Strategic retrieval of tutorial stories

Robin Burke*
Institute for the Learning Sciences, Northwestern University,
1890 Maple Ave., Evanston, IL 60201
burke@ils.nwu.edu

Abstract

This paper describes SPIEL, a system for retrieving and presenting tutorial stories for students who are using a social simulation to learn social skills. SPIEL's task is primarily retrieval, but it requires techniques from case-based reasoning to perform it. SPIEL's stories are stored in video form, which prevents the use of text-based processing or indexing. Instead of using a story's text, SPIEL uses complex structured indices intended to represent what the story is about.

SPIEL's tutoring task also differs from strict retrieval because it is not responding to a user's request for information. The system monitors the student's actions in a simulated world and brings up relevant stories in the course of that interaction. The standard notion of query does not apply. In its place, SPIEL has a system of storytelling strategies that define what kinds of stories are appropriate at what times. Each strategy makes its own comparison between story and situation, determinations that include but are not limited to assessments of similarity.

* 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, an Institute partner, and from IBM.

Introduction

One of the important ways an expert can help a novice is by telling stories and anecdotes (Schank, 1990; Orr, 1986). A computer system for storing and retrieving such stories for pedagogical uses is both a useful educational tool and an interesting challenge for both information retrieval (IR) and case-based reasoning (CBR) research.

Finding appropriate material to present to a user is a classic IR problem, but the IR model depends on user initiative: the user must recognize a knowledge need and seek to address it. A tutor, by contrast, must act as well as react: frequently, the student will not know that help is needed.

In this paper, I will describe SPIEL (Story Producer for InteractivE Learning), which is a testbed for exploring retrieval and presentation strategies for tutorial stories stored on video.

SPIEL differs from information retrieval and case retrieval systems in two important ways:

- SPIEL combines similarity, dissimilarity, and other types of comparisons to create strategic measures of a story's fitness for retrieval.
- Since SPIEL's stories are in video form, it cannot use text-based indexing methods standard in IR. It uses structured indices more complex than most CBR systems, because it is not storing cases that are themselves intelligible to the system.

SPIEL is designed to assist students who are learning social skills. It is embedded in an intelligent learning-by-doing architecture called **Gu**ided **S**ocial **S**imulation or GuSS. GuSS provides a social simulation in which

If you assume that Mrs. Swain will not participate in the advertising decision for Swain Roofing, you may be surprised. Here is a story in which an account executive made similar assumption and was wrong:
[Video: I went to this auto glass place one time where I had the biggest surprise. I walked in; it was 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 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.]
An assumption that a spouse will not participate in an advertising decision may be unrealistic. Think carefully about your assumptions.

Example 1. Telling the "Wife watching TV" story using the **Warn about assumption** storytelling strategy.

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. It allows the student to learn by doing.

Within GuSS, SPIEL is like an experienced pilot 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.

YELLO in action

Each YELLO scenario begins with the student in the office, receiving new accounts and contacting customers in the same way a real account executive would. One representative scenario is as follows:

- The student finds that the account is for "Swain Roofing" and the contact person is Mr. Ed Swain. The student calls the client and sets up an appointment.
- Upon arriving at the Swain's home, she meets Lucy Swain, Ed's wife and finds that Ed is not yet there.
- The student carries on small talk with Lucy but does not discover that Lucy "keeps the books" for the business and fails to gather any business information that Lucy could provide.
- Ed Swain arrives.

This is a good time to give some feedback to the student. She has misjudged Lucy's role in the business and missed an opportunity to get some business information from an important

You were hostile towards Mrs. Swain. Unfortunately, she has convinced Ed not to buy. Here is a story about a similar situation in which the salesperson used a different method and was successful.

[Video: I went to this auto glass place one time where I had the biggest surprise. I walked in; it was 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 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.]

You were hostile towards the client's spouse. In the future, you might consider being considerate when meeting members of the client's family.

Example 2. Telling the "Wife watching TV" story using the **Demonstrate alternative plan** storytelling strategy.

source. Without some feedback, the student might not realize that this opportunity existed.

One of the approximately 200 stories in SPIEL's library is relevant here. SPIEL lights up a button on the screen with the message: "A warning about something you just did." This tells the student that a story is available. If the student chooses to view the story, a video of the storyteller appears in a separate screen. Example 1 shows the text the student sees on the screen, along with a transcription of the video.

SPIEL precedes each story with an introductory paragraph called a *bridge* and summarizes each with a *coda* paragraph. The bridge explicitly connects the story with context of the student's activity; the coda brings the student back to the events of the simulation by suggesting some possible actions.

