Tailoring Retrieval to Support Case-Based Teaching*

Robin Burke
Department of Computer Science, University of Chicago
1100 E. 58th St., Ryerson Hall, Chicago, IL 60637
burke@cs.uchicago.edu (312) 702-4029

Alex Kass
Institute for the Learning Sciences, Northwestern University,
1890 Maple Ave., Evanston, IL 60201
kass@ils.nwu.edu (312) 491-3500

(To appear in Proceedings of the 13th Annual Conference on Artificial Intelligence, AAAI 1994.)

Abstract

This paper describes how a computer program can support learning by retrieving and presenting relevant stories drawn from a video case base. Although this is an information retrieval problem, it is not a problem that fits comfortably within the classical IR model (Salton & McGill, 1983) because in the classical model the computer system is too passive. The standard model of IR assumes that the user will take the initiative to formulate retrieval requests, but a teaching system must be able to initiate retrieval and formulate retrieval requests automatically. We describe a system, called SPIEL, that performs this type of retrieval, and discuss theoretical challenges addressed in implementing such a system. These challenges include the development of a representation language for indexing the system's video library, and the development of set of retrieval strategies and recognition knowledge that allow the system to locate educationally relevant stories.

1. Introduction

A real-world story told by an expert at an opportune moment can be an invaluable resource for training novices to perform complex tasks (Schank, 1990b). Students do not usually get the benefit of such stories because experts' time is scarce and expensive. However, interactive multi-media technology now makes it possible to store video clips of experts telling their important stories, and to present those clips under computer control. Realizing the educational potential of such stories requires an AI system that understands the educational purposes that a story can serve well enough to be able to recognize when a story should be told. In this paper we discuss a case-based teaching system that can retrieve tutorial stories stored on video as advice for students who are engaged in a learning-by-doing environment. The system, called SPIEL (Story Producer for InteractivE Learning), 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 (Kass et al., 1994) that provides a social simulation in which students can safely practice social skills, such as those required by diplomacy or business.

2. IR, CBR, and Case-Based Teaching

Selecting an instructive story from a video case base is a kind of information retrieval problem, but it is not one that fits comfortably within the classical model of information retrieval (Salton & McGill, 1983) because in the classical model the computer system is too passive. In the standard model of IR, the computer responds to well-formulated retrieval requests issued by a user. This model is appropriate only when intended users can be expected understand their information needs well enough to initiate requests at appropriate times and to formulate those requests correctly. Experts may sometimes understand their information needs well enough to fit this model, but non-experts learning a new, complex task rarely do. Systems intended to teach students to perform a complex task must, therefore, be more than passive data retrievers. They must be able to initiate retrieval and formulate retrieval requests automatically.

Active retrieval, including automatic formation of retrieval cues, has been investigated most extensively in the context of case-based reasoning (CBR) research (Kolodner, 1993). For instance, a case-based planner, such as CHEF (Hammond, 1989a) attempts to solve new problems by retrieving and adapting stored plans, and must be able to extract features of new problems and use those features to retrieve a suitable plan from its library.

The *case-based teaching architecture* (Schank, 1990a; Edelson, 1993; Burke, 1993) is a framework for building computer-based learning environments that present cases (such as experts' stories) to a student who is engaged in a complex task within a computer environment. Case-based teachers are not really case-based *reasoners* in the full sense, because they do not need to manipulate the internal content of their cases, they just need to retrieve and present them at the appropriate time. However, case-based teachers and case-based reasoners have in common the need to initiate the retrieval process and form retrieval cues. Case-based teachers are a good platform from which

* This work is supported in part by the Defense Advanced Research Projects Agency, monitored by the Air Force Office of Scientific Research and the Office of Naval Research. 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.

to study the retrieval problem because they isolate the retrieval problem from the other hard CBR problems, such as evaluation and adaptation.

We claim that the need for precision in educational case retrieval demands a system that understands the uses to which retrieved material will be put. In other words, just as an effective plan retriever must be based on a theory of goals and planning, an effective retriever of tutorial cases must be based on a theory the educational purposes that the retrieved cases can serve. Three types of theories are need; we will touch on each in this paper.

