Retrieval strategies for tutorial stories*

Robin Burke
Institute for the Learning Sciences
1890 Maple Ave., Evanston IL 60626, USA
burke@ils.nwu.edu

Abstract

Retrieving stories to present to students is a challenging application of case retrieval. This paper describes SPIEL, a system that retrieves tutorial stories, stored on video, for educational purposes. Although CBR methods are employed in SPIEL, its task requires a different emphasis than typically found in problem-solving CBR systems. One of the most significant of these is the centrality of multi-purpose retrieval in educational storytelling. SPIEL has a set of storytelling strategies, corresponding to different educational roles that stories can play, such as providing counter-examples or projecting possible results. To find stories that can fill these roles, the system uses a variety of comparisons including similarity, dissimilarity, and other relations. This paper describes three of these strategies in detail, showing how the strategies function in retrieval, what kinds of knowledge they use, and how they make use of SPIEL's indices.

Introduction

Is case retrieval a quick, associative process or is it a strategic, inferential one? Case retrieval can be viewed as a simple process that gathers raw material for the rest of the case-based reasoning process to use (Waltz 1989). However, there are many uses for cases that require highly specific comparisons driven by strategic considerations. I have been investigating strategic case retrieval for educational purposes by building a program, SPIEL (Story Producer for InteractivE Learning), to approximate what an instructor does in recalling a pedagogically-appropriate story to tell to a student.

Tutorial storytelling involves using stories in a variety of educational roles. They can project possible results of students' actions, provide counter-examples, and give suggestions, to name just a few possibilities. This paper describes in detail three of the retrieval strategies that are used to produce SPIEL's storytelling behavior. Each involves a comparison between a story and the situation in which it might be told, a comparison that differs according to the needs of the educational context. SPIEL's retrieval strategies select different subsets of features from a story's index and use a variety of measurements of fitness including similarity, dissimilarity and other kinds of relations.

What SPIEL does

SPIEL is designed to assist students who are learning social skills. It is embedded in an intelligent learning-by-doing architecture called **Guided Social Simulation** or **GuSS**. GuSS provides a social simulation in which students can safely practice social skills, such as those required by diplomacy or business. Currently, we are using this architecture to develop an application, YELLO, for teaching employees of Ameritech Publishing the fine points of selling Yellow Pages advertising. The goal in YELLO, and other GuSS applications, is to accomplish for the social environment what the flight simulator accomplishes for the physical environment of the cockpit, letting the student learn by doing. For more about GuSS and its precursor ESS, see (Kass et al. 1993), (Kass & Blevis 1991) and (Blevis & Kass 1991).

Within GuSS, SPIEL is like an experienced instructor watching over the student's shoulder. It monitors the simulation and presents stories from its library when they are relevant to the student's situation. Telling stories in the context of simulation is a particularly useful way to connect the student with an expert's experience. Stories help bring the simulation to life, and the student's activity in the simulation helps make the stories comprehensible.

SPIEL's stories are video clips of practitioners telling anecdotes about their own on-the-job experiences. SPIEL's knowledge of its stories comes from manually-constructed indices entered into the system to describe each story. The system currently has about 180 such stories gathered from interviews with Yellow Pages account executives. An earlier version of the program, described in (Burke & Kass 1992), contained 170 stories about selling consulting services.

* This work is supported in part by the Defense Advanced Research Projects Agency, monitored by the Air Force Office of Scientific Research under contract F49620-88-C-0058 and the Office of Naval Research under contract N00014-90-J-4117, by ONR under contract N00014-J-1987, and by the AFOSR under contract AFOSR-89-0493. The Institute for the Learning Sciences was established in 1989 with the support of Andersen Consulting, part of The Arthur Andersen Worldwide Organization. The Institute receives additional support from Ameritech and Northwest Water, Institute partners, and from IBM.

	Standard problem-solving	Story retrieval in SPIEL
Cue composition	Before retrieval	Incremental
Retrieval criteria	Solves a similar problem	Makes an educational point
Case structure	Represents problem solution	Video clip
Mandatory retrieval	Yes	No
Case evaluation	Yes	No
Between-case competition	Yes	No

Table 1. Differences between case retrieval in problem-solving CBR and story retrieval in SPIEL

Story retrieval vs. case retrieval

Educational tasks emphasize different requirements for case retrieval than most problem-solving tasks. Table 1 summarizes some of the most important differences in emphasis.

