Chapter 5

Collective Memory

The cognitive models and architectures discussed in chapter 4 all included memory systems that provide both long-term storage of information and a context in which that information is processed. It seems self-evident that any system that attempts to model cognition as an information processing activity must include components or facilities that provide these two functions. If no long-term memory were included in the system, only current information would be available for processing, severely limiting any form of cognition that could occur in that system. If no working context were included, processing would have to occur within the input/output streams of the information flow. Although some sort of filter architecture might be possible, that architecture could not solve problems that require context sensitive operations or relating bits of information that are not adjacent in the input/output stream. Consequently, I assume that for a collaborative group to function as a coherent, intelligent system, that system must include components that provide long-term storage of information and contexts in which conceptual processing can take place. I refer to these components as the group's *collective memory*.

We might expect the collective memory to resemble human memory in broad outline. For example, it seems reasonable to assume that separate subsystems will be used for long-term storage and for active processing, that information will move back and forth between the two, and that individual operations will be responsible for retrieving concepts, changing their form and content, and storing the result. Otherwise, it is hard to imagine how groups can build complex artifacts. However, we should also expect the abstraction I am calling the collective memory to differ from human memory. At the very least, it must take into account the multiple working memories of the individuals who comprise the group. Consequently, the strategy I follow in this chapter is to define a construct that is recognizable as a *plausible* analog or extrapolation of human memory, which provides

similar function, but is not necessarily identical to the corresponding components in conventional human memory system(s).

The information flow model introduced in chapter 2 included two basic types of information: tangible and intangible. Because we are considering collective intelligence and human cognition from an information processing perspective, we should consider the collective memory in relation to these types. That is, we must consider how tangible knowledge is stored over long periods of time and in which contexts it is accessed and processed. This results in memory subsystem for tangible knowledge. We must also consider similar issues for intangible knowledge, resulting in a second subsystem for that type of information. Together these two memory subsystems comprise the collective memory for a collaborative group.

Throughout this and the discussions that follow, it is important to recognize that we can view a group both as the collection of individuals who comprise it and as an entity in its own right. Consequently, when I speak of a group's collective memory, I am referring to a collective memory for that entity. It is related to the individual memory systems of the group's members, but it also functions with respect to the group as a whole.

The discussion is divided into two main parts. The first is concerned with components for storing and working with tangible forms of information, the second with comparable components for intangible knowledge. The chapter ends with a brief discussion of issues for further research. Because the focus is on collective intelligence in computer-based collaboration, I assume that groups are using a computer system, such as the ABC system described in chapter 3, to support their activities.

Tangible Knowledge

In this section, I focus on *tangible* knowledge and on constructs that provide functions similar to those provided by long-term, working memory, and extended memory in individuals. Together, these components can be regarded as the group's collective memory for tangible knowledge.

Long-Term Memory

Let's look first at long-term memory. One must keep in mind that the discussion is concerned with computer-based collaborative groups and with collective intelligence as a form of computer-mediated behavior. Because I am assuming that groups do most if not all of their tangible work in direct relation to a computer system, such as ABC, when we consider the long-term storage of the group's tangible knowledge, we can do so with regard to how that information is stored in the computer system.

I refer to the body of tangible knowledge that a group works with as the *artifact*, as suggested in Fig. 5.1. Within the ABC system, it is stored as a hypermedia graph structure. It is a straightforward extrapolation to regard the artifact as a form of long-term memory for the group. To see why this is so, we can examine the artifact with respect to important characteristics and functions associated with conventional long-term memory.

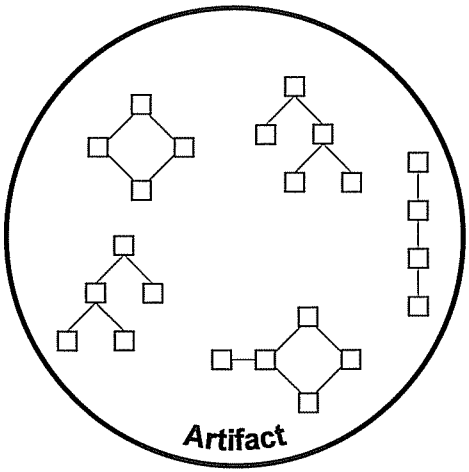

Fig. 5.1. The Artifact viewed as collective long-term memory for a group's tangible knowledge.

