pans and zooms through this space by rubber-banding a velocity vector in the view window. Users can also navigate by right-clicking to focus on individual regions, which will then be smoothly centered in the user's view. Users can gain an overview of the space by zooming out to examine all of the stickies, then zoom in on a region containing stickies of interest, and finally view the content of individual stickies by right-clicking on them. Users may also search the full text of all stickies in the space by keyword.

When a user adds a new sticky to the space, it appears on the user's ScratchPad, a holding area distinct from the multiscale view. The user can then navigate

[2] Although automatic or assisted placement of new items into the space is possible, we do not explore those possibilities here.

to the location in space he wants to place the sticky, pick it up from the ScratchPad, and drop it into position. Users may reposition their own stickies at a later time by dragging them with the mouse.

Awareness of other users is provided by a list of users currently logged in and by immediate notification of others' actions through real-time updates to the objects in the space (both their contents and their position). Plexus also supports casual interaction among simultaneous users through real-time chat.

To tailor Plexus to provide a particular kind of group memory, stickies and links can be annotated with a type. A type is simply a way to give a meaningful category to stickies and links, and to structure the kinds of contributions that can be made (in a method similar to gIBIS (Begeman, 1988)). In OfficeHours, our application of Plexus, types are assigned automatically according to the permitted discourse moves in the space (for example, a response to a Question sticky is automatically an Answer sticky).

Each Plexus space has a *czar*, a distinguished member with special privileges. The czar may architect the navigation features of the space, and remove inappropriate stickies and links, among other privileges.

### Implementation

Plexus is implemented as a Java client application that stores its content in a remote object-relational database. The chat and object update functionality is provided through a separate, multi-threaded message server.

## OfficeHours

To illustrate Plexus' use as a group memory, we have adapted it to create OfficeHours, a group memory of students' questions and answers about an undergraduate Artificial Intelligence course at the authors' institution.[3]

We chose this example because we believe an archive of answered questions provides a rich source of knowledge that should be shared among students and preserved in a group memory. Students benefit because they can not only have their own questions answered, but can gain additional insight into the course material by reviewing others' questions and answers. Teaching assistants (TAs) benefit by being able to respond to a particular question once, in the hopes that a student with the same question will browse the memory first for the answer. TAs can also see at a glance which topics have prompted the most questions.

The OfficeHours space contains stickies of two types: questions and answers. In addition, untyped stickies may be posted by the czar (the course TA) with additional notes or URLs of course material. Students contribute to the space by posting new questions or answering other students' questions. The TA may also remove stickies and alter the navigation features (landmarks, paths, and regions) of the space.

**Asking Questions and Posting Answers.** Students can post questions in two ways. They can post an unlinked question by clicking the "Ask a Question" button on the main interface or post a question about a particular sticky, in which case a link is created between the new question and the sticky. The identity of a question's author is hidden from other students, so that students will not feel uncomfortable about posting a question.

Answers are created by clicking on the "Answer this Question" button while viewing a question's content. A link is automatically created between the answer and its question. A question may have multiple answers linked to it.

**Navigation Elements.** The navigation elements of OfficeHours – landmarks, regions, and a path – are shown in Figure 3. They were designed to give a cognitively coherent picture of the intellectual territory of the course and to guide the placement of questions and answers. The four regions in the space correspond to broad approaches to Artificial Intelligence. Within each region, landmarks indicate important concepts that fall under each approach. (For example, the "Rule-Based Systems" landmark is located in the "Reasoning" region of the space.) A path, indicated by the heavier arrows, visits the landmarks roughly in the order the topics were covered in the course. A first-time visitor to the space can use this path to tour the major concepts of the course.

In Figure 3, the TA has posted questions and answers from his archive of email correspondence with students. As a further visual cue, the link between each question and answer is oriented towards the landmark for the concept most relevant to that question. A quick glance at the space reveals that "Rule-Based Systems," the fourth landmark on the path, has prompted the most questions.

A number of visualization cues support the presentation of the stickies in the information space (Table 1). The stickies are color-coded by the identity of the author, and a symbol is used inside the sticky to identify the type of the sticky. A student's own stickies are colored green, while those of other students are colored blue. The TA's stickies are all colored orange. This color scheme was designed so that students could quickly locate their own postings, and see if a question has been answered by the TA.

Landmarks are larger than and visually distinct from stickies. Unanswered questions have a white square behind them, so that the TA can quickly scan the space for questions awaiting answers.

### Evaluation

In a two-week trial, twenty-six students from the course were invited to try OfficeHours and contribute new

[3]By adapting Plexus, we mean implementing dialog boxes and graphical widgets suitable for OfficeHours. The interface described previously remains the same.