Cue composition: One of the major differences between SPIEL and the standard problem-solving CBR model (Kolodner & Jona 1992) is the way retrieval cues are put together. CBR systems create a retrieval cue by analyzing a statement of the problem to be solved. SPIEL retrieves its stories based on a continuous stream of actions by the student and the GuSS simulation. Any event in the history of the interaction could be relevant to the retrieval of a tutorial story. So, SPIEL has to build its retrieval cues incrementally throughout this on-going process.

Mandatory retrieval: A case-based problem solver must retrieve something, otherwise there will be no basis for building a solution. If what is retrieved is not a directly-applicable solution, it can be adapted. However, in teaching, there are many student states for which there will be no appropriate story to tell. Because it cannot adapt its stories, SPIEL has to retrieve only closely-relevant ones. A story that is far afield from the student's immediate concerns will be confusing. If SPIEL cannot find a very good story, it is better off waiting for the student to do something else. Usually, there is an appropriate story about once every 10-20 student actions.

Case evaluation: In the standard CBR model, the retrieval step uses indices to retrieve a set of candidate cases. The candidate cases themselves are then examined and evaluated, and the best case is chosen as the basis for a problem solution. The evaluation step enables the system to reject inappropriately retrieved cases. SPIEL has no language capacity with which to understand or evaluate the video clips it retrieves. It must retrieve conservatively, because whatever is retrieved will be presented to the student.

Between-case competition: Educational stories are not in competition to be the single right answer, as in many retrieval models, such as (Thagard et al. 1990). If there are stories from experts with opposing viewpoints on the student's situation, SPIEL needs to show the student the whole range of opinion by making all retrieved stories available.

These differences have three notable consequences for the design of a tutorial case retriever.

Because SPIEL does not have access to the content of what it is retrieving, the index has to contain everything that the system will need to decide to tell the story. It needs indices that are more detailed and complex than those typically used in case retrieval.

There is an implicit theory of problem solving in the similarity metrics found in problem-solving CBR: the more similar the input problem is to the problem solved by the case, the more likely it is that the solutions will also be similar. SPIEL's retrieval strategies also embody a theory of good cases for the educational context. They involve different kinds of judgments from those found in similarity metrics.

In SPIEL, a tutorial opportunity is defined as a situation for which there is a relevant story to tell. It does a storyteller no good to recognize a good time to tell a story it doesn't have. A similar insight was behind the design of ANON (Owens 1990), which used its case base to determine what features to search for in the input problem. SPIEL uses its story base in a similar way to guide the search of the events in the simulation.

The rest of this paper focuses mainly on retrieval strategies, but I touch briefly on the indexing and implementation issues to put the strategies in context.

Implementing storytelling strategies

One way to think about the problem an educational storyteller faces is to think of each story as a possible lesson and each strategy as a way to teach it. Storytelling strategies indicate the conditions under which the goal of telling a story can be achieved. Since a storyteller has many stories, any one of which could be relevant at any time, it is useful to think of it as an opportunistic system.

However, educational storytelling is not as open-ended as the general case of opportunism whose complexities were laid out in (Birnbaum 1986). SPIEL's storytelling strategies compare stories about an activity to specific contexts of student action. They seek coherence and relevance, not distant analogies. It is possible to describe concretely what an opportunity to tell a story using a particular strategy would look like.

SPIEL also has a strong advantage over the problem of opportunism in that it is embedded in a simulation. No truly novel actions can occur in GuSS since the student is constrained by the program's interface and the simulation operates in known ways. SPIEL knows with certainty what actions can and cannot happen in its world.