First, *long-term memory* is a permanent store. The artifact meets this criterion in a straightforward way. In ABC, information is stored in the graph structure and in the contents of nodes. This store is permanent, as much as any physical representation is permanent.

Second, long-term memory is frequently modeled as a semantic or associative network. An artifact built and maintained in ABC can be set up to resemble either type of graph structure. In both data models, nodes represent concepts and links represent relationships between concepts. Representing concepts as nodes with descriptive labels is straightforward. To represent relationships between concepts, ABC provides several options. Links can denote explicit semantic or associative relationships. Because all objects stored in ABC can carry attributes, attributes attached to links can define the link types required for a semantic network. For an associative network, no type scheme is needed. Consequently, ABC links can model either a semantic or an associative network. In its node/content relationship, ABC provides a second, hierarchical relationship between a superordinate concept and a cluster or a structure of subordinate concepts. Thus, the hierarchical properties of both memory models can be represented by either explicit tree structures or by the node/content relationship. A third type of relationship provided by ABC, called a hyperlink, can denote secondary semantic associations and connotations between concepts that are distant to one another in the main node/content hierarchy. Thus, a group's artifact can be implemented within an ABC-like data model so that it resembles either a semantic or an associative network and includes hierarchical chunking. Hence, an ABC artifact resembles the long-term memory components found in IPS models and architectures.

Third, information can be encoded, stored, and later retrieved from long-term memory. Because the artifact is dynamic, new concepts and new relationships can be added to it. In the ABC system, this means that new (labeled) nodes and new relationships — in the form of structural links, hyperlinks, or node/content relationships — can be added to the artifact. This is done by using special programs, called browsers, to work with the graph structure of the artifact or by using applications, such as text editors or drawing programs, to encode abstract concepts (represented by labeled nodes) as blocks of text, diagrams, or other forms. These tools, then, allow users to encode and store new concepts in the artifact. To retrieve information from this store, ABC users work with these same browsers to move from one concept or context to another, either by following

hyperlinks (representing semantic associations) or by successively opening browsers and/or applications on the contents of nodes. Currently, ABC does not include a content search function, as would be expected in conventional models of long-term memory; however, such a function could be added. Thus, because information can be encoded and stored in an ABC-maintained artifact and later retrieved from it, the artifact satisfies this third requirement for a long-term memory.

Fourth, information retrieved from long-term memory is activated and processed in a working memory component. I claim this property here but delay justifying it until the next section.

Thus, the artifact, as maintained in an ABC-like collaboration system, is a large, permanent store that can be structured as a semantic or associative network, into which encoded information can be inserted and from which specific concepts can be retrieved or activated. Consequently, it can be viewed as a form of long-term memory for tangible knowledge.

Working Memory

Working memory is the context in which information is activated and where conceptual processing takes place with regard to that information. In this section, I will focus on the component as a whole and defer discussion of the specific processes that operate within it until the next chapter.

As suggested in the preceding section, ABC browsers and applications can be viewed as a form of working memory with respect to information stored in the artifact. In IPS models, working memory has been described both as a separate cache that is loaded and unloaded with respect to long-term memory and as the currently activated parts of long-term memory. Researchers frequently shift between these two views as needed by context or for convenience. To simplify the discussion, I generally assume the cache version, although ABC tools can support either view. For example, a browser or application, as shown in Fig. 5.2, can be opened on any given subgraph or file of data that is part of the artifact. It presents the user with a view of the data and with functions for adding, deleting, and changing that

information. Thus, that segment of the artifact can be said to be *activated* in the browsers and applications because it can be consciously attended to and modified by users of ABC.

If one prefers, the relationship between long-term memory and working memory can be viewed as one of retrieving and storing, rather than activating. That is, one can view the segment of the artifact presented by a browser or application as having been retrieved from the artifact. The user can then perform various operations on this information using the functions provided by the browsers and applications. When the user has finished his or her work on this segment, the (possibly modified) contents are returned to the artifact for permanent storage.

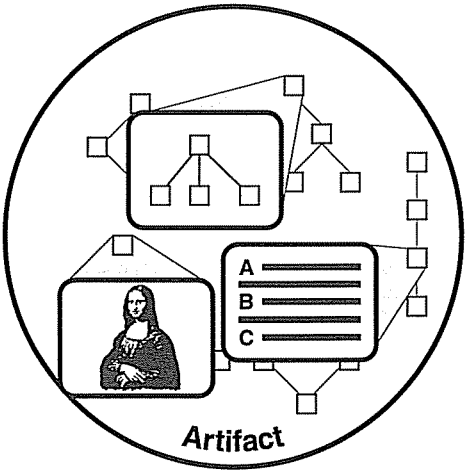

Fig. 5.2. Browsers and applications viewed as working memory for a group's tangible knowledge.

