

COLLAGEN

Applying Collaborative Discourse Theory to Human-Computer Interaction

Charles Rich, Candace L. Sidner, and Neal Lesh

- We describe an approach to intelligent user interfaces, based on the idea of making the computer a collaborator, and an application-independent technology for implementing such interfaces.

What properties of a user interface would make you want to call it intelligent? For us, any interface that is called intelligent should at least be able to answer the six types of questions from users shown in figure 1. Being able to ask and answer these kinds of questions implies a flexible and adaptable division of labor between the human and the computer in the interaction process. Unlike most current interfaces, an intelligent user interface should be able to guide and support you when you make a mistake or if you don't know how to use the system well.

What we are suggesting here is a paradigm shift. As an analogy, consider the introduction of the undo button. This one button fundamentally changed the experience of using interactive systems by removing the fear of making accidental mistakes. Users today expect every interactive system to have an undo button and are justifiably annoyed when they can't find it. By analogy, to focus on just one of the question types in figure 1, what we are saying is that every user interface should have a "What should I do next?" button.

Note that we are not saying that each of the questions in figure 1 must literally be a separate

button. The mechanisms for asking and answering these questions could be spoken or typed using natural (or artificial) language, adaptive menus, simple buttons, or some combination of these. We have experimented with all these mechanisms in the various prototype systems described later.

Finally, some readers might object that answering the question types in figure 1 should be thought of as a function of the application rather than the interface. Rather than getting into a unproductive semantic argument about the boundary between these two terms, we prefer instead to focus on what we believe is the real issue, namely, whether this characterization of intelligent user interfaces can lead to the development of a reusable middleware layer that makes it easy to incorporate these capabilities into diverse systems.

Again, there is a relevant historical analogy. A key to the success of so-called WIMP (windows, icons, menus, and pointers) interfaces has been the development of widely used middleware packages, such as MOTIF and SWING. These middleware packages embody generally useful graphic presentation and interaction conventions, such as tool bars, scroll bars, and check boxes. We believe that the next goal in user interface middleware should be to codify techniques for supporting communication about users' task structure and process, as suggested by the question types in figure 1. This article describes a system, called COLLAGEN, which is the first step in this direction.

Who should/can/will do ____ ?
 What should I/we do next ?
 Where am/was I ?
 When did I/you/we do ____ ?
 Why did you/we (not) do ____ ?
 How do/did I/we/you do ____ ?

Figure 1. Six Questions for an Intelligent Interface.

Adapted from the news reporter's "five Ws." The blanks are filled in with application-specific terms, ranging from high-level goals, such as "prepare a market survey" or "shut down the power plant," to primitive actions, such as "underline this word" or "close valve 17."

Collaboration

What does all this have to do with the "collaborative discourse theory" in the title of this article? The goal of developing generic support for communicating about the user's task structure cannot, we feel, be achieved by taking an engineering approach focused directly on the questions in figure 1. We therefore started this research by looking for an appropriate theoretical foundation, which we found in the concept of collaboration.

Collaboration is a process in which two or more participants coordinate their actions toward achieving shared goals. Most collaboration between humans involves communication. *Discourse* is a technical term for an extended communication between two or more participants in a shared context, such as a collaboration. *Collaborative discourse theory* (see sidebar) thus refers to a body of empirical and computational research about how people collaborate. Essentially, what we have done in this project is apply a theory of human-human interaction to human-computer interaction.

In particular, we have taken the approach of adding a collaborative interface agent (figure 2) to a conventional direct-manipulation graphic user interface. The name of our software system, COLLAGEN (for COLLABORATIVE AGENT), is derived from this approach. (Collagen is also a fibrous protein that is the chief constituent of connective tissue in vertebrates.)

The interface agent approach mimics the relationships that typically hold when two humans collaborate on a task involving a shared artifact, such as two mechanics working

on a car engine together or two computer users working on a spreadsheet together.

In a sense, our approach is a very literal-minded way of applying collaborative discourse theory to human-computer interaction. We have simply substituted a software agent for one of the two humans that would appear in figure 2 if it were a picture of human-human collaboration. There might be other ways of applying the same theory without introducing the concept of an agent as separate from the application (Ortiz and Grosz 2000), but we have not pursued these research directions.

Notice that the software agent in figure 2 is able both to communicate with and observe the actions of the user and vice versa. Among other things, collaboration requires knowing when a particular action has been done. In COLLAGEN, this collaboration can occur two ways: (1) a reporting communication ("I have done *x*") or (2) direct observation. Another symmetrical aspect of the figure is that both the user and the agent can interact with the application program.

There are many complex engineering issues regarding implementing the observation and interaction arrows in figure 2, which are beyond the scope of this article (Lieberman 1998). In all our prototypes, the application program has provided a program interface (API) for performing and reporting primitive actions and querying the application state. Communication between the user and the agent has variously been implemented using speech recognition and generation, text, and menus.

Outline of the Article

The remainder of this article lays out the work we have done in more detail, starting with snapshots of four interface agents built by us and our collaborators in different domains using COLLAGEN. Following these examples comes a description of the technical heart of our system, which is the representation of the discourse state and the algorithm for updating it as an interaction progresses. Next, we discuss another key technical contribution, namely, how COLLAGEN uses plan recognition in a collaborative setting. Finally, we present the overall system architecture of COLLAGEN, emphasizing the application-specific versus application-independent components. We conclude with a brief discussion of related and future work.

Application Examples

This section briefly describes four interface agents built by us and our collaborators in four different application domains using COLLAGEN.

Theory

Grosz and Sidner (1986) proposed a tripartite framework for modeling task-oriented discourse structure. The first (*intentional*) component records the beliefs and intentions of the discourse participants regarding the tasks and subtasks (purposes) to be performed. The second (*attentional*) component captures the changing focus of attention in a discourse using a stack of “focus spaces” organized around the discourse purposes. As a discourse progresses, focus spaces are pushed onto, and popped off of, this stack. The third (*linguistic*) component consists of the contiguous sequences of utterances, called *segments*, which contribute to a particular purpose.

Grosz and Sidner (1990) extended this basic framework with the introduction of SHAREDPLANS, which are a formalization of the collaborative aspects of a conversation. The SHAREDPLAN formalism models how intentions and mutual beliefs about shared goals accumulate during a collaboration. Grosz and Kraus (1996) provided a comprehensive axiomatization of SHAREDPLANS, including extending it to groups of collaborators.