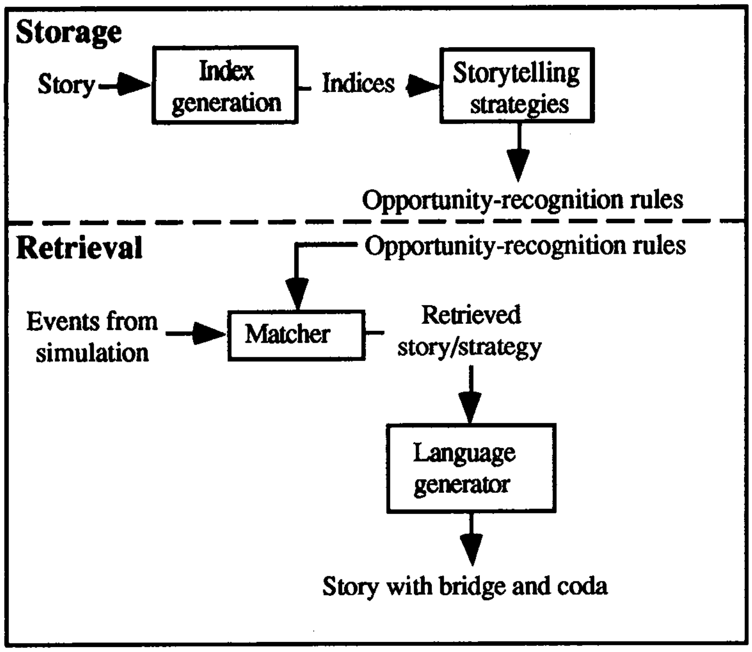

Figure 1. SPIEL's processing

SPIEL's problem of opportunism is therefore much simpler than the general one. The simulated world provides a limited space in which events can occur; storytelling strategies single out precise areas of the space that constitute opportunities. These properties enable SPIEL to use what Birnbaum calls the "elaborate and index" model of opportunism:

> ...spend some effort, when a goal is formed, to determine a number of situations in which it might be easily satisfied...and then index the goal in terms of all the features that might arise in such situations. (pg. 146)

SPIEL works from its database of stories to determine what is an opportunity for intervention. Its processing can be divided into two phases: storage time, when new stories are put into the system and the system considers how it might tell them; and retrieval time, when a student interacts with the GuSS system and and SPIEL watches for opportunities to give tutorial feedback. Figure 1 shows these phases.

At storage time:

1. Indices are attached to each story in the database.
2. SPIEL's storytelling strategies are applied to each index. If the strategy is applicable to the index, a set of rules is generated that will recognize an opportunity to tell that story using the storytelling strategy.

At retrieval time:

1. Opportunity-recognition rules are matched against the state of the simulation.
2. When the rules for a particular combination of story and strategy match successfully, the story is retrieved.
3. After retrieval, natural language bridge and coda are generated to integrate the story into the student's current context.

Indices for educational stories

SPIEL's indices are created manually. Since the stories are in video form, automatic processing would entail speech (and possibly gesture) recognition as well as natural language understanding. SPIEL's design therefore calls for a human indexer to watch each story being told and use an indexing tool to compose indices that capture interpretations of the story's meaning.

These interpretations have the general form, "X believed Y, but actually Z," which is a form of *anomaly*. An anomaly is a failure of expectation that requires explanation (Schank 1982). Typically, anomalous occurrences are what make stories interesting and useful, and they are a natural way to summarize what a story is about. Anomalies are especially important in stories about social activity since students are learning what expectations they should have and how to address expectation failures.

The anomaly forms the core of SPIEL's indexing representation. Consider the story transcribed here whose index is shown in Figure 2.

> I went to this auto glass place one time where I had the biggest surprise. I walked in; it was a big, burly man; he talked about auto glass. So we were working on a display ad for him.
>
> It was kind of a rinky-dink shop and there was a TV playing and a lady there watching the TV. It was a soap opera in the afternoon. I talked to the man a lot but yet the woman seemed to be listening, she was

