

Using a Model of Collaborative Dialogue to Teach Procedural Tasks

Jeff Rickel,¹ Neal Lesh,² Charles Rich,² Candace L. Sidner² and Abigail Gertner³

¹ USC Information Sciences Institute, 4676 Admiralty Way, Marina del Rey, CA, 90292
rickel@isi.edu, <http://www.isi.edu/isd/rickel>

² Mitsubishi Electric Research Laboratories, 201 Broadway, Cambridge, MA, 02139
lesh,rich,sidner@merl.com, <http://www.merl.com/projects/collagen>

³ MITRE Corporation, 202 Burlington Road, Bedford, MA, 01730
gertner@mitre.org, <http://www.mitre.org/resources/centers/it/g068>

Abstract

Previous research on building intelligent tutoring systems has not leveraged general models of collaborative discourse, even though tutoring is an inherently collaborative and often discourse-based activity. Similarly, previous research on collaborative discourse theory has rarely addressed tutorial issues, even though teaching and learning are crucial components of collaboration. We help bridge the gap between these two related research threads by presenting a tutorial agent, called Paco, that we built using an application-independent collaboration manager, called Collagen. Our primary contribution is to show how a variety of tutorial behaviors can be expressed as rules for generating candidate discourse acts in the framework of collaborative discourse theory.

1 Introduction

Our research objective is to develop computer tutors that collaborate with students on tasks in simulated environments. Towards this end, we seek to integrate two separate but related research threads: intelligent tutoring systems (ITS) and collaborative dialogue systems (CDS). Research on ITS (e.g., [1, 21, 23]) focuses on computer tutors that adapt to individual students based on the target knowledge the student is expected to learn and the presumed state of the student's current knowledge. Research on CDS (e.g., [8, 13, 22]), with an equally long history, focuses on computational models of human dialogue for collaborative tasks.

Unfortunately, there has been a surprising lack of cross-fertilization between these two research areas. Work on tutorial dialogue for intelligent tutoring systems (e.g., [3, 14, 24]) has not leveraged general models of collaborative dialogue. Similarly, research on collaborative dialogues has focused on modeling conversations between peers or between an expert and novice, but has rarely addressed tutorial issues.

To help integrate ITS and CDS, we developed a tutorial agent in Collagen [17], a middleware system based on a long line of research on collaborative discourse [8, 6, 7, 5, 13]. Collagen maintains a model of the discourse state shared by the user (e.g., student) and the computer agent (e.g., tutor). The discourse state includes information about the current focus of attention and the collaborators' mutually believed plans. Agents constructed using Collagen use the discourse state to generate an agenda of candidate *discourse acts*, including both "physical" actions and utterances, and then choose one to perform or utter.

Our tutorial agent, Paco (Pedagogical Agent for Collagen), teaches students procedural tasks in simulated environments, building on ideas from earlier tutoring systems [18, 19]. While Paco can engage in slightly more sophisticated conversations than previous such tutors, our primary contribution is to show how a variety of tutorial behaviors can be expressed as rules for generating candidate discourse acts in Collagen. Translating behaviors developed in ITS into the framework of CDS is a first step towards building tutoring agents that can leverage advances in collaborative discourse theory. Also, since Paco is domain-independent, its tutorial actions can be added to the set of candidate discourse acts of any agent built with Collagen, allowing such agents to tutor in addition to their normal role as assistants. Finally, a third goal of this work is to report on Collagen’s value for building tutorial agents, both in terms of the theory it reflects and the software architecture it supports.

2 Pedagogical Approach

We designed Paco to support simulation-based training, in which students learn tasks by performing them in a simulation of the real work environment. (Of course, if the target work environment is actually a software application, that application can serve as the simulator.) The computer tutor’s instruction and assistance are situated in the performance of domain tasks in the simulated world. That is, the tutor chooses a scenario (task to perform starting from a particular simulation state), works through it with the student, and then repeats until all scenarios have been mastered.

Our pedagogical approach is based on the apprenticeship model of learning [2], which requires two capabilities. First, the tutor must be able to perform and explain the task. Second, it must be able to monitor the student as she performs the task, providing assistance when needed as well as critique or positive feedback when appropriate. As the student gains proficiency, the assistance provided should decrease. Ideally, students should learn to flexibly apply well-defined procedures in a variety of situations.