This example shows the use of a *storytelling strategy* called **Warn about assumption**. By observing the conversation with Lucy, the system can gather evidence that the student assumes her interaction with Lucy will not be an important part of the sales call. It tells the story to show how a similar assumption has been found inaccurate in a similar situation in the past.

Storytelling strategies enable a storyteller to make a variety of different points with the same material. A teacher could use the "Wife watching TV" story in a different way: suppose the student were hostile and dismissive towards Lucy during the sales call, resulting in a poor outcome. This would be a good time to tell the story as a counter-example, as shown in Example 2. Here the story shows how another salesperson succeeded where the student failed.

These examples illustrate the synergy between simulation and storytelling. Without a story to provide the impetus to examine the situation, the student might never realize what opportunities had been missed. Without the active engagement provided by the simulation, however, the student might lack the motivation and context to listen, understand and remember the story.

How stories teach

The primary challenge for a tutorial storyteller is to tell the right story at the right time. I have approached this problem by analogy to case-based reasoning. Like case-based problem-solving systems, such as CHEF (Hammond, 1986), SPIEL must locate knowledge structures in its memory. CHEF retrieved recipes using indices that describe the types of dishes the recipes produce: the cooking goals they satisfy. SPIEL's goal is to tell its stories in instructive ways. Its retrieval system is therefore based on the kinds of learning that stories can produce.

When students learn by performing an activity, it is frequently because of failures: unsuccessful results or unmet expectations. Failure-driven learning theory (Schank, 1982) holds that there is a central cycle in performance learning. The learner makes predictions about results of his or her own actions or the behavior of others, and tests these predictions by observing what happens. If all goes well, these expectations feed into decisions about the next action or prediction. However, if the expectations fail, the learner has failed to understand something about the world. The problem helps focus attention on exactly what may be wrong in the student's thinking about the world. To prevent future failures, the student must explain how this one occurred in sufficient detail to allow the problem to be fixed.

Explanation is often the most difficult aspect of the prediction-observation-failure-explanation cycle. The student has to identify why the prediction was wrong and what to do to prevent future errors. Stories can help this kind of learning in two ways. If the student has failed in some task, the story base can recall a story that helps the student explain the failure and recover from it. If the student is heading in the wrong direction without being aware of it, a story can point out the problem by showing where the direction is leading. It can spur learning by getting the student to recognize a failure.

In order for the story retriever to address the student's acutal and possible failures, it has to be able to recognize the student's mistakes. In intelligent tutoring systems research, the standard approach is to create a model of the student's thinking and an expert's thinking. The differences between the models constitute the things the student needs to learn.

There are a number of reasons why this is not feasible for SPIEL. No one "correct" model of expert social knowledge exists. People's understanding of social situations is for the most part beyond the capability of existing AI models; it is large and idiosyncratic, tied to culture and personal experience. Even if such modeling were possible, the knowledge necessary for building a model would have to be laboriously extracted from each student.

A more subtle problem is that an expert model would, in some sense, be redundant with the knowledge embedded in the stories that the system has. If SPIEL identifies a particular action as dangerous, it is because it has a story that shows that the action can lead to disaster. Stories themselves serve as the basis for assessing what is a possible failure.

Finding storytelling opportunities

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. The storyteller wants to communicate as many stories as possible; 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.

The general problem of opportunism, locating and responding to opportunities for the pursuit of multiple goals, is a difficult unsolved problem in artificial intelligence. Birnbaum