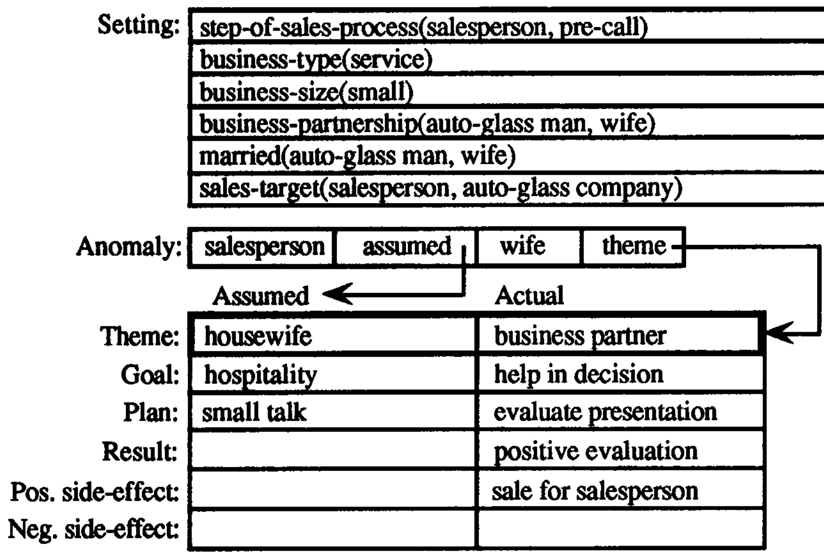

Figure 2. Index for "Wife watching TV" story

asking a couple of questions. She talked about the soap opera a little bit and about the weather.

It turns out that after he and I worked on the ad, he gave it to her to approve. It turns out that after I brought it back to approve, she approved the actual dollar amount. He was there to tell me about the business, but his wife was there to hand over the check.

So if I had ignored her or had not given her the time of day or the respect that she was deserved, I wouldn't have made that sale. It's important when you walk in, to really listen to everyone and to really pay attention to whatever is going on that you see.

The index contains

- the anomaly in the story, which can be phrased as "the salesperson assumed the wife would have the role of housewife, but actually she was a business partner."
- the setting, the story's position within the overall social task, including a representation of the social relationships between the actors in the story, and
- intentional chains surrounding and explaining the anomalous occurrences.

SPIEL's indices are considerably more complex than those typically found in case-based reasoning systems. The fact that SPIEL's cases are video clips is part of the reason. Its indices have to say everything that SPIEL needs to know about its stories since the stories themselves cannot be evaluated. Another reason that SPIEL's indices show complexity is the task of educational storytelling itself. SPIEL has a variety of reasons for telling stories, each of which requires a slightly different perspective on the index. SPIEL needs complex indices to meet the demands of a variety of strategies. The three strategies I describe next show some of these uses. The others can be found in (Burke in preparation).

The Warn about plan storytelling strategy

A storytelling strategy is a way of using a story to teach. Consider Example 1, which shows the **Warn about plan** strategy in action. The student is pressing for a large ad campaign, much larger than the client needs or can really afford. The customer is an evasive type, and does not immediately reject the idea: instead, he stalls. The student could lose a lot of time in a futile effort to make this sale. At this point, the storyteller intervenes with a story about an analogous situation where, instead of stalling, the customer rejected the ad program and rebuked the salesperson. Telling the story at this point helps the student identify the problem before going too far along in this direction.

An accurate assessment of how far to let the student go would require a great deal of knowledge about the scenario the student is engaged in, more knowledge than is available to SPIEL. GuSS's open-ended simulation design does not allow any quick read-out of what outcomes are likely and how far away the student is from obtaining them. SPIEL would have to perform a complete envisionment of possible simulation outcomes to find out, more processing than can be accomplished in the midst of student interaction. Instead, SPIEL uses prototypical knowledge about inconclusive outcomes. For SPIEL to tell the "Would you buy this ad?" story in the example, it must know that overcoming the customer's stall in this kind of case is time-consuming and not particularly educational.

At storage time, each storytelling strategy selects stories that are compatible with its educational role. A story with a good outcome could not be sensibly used with **Warn about plan**, for example. The second task of the storytelling strategy is to generate, for each compatible story, a Recognition Condition Description (RCD), a representation that describes a situation in which the story

- The student doesn't gather very much information at the pre-call stage.
- Back at the office, the student prepares a very large ad campaign.
- When the student presents this to the client, the client says "You know, I really have to talk to Ed about this..." and is inwardly very doubtful of the value of such a large expenditure.

Storyteller: A story about a failure in a situation similar to yours...

You made a recommendation that was much larger than Dave's expectations. Here is a story in which doing that led to problems:

"I remember my first year 1970. I was on my first Yellow Page sales canvass. In those days, they didn't give you a lot of time to show ability and I wasn't doing very well. My manager told me that I had one week to start producing or they were going to let me go.