Thus, ABC browsers and applications provide contexts where the contents of long-term memory can be consciously attended to and processed, and they can accommodate the two complementary views of working memory found in IPS models. Consequently, we can regard them as a form of working memory for tangible knowledge.

Extended Memory

Extended memory is an external physical representation of information that functions as an extension of an individual's working memory. Often solving a problem, such as an arithmetic problem, is closely related to producing a sequence of intermediate products represented on paper. Or, when planning a document, many writers will jot down ideas on post-it notes during a brainstorming phase to relieve the burden on working memory and later organize those ideas into some sort of plan for the document. Both the post-it notes and the document plan can be viewed as forms of extended memory. ABC browsers and applications can also function as forms of extended memory for knowledge-construction tasks. For example, a graph browser can enable a user to represent and work with 40 or 50 concepts, as opposed to the four or five that can be dealt with in human working memory. Consequently, these tools function both as a form of collective working memory and as an extended memory for intangible knowledge.

Thus, we have identified constructs that are recognizable as approximations or extrapolations for the three memory systems found in cognitive models and architectures — long-term memory, working memory, and extended memory. All are used by collaborative groups to work with tangible information. It is, therefore, plausible to view these components as comprising a collective memory for tangible knowledge.

Intangible Knowledge

Not all of the information a group works with is tangible and, thus, stored in the artifact. A great deal remains *intangible,* carried in the heads of the individuals who comprise the group. In this section, I identify constructs that function as a memory system for tangible knowledge. It, too, consists of separate components for long-term, working, and extended memory, which, together, can be viewed as a collective memory for intangible knowledge.

Long-Term Memory

Collaborative groups incorporate two forms of intangible knowledge: *private* and *shared.* Individuals are frequently added to a team because they have expertise or specialized knowledge not shared by all members of the group. Their *private knowledge* informs their actions and is used by them as they work on their parts of the artifact. Thus, it is a group resource, but it is not part of the group's collective knowledge because it is not immediately accessible to the group as a whole. In some cases, portions of this knowledge will be encoded into products, become part of the artifact, and thereby become part of the collective long-term memory for tangible knowledge. It is the residue that I want to look at here — knowledge that remains intangible but is no longer private.

Shared intangible knowledge is the intersection of the different collections of information stored in the individual long-term memories of the group members. It is developed and maintained by them in two ways. First, an individual member of the group who possesses private knowledge may transfer portions of his or her knowledge to the group as a whole. For example, one group member may brief the rest of the group on a technical topic. Second, the group as a whole may collectively build a segment of shared knowledge. For example, they may analyze a problem together and, in the course of their discussion, construct a solution. But, regardless of how it is derived, the body of intangible knowledge that is accessible to all members of a group can be regarded as a form of *collective long-term memory.*

Figure 5.3 illustrates the relationships among the artifact and both shared and private intangible knowledge. The artifact is shown in the center. It is surrounded by a layer of shared intangible knowledge, held in common by all members of the group; it is shown as the lightly shaded area in the figure. Around shared knowledge are the more extensive bodies of privately intangible knowledge, shown as darkly shaded areas, known to individual members of the group. This entire construct is dynamic, as new knowledge is constructed and new information is brought into the group, but it represents the total body of information that is available for use by the group.