Most recently, Lochbaum (1998) developed an algorithm for discourse interpretation using SHAREDPLANS and the tripartite model of discourse. This algorithm predicts how conversants follow the flow of a conversation based on their understanding of each

other’s intentions and beliefs.

References

- Grosz, B. J., and Sidner, C. L. 1986. Attention, Intentions, and the Structure of Discourse. *Computational Linguistics* 12(3): 175–204.
- Grosz, B. J., and Sidner, C. L. 1990. Plans for Discourse. In *Intentions and Communication*, eds. P. R. Cohen, J. L. Morgan, and M. E. Pollack, 417–444. Cambridge, Mass.: MIT Press.
- Grosz, B. J., and Kraus, S. 1996. Collaborative Plans for Complex Group Action. *Artificial Intelligence* 86(2): 269–357.
- Lochbaum, K. E. 1998. A Collaborative Planning Model of Intentional Structure. *Computational Linguistics* 24(4): 525–572.

We have also built agents for air travel planning (Rich and Sidner 1998) and e-mail (Gruen et al. 1999). All these agents are currently research prototypes.

Figures 3 through 6 show a screen image and a sample interaction for each example agent. Instances of the question types in figure 1 are underlined in the sample interactions.

Each screen image in figures 3 through 6 consists of a large application-specific window, which both the user and the agent can use to manipulate the application state, and two smaller windows, labeled Agent and User, which are used for communication between the user and the agent. Two of the example agents (figures 3 and 6) communicate using speech-recognition and -generation technology; the other two allow the user to construct utterances using hierarchical menus dynamically generated based on the current collaboration state.

The agent in figure 3 helps a user set up and program a video cassette recorder (VCR). The image in the figure, which a real user would see on his/her television screen, includes the VCR itself so that the agent can point at parts of the VCR during explanations (see line 12). In the first part of the VCR agent transcript, the agent helps the user eliminate the annoying blinking 12:00 that is so common on VCR clocks. Later on, the agent walks the user through the task of connecting a camcorder to the VCR.

The agent in figure 4 was developed in collaboration with the Industrial Electronics and


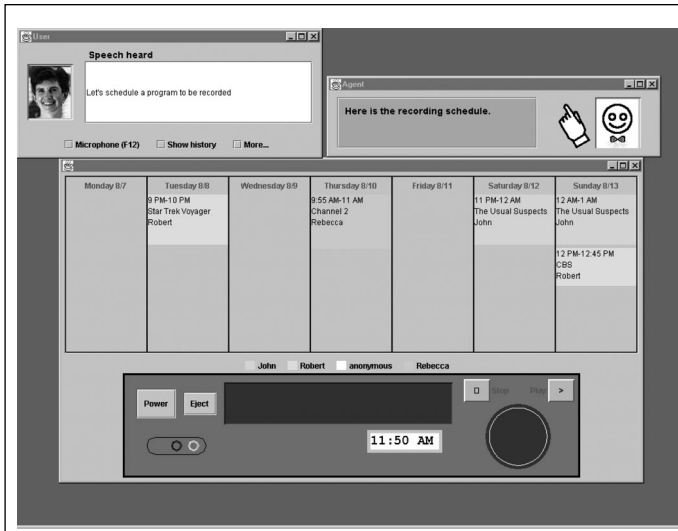


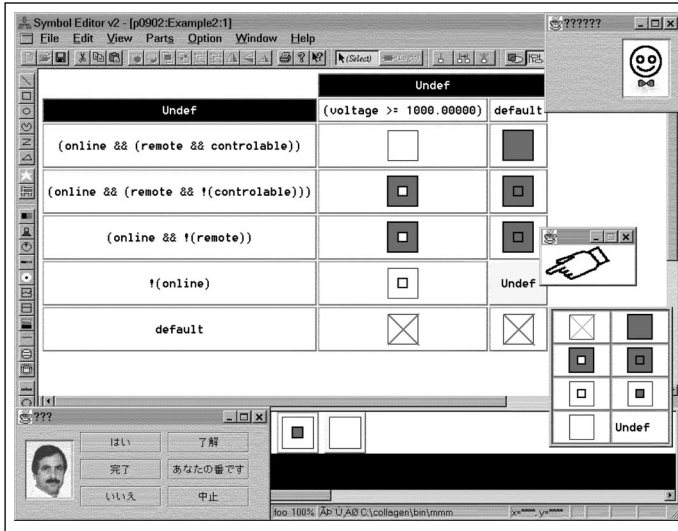
Figure 2. Collaborative Interface Agent.

Systems Laboratory of Mitsubishi Electric in Japan. The application program in this case is a sophisticated graphic interface development tool, called the SYMBOL EDITOR. Like many such tools, the SYMBOL EDITOR is difficult for novice users because there are too many possible things to do at any moment, and the system itself gives no guidance regarding what to do




- 1 User says "Hello."
- 2 Agent says "Hello, what's your name?"
- 3 User says "Candy."
- 4 Agent says "Hello, Candy!"
- 5 User says "Let's set the clock."
- 6 Agent says "What is the time?"
- 7 User says "The time is 11:00 AM."
- 8 Agent sets the clock to 11:00 AM.
- ...
- 9 User says "How do I play back a video from my camcorder?"
- 10 Agent says "First, plug in the camcorder cables."
- 11 User says "How do I plug in the camcorder cables?"
- 12 Agent points to where you plug in the camcorder cables.
- 13 Agent says "The red cable plugs into the red terminal, the black cable plugs into the black terminal, and the yellow cable plugs into the yellow terminal. Make sure the other ends of the cables plug into the corresponding colors on the camcorder."
- 14 User says "What next?"
- 15 Agent says "Now press Play on your camcorder. You should see a picture on the TV screen."
- 16 User says "Ok."
- 17 Agent says "I don't see a picture. Make sure the camcorder is plugged in or has fresh batteries, and is turned on."
- 18 User says "Ok."
- 19 Agent says "Good. It's working."

Figure 3. Example Video Cassette Recorder Agent Built with COLLAGEN.