"I called on a graphic artist in Indianapolis. He had a 2HS, a one-inch ad. I walked in, asked two minor questions, and I laid down a quarter-page piece of spec in front of this man and told him he needed this ad. The man looked at me and he said, 'Would you buy this ad?' He turned it right back to me. I didn't know what to say. It shocked me. Finally, I said, 'No, I wouldn't.' He didn't need a quarter page; a one-inch ad is what he needed. The gentleman proceeded to give it to me, up one side and down the other. He told me that I was there for my own greedy interests, trying to make commissions instead of taking care of his advertising and caring about him. He said, 'If you're ever going to make it in this business, you'd better start paying attention to your customers.'

"I walked out of that call a different salesman because I realized then that the only way to sell Yellow Pages is to sell what the customer needs, not what I need. Learning that lesson turned my sales career around. I started making sales, and by the end of the canvass, I was one of the top producers."

You made a recommendation that is much larger than the client's expectations. That might not be a good idea.

Example 1. Telling "Would you buy this ad?" using the **Warn about plan** strategy

could be told using the strategy, a tutorial opportunity provided by the story.

Computing the RCD involves making inferences based on the story's index, essentially asking the question "What would have to happen for this story to be good to tell?" The answer to this question is different for every storytelling strategy. For **Warn about plan**, the situation should be pretty much the same as in the story: the student should have overshot the mark by designing a large ad program. The difference is that the student should be unlikely to discover this error. Immediate rejection, which is the bad outcome of the story, is not what should be looked for. Instead, the system needs an inconclusive outcome that bodes ill.

What occurs in Example 1 is one kind of outcome of this kind: if the person the student is talking to has authority to buy and that person defers the decision to someone else, this is usually an evasive maneuver, not motivated by a real need for consultation. Also, if the customer delays the decision: "I don't have time for this now. Let's talk about it next week." Neither of these objections are associated with any of the important educational goals of the system, so intervention is justified.

A description of the RCD that characterizes the tutorial opportunity would therefore look something like this:

WHEN the student is closing the sale and speaking to someone who is the decision maker,
LOOK FOR the student to present a very large ad program,
the client to have a negative belief about that program, and
the client to defer the decision to another,
or the client to put the decision off to another time,
THEN TELL "Would you buy this ad?"
AS a "Warn about plan" story.

This is the format of an RCD. It contains a trigger (the "when" part) that describes the conditions under which the opportunity becomes possible, usually a function of the stage of the sales process that the student is in, combined with characteristics of the other agents involved in that stage. If the triggering conditions are met, the system tries to gather evidence that similar intentions are at work in the simulation.

Warn about plan can also be used to show effects that would not appear within the scenario. For example, scare tactics may persuade a customer to buy once, but they hurt the client relationship and eventually result in lost business. This problem does not show up in YELLO since a scenario ends when the student makes one sale. A student who successfully uses this tactic in the simulation might think it is a good idea in general unless the storyteller can show an example to demonstrate otherwise.

Another use of the storytelling strategy is in the generation of the natural language texts that surround and explain the story. As shown in Example 1, there are three parts to the explanation: the headline, for the attention-getting initial statement; the bridge, the introductory paragraph explaining why the story has come up; and the coda, that closes the story presentation with a recommendation or evaluation. Each strategy has a set of natural language templates for these texts, which are filled in by generating appropriate phrases at retrieval time.

The Demonstrate risks storytelling strategy

There is tremendous variability in the social world. Approaches that have always worked may fail in a new situation for no apparent reason; unlikely strategies may fortuitously succeed. Both in GuSS's simulator and in real life, the learner gets to see only one outcome at a time. One important role that real-world stories can play is to illustrate the range of real experience by presenting counter-examples that contrast with what is happening to the student. Studies of apprenticeship situations indicate that experts often use stories in exactly this way (Lave & Wenger 1991).

The **Demonstrate risks** strategy shows that a successful result in the simulation is not always repeatable in the real world. Consider the situation where the student closes a sale, but continues conversing with the customer. This is risky, since it gives the customer an opportunity to reconsider. To show the risk, SPIEL tells about a salesperson who lost a sale through a careless remark made after a sale was closed.

This strategy resembles **Warn about plan** because it uses a story about a failure to warn the student. The difference is that the **Warn about plan** strategy is used when the simulation is not likely to give immediate feedback about the student's actions. The **Demonstrate risks** strategy looks for the simulation giving feedback, but of the wrong kind. It waits for the simulation to do something opposite from the outcome found in the story, and tells a story to show how the real world can be different from the simulation.