1. **A theory of indexing vocabulary:** a representation language that can be used to form labels for each story by encoding the features of story that may prove important for the system's educational purposes.
2. **A theory of educational relevance:** a taxonomy of the different kind of points that stories can be used to make, a theory of the features of stories that are important for making these kinds of points, and the features of a student's situation that are appropriateness conditions for making each kind of point.
3. **A theory of opportunity recognition:** To recognize storytelling opportunities, the tutor must be able to infer higher-level descriptions of the student’s actions from a stream of low-level event descriptions. It must draw reasonable conclusions about what the student is trying to accomplish, and relate its observations to the contents of the story base.

3. A Case-Presentation Example

The goal of GuSS applications is to accomplish for social environments what other simulators accomplish for physical environments, supporting learning by doing. Within GuSS, SPIEL is like an experienced practitioner 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. The following example illustrates the kind of interaction that SPIEL has with students using the YELLO program, in which the student’s task is to try to sell Yellow Pages advertising. SPIEL has 178 stories told by account executives with great experience in selling. Consider the following interaction in which one of these stories is retrieved:

Student: So, we're going to go ahead with the 1/4 page ad with color?
Customer: OK.
Student: Just sign right here.
Customer: [signs]
Student: I think the color is really going to attract people to your ad.
Customer: I sure like the way it looks.
Student: Ask your customers what they think. I'll bet you'll find its an attention-getter.
Customer: See you next year.
Student: See you. [leaves]

This is a sales success, but SPIEL has a cautionary story to tell. Every moment the student remains after the close of the sale gives the customer an opportunity to retract his buying decision. Although that did not happen here, it is a possibility that the student should be aware of.

SPIEL intervenes by first signaling to the student that it has a story available. Then, if the student expresses an interest in hearing the story, it explains why the story is relevant, shows a video in which an experienced salesperson relates the story as a personal experience, and finally sums up the main point of that story as it applies in the current situation (see Figure 1):

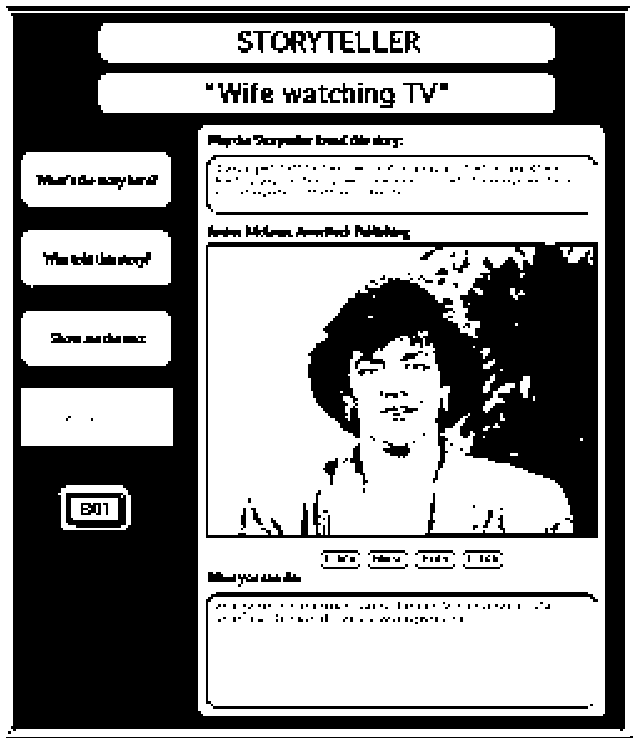

Figure 1. Example of story presentation window

Headline: A story showing the risks of your approach…

Bridge: You kept talking to the client after the sale was closed. Nothing bad happened but here's a story in which doing that led to problems:

Video: I was in the South Bend/Mishawaka area. This was my first or second year. I was calling on a swimming pool contractor. He had a quarter page in South Bend. I was proposing a quarter page in another directory. It was sitting at his kitchen table. And the guy was hesitating; he didn't know... So, after a few more attempts, he says to me "OK, I'll go with the other quarter page." He bought it. I pushed the order form in front of him. He signed it. It's done.