- 1 Agent says "Please list the other (non-Boolean) attributes used in the columns of the logic table."
- 2 User says "The attributes are: voltage."
- 3 Agent creates a new user-defined attribute.
- 4 Agent enters "voltage" as the name of the new attribute.
- 5 Agent says "Please select the data type of the new attribute."
- 6 User selects float as the data type of the new attribute.
- 7 User clicks ok to accept the definition of voltage.
- 8 Agent says "Please edit the initial setting for the voltage attribute (if necessary) and then click ok."
- 9 User clicks ok to accept 0.0 as the initial setting for voltage.
- 10 Agent shows the expression editor for column 1.
- 11 Agent enters (voltage > 0.0) as the title for column 1.
- 12 Agent says "Please define an expression for column 1."
...[user goes to lunch]...
- 14 User says "Where are we?"
- 15 Agent says "You are defining an expression for column 1, as part of making a logic table."
- 16 User enters (voltage >= 1000.0) as the title for column 1.
- 17 Agent shows the logical operator window.
- 18 Agent selects the >= operator.
- 19 Agent drags and drops >= into expression for column 1.
- 20 Agent selects the voltage attribute.
- 21 Agent drags and drops voltage into expression for column 1.
- 22 Agent types "1000.0" into expression for column 1.

Figure 4. Example Symbol Editor Agent Built with COLLAGEN.



- 1 Agent says "Let's start the generator."
- 2 User says "Ok."
- 3 Agent says "You take it from here."
- 4 User presses the alarm reset button.
- 5 User says "What next?"
- 6 Agent says "Please press the gas turbine generator on button."
- 7 Agent points to the gas turbine generator on button.
- 8 User presses the gas turbine generator button.
- 9 Agent says "Good. We have successfully started the generator."

Figure 5. Example Gas Turbine Agent Built with COLLAGEN.

next. Our agent guides a user through the process of achieving a typical task using the SYMBOL EDITOR, automatically performing many of the tedious subtasks along the way.

The agent in figure 5 was developed in collaboration with the Information Sciences Institute of the University of Southern California (Rickel et al. 2001). This agent teaches a student user how to operate a gas turbine engine and generator configuration using a simple software simulation. The first time the agent teaches a new task or subtask, it walks the student through all the required steps. If a task has already been performed once, however, the agent tells the student to “take it from here” (line 3). If the student later asks for help (line 5), the agent will describe just the next step to be performed.

The gas turbine agent is part of a larger effort, which also involves the MITRE Corporation (Gertner, Cheikes, and Haverty 2000), to incorporate application-independent tutorial strategies into COLLAGEN. Teaching and assisting are best thought of as points on a spectrum of collaboration (Davies et al. 2001) rather than as separate capabilities.

Finally, Figure 6 shows an agent being developed at the Delft University of Technology to help people program a home thermostat (Keyson et al. 2000). The transcript here illustrates only a simple interaction with the agent. This agent will eventually be able to help people analyze their behavior patterns and construct complicated heating and cooling schedules to conserve energy. This work is part of a larger research project at Delft to add intelligence to products.

Discourse State

Participants in a collaboration derive benefit by pooling their talents and resources to achieve common goals. However, collaboration also has its costs. When people collaborate, they must usually communicate and expend mental effort to ensure that their actions are coordinated. In particular, each participant must maintain some sort of mental model of the status of the collaborative tasks and the conversation about them. We call this model the *discourse state*.

Among other things, the discourse state tracks the beliefs and intentions of all the participants in a collaboration and provides a focus-of-attention mechanism for tracking shifts in the task and conversational context. All this information is used by an individual to help understand how the actions and utterances of the other participants contribute to the common goals.

The screenshot shows a window titled "Thermostat View". On the left, a list of rooms includes "living room", "kitchen", "bathroom", "bedroom dave", "bedroom lisa", "bedroom parents", "attic", and "study". A circular dial in the center is set to "14". On the right, there is an "Agent" window with a hand icon, a smiley face, and a slider between "(passive)" and "(active)". Below the room list is a "User" dialog box with a portrait of a man and buttons for "Yes", "Ok", "Accomplished", "What next?", "No", and "Stop". There are also checkboxes for "Show more things to say ..." and "Show history and context". A clock in the bottom right corner shows "12:53".

1 Agent says "Let's set the temperature in a room."
 2 User selects the bathroom.
 3 User says "What next?"
 4 Agent says "What should the temperature setting be?"
 5 User says "18 degrees Centigrade."
 6 Agent sets the temperature in the bedroom to 18 degrees.

Figure 6. Example Thermostat Agent Built with COLLAGEN.

To turn a computer agent into a collaborator, we needed a formal representation of a discourse state and an algorithm for updating it. The discourse state representation currently used in COLLAGEN, illustrated in figure 7, is a partial implementation of Grosz and Sidner's theory of collaborative discourse (see sidebar); the update algorithm is described in the next section.

COLLAGEN's discourse state consists of a stack of goals, called the *focus stack* (which will soon become a stack of focus spaces to better correspond with the theory), and a *plan tree* for each goal on the stack. The top goal on the focus stack is the "current purpose" of the discourse. A plan tree in COLLAGEN is an (incomplete) encoding of a partial SHAREDPLAN between the user and the agent. For example, Figure 7 shows the focus stack and plan tree immediately following the discourse events numbered 1 through 3 on the right side of the figure.

Segmented Interaction History

The annotated, indented execution trace on the right side of figure 7, called a *segmented*