This approach to retrieving counter-examples differs from existing CBR systems that use contrasting examples, such as HYPO (Ashley & Rissland 1987). HYPO retrieves cases based on similarity along certain dimensions and builds a structure, the claim lattice, of the retrieved cases, enabling it to identify contrasts along other dimensions. SPIEL uses its knowledge of contrasts in the retrieval process itself, so that it only retrieves those cases that make the needed educational point.

To arrive at the recognition condition description for such a tutorial opportunity, SPIEL must look for a result that would be opposite from what occurs in the story. The story described above showed a salesperson losing a sale during conversation after the close of a sale. SPIEL uses an opposite-finding inference mechanism to identify a outcome which is opposite from the one shown in the story. The story is a good counter-example if the student does a similar action, but does not lose anything.

The "Warn about assumption" strategy

Newcomers to a social domain may inappropriately transfer expectations from the rest of their social lives. A new salesperson may think, for example, that a friendly, talkative customer is more likely to buy than a "get down to business" type, when in many cases, the opposite is true. Since the indices used for stories incorporate the difference between a person's view of the world and how the world actually turned out to be, SPIEL is well poised to help students by pointing out their unrealistic expectations.

If the program has evidence that the student has a particular assumption, the **Warn about assumption** strategy calls for it to tell a story about a time when a similar assumption was wrong. Suppose the student has an opportunity to gather information about a client's business from the client's spouse, but does not take that opportunity. It is reasonable to infer that the student assumes that the spouse does not have an important role in the business. SPIEL can tell the "Wife watching TV" story at this point, showing a case where a similar assumption proved wrong.

I call this type of strategy is a *perspective-oriented* strategy, because what is important is the contrasting perspective in a story. Recognizing stories that are relevant in this way is more difficult than recognizing stories that contain similar actions. Perspective-oriented strategies need knowledge about how certain actions are characteristic indicators of the student's mental state. Students do not always act in ways that clearly indicate their beliefs, but if they do, the system should be prepared to respond. In this story, the salesperson assumes, in an information-gathering context, that the client's spouse will not have a business role. What actions would indicate such an assumption on the student's part? In information gathering, the primary task is to ask questions about the client's business. If the student fails to ask such questions of someone when given the opportunity, this is a good indication of an assumption on the student's part that there is no information to gather. Another possible indicator would be a deliberate act on the student's part to exclude the spouse from the sales presentation.

Conclusion

The three storytelling strategies described in this paper and summarized in Table 2 show some of the demands that the task of educational storytelling places on a case retrieval system. The retriever must not only respond to crucial similarities between a story and the situation in which it is told, but also differences, such as opposite outcomes as called for by the **Demonstrate risks** strategy, and

Storytelling strategy	Summary	Story is about:	Tutorial opportunity
Warn about plan	Tell a story about an unsuccessful plan when the student has begun executing a similar plan.	The negative outcome of a plan.	Look for similar setting, similar goal and similar plan. Look for a negative, but inconclusive outcome.
Demonstrate risks	Tell a story about a negative result of a particular plan when the student has just executed a similar plan but had success.	The negative outcome of a plan.	Look for similar setting, similar goal and similar plan. Look for an opposite outcome.
Warn about assumption	Tell a story about an erroneous assumption that someone made when the student appears to have made the same assumption.	An assumption that didn't hold.	Look for similar setting. Look for actions that are indicative of the assumption.

Table 2. Summary of the three storytelling strategies

evidential relations, used by **Warn about assumption** and other perspective-oriented strategies that look for actions that indicate beliefs.

The design of SPIEL is a response to these demands. It uses structured, strategic comparisons between stories and the situations in which they are told. These retrieval strategies are an explicit counterpart to what is implicit in the similarity metrics used in standard problem-solving case retrieval: a theory of what cases are useful for the task. Extensions of the standard problem-solving paradigm (Kolodner, 1989) and new applications for case-based reasoning technology, such as education, entail new notions of utility, and with them, new metrics that combine similarity judgments and other measures of fitness.

Putting such strategies to work need not render case retrieval prohibitively inefficient. Since strategies are implemented as procedures that operate at storage time, SPIEL performs a minimum of inference at retrieval time, yet still remains sensitive to the strategic considerations involved in recognizing tutorial opportunities.