As I'm putting my stuff together, I made this comment that cost me the quarter page. I said, as I'm packing up, "I'm sure that you're going to get a lot of phone calls from that new ad." He looked at me and he said, "You know, I don't need any more phone calls. I'll tell you what, let's not do that this year, maybe next." I talked myself out of a quarter page. I've never done it since. I walked out. There was nothing I could say. I had it and I lost it. All I had to say was "Thank you very much Joe. See you next year." But I didn't. I had to tell him about the calls, which I'd already done twenty times.[1]

Coda: Nothing bad happened to you because you kept talking to the client after the sale was closed, but sometimes the client changes his mind.

In this example, the storyteller augments the student's simulated experience in an important way. Without the "Talked myself out of a sale" story, the student, who was successful, might never realize the risks inherent in remaining after the sale. The story arrives just at the time when it is most relevant, after the risk is past and the student thinks all went well, and it is exactly on point as a counterexample: it shows a situation in which the same tactic had an opposite outcome.

3.1. How the system works

SPIEL operates in storage and retrieval phases. The storage phase prepares the system for storytelling. The storyteller examines its stories in the light of its set of *storytelling strategies*, its library of different ways that stories can be used to teach, and determines what kinds of situations could arise that would constitute opportunities to tell each story. In the retrieval phase, the system tries to recognize those opportunities as the student is performing the task. This design is similar to approaches to "opportunistic memory" architectures proposed for opportunistic planning (Hammond, 1989b). A schematic picture of the phases is shown in Figure 2.

The storage phase begins with the manual construction of indices. An indexer watches the story on video and uses an indexing tool to construct labels representing what it is about. Indices are then processed using storytelling strategies, representations of the tutorial purposes stories can serve. SPIEL determines, for each combination of story and strategy, how and when it might tell the story using the strategy, thus characterizing the storytelling opportunities that the story affords. The rule generator converts descriptions of these opportunities into procedures for recognizing them, in the form of rules that are compatible with the learning-by-doing environment.

[1] SPIEL's stories are video clips from videotaped interviews with experienced Yellow Pages account executives. The surrounding explanations are generated by the program.

The retrieval phase of SPIEL is implemented by a rule-based system that interfaces with the GuSS learning environment. The rules created during the storage phase are matched against on-going events in the simulated world, as the student takes actions and the simulated characters respond. Successful recognition of an opportunity causes a story and its associated strategy to be retrieved. The rest of this paper discusses some of the details of the implementation of SPIEL's retrieval architecture, starting with the indices for stories. For brevity, we have omitted the discussion of SPIEL’s simple natural language generation.

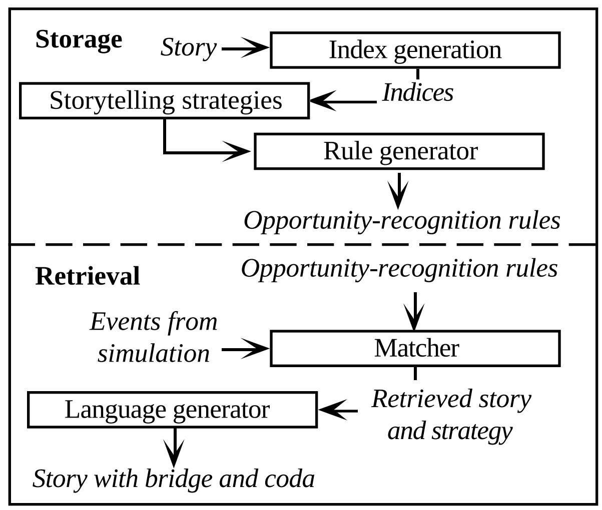

Figure 2. How SPIEL works

4. Indexing stories

In traditional IR, indices are keywords. The meaning of a document is approximated from words found in the text. Because SPIEL's stories are in video form, text ceases to become a "cheap feature." Indices must be entered manually. SPIEL makes use of structured indices that are more precise representations of the story's point or points.