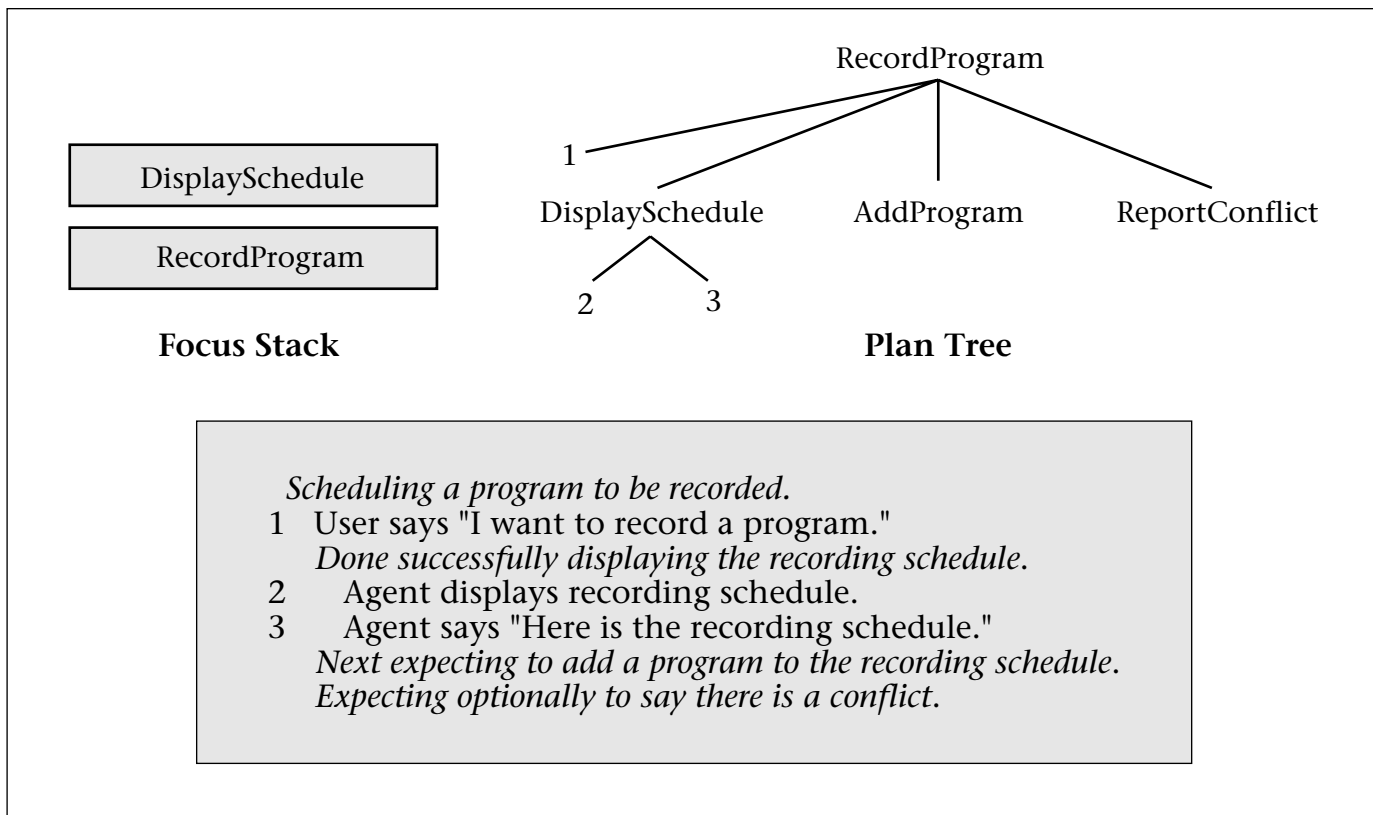


Figure 7. Example Discourse State and Segmented Interaction History for Video Cassette Recorder Agent.

interaction history, is a compact, textual representation of the past, present, and future states of the discourse. We originally developed this representation to help us debug agents and COLLAGEN itself, but we have also experimented with using it to help users visualize what is going on in a collaboration (see discussion of history-based transformations in Rich and Sidner [1998]).

The numbered lines in a segmented interaction history are simply a log of the agent's and user's utterances and primitive actions. The italic lines and indentation reflect COLLAGEN's interpretation of these events. Specifically, each level of indentation defines a segment (see sidebar) whose purpose is specified by the italicized line that precedes it. For example, the purpose of the top-level segment in figure 4 is scheduling a program to be recorded.

Unachieved purposes that are currently on the focus stack are annotated using the present tense, such as *scheduling*, whereas completed purposes use the past tense, such as *done*. (Note in figure 7 that a goal is not popped off the stack as soon as it is completed because it might continue to be the topic of conversation, for example, to discuss whether it was successful.)

Finally, the italic lines at the end of each seg-

ment, which include the keyword *expecting*, indicate the steps in the current plan for the segment's purpose that have not yet been executed. The steps which are "live" with respect to the plan's ordering constraints and preconditions have the added keyword *next*.

Discourse Interpretation

COLLAGEN updates its discourse state after every utterance or primitive action by the user or agent using Lochbaum's discourse-interpretation algorithm (see sidebar), with extensions to include plan recognition (see next section) and unexpected focus shifts (Lesh, Rich, and Sidner 2001).

According to Lochbaum, each discourse event is explained as either (1) starting a new segment whose purpose contributes to the current purpose (and thus pushing a new purpose on the focus stack), (2) continuing the current segment by contributing to the current purpose, or (3) completing the current purpose (and, thus, eventually popping the focus stack).

An utterance or action contributes to a purpose if it either (1) directly achieves the purpose, (2) is a step in a recipe for achieving the purpose, (3) identifies the recipe to be used to achieve the

purpose, (4) identifies who should perform the purpose or a step in the recipe, or (5) identifies a parameter of the purpose or a step in the recipe. These last three conditions are what Lochbaum calls “knowledge preconditions.”

A *recipe* is a goal-decomposition method (part of a task model). COLLAGEN’s recipe definition language (see figure 8) supports partially ordered steps, parameters, constraints, preconditions, postconditions, and alternative goal decompositions.

Our implementation of the discourse-interpretation algorithm described earlier requires utterances to be represented in Sidner’s (1994) artificial discourse language. For our speech-based agents, we used standard natural-language processing techniques to compute this representation from the user’s spoken input. Our menu-based systems construct utterances in the artificial discourse language directly.

Discourse Generation

To illustrate how COLLAGEN’s discourse state is used to generate, as well as interpret, discourse behavior, we briefly describe here how the VCR agent produces the underlined utterance on line 5 in figure 9, which continues the interaction in figure 7.

The discourse-generation algorithm in COLLAGEN is essentially the inverse of discourse interpretation. Based on the current discourse state, it produces a prioritized list, called the *agenda*, of (partially or totally specified) utterances and actions that would contribute to the current discourse purpose according to cases 1 through 5 described earlier. For example, for the discourse state in figure 7, the first item on the agenda is an utterance asking for the identity of the program parameter of the AddProgram step of the plan for RecordProgram.

In general, an agent can use any application-specific logic it wants to decide on its next action or utterance. In most cases, however, an agent can simply execute the first item on the agenda generated by COLLAGEN, which is what the VCR agent does in this example. This utterance starts a new segment, which is then completed by the user’s answer on line 6.