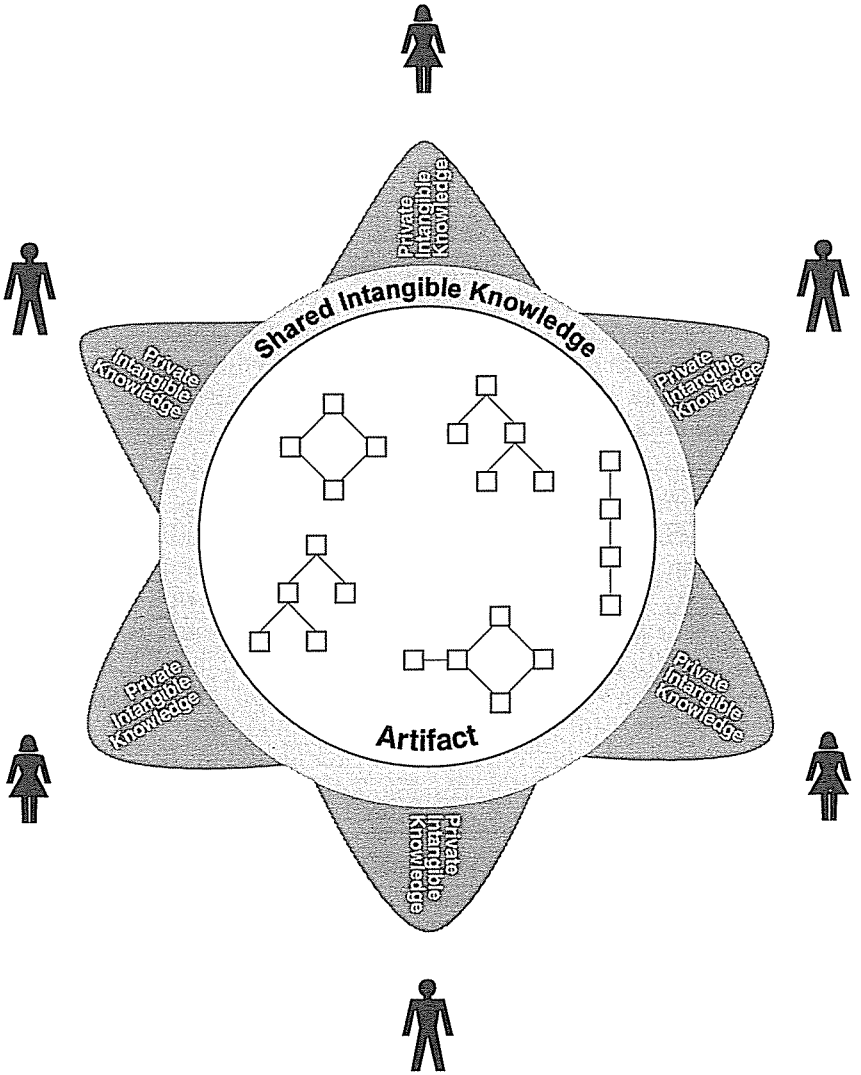

Fig. 5.3. Artifact surrounded by a group's shared and private intangible knowledge.

Working Memory

The contexts in which intangible knowledge is activated and processed by the group form a *collective working memory* for this type of information. It is there that shared knowledge is built and there that it is refreshed and calibrated to maintain it within limits sufficient for it to function as a collective resource. The contexts in which this occurs are the social and intellectual situations in which a group collectively attends to particular segments of intangible knowledge. These situations range from chance discussions in the hallway to formal design reviews. Thus, we can regard activities such as these as a form of working memory for intangible knowledge, albeit a very abstract one.

The literal mechanism through which this abstract working memory is realized is the individual working memories of the individuals participating in the activity. From this perspective, the collective working memory is the union of the conventional working memories of the individuals taking part an event where shared intangible knowledge is developed. But that definition is not sufficient, because shared knowledge cannot be built in isolation. There must be communication, at the least, and, in most cases, knowledge must be collectively built and maintained by the group if it is to function as shared knowledge. Thus, a group's collective working memory for intangible knowledge consists of both the individual working memories of its members and the situations in which they activate their respective versions of shared knowledge and collectively attend to and/or modify it.

One of the most common and most important of these situations is the meeting. To see in what way it makes sense to regard a meeting as a form of working memory, let's look briefly at the way information is processed in those situations. Working memory provides the context for three basic kinds of processes: activation/retrieval, conceptual processing, and encoding/storage. Activation takes place in meetings as groups reconstruct shared knowledge. During discussions, groups frequently recall earlier discussions, decisions, or factual data to refresh and calibrate this knowledge or to incorporate it into new conceptual structures or decisions. The mechanisms for doing this are conventional. As topics arise during discussion, they key access functions, activating portions of members' respective long-term memories. If the topic is known to all, the activated structures are

approximately the same, although not identical. However, if the issue is important, and hence memorable, the difference among versions are likely to be small. Thus, when I speak of shared knowledge, I am referring to an abstraction or an ideal that does not literally exist; what many groups do achieve, however, is an approximation that is sufficient to enable them to work coherently. Determining the degree of similarity among the different versions required for knowledge to be considered shared is a topic for future research.