The points of SPIEL’s first-person anecdotes often comes from an expectation violation, a event that surprises the narrator or someone else in the story (Schank, 1982; Schank et al., 1990). In the “Talked myself out of a sale” story, the salesperson is surprised to find his attempt at reassuring the customer backfires into a loss of the sale. These *anomalies* make the stories interesting and they are also the key to the stories’ relevance in educational situations. An indexing representation for such stories must therefore a language for expressing anomalies. The general form of an anomaly can be stated as follows:

> Actor X had an expectation that Y would happen,
> but actually Z happened.

A statement of the anomaly from the "Talked myself out of a sale" story is

The salesperson wanted to reassure the client about his purchase but actually the salesperson lost the sale.

This anomaly captures one important aspect of what the story is about. A story usually says (or implies) more than this, including why the expectation came about and why it failed. The salesperson obviously thought that talking about the increased number of calls the client would get would be reassuring, and thought that reassuring the client was an important goal. The expectation failed because the time spent after the sale gave the customer time to reconsider. Figure 3 is an outline of the “Talked myself out of a sale” index. The index contains the anomaly (shown in bold) and the explanation of the anomaly. The anomaly highlights the feature that is different between the hoped-for case and the actual case.

Viewer:	**Salesperson**	
Perspective:	**wanted**	**actual**
Agent:	**salesperson**	**salesperson**
Goal:	get customer confidence	make sale
Plan:	remain after sale	remain after sale
Result:	**get confidence**	**failure to make sale**
Neg. side-eff.		customer not confident

Figure 3. Index for “Talked myself out of a sale."

5. Determining relevance

SPIEL's ability to tailor its retrieval to the task of education comes from its storytelling strategies, its knowledge of the conditions under which a story with particular characteristics will be relevant to a student. In the example above, the "Talked myself out of a sale" story was told using the **Demonstrate risks** strategy. This strategy selects stories about the failure of a particular course of action and tells them when the student has executed a similar course of action but had success. The story is relevant because it contradicts the evidence of the simulation with respect to the student's recent success. This is an educationally-appropriate way to intervene since, without the story, the student might not see the risks identified in the story.

5.1. SPIEL's storytelling strategies

There are two main sources of relevance for stories in a learning-by-doing environment: (1) they can point out errors that the student may be making when the environment does not provide sufficient feedback for the student, and (2) they can help the student explain expectation failures when the environment does not give all the information needed to construct an explanation. SPIEL has thirteen storytelling strategies, which fall into four categories based on their relationship between the stories they present and the situations in which they present them. The strategies are described in detail in (Burke, 1993).

Strategies that show the student alternatives. These strategies present stories whose outcomes differ substantially from what the student has achieved in the simulated world. **Demonstrate risks** falls in this category. Confronting students with stories about alternative possibilities gets them to question their expectations about the simulated world.

Strategies that critique the student's expectations. By showing examples of situations where people in a similar role had preconceptions that were incorrect, these strategies help students transfer expectations from their everyday social lives to the specialized social environment they are learning about.

Strategies that project possible results of actions the student is taking. Normally, students receive immediate feedback from the execution of their plans. If they do not get such feedback, projection strategies call for the storyteller to provide examples that show possible outcomes.

Strategies that explain the perspectives of other people to the student. These strategies recall stories that explain why people acted as they did in real world situations, making the unexpected actions of others in the simulation comprehensible.

These categories cover the most important ways that stories can be relevant to students using GuSS. They concentrate especially on plans and outcomes, which is consonant with the emphasis in GuSS on learning-by-doing. Students are largely engaged in planning and acting in the social environment. In a different kind of task, such as a design task, students would focus on different kinds of expectations, such as how a design feature achieves a given function and how it will interact with other features. The case-based teaching system Creanimate (Edelson, 1993) has retrieval strategies organized around the problem of design.

5.2. Demonstrate risks: A storytelling strategy

Whenever one confronts a complex problem, it is important to know the range of possible outcomes: What is the best and worst that could happen? Both in real life and in GuSS's simulator, the student gets to see only one outcome at a time. Stories that show a range of alternatives provide necessary balance. Alternative-finding storytelling strategies are designed to bring to the student's attention stories that differ in significant ways from the student's experience in the simulated world.

The idea behind **Demonstrate risks** is to tell a story about a failure when the student has experienced success.