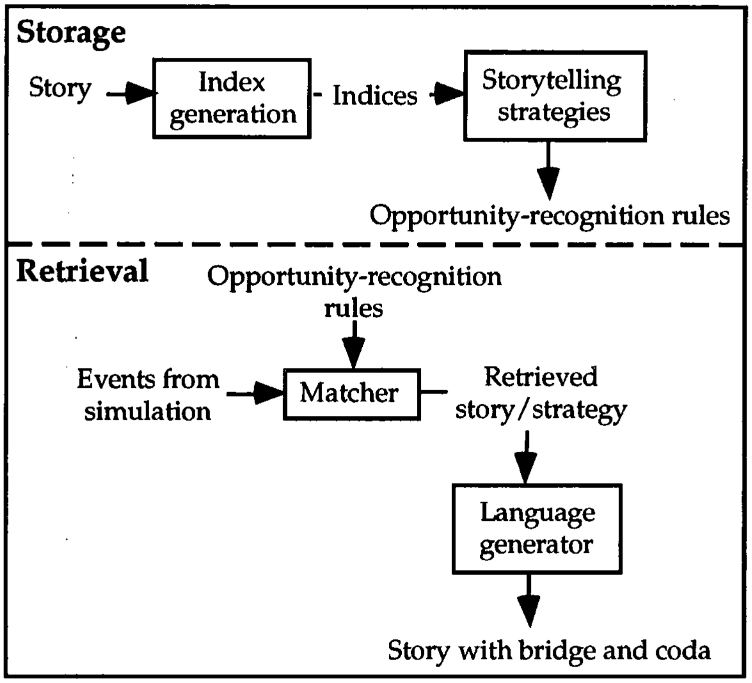

Figure 1. SPIEL's processing

(1986) argues that unbounded inference can be required to recognize a novel opportunity to pursue a goal. However, the educational storytelling problem in SPIEL is a special case of opportunism, one that is amenable to solution.

One difficulty associated with noticing real-world opportunities is that it can be difficult to predict how they can arise. Birnbaum uses revenge as an example. There is no single state of the world that can be described as constituting a "revenge opportunity." Such an opportunity will always be a function of the desires of the revengee and the actions available to the revenger.

Educational storytelling is not so open-ended. SPIEL's storytelling strategies compare stories about activity to specific contexts of student action. They seek coherence and relevance, and do not look for distant analogies. It is possible to describe concretely what an opportunity would look like to tell a story using a particular strategy.

Perhaps SPIEL's strongest advantage over the problem of opportunism is 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.

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." (page 146)

This model fails when it is too difficult to determine all of the situations in which a goal can be satisfied or to enumerate the features of such situations, but these are conditions that SPIEL can meet. It is possible to do some processing, when a story is stored in memory, to determine the situations in which it could be told and then prepare to recognize such situations when they arise.

SPIEL works from its database of stories to determine what should count as an opportunity for intervention. Its processing can be divided into two stages: storage time, when new stories are put into the system; and retrieval time, when a student interacts with a GuSS system and gets 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 implemented by recognition condition generators that 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.
3. Optimizations are performed on the final opportunity-recognition rule set.

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.

In the following sections, I will describe SPIEL's indexing method and the storytelling strategies. A discussion of the system of opportunity-recognition rules requires an understanding of the internal operation of the GuSS simulation and is beyond the scope of this paper.[1]

Indexing

SPIEL uses video to communicate its stories vividly and memorably. However, this medium is much more demanding of an indexing scheme than text. Video does not support text-based processing or indexing. To automate indexing would require speech (and possibly gesture) recognition as well as natural language understanding, technical challenges well beyond our intent in designing a learning environment.

In SPIEL, a human indexer watches stories, and uses an indexing tool to create indices. The indexer is expected to know enough about the domain to understand the meaning of stories, but does not need to know how those stories will appear in an educational context. The indexer constructs indices that capture different interpretations of a story's meaning. SPIEL's uses these characterizations of stories to perform its educational task.

The interpretation of the "Wife watching TV" story that is at work in Example 1 goes something like this:

> "The salesperson assumed the wife would be a housewife, but actually she was a business partner."

This is an *anomaly*, a failure of expectation that requires explanation (Schank, 1982). Typically, anomalous occurrences are what make stories interesting and useful. They are especially important in stories about social activity since students in such areas need to

[1]GuSS's system of decision rules is covered in (Kass, et al., 1992). The implementation of SPIEL's opportunity-recognition rules within that system is discussed in (Burke, in preparation).

Setting:	step-of-sales-process(salesperson, presentation)
	business-type(service)
	business-size(small)
	business-partnership(auto-glass man, wife)
	married(auto-glass man, wife)
	sales-target(salesperson, auto-glass company)

Anomaly:	salesperson	assumed	wife	theme

	Assumed	Actual
Theme:	housewife	business partner
Goal:	hospitality	help in decision
Plan:	small talk	evaluate presentation
Result:		positive evaluation
Pos. side-effect:		sale for salesperson
Neg. side-effect:		

Figure 2. Index for "Wife watching TV."

learn what to expect and how to explain and address their failures. In stories about social activity, anomalies usually take the general form, "X believed Y, but actually Z."

SPIEL's indices are considerably more complex than those typically found in case-based reasoning systems. There are two reasons for this complexity.

One reason has to do with the medium of video. Because the retriever does not have access to the content of what it is retrieving, the index has to contain everything that the retriever will need. This is not true in CBR or IR systems. A CBR retriever retrieves cases, represented in terms that the system uses for problem-solving. If necessary, it can use the contents of the case to make decisions during retrieval. Similarly, in information retrieval systems, the textual contents of indexed documents can be employed in retrieval processing. SPIEL cannot examine the "text" of its video stories, so indices must compensate.