Figure 1 shows an example dialogue with our current implementation of Paco that illustrates some of the key features we support. Paco is teaching the student how to operate the gas turbine engines that propel naval ships. Paco has previously worked through a simple scenario in which the student engaged one of the turbine engines. Now, Paco is going to teach the same procedure under slightly more complicated conditions: (1) a high vibration alarm has occurred on the gas turbine generator, shutting the generator down, so the student will have to reset the alarm before starting the generator; and (2) a second engine is already running, so the student will have to stop it before starting up the desired engine. The remainder of the paper will use this example dialogue to illustrate aspects of our design.

If there were no overlap among tasks and scenarios, Paco could be implemented in an obvious way: the tutor would first demonstrate the entire task, then repeatedly let the student practice the task, providing assistance where necessary. However, different tasks often share common subtasks or actions, and different scenarios often require variants of the same task. Therefore, at any moment, a student’s level of mastery may differ across the different parts of a task. For example, a new scenario may require branches of a task that the student has not yet seen (e.g., lines 8-12 and 18-35 in the example dialogue) while also requiring steps and subtasks that have been mastered already.

To address this issue, Paco uses a student model to dynamically interleave demonstration and coached practice, using the approach introduced by Rickel [18]. As the student and Paco progress through a task, Collagen will repeatedly identify the set of valid next steps in the plan to solve the current task. Paco consults the student model to see whether the student has sufficient knowledge to choose the next step. If so, it will expect the student to take the next step, and will provide assistance

only if the student requests it or makes a mistake. If not, Paco will intervene and teach the student what to do next (e.g., lines 8-12 and 18-35). Thus, as Paco and the student work through tasks, initiative will pass back and forth between them based on the student's prior experience. Whenever Paco decides that the initiative should shift, it will let the student know through verbal comments (e.g., "You take it from here").

Paco represents the procedures it will teach using Collagen's declarative language for domain-specific procedural knowledge. This knowledge serves as a model of how domain tasks should be performed. Each task is associated with one or more *recipes* (i.e., procedures for performing the task). Each recipe consists of several elements drawn from a relatively standard plan representation. First, it includes a set of steps, each of which is either a primitive action (e.g., press a button) or a composite action (i.e., a subtask). Composite actions give tasks a hierarchical structure. Second, there may be ordering constraints among the steps; these constraints define a partial order over the steps. Third, a task and its steps can have parameters, and a recipe can specify constraints (bindings) among the parameters of a task and its steps. Finally, steps can have preconditions (to allow Collagen to determine whether a step can be performed in the current state) and postconditions (to determine whether the effects of a step have been achieved).

3 Collagen as a Foundation for Teaching Procedural Tasks

Collagen's main value for building tutoring systems is that it provides a general model of collaborative dialogue based on well-established principles from computational linguistics. The model includes two main parts: (1) a representation of discourse state and (2) a discourse interpretation algorithm that uses plan recognition to update the discourse state given the actions and utterances of the user and agent. Previous tutoring systems for procedural tasks do not include dialogue managers with the same level of generality.

Based on the work of Grosz and Sidner [6], Collagen partitions the discourse state into three interrelated components: the linguistic structure, the attentional state, and the intentional structure. The linguistic structure, implemented as a *segmented interaction history* in Collagen, groups the dialogue history into a hierarchy of discourse segments. Each segment is a contiguous sequence of actions and utterances that contribute to some *purpose* (e.g., performing a task or subtask). For example, Figure 2 shows the segmented interaction history for a portion of the example dialogue.

The attentional state, i.e., what the user and agent are talking about and/or working on *now*, is represented by a stack of discourse purposes called the focus stack [6]. When a new discourse segment is begun, its purpose is pushed onto the stack. When a discourse segment is completed or discontinued, its purpose is popped off the stack. The stack mechanism is crucial for the proper treatment of interruptions. Additionally, the attentional state maintained by Collagen includes an extension to the original model of Grosz and Sidner to capture which participant holds the conversational initiative. This allows Paco to decide when to explicitly pass the initiative to the student (e.g., "You take it from here").

An explicit representation of attentional state is important in tutoring systems so that the tutor does not confuse the student with unexpected shifts in focus. For example, if the student begins work on a subtask and then asks for help, the tutor's advice should relate to that subtask, even if other actions outside that subtask could be done next. For tutoring systems that follow a strict sequence through domain tasks, a focus stack may not be needed. However, to support mixed-initiative dialogues where the student and tutor have more freedom to choose the execution order for tasks and initiate topics for discussion, a focus stack helps ensure the coherence of the tutor's responses.