Plan Recognition

Plan recognition (Kautz and Allen 1986) is the process of inferring intentions from actions. Plan recognition has often been proposed for improving user interfaces or to facilitate intelligent help features. Typically, the computer watches “over the shoulder” of the user and jumps in with advice or assistance when it thinks it has enough information.

```
public recipe RecordRecipe achieves RecordProgram {
    step DisplaySchedule display;
    step AddProgram add;
    optional step ReportConflict report;
    constraints {
        display precedes add;
        add precedes report;
        add.program == achieves.program;
        report.program == achieves.program;
        report.conflict == add.conflict;
    }
}
```

Figure 8. Example Recipe in Video Cassette Recorder Task Model.

Definition of the recipe used in figure 7 to decompose the nonprimitive RecordProgram goal into primitive and nonprimitive steps. COLLAGEN task models are defined in an extension of the JAVA language, which is automatically processed to create JAVA class definitions for recipes and act types.

Scheduling a program to be recorded.

- 1 User says "I want to record a program."
Done successfully displaying the recording schedule.
- 2 Agent displays recording schedule.
- 3 Agent says "Here is the recording schedule."
- 4 User says "Ok." *Done identifying the program to be recorded.*
- 5 Agent says "What is the program to be recorded?"
- 6 User says "Record 'The X-Files!'."
Next expecting to add a program to the recording schedule.
Expecting optionally to say there is a conflict.

Figure 9. Continuing the Interaction in Figure 7.

In contrast, our main motivation for adding plan recognition to COLLAGEN was to reduce the amount of communication required to maintain a mutual understanding between the user and the agent of their shared plans in a collaborative setting (Lesh, Rich, and Sidner 1999). Without plan recognition, COLLAGEN’s discourse-interpretation algorithm onerously required the user to announce each goal before performing a primitive action that contributed to it.

Although plan recognition is a well-known

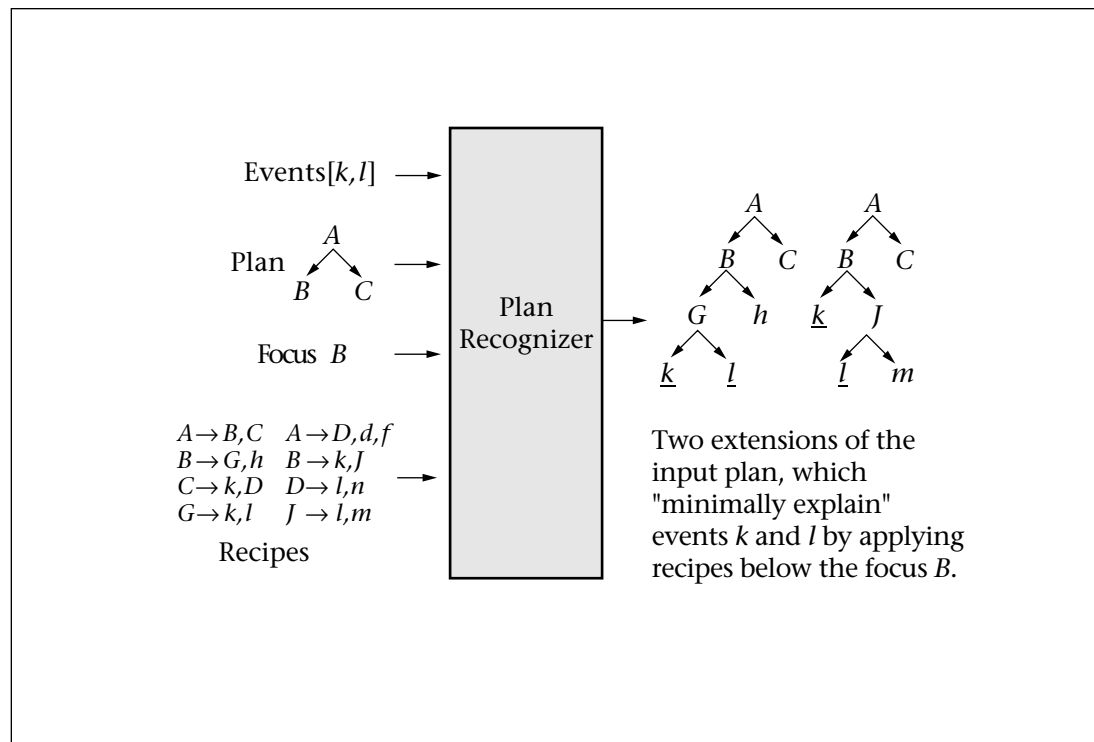


Figure 10. Plan Recognizer Input and Output in a Collaborative Setting.

feature of human collaboration, it has proven difficult to incorporate into practical computer systems because of its inherent intractability in the general case. We exploit three properties of the collaborative setting to make our use of plan recognition tractable. The first property is the *focus of attention*, which limits the search required for possible plans.

The second property of collaboration we exploit is the interleaving of developing, communicating about, and executing plans, which means that our plan recognizer typically operates only on partially elaborated hierarchical plans. Unlike the “classical” definition of plan recognition, which requires reasoning over complete and correct plans, our recognizer is only required to incrementally extend a given plan.

Third, it is quite natural in the context of a collaboration to ask for clarification, either because of inherent ambiguity or simply because the computation required to understand an action is beyond a participant’s abilities. We use clarification to ensure that the number of actions the plan recognizer must interpret will always be small.

Figure 10 illustrates roughly how plan recognition works in COLLAGEN. Suppose the user performs action k . Given the root plan (for example, A) for the current discourse purpose (for example, B) and a set of recipes, the plan

recognizer determines the set of minimal extensions to the plan that are consistent with the recipes and include the user performing k . If there is exactly one such extension, the extended plan becomes part of the new discourse state. If there is more than one possible extension, action k is held and reinterpreted along with the next event, which can disambiguate the interpretation (which l does not), and so on. The next event might in fact be a clarification.

Our algorithm also computes essentially the same recognition if the user does not actually perform an action but only proposes it, as in, “Let’s achieve G .” Another important, but subtle, point is that COLLAGEN applies plan recognition to both user and agent utterances and actions in order to correctly maintain a model of what is mutually believed.

System Architecture

Figure 11 summarizes the technical portion of this article by showing how all the pieces described earlier fit together in the architecture of a collaborative system built with COLLAGEN. This figure is essentially an expansion of figure 2, showing how COLLAGEN mediates the interaction between the user and the agent. COLLAGEN is implemented using JAVA BEANS, which makes it easy to modify and extend this architecture.