The other reason that SPIEL's indices show complexity is the task of educational storytelling itself. Most case-based reasoning systems are engaged in a single task, usually problem-solving. Cases serve only one purpose: to supply candidate solutions. SPIEL has a variety of reasons for telling stories; each of its storytelling strategies uses the index in a slightly different way. SPIEL needs complex indices to meet the demands of a variety of strategies.

Figure 2 shows an example of an index for "Wife watching TV". As well as the anomaly, the index describes

- the physical, temporal and social setting, placing the story within the overall task and within a network of social relationships.
- intentional chains that link plans of action with the goals and themes that give rise to them and the results and side-effects that

follow from them, thereby explaining the anomalous parts of the story.[2]

Storytelling strategies

A tutorial storyteller needs indices to know what its stories are about. It also needs to know how stories can be usefully employed. For maximum pedagogical impact, a story must appear precisely at the time that the student is prepared to learn from it. For example, when the student has acted on the basis of an incorrect assumption as in Example 1, it is a good time to tell a story showing that such expectations are often in error.

Each storytelling strategy is a domain-independent description of what kind of stories can be employed in what circumstances. A strategy can be described by a simple conditional statement: "Tell a story about X when Y."

For example, the **Warn about assumptions** strategy calls for the storyteller to

> Tell a story about an assumption that someone made which did not hold when the student appears to have a similar assumption.

This allows the storyteller to correct the student's mistaken assumptions, if it has a story that shows how that assumption can fail.

SPIEL has two other strategies for dealing with erroneous beliefs that students may have. Students may have unrealistic hopes about what they can achieve. For example, a student might try to sell an ad that is much larger than what the client needs. Given an appropriate story, SPIEL could use the **Warn about optimism** strategy:

> Tell a story about a desire that someone had which was not realized when the student appears to have a similar desire.

The opposite problem arises if a student is too timid. Aiming too low is just as disastrous for a salesperson as aiming too high. For a student that shows such fears, SPIEL uses the **Warn about pessimism** strategy:

> Tell a story about a fear that someone had which was not realized when the student appears to have that fear.

Stories are intended to bring the real world of practice into the simulation, which is by nature artificial and simplified. One of the important benefits that stories can bring is to show the ways in which the real world is different from the simulation. SPIEL has a story about a salesperson who continued conversing with a cleint after making a sale, and had the client change his mind. If the student hangs around after a sale is closed, SPIEL can tell this story. It shows the student that an action that works in the simulated world may be disastrous on the job. This the **Demonstrate risks** strategy.

> Tell a story about a negative result of a particular course of action when the student has just executed a similar course of action but had success.

The same logic holds when there is a story about a success that ensues from a particular action, contrary to what the simulation shows. **Demonstrate opportunities** is the name of the strategy that presents this kind of story.

> Tell a story about a positive result of a particular course of action when the student has just executed a similar course of action but had a poor result.

Stories can also teach students particular techniques and approaches. SPIEL takes advantage of the failure-based learning cycle to do this. If it has a success story, one way to tell it is to wait for the student to fail in a similar situation and then tell the story to show a different way to approach the same situation.

---

[2]This index structure is derived from the Universal Indexing Frame, reported in (Schank, et al., 1990). See also (Domeshek, 1992).

This is the **Demonstrate alternative plan** strategy from Example 2.

> Tell a story about a successful plan to achieve a particular goal when the student has executed a significantly-different plan and failed to achieve the goal.

One of the important features of a learning environment is that students should get quick feedback about how well they are doing. It is not always realistic to do so in a simulation, since real salespeople often do not get such feedback. However, SPIEL can intervene to give feedback when the simulation is not likely to.

For example, a student might use high-pressure fear tactics to win a sale. In real life, this endangers the salesperson's chances of selling to the customer next time around. This consequence would not be apparent in the YELLO simulation, however, since the encounter with a client ends after one sale. A story about the bad long-term consequences of such tactics would be told by the **Warn about plan** strategy.

> Tell a story about an unsuccessful plan when the student has begun executing a similar plan and will not get immediate feedback about it.

The same principle holds for successful sales approaches with the **Reinforce plan** strategy:

> Tell a story about a successful plan to achieve a particular goal when the student has just started to execute a similar plan and will not get immediate feedback about it.

In a social domain, knowing what others are doing is just as important as knowing what to do oneself. **Explain other's plan** is a strategy for helping the student understand the actions of other agents in the simulated world.