- (1) **Paco** sets the state of the simulator.
- (2) **Paco** “I’ve reset the simulation so that we can try engaging an engine under slightly different conditions.”
- (3) **Paco** “Let’s engage engine one.”
- (4) **Student** “Ok.”
- (5) **Paco** “You take it from here.”
- (6) **Student** “Let’s start the generator.”
- (7) **Paco** “Good.”
- (8) **Paco** (Pointing) “Now, notice that the alarm light is on.”
- (9) **Paco** “Before starting a generator, always make sure that the alarm light is off so that you’ll know when a new alarm appears.”
- (10) **Paco** (Pointing) “Press the alarm reset button on the generator in order to extinguish the alarm.”
- (11) **Student** presses the alarm reset button on the generator.
- (12) **Paco** “Good.”
- (13) **Student** “Ok.”
- (14) **Paco** “You take it from here.”
- (15) **Student** presses the on button on the generator.
- (16) **Paco** “Great.”
- (17) **Paco** “We succeeded in starting the generator.”
- (18) **Paco** (Pointing) “Now, notice that engine two is running.”
- (19) **Paco** “Before starting an engine, always make sure that the other engine is stopped.”
- (20) **Paco** “Let’s stop engine two.”
- (21) **Paco** (Pointing) “First, press the disengage button on engine two in order to disengage the clutch.”
- (22) **Student** presses the disengage button on engine two.
- (23) **Paco** “Right.”
- (24) **Paco** “The next step of stopping engine two is to close the fuel valves on engine two.”
- (25) **Student** “Ok.”
- (26) **Paco** “You take it from here.”
- (27) **Student** closes fuel valve one on engine two.
- (28) **Paco** “Good.”
- (29) **Student** closes fuel valve two on engine two.
- (30) **Paco** “Good job.”
- (31) **Paco** (Pointing) “The next step is to press the off button on engine two.”
- (32) **Student** presses the off button on engine two.
- (33) **Paco** “Great.”
- (34) **Paco** (Pointing) “The indicator light tells us that engine two is stopped.”
- (35) **Paco** “So, we succeeded in stopping engine two.”
- (36) **Student** “Ok.”
- (37) **Paco** “You take it from here.”
- (38) **Student** opens fuel valve one on engine one.
- (39) **Paco** “Good.”
- (40) **Student** opens the bleed valve on engine one.
- (41) **Paco** “That was a correct action, but I wasn’t expecting you to do it now.”
- (42) **Paco** “Let’s go back to opening the fuel valves on engine one.”
- (43) **Student** opens fuel valve two on engine one.
- (44) **Paco** “Great.”
- (45) **Student** presses the engage button on engine one.
- (46) **Paco** “Whoops, you can’t do that yet.”
- (47) **Student** “What next?”
- (48) **Paco** (Pointing) “Press the on button on engine one.”
- (49) **Student** presses the on button on engine one.
- (50) **Paco** “Good.”
- (51) **Student** “I think I should set the throttle speed.”
- (52) **Paco** “Right.”
- (53) **Student** “What should the speed be?”
- (54) **Paco** “The speed should be stop.”
- (55) **Student** sets the throttle speed to stop.
- (56) **Paco** “Good.”
- (57) **Student** presses the engage button on engine one.
- (58) **Paco** “Good job.”
- (59) **Paco** “We succeeded in engaging engine one.”

Figure 1: An example dialogue in which Paco teaches steps in the current task (e.g., lines 21, 31), teaches that subtasks are finished (e.g., lines 17, 35), provides help when requested (e.g., line 48), passes initiative to the student when she should know what to do next (e.g., lines 5, 14, 26), provides positive feedback (e.g., lines 7, 16, 23), diagnoses and critiques incorrect actions (e.g., line 46), explains why steps needs to be done (e.g., line 9), discusses postconditions for actions (e.g., lines 8, 34), and helps the student stay focused on the current subtask (e.g., line 42).

Engaging engine one

3 Paco says "Let's engage engine one."
4 Student says "Ok."
5 Paco says "You take it from here."
Done successfully starting the generator:

6 Student says "Let's start the generator."
7 Paco says "Good."
Done user successfully pressing the alarm reset button on the generator:

8 Paco says "Now, notice that the alarm light is on."
9 Paco says "Before starting a generator, always make
sure that the alarm light is off so that
you'll know when a new alarm appears."
10 Paco says "Press the alarm reset button on the
generator in order to extinguish the alarm."
11 Student presses the alarm reset button on the generator.
12 Paco says "Good."
13 Student says "Ok."
14 Paco says "You take it from here."
15 Student presses the on button on the generator.
16 Paco says "Great."
17 Paco says "We succeeded in starting the generator."