Acknowledgments

The GuSS project has prospered under the guidance of Roger Schank, Alex Kass and Eli Blevis. I would like to acknowledge their contribution to the ideas in this paper. GuSS is the product of many minds and efforts including those of Mark Ashba, Greg Downey, Debra Jenkins, Smadar Kedar, Jarret Knyal, Maria Leone, Charles Lewis, John Miller, Thomas Murray, Michelle Saunders, Jeannemarie Sierant, and Mary Williamson. The sales and sales training organizations of Ameritech Publishing have also contributed to this research in many ways.

References

Ashley, K., and Rissland, E. 1987. Compare and Contrast, A Test of Expertise. In Proceedings of the 6th National Conference on Artificial Intelligence. AAAI.

Birnbaum, L. A. 1986. Integrated Processing in Planning and Understanding. PhD thesis, Dept. of Computer Science, Yale University.

Blevis, E. and Kass, A. 1991. Teaching by Means of Social Simulation. In Proceedings of the International Conference on the Learning Sciences. AACE.

Burke, R. 1992. Knowledge acquisition and education: a case for stories. In Proceedings of the Symposium on Cognitive Aspects of Knowledge Acquisition. AAAI.

Burke, R. in preparation. Representation, Storage and Retrieval of Stories in a Guided Social Simulation. PhD thesis, Dept. of EECS, Northwestern University.

Burke, R. and Kass, A. 1992. Integrating Case Presentation with Simulation-based Learning-by-doing. In Proceedings of 14th Annual Conference of the Cognitive Science Society, 629-634. Hillsdale, NJ: Lawrence Erlbaum Assoc.

Domeshek, E. 1992. Do the Right Thing: A Component Theory for Indexing Stories as Social Advice. PhD thesis, Dept. of Computer Science, Yale University. Issued as Technical Report No. 26, Institute for the Learning Sciences.

Edelson, D. 1993 Learning from Stories: Indexing, Reminding and Questioning in a Case-based Teaching System. PhD thesis, Dept. of EECS, Northwestern University.

Gentner, D. and Forbus, K. 1991. MAC/FAC: A Model of Similarity-Based Retrieval. In Proceedings of 13th Annual Conference of the Cognitive Science Society. Hillsdale, NJ: Lawrence Erlbaum Assoc.

Hammond, K. 1989. *Case-based Planning: Viewing Planning as a Memory Task.* Boston: Academic Press.

Kass, A. and Blevis, E. 1991. Learning Through Experience: An Intelligent Learning-by-Doing Environment for Business Consultants. In Proceedings of the Conference on Intelligent Computer-Aided Training. NASA.

Kass, A.; Burke, R.; Blevis, E.; and Williamson, M. 1993. Constructing Learning Environments for Complex Social Skills. *Journal of the Learning Sciences.* To appear.

Kolodner, J. 1989. Judging Which is the 'Best' Case for a Case-based Reasoner. In Proceedings of the Case-Based Reasoning Workshop. DARPA/ISTO.

Kolodner, J. and Jona M. 1992. Case-based Reasoning. In *Encyclopedia of Artificial Intelligence (2nd ed.),* S. Shaprio (ed.). New York: Wiley.

Lave, J. and Wenger, E. 1991. *Situated Learning: legitimate peripheral participation.* New York: Cambridge University Press.

Owens, C. 1989. Integrating Feature Extraction and Memory Search. In Proceedings of the 11th Annual Conference of the Cognitive Science Society. Hillsdale, NJ: Lawrence Erlbaum Assoc.

Schank, R. C. 1982. *Dynamic Memory: A Theory of Learning in Computers and People.* New York: Cambridge University Press.

Schank, R. C. 1991. *Tell Me a Story.* New York: Charles Scribner's Sons.

Stevens, S. 1989. Intelligent Interactive Video Simulation of a Code Inspection. *Communications of the ACM,* 32(7), 832-843.

Thagard, P.; Holyoak, K.; Nelson, G.; and Gochfeld, D. 1990. Analog retrieval by constraint satisfaction. *Artificial Intelligence,* 46(3): 259-310.

Waltz, D. L. 1989. Is Indexing Used for Retrieval? In Proceedings of the Case-Based Reasoning Workshop, DARPA/ISTO.

Williams, S. M. 1993. Putting Case-Based Instruction into Context: Examples from Legal, Business, and Medical Education. *Journal of the Learning Sciences.* To appear.