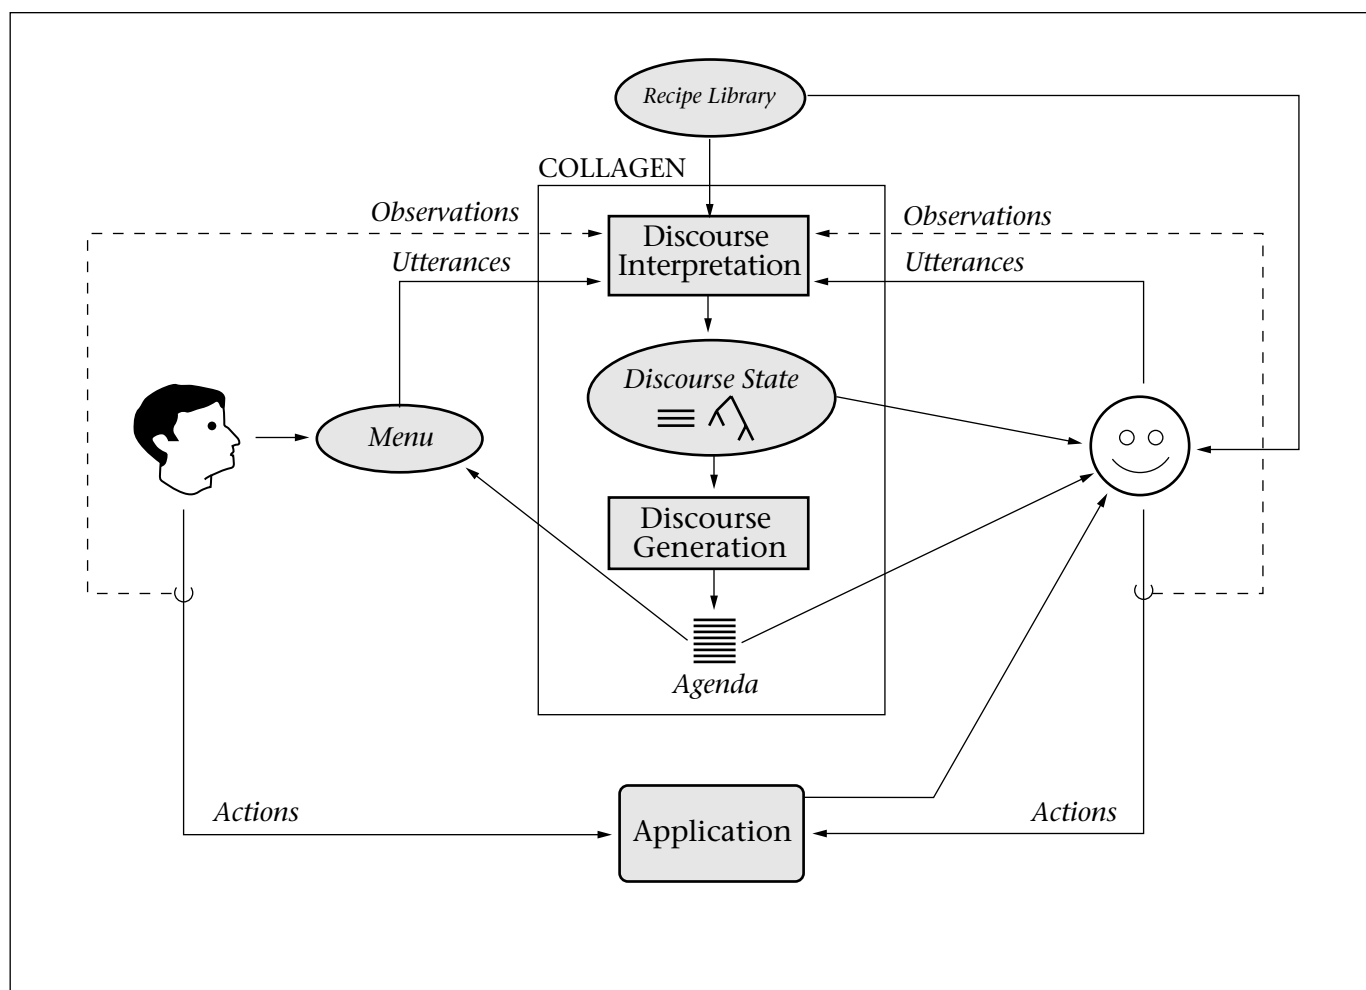


Figure 11. Architecture of a Collaborative System Built with COLLAGEN.

The best way to understand the basic execution cycle in figure 8 is to start with the arrival of an utterance or an observed action (from either the user or the agent) at the discourse-interpretation module at the top center of the diagram. The discourse-interpretation algorithm (including plan recognition) updates the discourse state as described earlier, which then causes a new agenda to be computed by the discourse-generation module. In the simplest case, the agent responds by selecting and executing an entry in the new agenda (which can be either an utterance or an action), which provides new input to discourse interpretation.

In a system without natural language understanding, a subset of the agenda is also presented to the user in the form of a menu of customizable utterances. In effect, this menu is a way of using expectations generated by the collaborative context to replace natural language understanding. Because this is a mixed-initia-

tive architecture, the user can, at any time, produce an utterance (for example, by selecting from this menu) or perform an application action (for example, by clicking on an icon), which provides new input to discourse interpretation.

In this simple story, the only application-specific components an agent developer needs to provide are the recipe library and an API through which application actions can be performed and observed (for an application-independent approach to this API, see Cheikes et al. [1999]). Given these components, COLLAGEN is a *turnkey technology*—default implementations are provided for all the other needed components and graphic interfaces, including a default agent that always selects the first item on the agenda.

In each of the four example applications (figures 3 through 6), however, a small amount (for example, several pages) of additional application-specific code was required to achieve the

desired agent behavior. As the arrows incoming to the agent in figure 11 indicate, this application-specific agent code typically queries the application and discourse states and (less often) the recipe library. An agent developer is free, of course, to use arbitrarily complex application-specific and generic techniques, such as a theorem proving and first-principles planning, to determine the agent's response to a given situation.