Figure 2: The segmented interaction history for a portion of the dialogue in Figure 1.

While the linguistic structure and attentional state closely reflect the actual temporal order of actions and utterances in the dialogue, the intentional structure represents the decisions that have been made as a result of those actions and utterances, independent of their order. Collagen represents the intentional structure as *plan trees*, which are a partial implementation of SharedPlans [7, 5]. Nodes in the tree represent mutually agreed upon intentions (e.g., to perform a task), and the tree structure represents the subgoal relationships among these intentions. Plan trees also record other types of decisions, such as whether a recipe has been chosen for a task, whether any of its parameters have been determined, and who is responsible for performing the task (e.g., student, agent, or both).

The heart of Collagen is the discourse interpretation algorithm, which specifies how to update the discourse state given a new action or utterance by either the user or agent. Its objective is to determine how the current act contributes to the collaboration. For example, the act could contribute to the current discourse segment's purpose (DSP) by directly achieving it (e.g., pressing a button when that action is the current DSP), proposing how it can be achieved (i.e., suggesting a recipe), proposing or performing a step in its recipe, or proposing a value for one of its unspecified parameters. Collagen extends Lochbaum's discourse interpretation algorithm [13] with plan recognition, which can recognize when an act contributes to a DSP through one or more implicit acts [12].

Collagen's discourse interpretation algorithm proceeds as follows. If the current act contributes to the current DSP, it is added to the segment and the plan tree is updated accordingly. If not, Collagen searches up through the plan tree to see if an act contributes to any other action in the plan; if so, and if the act is a valid next step, it represents a shift in focus. Collagen pops all purposes off the stack that are not parents of the matched step, then pushes any necessary purposes on until the act is in focus. Finally, if nothing in the plan tree matches the current act, it is treated as an interruption and pushed onto the stack without popping anything.

Collagen has recently been extended to perform "near-miss" plan recognition if it cannot find a correct interpretation of an act. It systematically searches for extensions to the plan tree that would explain the current act if some constraint were relaxed. For example, it can recognize acts that would violate an ordering constraint, unnecessarily repeat a step that was already performed, or perform a step that should be skipped because its effects are already satisfied. Thus, near-miss




Figure 3: Paco's Architecture

plan recognition attempts to find plausible interpretations of student errors, providing a domain-independent capability for student diagnosis. It is also extensible, allowing a domain author to define new types of errors or even add explicit buggy recipes. Additionally, Collagen is being extended to use causal information in recipes to repair plans after an incorrect action, or external event, occurs.

3.1 Architecture

Figure 3 shows how Paco fits into the general Collagen architecture. The three software components in this architecture are the simulator, Collagen, and the agent (e.g., Paco). Collagen makes very few assumptions about the simulator. Primarily, it assumes that the user (e.g., student) and agent (e.g., Paco) can both perform domain actions (e.g., open a fuel valve) and can observe the actions taken by each other. Collagen provides an API for such event messages, so that it will be able to interpret them. Collagen makes no assumptions about the simulator's user interface. The simulator can, however, optionally specify a screen location for domain actions, which allows the agent to use a pointing hand to draw the user's attention to an object or indicate that the agent is performing an action.

Collagen represents utterances using an artificial discourse language derived from earlier work by Sidner [20]. The language is intended to include the types of utterances that people use when collaborating on tasks. Currently, Collagen's language includes utterance types for agreeing ("yes" and "OK") and disagreeing ("no"), proposing a task or action (e.g., "Let's engage engine one"), indicating when a task has been accomplished (e.g., "We succeeded in stopping engine two"), abandoning a task, asking about or proposing the value of a parameter to a task or action, asking or proposing how a task should be accomplished, and asking what should be done next ("What next?"). Current work is extending Collagen's language to include additional elements from Sidner's language, especially to support negotiation about task decisions.

To bypass natural language understanding issues, Collagen provides a window to allow the user to construct utterances and to display the agent's utterances. In both windows, it converts its internal discourse language into English (or other language) strings, using a combination of domain-independent and (optional) domain-specific text templates. In the user window, users construct utterances by selecting from a menu of utterances and utterance types, and they can modify any utterance by selecting any phrase within it (representing a field in the original text template) and choosing a replacement phrase. Optionally, Collagen can also use speech recognition software to

allow the user to speak these utterances rather than creating them through the GUI, and it can use speech synthesis software to allow the agent to speak its utterances.