> Tell a story about a plan that the student might not know about when the student has just observed someone in the simulation execute a similar plan.

For example, SPIEL has a story about a client who felt it was his right to verbally abuse salespeople. If the student encounters an abusive client, SPIEL can use the story to show that some customers are just like that.

One interesting consequence of this set of storytelling strategies is that SPIEL cannot restrict itself to performing a simple assessment of similarity when retrieving. In Example 2, the storyteller presents a story that is in some ways similar to the student's situation, but it shows someone handling that situation quite the opposite from the student's approach: being friendly instead of hostile. The **Demonstrate alternative plan** strategy calls for an examination of specific areas of similarity and dissimilarity in the relation between story and situation.

It should also be noted that these storytelling strategies are not specific to selling or to business. They are useful for presenting stories about any social activity to students who are practicing their social skills.

Conclusion

SPIEL is a specialized retrieval system. It is designed for the task of presenting pedagogical stories to students who are learning by interacting with a simulated environment. Still, its construction has some broader lessons for research at the boundaries of case-based reasoning and information retrieval.

As video becomes a more common source of information, the appeal of strictly text-based approaches to indexing will diminish. SPIEL is an experiment in designing indices that represent what stories are about richly enough so that they can be employed in a variety of ways. Such indices are by necessity complex, but with an indexing tool that supplies the format and the representational vocabulary, a story can be indexed in a few minutes.

Building retrieval systems for specific tasks helps illuminate the variety of uses to which retrieved material can be put. A tutor needs to make its interventions directly relevant to the student's actions. SPIEL's nine storytelling strategies describe different ways that stories can be made relevant. They incorporate similarity judgments, but also are sensitive to important differences and other relations.

SPIEL is therefore in a quite different position from both a CBR retriever and IR system. In CBR, a case is considered worth applying to the current problem if it solves a similar problem. In IR, a text is worth retrieving and presenting to the user if it is similar, along some metric, to the user's retrieval request. In a tutorial situation, similarity remains important, but it is not the only kind of relation that is needed by SPIEL's storytelling strategies.

References

Birnbaum, Lawrence A. (1986) Integrated Processing in Planning and Understanding. PhD Thesis, Yale University.

Burke, Robin (1992). Knowledge acquisition and education: a case for stories. In Proceedings of the AAAI Symposium on Cognitive Aspects of Knowledge Acquisition, (pp. 39-46). Stanford University: AAAI.

Burke, Robin, & Alex Kass (1992). Integrating Case Presentation with Simulation-based Learning-by-doing. In Fourteenth Annual Conference of the Cognitive Science Society, (pp. 629-634). Bloomington, IN: Lawrence Erlbaum Assoc.

Burke, Robin (1993). Intelligent retrieval of the video stories in a social simulation. In Educational Multimedia and Hypermedia Annual. To appear.

Burke, Robin (in preparation) Representation, Storage and Retrieval of Stories in a Guided Social Simulation. PhD thesis, Northwestern University.

Domeshek, Eric A. (1992) Do the Right Thing: A Component Theory for Indexing Stories as Social Advice. PhD thesis, Yale University.

Hammond, Kristian J. (1986) Case-based Planning: An Integrated Theory of Planning, Learning and Memory. PhD thesis, Yale University.

Kass, Alex, Robin Burke, Eli Blevis, & Mary Williamson (1992). The GuSS Project: Integrating Instruction and Practice through Guided Social Simulation (Technical Report No. 34). Institute for the Learning Sciences.

Kolodner, Janet L., & Menachem Y. Jona (1991). Case-based Reasoning: An Overview (Technical Report No. 15). Institute for the Learning Sciences.

Lave, Jean, & Etienne Wenger (1991). Situated Learning: legitimate peripheral participation. Cambridge University Press.

Orr, Julian (1986). Narratives at work: Story telling as cooperative diagnostic activity. In Proceedings of the Conference on Computer-Supported Cooperative Work. Austin, TX.

Salton, Gerard, & Michael J. McGill (1983). Introduction to modern information retrieval. New York: McGraw-Hill.

Schank, Roger C. (1982). Dynamic Memory: A Theory of Learning in Computers and People. Cambridge University Press.

Schank, Roger C., et al. (1990). A Content Theory of Memory Indexing (Technical Report No. 1). Institute for the Learning Sciences.

Schank, Roger C. (1990). Teaching Architectures (Technical Report No. 3). Institute for the Learning Sciences.

Williams, Susan M. (1991). Putting Case-Based Instruction into Context: Examples from Legal, Business, and Medical Education (Technical Report) Learning Technology Center, Vanderbilt University.