Related Work

This work lies at the intersection of many threads of related research in AI, computational linguistics, and user interfaces. We believe it is unique, however, in its combination of theoretical elements and implemented technology. Other theoretical models of collaboration (Levesque, Cohen, and Nunes 1990) do not integrate the intentional, attentional, and linguistic aspects of collaborative discourse, as SHAREDPLAN theory does. Our incomplete implementation of SHAREDPLAN theory in COLLAGEN does not, however, deal with the many significant issues in a collaborative system with more than two participants (Tambe 1997).

There has been much related work on implementing collaborative dialogues in the context of specific applications, based either on discourse planning techniques (Ahn et al. 1995; Allen et al. 1996; Chu-Carroll and Carberry 1995; Stein, Gulla, and Thiel 1999) or rational agency with principles of cooperation (Sadek and De Mori 1997). None of these research efforts, however, have produced software that is reusable to the same degree as COLLAGEN. In terms of reusability across domains, a notable exception is the VERBMOBIL project,¹ which concentrates on linguistic issues in discourse processing without an explicit model of collaboration.

Finally, a wide range of interface agents (Maes 1994) continue to be developed, which have some linguistic and collaborative capabilities, without any general underlying theoretical foundation.

Future Work

We are currently extending our plan-recognition algorithm to automatically detect certain classes of so-called "near miss" errors, such as performing an action out of order or performing the right action with the wrong parameter. The basic idea is that if there is no explanation for an action given the correct recipe library, the plan recognizer should incrementally relax the constraints on recipes before giving up and

declaring the action to be unrelated to the current activity. This function is particularly useful in tutoring applications of COLLAGEN.

We have, in fact, recently become increasingly interested in tutoring and training applications (Gertner, Cheikes, and Haverty 2000; Rickel et al. 2001), which has revealed some implicit biases in how COLLAGEN currently operates. For example, COLLAGEN's default agent will always, if possible, itself perform the next action that contributes to the current goal. This behavior is not, however, always appropriate in a tutoring situation, where the real goal is not to get the task done but for the student to learn how to do the task. We are therefore exploring various generalizations and extensions to COLLAGEN to better support the full spectrum of collaboration (Davies et al. 2001), such as creating explicit tutorial goals and recipes and recipes that encode "worked examples."

We are also working on two substantial extensions to the theory underlying COLLAGEN. First, we are adding an element to the attentional component (see sidebar) to track which participant is currently in control of the conversation. The basic idea is that when a segment is completed, the default control (initiative) goes to the participant who initiated the enclosing segment.

Second, we are beginning to codify the negotiation strategies used in collaborative discourse. These strategies are different from the negotiation strategies used in disputations (Kraus, Sycara, and Evanchik 1998). For example, when the user rejects an agent's proposal (or vice versa), the agent and the user should be able to enter into a subdialogue in which their respective reasons for and against the proposal are discussed.

Note

1. Verbmobil, 2000, <http://verbmobil.dfki.de>.

References

- Ahn, R.; Bunt, H.; Benn, R.; Borghuis, C.; and Van Overveld, A. 1994. The DENKArchitecture: A Fundamental Approach to User Interfaces. *Artificial Intelligence Review* 8:431–445.
- Allen, J.; Miller, B.; Ringger, E.; and Sikorski, T. 1996. A Robust System for Natural Spoken Dialogue. In *Proceedings of the Thirty-Fourth Annual Meeting of the Association for Computational Linguistics*, 62–70. San Francisco, Calif.: Morgan Kaufmann.
- Cheikes, B.; Geier, M.; Hyland, R.; Linton, F.; Riffe, A.; Rodi, L.; and Schaefer, H. 1999. Embedded Training for Complex Information Systems. *International Journal of Artificial Intelligence in Education* 10:314–334.
- Chu-Carroll, J., and Carberry, S. 1995. Response Generation in Collaborative Negotiation. In *Proceedings of the Thirty-Third Annual Meeting of the Association for*

Computational Linguistics, 136–143. San Francisco, Calif.: Morgan Kaufmann.

Davies, J.; Lesh, N.; Rich, C.; Sidner, C.; Gertner, C.; and Rickel, J. 2001. Incorporating Tutorial Strategies into an Intelligent Assistant. In Proceedings of the International Conference on Intelligent User Interfaces, 53–56. New York: Association of Computing Machinery.

Gertner, A.; Cheikes, B.; and Haverty, L. 2000. Dialogue Management for Embedded Training. In AAAI Symposium on Building Dialogue Systems for Tutorial Applications, 10–13. Technical Report FS-00-01. Menlo Park, Calif.: AAAI Press.

Gruen, D.; Sidner, C.; Boettner, C.; and Rich, C. 1999. A Collaborative Assistant for E-Mail. In Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems, Extended Abstracts, 196–197. New York: Association of Computing Machinery.

Kautz, H. A., and Allen, J. F. 1986. Generalized Plan Recognition. In Proceedings of the Fifth National Conference on Artificial Intelligence, 32–37. Menlo Park, Calif.: American Association for Artificial Intelligence.

Keyson, D.; de Hoogh, M.; Freudenthal, A.; and Vermeeren, A. 2000. The Intelligent Thermostat: A Mixed-Initiative User Interface. In Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems, Extended Abstracts, 59–60. New York: Association of Computing Machinery.

Kraus, S.; Sycara, K.; and Evanchik, A. 1998. Argumentation in Negotiation: A Formal Model and Implementation. *Artificial Intelligence* 104(1–2): 1–69.

Lesh, N.; Rich, C.; and Sidner, C. 2001. Collaborating with Focused and Unfocused Users under Imperfect Communication. In *Proceedings of the Ninth International Conference on User Modeling*, 64–73. New York: Springer-Verlag.

Lesh, N.; Rich, C.; and Sidner, C. 1999. Using Plan Recognition in Human-Computer Collaboration. In *Proceedings of the Seventh International Conference on User Modeling*, 23–32. New York: Springer-Verlag.

Levesque, H. J.; Cohen, P. R.; and Nunes, J. H. T. 1990. On Acting Together. In Proceedings of the Eighth National Conference on Artificial Intelligence, 94–99. Menlo Park, Calif.: American Association for Artificial Intelligence.

Lieberman, H. 1998. Integrating User Interface Agents with Conventional Applications. In Proceedings of the International Conference on Intelligent User Interfaces, 39–46. New York: Association of Computing Machinery.

Maes, P. 1994. Agents That Reduce Work and Information Overload. *Communications of the ACM* 37(17): 30–40.