4 Tutorial Behaviors as Collaborative Discourse Acts

Table 1 is a summary of our progress in integrating ITS and CDS: it lays out in detail how Paco’s tutorial behaviors are generated from Collagen’s discourse state representation and Paco’s student model. The first column of the table is a ranked list of the tutorial act types. The second column describes procedures that generate zero or more instances of each act type from the current discourse state and student model. When it is Paco’s turn, it constructs a prioritized agenda by evaluating the procedures for each act type and then selects the highest ranked act in this agenda.¹ The third column of the table shows the semantics of each act type in Sidner’s [20] artificial discourse language, which determines how the act will be interpreted by Collagen’s discourse interpretation algorithm. Several of the act types have subcases, shown in the fourth column, which share the same basic semantics, but differ in how they are rendered into English (fifth column).

Paco uses several elements of the discourse state to generate its discourse acts including the focus of attention, the initiative, and plan trees. The focus of attention is used, for example, to avoid teaching a step unless its purpose is in focus. The focus stack also indicates when the student has interrupted the current task, which causes Paco to generate a discourse act which would end the current interruption. In addition to the shared focus maintained by Collagen, Paco also maintains a *private focus* because it prefers to finish teaching an action before moving on. If the student starts working on another part of the plan (thus popping the current focus from the shared focus stack) while there are still legal steps within Paco’s private focus (e.g., line 40 in Figure 1), then Paco will add a Correct Focus action (e.g., line 42) to the agenda. Paco might choose to execute a higher-ranked element on the agenda first (e.g., line 41) but will re-generate the Correct Focus action unless the student returns to the previous subtask by herself.

The various conditions for generating discourse acts are easy to compute given the data structures maintained by Collagen. For example, several of the acts operate on the *valid next actions*, which refers to the plan steps that can be executed next based on precondition and ordering constraints.² Collagen computes this information during discourse interpretation. Additionally, Collagen’s near-miss recognition computes the conditions needed to generate the various subcases of Negative Feedback (e.g., line 46). Finally, when the student asks for help (e.g., line 47) this pushes a discourse purpose of helping the student onto the stack which remains there until the agent provides the help (e.g., line 48).

Using the generic capabilities of Collagen to record information about a user, Paco maintains a simple overlay model [4] that records, for each step in a recipe, whether the student has been exposed to it. In Table 1, the condition “the student knows step ω ” means that the student has been taught this step before. The condition “student knows step ω needs to be done” means the student has been taught all the steps that connect ω to the root of the current plan. Finally, Paco’s student model also records which actions the student has been told that she has completed (e.g., line 17). The condition “the student knows when ω is complete” means that the tutor has told the student when ω was complete, at least once before.

We use Collagen’s generic representation for recipes to store domain-specific knowledge about why actions need to be performed. That is, Paco’s domain knowledge includes recipes that achieve the subgoal of explaining why an action, or more specifically a step of a recipe, should be performed.

¹An agent that was more of an assistant might also include acts in Collagen’s default agenda in its ranking.

²Paco also uses information about the preferred order of executing actions to determine which actions to teach.

Tutorial act type	Add instance to agenda for ...	Semantics	Subcases (if any)	Example gloss
Positive feedback (rank 1)	the user's most recent action α if it was, or proposed, a valid next action and has not yet received feedback	$accept(should(\alpha))$	α finished subtask	Great job.
			α wasn't proposed by tutor	Nice.
			α caused unnecessary focus shift	That was a correct action, but I wasn't expecting you to do it now.
			α finished top-level goal	We're done with this scenario.
			<i>none of above</i>	Good.
Negative feedback (rank 1)	the user's most recent action α if it was, or proposed, an invalid next action and has not yet received feedback	$reject(should(\alpha))$	α was already done	Whoops, you already did that.
			α 's purpose was already achieved	Whoops, you didn't need to do that.
			α has an unsatisfied precondition	Whoops, you can't do that yet.
			executing α violates an ordering constraint	Whoops, it's too soon to do that.
End interruption (rank 2)	each step ω that is an unstopped interruption on the focus stack	$propose(-should(\omega))$	ω has known purpose	Let's stop closing the fuel valves.
			ω has unknown purpose	That is not relevant to our current task.
Teach complete (rank 3)	each non-primitive ω in the current plan <i>s.t.</i> ω is complete and the student does not know when ω is complete	$propose(achieved(\omega))$		We succeeded in closing the fuel valves.
Correct Focus (rank 4)	step ω if it is the tutor's private focus but not the action on top of the focus stack	$propose(should(\omega))$		Let's return to opening the fuel valves.
Give initiative (rank 5)	any valid next plan step ω that the student knows needs to be done, if the tutor has initiative and the student has not requested help	$propose(initiative = user)$	tutor has just proposed ω	Go ahead.
			tutor has not just proposed ω	You take it from here.
Explain Why (rank 6)	every plan step ω that is teachable (see Teach Step) and is currently unexplained and has an explanation recipe	<i>first step of explanation recipe</i>		Before starting an engine, always make sure that the other engine is stopped.
Teach step (rank 7)	every valid next plan step ω that the student does not know and whose parent is in focus	$propose(should(\omega))$	ω is primitive	Now, you should press the on button.
			ω is non-primitive	The next step of engaging the engine is to open the fuel valves.
Remind step (rank 8)	every valid next plan step ω that the student knows and whose parent is in focus	$propose(should(\omega))$		You need to press the on button.
Propose new scenario (rank 8)	purpose ω , if the current plan is complete, where ω is the next task to work on	$propose(should(\omega))$		Let's try another scenario. Let's engage engine one.
Shift Focus (rank 8)	every plan step ω that is not currently on top of the focus stack and the student knows has to be done and has a child c that is a valid next plan step and c is not known by the student	$propose(should(\omega))$		Let's open the fuel valves.