One reason this is important is that no simulation is perfect. If the simulation is not accurately reproducing the circumstances encountered in the story, the story points out possible differences between the simulation and the real world of practice. The example in section 2.1 shows **Demonstrate risks** in action in YELLO. In general, this strategy needs to see that the student is in a situation similar to that in the story, but has achieved an outcome that is very different: the student has succeeded where the story shows a failure. The strategy calls on SPIEL to predict an outcome that would be "very different" from losing the sale, different enough that it makes sense to tell the story as an alternative. Such opposite-finding inference occurs in each alternative-finding strategy.

An opposite outcome is not something that contradicts the original in every respect. It does not make sense to say that the opposite of achieving X is failing to achieve Y, where X and Y are opposite goals. If anything, these would be similar outcomes. An opposite result is an opposite impact on the same goal. This is a simple example, but more complex cases of the opposite computation, such as that in the strategy **Demonstrate alternative plan**, where an opposite plan must be found, require involved opposite-finding inference. The details of the opposite-finding algorithm and other inference mechanisms in SPIEL can be found in (Burke, 1993).

6. Implementing recognition

Indices represent what a story is about. Strategies represent how, in general, stories can be made relevant to a student. There is one final link that a storytelling tutor must make, which is the connection between students' actions and its notion of a relevant story. As shown in the examples, this involves a variety of considerations including social knowledge, such as the scope of a conversation, and practical details that are a function of the simulation environment itself.

The tutor needs a general understanding of the task: that selling involves sub-goals like gathering information, constructing presentations, making sales pitches, and answering objections, for example. At a more concrete level, SPIEL has to know particular details of how the selling task is achieved by the student in the YELLO application. It has to know which kinds of answers from a customer constitute substantive information and which are evasions, what actions constitute the construction of a prospective ad, and so forth. SPIEL is operating in a environment with limited scope. It does not have to be concerned about every possible way a student can ask a question; it need only recognize those choices that GuSS actually permits.

6.1. Applying storytelling strategies

How would an opportunity to tell the "Talked myself out of a sale" story be recognized by SPIEL? The tutor would look for (1) the student to successfully complete the close of a sale, but (2) remain to talk with the customer, and then (3) finish the sales call without losing the sale. Then it would (4) tell the "Talked myself out of a sale" story to point out the possibly risky tactic the student had used..

We call the sequence of conditions for recognizing a storytelling opportunity a *recognition condition description* or RCD. The RCD describes what situation the student would have to encounter for the story to be relevant. Recognition condition descriptions have three parts: (1) a context in which the story would be worth telling, (2) conditions that indicate that the story is relevant in that context, and (3) the story itself and the manner in which it should be presented. These three parts can be thought of as filling roles in the following abstract rule form:

When *trigger*, **Look for** *evidence*, **Then** *presentation*.

To create the RCD, a storytelling strategy must reason about how concepts in the index or those inferred from the index may be manifested in observable actions in the simulation. SPIEL has a library of manifestation rules designed to predict characteristic actions that correspond to internal states.

For an example, the RCD generated for the story "Talked myself out of a sale" when told using the **Demonstrate risks** strategy looks like this:

WHEN the student is in the closing stage of the sales call and talking to the buyer,
LOOK FOR the student to make the sale, continue conversing with the customer, and not to lose the sale,
THEN TELL "Talked myself out of a sale" AS a "Demonstrate risks" story.

6.2. From description to mechanism

Creation of the trigger and evidence conditions completes the construction of an RCD. The RCD represents the recognition conditions for a tutorial opportunity describing a good time to tell a particular story using a particular strategy. The next part of the storage-time processing of SPIEL is rule generation, the construction of procedures capable of recognizing storytelling opportunities.

The GuSS simulation is organized around a production system interpreter. Each of the characters in the social simulation is implemented by a production system, like those found in expert systems. To fit into this environment most effectively, the retrieval-time component of SPIEL is also implemented as a production system. The task of the rule generation component of SPIEL is to produce production rules capable of recognizing the tutorial opportunities that have been created by applying storytelling strategies to indices.