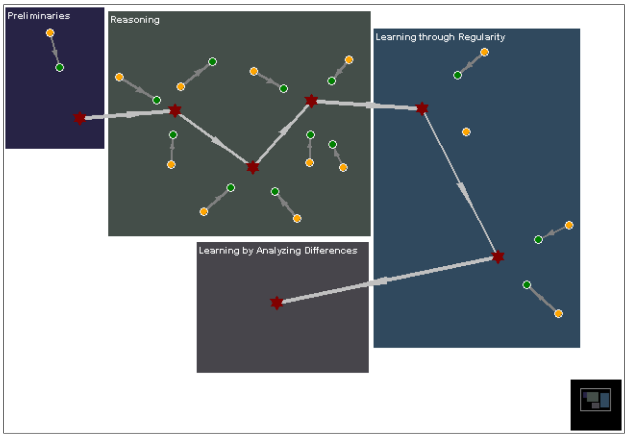


Figure 3: Navigation features in OfficeHours. The background color has been altered to better distinguish the regions.

| | |
|---|---|
|  | A question |
|  | An unanswered question |
|  | An answer |
|  | An untyped sticky |
|  | A landmark |
| 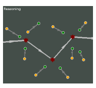 | A named, colored region |
| 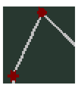 | A path connecting multiple landmarks |

Table 1: Visualization cues in OfficeHours.

questions. The short duration of the trial and its coincidence with the institute's final exam period limited student participation. However, several students used the interface to browse the initial set of questions and answers, and reacted favorably to OfficeHours in email communication to the authors. A longer-term study with students from another course is planned for the upcoming term.

## Related Work

Other systems have provided the function of a group memory, but without visualization of the memory's contents. Answer Garden (Ackerman 1998), in particular, archives frequently asked questions and answers for a group in a categorized list, and provides a hypertext interface to the archive. Its task is most similar to that of OfficeHours, and adding a spatial interface to Answer Garden would be another test of the suitability of the spatial metaphor for this task.

Another group memory system, TeamInfo (Berlin *et al.* 1993), allows users to save keyword-annotated email messages into a group hypertext. A concern that arose in that work was that relevant messages would be saved and searched under different sets of keywords by different users. We believe a navigable, visual space can help to alleviate this problem by making clear the conceptual organization of the space, guiding the placement of items, and by allowing users to browse a space as well as search by keywords.

gIBIS graphically displays a group hypertext structured specifically around the IBIS model of design deliberation (Conklin & Begeman 1988). The success of gIBIS in capturing design rationale and graphically promoting new contributions has made it a useful model for our work. We want to fully explore the potential of visual, spatial representations for this and other kinds of group memories.

Other work addresses the design issues and difficulties inherent in integrating collaborative applications into a particular workplace setting (Zimmermann & Selvin 1997; Rogers 1994). We do not address these concerns here, as we are mainly focused on novel interfaces for such applications.

## Conclusion and Future Work

This project is intended to assess the advantages of a spatial, visual interface to a group memory. We believe such an interface has several advantages for its users, by presenting them with a visually coherent representation of the memory that assists navigation, guides contributions, and promotes awareness of other users.

Future work will focus on evaluating this type of interface for group memories. A longer-term trial of OfficeHours will take place in the upcoming term. Also, the availability of many lists of frequently asked questions (FAQs) on the World-Wide Web suggests a comparison of text-based and spatial navigation of questions and answers, to verify the advantages of the spatial, visual approach to group memory.

Finally, we would like to investigate what other kinds of group memories can be cast into a spatial rubric, such as spaces that support the on-line discussions of a real-world reading group.

**Acknowledgements.** The authors would like to thank Randy Davis for helpful comments and discussion throughout the project, Patrick Winston for many insights into the information architecture of OfficeHours, Michael Oltmans and Christine Alvarado for helpful comments on the Plexus interface, and the support of the NDSEG Graduate Fellowship Program and DARPA Grant N0014-94-1-1090.

## References

Ackerman, M. S. 1998. Augmenting organizational memory: a field study of Answer Garden. *ACM Transactions on Information Systems* 16(3):203–224.

Bederson, B. 1994. Pad++: Advances in multiscale interfaces. In *Proceedings of CHI '94*, Conference Companion, 315–316.

Berlin, L. M.; Jeffries, R.; O'Day, V. L.; Paepcke, A.; and Wharton, C. 1993. Where did you put it? issues in the design and use of a group memory. In *CHI '93. Conference proceedings on Human factors in computing systems*, 23–30.

Chalmers, M.; Ingram, R.; and Franger, C. 1996. Adding imageability features to information displays. In *ACM Symposium on User Interface Science and Technology*, 33–39.

Conklin, J., and Begeman, M. L. 1988. gIBIS: A hypertext tool for exploratory policy discussion. *ACM Transactions on Information Systems* 6(4):303–331.

Foltz, M. A. 1998. Designing navigable information spaces. Master's thesis, Massachusetts Institute of Technology.

Lokuge, I.; Gilbert, S. A.; and Richards, W. 1996. Structuring information with mental models: A tour of Boston. In *Proceedings of CHI '96*.

Lynch, K. 1960. *The Image of the City*. Cambridge, Massachusetts: The Technology Press and the Harvard University Press.

Rogers, Y. 1994. Exploring obstacles: Integrating CSCW in evolving organizations. In *Proceedings of Computer-Supported Cooperative Work '94*, 67–77.

Roseman, M., and Greenberg, S. 1996. TeamRooms: network places for collaboration. In *Proceedings of CSCW '96*, 325–333.

Zimmermann, B., and Selvin, A. M. 1997. A framework for assessing group memory approaches for software design projects. In *Proceedings of the Conference on Designing Interactive Systems*, 417–426.