Table 1: Tutorial discourse acts

Typically, these recipes are composed of one or more utterances of text written by a domain expert, but in principle, explanation recipes can contain any type of primitive or abstract actions.³ Whenever Paco generates a candidate discourse act to teach a step, it also checks to see if an explanation recipe exists for that step. If so, and if the step has not already been explained, Paco generates a candidate discourse act of executing the first step of the recipe.

The conditions for generating discourse acts represent necessary, but not sufficient, conditions for Paco to perform the act. An advantage of making explicit all necessary conditions for a discourse act is to make it easier to extend Paco with new discourse acts or extend other agents with the ability to perform Paco's tutorial actions. However, this approach leaves open the question of how to choose which act to perform. Paco chooses which act to perform based on the rankings of the discourse acts, given in the first column of Table 1. For example, Paco prefers to give initiative when the student knows what to do next rather than teach or remind her what to do next. We hypothesize that different rankings or other methods for choosing an act from the agenda will produce different tutoring styles.

5 Discussion

To facilitate comparison between Paco and other tutorial dialogue systems, the following list outlines some of the main dimensions along which such systems can be compared, categorizes Paco along these dimensions, and provides some of the motivations and trade-offs involved in our design choices:

- Our work focuses on the pragmatics of natural language understanding, i.e., the use of a discourse interpretation algorithm and a rich representation of discourse state. Our claim is that tutorial dialogues will be more natural for students if computer tutors follow the principles of human collaborative dialogues, on which much research in computational linguistics has focused. We do not yet have strong evidence to substantiate this claim, but investigating that hypothesis is the primary focus of our work.
- Paco performs relatively sophisticated domain reasoning, based on the application of Collagen's domain-independent algorithms to a domain-specific task model (recipe library). Specifically, Collagen decides which domain actions can be done next based on its recipe library, its knowledge of which actions and utterances have been performed so far, and its knowledge of the current simulation state. Its reasoning does not yet include a full planner, as can be found in Rickel and Johnson's Steve tutor [19], but we recognize that such planning capabilities are important in many domains, and we are currently extending Collagen in that direction. Collagen's advantage over Steve is that it requires less domain knowledge (specifically, it does not require causal links among task steps), but this limits its ability to recover from some student errors (e.g., that would require repeating earlier actions) and to recognize when some steps can be skipped (e.g., because they only establish preconditions for later steps whose postconditions are already satisfied).
- Paco currently uses a simple overlay student model. The student model is crucial for Paco's approach to interleaving demonstration and coached practice. We do not currently use a bug library, which would allow Paco to recognize common errors and provide more specific feedback aimed directly at those errors, but Collagen's near-miss capability is capable of exploiting such knowledge if it is provided.