Ortiz, C. L., and Grosz, B. J. 2001. Interpreting Information Requests in Context: A Collaborative Web Interface for Distance Learning. *Autonomous Agents and Multi-Agent Systems*. Forthcoming.

Rich, C., and Sidner, C. 1998. COLLAGEN: A Collaboration Manager for Software Interface Agents. *User Modeling and User-Adapted Interaction* 8(3–4): 315–350.

Rickel, J.; Lesh, N.; Rich, C.; Sidner, C.; and Gertner, A. 2001. Building a Bridge between Intelligent Tutoring and Collaborative Dialogue Systems. Paper presented at the Tenth International Conference on Artificial Intelligence in Education, 19–23 May, San Antonio, Texas.

Sadek, D., and De Mori, R. 1997. Dialogue Systems. In *Spoken Dialogues with Computers*, ed. R. D. Mori. San Diego, Calif.: Academic.

Sidner, C. L. 1994. An Artificial Discourse Language for Collaborative Negotiation. In Proceedings of the Twelfth National Conference on Artificial Intelligence, 814–819. Menlo Park, Calif.: American Association for Artificial Intelligence.

Stein, A.; Gulla, J. A.; and Thiel, U. 1999. User-Tailored Planning of Mixed-Initiative Information-Seeking Dialogues. *User Modeling and User-Adapted Interaction* 9(1–2): 133–166.

Tambe, M. 1997. Toward Flexible Teamwork. *Journal of Artificial Intelligence Research* 7:83–124.



Charles Rich is a senior research scientist at Mitsubishi Electric Research Laboratories (MERL) in Cambridge, Massachusetts. He received his Ph.D. from the Massachusetts Institute of Technology (MIT) in 1980. The thread connecting all Rich's research has been to make interacting with a computer more like interacting with a person. As a founder and director of the PROGRAMMER'S APPRENTICE project at the MIT Artificial Intelligence Laboratory in the 1980s, he pioneered research on intelligent assistants for software engineers. Rich joined MERL in 1991 as a founding member of the research staff. He is a fellow and past councilor of the American Association for Artificial Intelligence and was cochair of AAAI-98. His e-mail address is rich@merl.com.



Candace L. Sidner is a senior research scientist at Mitsubishi Electric Research Laboratories (MERL) in Cambridge, Massachusetts. She received her Ph.D. from the Massachusetts Institute of Technology in 1979. Sidner has researched many aspects of user interfaces, especially those involving speech and natural language understanding and human-computer collaboration. Before coming to MERL, she was a member of the research staff at Bolt Beranek Newman, Inc., Digital Equipment Corp., and Lotus Development Corp. and a visiting scientist at Harvard University. She is a fellow and past councilor of the American Association for Artificial Intelligence, past president of the Association for Computational Linguistics, and cochair of the 2001 International Conference on Intelligent User Interfaces. Her e-mail address is sidner@merl.com.



Neal Lesh is a research scientist at Mitsubishi Electric Research Laboratories in Cambridge, Massachusetts. He received his Ph.D. from the University of Washington in 1998. His research interests of late lie primarily in human-computer collaborative problem solving. He has also worked in the areas of interface agents, interactive optimization, goal recognition, data mining, and simulation-based inference. After completing his thesis on scalable and adaptive goal recognition at the University of Washington with Oren Etzioni, Lesh was a postdoc at the University of Rochester with James Allen. His e-mail address is lesh@merl.com.

AI Proceedings from AAAI Press

ISMB-2000—San Diego, California

Proceedings of the Eighth International Conference on Intelligent Systems for Molecular Biology

Edited by Russ Altman, Timothy Bailey, Philip Bourne, Michael Gribskov, Thomas Lengauer, Ilya Shindyalov, Lynn Ten Eyck, and Helge Weissig

The ISMB conference series provides a general forum for disseminating the latest developments in bioinformatics. ISMB is a multidisciplinary conference that brings together scientists from computer science, mathematics, molecular biology, and statistics. Its scope includes the development and application of advanced computational methods for biological problems.

ISBN 1-57735-115-0 436 pp., index. \$60.00 softcover

FLAIRS-2000—Orlando, Florida

Proceedings of the Thirteenth International Florida Artificial Intelligence Research Society Conference

Edited by Jim Etheredge and Bill Manaris

This year's Florida AI Research Society Conference proceedings covers a wide range of topics, including applications, computer vision, evolutionary computation, intelligent agents, knowledge-based systems, learning and AI, logic and AI, logic programming, natural language processing, planning, AI applied to spacecraft autonomy, evolutionary computation, intelligent tutoring systems, knowledge management, neural network applications, parallel and distributed reasoning, reasoning about function, spatiotemporal reasoning, uncertain reasoning, and verification, validation, and knowledge-base refinement.

ISBN 1-57735-113-4 700 pp., illus., index, \$60.00 softcover

AIPS-2000—Breckenridge, Colorado

Proceedings of the Fifth International Conference on Artificial Intelligence Planning and Scheduling Systems

Edited by Steve Chien, Subbarao Kambhampati, & Craig A. Knoblock

The 2000 proceedings papers span a wide range of topics including algorithms for planning and scheduling in a variety of environments, planner implementation, and applications of planning and scheduling research.

ISBN 1-57735-111-8 416 pp., index, \$50.00 softcover

KDD-98—New York, New York

Proceedings of the Fourth International Conference on Knowledge Discovery and Data Mining

Edited by Rakesh Agrawal, Paul Stolorz, and Gregory Piatetsky-Shapiro

Paper topics in this proceedings include theory and foundational issues in KDD, data mining methods and algorithms, database-centric data mining, and KDD process and human interaction applications.

ISBN 1-57735-070-7 392 pp., index, \$65.00 softcover

IAAI-01—Seattle, Washington

The Thirteenth Annual Conference on Innovative Applications of Artificial Intelligence

Edited by Hayam Hirsh and Steve Chien

Books in the Innovative Applications of Artificial Intelligence (IAAI) series report on the nature and range of real-world problems that AI technology can address successfully today. All of the contributions in these books have a practical slant, showing how AI has been successfully applied to a wide spectrum of domains and tasks.

ISBN 1-57735-134-7 202 pp., index. \$30.00 softcover

Ordering Information

To order call: (650) 328-3123 or fax to (650) 321-4457.
AAAI Press • 445 Burgess Drive • Menlo Park, CA 94025
<http://www.aaai.org>
MasterCard, VISA, and American Express accepted.