As discussion within a meeting progresses, the conceptual structures currently activated in the working memories of individual participants changes. People add new ideas, argue against a previous point, forge new relationships between old ideas, and so on. Thus, a form of collective conceptual processing takes place.

Eventually, new information structures that evolve during meetings may be encoded and stored in the collective long-term memory — the union of individual long-term memories of the group members. This is more likely to be the case if the new construct is considered "important." New information that is shared in the moment but not regarded as important is likely to be forgotten by some or all of the participants and, hence, remains private relative to the group as a whole.

Thus, meetings provide a context in which shared knowledge can be activated, developed, and selectively encoded and stored in the collective long-term memory. Consequently, meetings can be viewed as a form of working memory where shared intangible knowledge is processed. Other forms of intellectual and social interaction serve a similar function; perhaps future research will identify them and examine them more closely.

Extended Memory

In chapter 2, we saw that during meetings groups often use display devices such as a whiteboard to list points, to draw diagrams, or to write a sentence or a source code statement. Doing so makes visible the abstract concepts or structures being discussed. Consequently, such devices can help members remove extraneous differences in their respective understandings of what they are discussing, thereby contributing to the group's development of a common understanding of an idea. Consequently, ephemeral products

such as these can be viewed as forms of *collective extended memory* with respect to a group's developing shared intangible knowledge.

Looking back at the chapter as a whole, I have identified two sets of constructs that function within collaborative groups as long-term, working, and extended memories. The first was concerned with the tangible knowledge a group works with and, hence, can be viewed as a collective memory for that type of knowledge. The second was concerned with intangible knowledge and, hence, can be viewed as a collective memory for intangible knowledge. Together, they constitute the collective memory for the group, in which all forms of information to which it has access as a collective resource are stored and processed.

Issues for Research

Several issues or questions for potential research arise from the preceding discussion.

- *What is the structure of a group's tangible long-term memory?*

Because it is tangible, the group's long-term memory, unlike human long-term memory, can be examined directly. What does it look like? What are its structural characteristics? Will groups build conceptual structures that resemble those thought to exist in human long-term memory (e.g., semantic or associative networks) or will they be structured differently? Will they build larger or smaller hierarchical components, perhaps influenced by the larger capacity of the collective working memory? To what extent will the total structure be kept internally consistent? How extensively will groups link "distant" concepts with one another through explicit semantic associations (such as hyperlinks)? Do differences in the degree of connectivity translate into greater or lesser coherence in the work of the group?

- *What is its pattern of growth and use?*

Because the collective long-term memory is dynamic, what will be its pattern of growth? Will it grow at a uniform rate for the duration of a project, grow by fits and starts, or grow quickly at first and then gradually decrease in rate? How will members of the group

use it and participate in its development? Will material be copied and reused? Will individual members dwell in particular regions or will they range widely over the whole structure? Will work on one section lead to changes in another? Will internal consistency decrease with duration and/or size of the project? What problems will inconsistencies cause and how significant will they be?

- *How can we observe and characterize intangible long-term memory?*

Because the intangible long-term memory is maintained in individual minds, it is both distributed and unobservable. Are existing methodologies for studying it adequate? If not, which tools and methods can we develop so that we can "see" it and study it in close detail? Once we have the requisite tools, what will be its structure and form?

- *How is it developed, calibrated, and used?*

How will intangible knowledge grow and develop over the course of a project? Because it is more amorphous than its tangible counterpart, how will group members keep their respective versions consistent? How similar must these versions be to be considered shared? Will groups engage specific calibrating activities or will they allow portions of it to dissipate, resulting in a form of collective "forgetting"? How is intangible knowledge used by the group? How important is it?

- *What is the relationship between the tangible and intangible memory systems?*

How do the tangible and intangible memory systems relate to one another? Are there regular and observable patterns of development, where private knowledge becomes shared and, in turn, influences development of the artifact? Does the reverse process occur? What are the relative sizes of the two memory systems? Is this ratio consistent across groups, tasks, project sizes, durations, and so on? Do the two memories grow at similar rates? Does it matter?

- *What is the relationship between these systems and the strategic behavior of a group?*

How is the overall strategy of a group reflected in the growth and development of the two long-term memory systems? For example, a

group that works by consensus might be expected to develop a large body of intangible knowledge in parallel with the tangible structure it is building, whereas a group that works largely by partitioning its work might develop a relatively smaller body of shared knowledge. Is this true? Can groups be trained or influenced in the ways they develop the collective memory and, if so, will it make a difference in their productivity?