³Collagen's facilities for executing recipes in a collaborative setting can be used to complete the explanation

- Text and gesture generation are currently relatively simple in Paco, but this is only a matter of research focus. Collagen uses text templates for text generation, and it uses a pointing hand to direct the student's attention to elements of the simulator. We believe this approach suffices for simple 2D simulations. However, more sophisticated text generation would certainly improve Paco, and elsewhere we have elaborated on the costs and benefits of more fully embodied pedagogical agents [10]. As for graphics, we assume that the simulator will provide appropriate graphics for the simulated world, and some additional graphics could be useful in helping students understand the inner workings of equipment they are learning to operate [9], but this has not been a focus of our work.
- Paco does not include a conventional dialogue planning module, i.e., an explicit search for a sequence of utterances that will achieve a desired mental state in the student. The agent's utterances are selected (using a simple priority scheme) from the candidate discourse acts that follow naturally from the current discourse state. We are interested in investigating more sophisticated dialogue planning, but we have no strong evidence yet that it will be required for teaching procedural tasks. One intermediate position that we are currently exploring is the use of "tutorial recipes," which can be viewed as cached dialogue plans. Collagen can use such recipes to guide its interaction with students using the same mechanisms by which it uses domain recipes. Also, Collagen's plan trees can be viewed as plan-like structures that encode expectations for future utterances and actions that will complete the current task, including ordering constraints and subgoal relationships among these discourse acts. Thus, while Paco does not plan its dialogue acts in a traditional sense, its plan trees play a similar role.
- Paco does not allow free-form student utterances, so it does not include any parsing or semantic interpretation of sentences. Instead, the student constructs utterances through a GUI. This is mainly because we are focusing on the dialogue manager; we are not making any claims about the utility of natural language understanding for teaching procedural tasks. However, we do believe that a GUI will be adequate for teaching many procedural tasks, although full natural language understanding would certainly be better if it could be achieved.

The use of Collagen as a dialogue manager for a tutorial system, as an alternative to building such a system from scratch, also presents some trade-offs. To connect an application and agent to Collagen, one must make several commitments. First, one must write a software module that maps application events into Collagen discourse acts and vice versa. However, a similar module is required to connect any tutor to an external simulator, and Collagen provides a nice interface for making such connections. Second, Collagen requires a recipe library that encodes domain task knowledge, but, again, something similar will be required for any intelligent tutoring system for procedural tasks. One important commitment is that the domain task knowledge must be expressed in Collagen's recipe library representation, as opposed to having the freedom to express it procedurally (e.g., as production rules) or through a custom declarative language. The biggest disadvantage this poses is that Collagen may not exploit some types of knowledge (e.g., causal links or temporal constraints) that are important in a domain, or its semantics (e.g., the definition and implications of ordering constraints) may not be appropriate for some domains. Similarly, one must map all student and tutor utterances into Collagen's act types, although this may not be a serious limitation since Collagen allows new act types to be added. The benefit of providing a recipe library and mapping to Collagen's act types is that it maintains the discourse state based on principles from collaborative discourse theory, and it includes both normal and near-miss plan recognition.

We are interested in several areas of future work. Paco thus far has been primarily a reimplementation (on a new foundation) of fairly standard ITS behaviors. As the next step, we plan to better

leverage Collagen's rich discourse state representation to implement aspects of tutorial dialogue that have not been treated in a fully general way in previous ITS work. We are also interested in broadening the types of tutorial discourse acts we consider to include those used in recent analyses of human tutorial dialogues [11, 15, 16], and we are especially interested in exploring the relationship of "hinting" strategies to collaborative discourse theory. Some of these issues may require integrating information in Paco's student model into Collagen's discourse interpretation algorithm. Finally, we would like to experimentally evaluate Paco's ability to teach procedural tasks.

6 Conclusion

In conclusion, we believe that building Paco has been a demonstration of successful cross-fertilization between research in intelligent tutoring and collaborative dialogue systems in at least three respects. First, we showed how a variety of tutorial behaviors can be expressed as rules for generating candidate discourse acts in the framework of CDS. This allows us to immediately apply many notions from CDS in our tutorial agents.

Second, building Paco has given us the opportunity to evaluate the suitability of a particular piece of CDS technology, namely Collagen, for building ITS systems. Our experience has been that using Collagen as the starting point for implementing Paco was a great improvement over programming tutorial agents "from scratch," as we have done in the past. Also, using Collagen led us to design Paco as a composition of a generator of candidate discourse acts and a set of preferences for selecting from these acts. This approach makes it easier to understand, explain, and share tutorial behaviors.

Third, building a tutorial agent in Collagen has revealed some implicit biases in how Collagen operates. As a result, we are exploring various generalizations and extensions to Collagen to better support the full spectrum of collaboration.