The first step in rule generation is to make the recognition conditions concrete, to reason about how these conditions might arise and describe their occurrence in terms of actions that the student might actually take. To perform this elaboration, SPIEL needs knowledge about observability in the simulated world. It must know what kinds of actions are and are not available to the student and the simulated agents. The output of this step is an expanded RCD (or eRCD) representing the same recognition conditions as the original RCD, but placing them in terms of their concrete, observable manifestations.

From the expanded RCD, the system designs a set of rules that will recognize the situation it describes. The initial stage of rule design is the construction of a *rule specification*, a directed graph indicating what production rules are needed and how they will relate to each other. Each node in the graph is a rule that the generator needs to produce. The triggering conditions are implemented by two rules, one that looks for the individual who is the buyer, and one that looks for the student to begin the closing stage of the sales process. If the student gets the sale but then leaves immediately, the story is not relevant. So there is a deactivating rule that halts recognition in this case. There is a second deactivating rule that comes in if the customer does what the customer does in the story, rethink the buying decision. In this case, the student does not need to see the story because the simulation has amply demonstrated the folly of remaining after the sale.

In the final rule specification graph, each node contains a condition or set of conditions that a rule will have to recognize. Rules are produced by walking through this graph translating from the descriptions into simulation states that rules can directly test for. For example, a “customer signing contract” condition becomes a test for a signing event with the actor and object slots bound to certain values.

7. Current Status

Formative evaluation of an earlier version YELLO containing a subset of SPIEL's stories was performed using Yellow Pages account executives as subjects. We found that users reacted well to the stories, finding SPIEL's interventions interesting and relevant. (Kass, et al. 1994) describes some of these findings.

SPIEL is currently being adapted to teach personnel management. The most significant drawback in this conversion is the need for a large knowledge base of domain knowledge in order to apply storytelling strategies. This knowledge does not transfer from one domain to another and its acquisition forms a significant bottleneck for developing storytelling programs based on SPIEL. We are beginning to address this problem using knowledge-acquisition techniques to build an interactive tool that draws on the user's domain knowledge when applying storytelling strategies.

8. Conclusion

Currently users only benefit from on-line databases when (a) they realize they have an information need, and (b) they know how to formulate retrieval queries. We feel there is much greater potential inherent in on-line databases, such as the video case-base discussed in this paper. However, realizing that potential requires intelligent systems that can initiate retrieval, formulating their own retrieval requests, not just responding to retrieval requests from a user.

In building such a system, we have moved away from generic information retrieval: we found that what was required was a careful analysis of the *purposes* that retrieved information will serve. Such an analysis makes it possible to create retrieval systems like SPIEL that tailor retrieval by actively reasoning about those purposes.

References

Burke, R. D. 1993. Representation, Storage and Retrieval of Tutorial Stories in a Social Simulation. PhD Thesis, Northwestern University.

Edelson, D. C. 1993. Learning from Stories: Indexing, Reminding and Questioning in a Case-based Teaching System. PhD Thesis, Northwestern University.

Hammond, K. J. 1989a. *Case-based planning*. Boston: Academic Press.

Hammond, K. J. 1989b. Opportunistic Memory. In Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, 504-510. Menlo Park, CA: International Joint Conferences on Artificial Intelligence, Inc.

Kass, A., R. Burke, E. Blevis, and M. Williamson 1994. Constructing learning environments for complex social skills. *Journal of the Learning Sciences*. Forthcoming.

Kolodner, J. L. 1993. *Case-based reasoning*. San Mateo, CA: Morgan Kaufmann.

Salton, G., and M. J. McGill 1983. *Introduction to modern information retrieval*. New York: McGraw-Hill.

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

Schank, R. C. 1990a. Teaching Architectures, Technical Report, #3. Institute for the Learning Sciences.

Schank, R. C. 1990b. *Tell Me a Story: A New Look at Real and Artificial Memory*. New York: Charles Scribner's Sons.

Schank, R. C., R. Osgood, M. Brand, R. Burke, E. Domeshek, D. Edelson, W. Ferguson, M. Freed, M. Jona, B. Krulwich, E. Ohmaye, and L. Pryor 1990. A Content Theory of Memory Indexing, Technical Report, #2. Institute for the Learning Sciences.