References

- [1] J. R. Carbonell. AI in CAI: An artificial-intelligence approach to computer-assisted instruction. *IEEE Transactions on Man-Machine Systems*, 11(4):190–202, 1970.
- [2] A. Collins, J. S. Brown, and S. E. Newman. Cognitive apprenticeship: Teaching the crafts of reading, writing, and mathematics. In L. Resnick, editor, *Knowing, Learning, and Instruction: Essays in Honor of Robert Glaser*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1989.
- [3] R. K. Freedman. *Interaction of Discourse Planning, Instructional Planning and Dialogue Management in an Interactive Tutoring System*. PhD thesis, Northwestern University, 1996.
- [4] I. P. Goldstein. Overlays: A theory of modelling for computer-aided instruction. Artificial Intelligence Laboratory Memo 495, Massachusetts Institute of Technology, Cambridge, MA, 1977.
- [5] B. J. Grosz and S. Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269–357, 1996.
- [6] B. J. Grosz and C. L. Sidner. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3):175–204, 1986.
- [7] B. J. Grosz and C. L. Sidner. Plans for discourse. In P. Cohen, J. Morgan, and M. Pollack, editors, *Intentions in Communication*, chapter 20, pages 417–444. MIT Press, 1990.
- [8] B. J. Grosz [Deutsch]. The structure of task oriented dialogs. In *Proceedings of the IEEE Symposium on Speech Recognition*, Pittsburgh, PA, April 1974. Carnegie-Mellon University. Also available as Stanford Research Institute Technical Note 90, Menlo Park, CA.

- [9] J. D. Hollan, E. L. Hutchins, and L. Weitzman. Steamer: An interactive inspectable simulation-based training system. *AI Magazine*, 5(2):15–27, 1984.
- [10] W. L. Johnson, J. W. Rickel, and J. C. Lester. Animated pedagogical agents: Face-to-face interaction in interactive learning environments. *International Journal of Artificial Intelligence in Education*, 11:47–78, 2000.
- [11] S. Katz, G. O'Donnell, and H. Kay. An approach to analyzing the role and structure of reflective dialogue. *International Journal of Artificial Intelligence in Education*, 11:320–343, 2000.
- [12] N. Lesh, C. Rich, and C. L. Sidner. Using plan recognition in human-computer collaboration. In *Proceedings of the Seventh International Conference on User Modeling*, pages 23–32, Banff, Canada, 1999.
- [13] K. E. Lochbaum. A collaborative planning model of intentional structure. *Computational Linguistics*, 24(4):525–572, 1998.
- [14] N. K. Person, A. C. Graesser, R. J. Kreuz, V. Pomeroy, and the Tutoring Research Group. Simulating human tutor dialog moves in autotutor. *International Journal of Artificial Intelligence in Education*, 12, 2001. Forthcoming.
- [15] K. Porayska-Pomsta, C. Mellish, and H. Pain. Aspects of speech act categorisation: Towards generating teachers' language. *International Journal of Artificial Intelligence in Education*, 11:254–272, 2000.
- [16] A. Ravenscroft and R. M. Pilkington. Investigation by design: Developing dialogue models to support reasoning and conceptual change. *International Journal of Artificial Intelligence in Education*, 11:273–298, 2000.
- [17] C. Rich and C. L. Sidner. COLLAGEN: A collaboration manager for software interface agents. *User Modeling and User-Adapted Interaction*, 8(3-4):315–350, 1998.
- [18] J. Rickel. An intelligent tutoring framework for task-oriented domains. In *Proceedings of the International Conference on Intelligent Tutoring Systems*, pages 109–115, Montréal, Canada, June 1988. Université de Montréal.
- [19] J. Rickel and W. L. Johnson. Animated agents for procedural training in virtual reality: Perception, cognition, and motor control. *Applied Artificial Intelligence*, 13:343–382, 1999.
- [20] C. L. Sidner. An artificial discourse language for collaborative negotiation. In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94)*, pages 814–819, Menlo Park, CA, 1994. AAAI Press.
- [21] D. Sleeman and J. Brown, editors. *Intelligent Tutoring Systems*. Academic Press, 1982.
- [22] D. R. Traum. *A Computational Theory of Grounding in Natural Language Conversation*. PhD thesis, Department of Computer Science, University of Rochester, Rochester, NY, 1994.
- [23] E. Wenger. *Artificial Intelligence and Tutoring Systems*. Morgan Kaufmann, Los Altos, CA, 1987.
- [24] B. P. Woolf. *Context-Dependent Planning in a Machine Tutor*. PhD thesis, Department of Computer and Information Science, University of Massachusetts at Amherst, 1984.