

An Intelligent Tutoring System Approach to Teaching People How to Learn

Richard G. Feifer

Center for the Study of Evaluation -- UCLA Graduate School of Education
UCLA Artificial Intelligence Laboratory

ABSTRACT

Sherlock is an intelligent tutoring system designed to teach people to build simplified knowledge representations (graphic maps) to facilitate learning of a text. Previous attempts to automate instruction in graphic mapping have had problems because they attempted to diagnose a learner's misunderstandings by looking at a finished graphic map. Sherlock uses a knowledge-based approach to diagnose a learner's misunderstandings by looking at the knowledge and processes that lead to a learner's graphic map, rather than the completed map.

In Sherlock's model a semantic network is used to represent the knowledge in the text. A production system models the strategy for constructing a graphic map by initiating spreading activation on the semantic network, and interpreting the resulting activation patterns. In a limited evaluation Sherlock was able to correctly determine if a construction was appropriate 96% of the time.

INTRODUCTION

In this paper I examine the problem of teaching people to use a learning strategy called **graphic mapping**. In graphic mapping the learner is taught a simplified knowledge representation scheme. The learner then uses this scheme to pictorially represent textual material. Figure 1 contains a sample graphic map construction.

Researchers have shown that comprehension of text can be enhanced by having the learner construct a graphic map representing the text (Dansereau, 1978; Dansereau, Collins, McDonald, Holley, Garland, Diekhoff & Evans, 1979a,b; Anderson, 1979; Novak, Gowin & Johnson, 1983).

The decisions that a learner makes in constructing a map appear to be equally dependent on both the learner's strategy for building a graphic map and the learner's understanding of the text. Intelligent tutoring systems have been developed which attempt to model a learner's strategy knowledge or a learner's factual knowledge, but not both. Sherlock models the acquisition of both strategy and facts in an integrated manner (Feifer, Dyer, & Baker; 1988).




Figure 1: A Graphic Map

Even domains which seem to rely predominantly on one kind of knowledge involve, to at least some extent, both kinds of knowledge. Programming, for example, requires a knowledge of the syntax and commands of a computer language in addition to knowledge of how to build programs in that language. Thus, a tutoring system able to model the acquisition of both kinds of knowledge can more completely model the acquisition of knowledge in any domain.

SHERLOCK

The Sherlock tutoring environment provides the learner with three components: (a) a text to be represented pictorially, (b) a screen containing icons¹ representing concepts within the text, and (c) a set of gm-links² which the learner can use to connect the icons.

Sherlock has been programmed to facilitate learning of a text excerpted from an introductory business law text book describing consideration and its place in a contract. Sherlock offers the learner six gm-links: PART, IS-A, LEADS, EQUIV, PROP and NOT. Learners are instructed to find icons that they believe are related and then choose the gm-link that best represents the relationship. For example, if the learner believes that one icon represents a concept that is a generalization of a concept represented by a second icon, he should make an IS-A gm-link between the two icons.

Sherlock's Input and Output

Once the learner has specified a gm-link, Sherlock tries to evaluate it and provide meaningful feedback. What follows are excerpts of a transcript (Feifer, 1989) from actual interactions between Sherlock and a learner named Janey. The text in parentheses are descriptions of the interactions and the figures that illustrate the content of the screen at each point in the interaction. All other text are comments made by the learner.

"I was first looking at these icons here (Figure 2); instead I am going to look at the links (the list of link types) to try and link something together from these different links I have to choose from."




Figure 2: Janey - 1

"Now they are equivalent according to this. So I am going to use an equivalent link (Figure 3)."

"According to this paragraph a consideration act is a lawful act. So a lawful act is a legal act. So I am going to link them as equivalents because they are both legal acts."




Figure 3: Janey - 2

¹I use the term icon to refer to a box containing text that the learner can use to represent a conceptual entity.

²To avoid confusion between the graphic map that the learner constructs and Sherlock's internal knowledge representation, links used in a graphic map are referred to as gm-links.

(Sherlock requests information regarding her reasons for making the link by presenting her with the menu shown in Figure 4.)

Why did you do that?
The likelihood of a legal act being a consideration act is equal to the reverse
Every legal act is a consideration act
Every consideration act is a legal act
legal act and consideration act refer to the same concept
The likelihood of a legal act being a consideration act is not equal to the reverse
legal act is another name for consideration act
It is not as likely that a legal act is a consideration act
DONE
ABORT

Figure 4: Janey - 3

“In general I wouldn't say that a ‘consideration act’ is a ‘legal act’. In general. But this (the text) is telling me that it is. So I am answering these question in regards to the material here? Yes.”

(She chooses “Every consideration act is a legal act” and “It is not as likely that a legal act is a consideration act.”)

(Sherlock provides her with the plan feedback shown in Figure 5)

“Sounds like algebra, and I am just terrible at algebra.”

“I guess that is what ... I used, but it doesn't really make good sense. That is the plan I used, but it doesn't seem like it is a very good plan or that it is going to be equivalent.”

(Sherlock provides the comment shown in Figure 6)

“It means where I made the equivalent link? Huh. Let's check that out.”

You used the plan:
 if
 every x is a y
 it is not as likely that a y is a x
 then
 make an EQUIV link from x to y
 that is not a good plan

Figure 5: Janey - 4

At times like this I would make an IS-A link.

Figure 6: Janey - 5

(Sherlock asks her if she would like to remove the link. She answers "YES". In Figure 7 the link is removed.)




Figure 7: Janey - 6

"So, 'legal act' well... 'consideration act' would be... would be a subset of a 'legal act,' or vice versa. Or do I have to decide that? I guess I have to decide that. Huh, OK. I would say that a..."

"Well, according to this a 'consideration act' is a type of a 'legal act.' So I am going to connect 'consideration act.'" (Figure 8).




Figure 8: Janey - 7

"Well according to this a 'consideration act' is a 'legal act'. It is one type, it is a type of a 'legal act.' So it's got to be legal. So... it's some type, or kind of 'legal act.'"

(Sherlock provides the feedback shown in Figure 9)

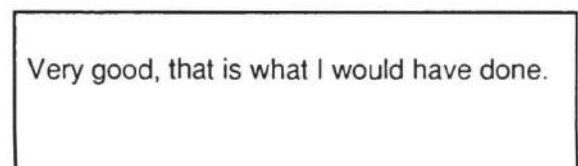


Figure 9: Janey - 8

In this sequence Sherlock determines that the learner seems to understand the relationship between a consideration act and a legal act, but has a bad strategy for representing that relationship. Sherlock decides that the problem is not with the facts used to choose the gm-link, but rather with the type of gm-link that those facts lead to.

Sherlock's Architecture

Sherlock uses the eight components shown in Figure 10.

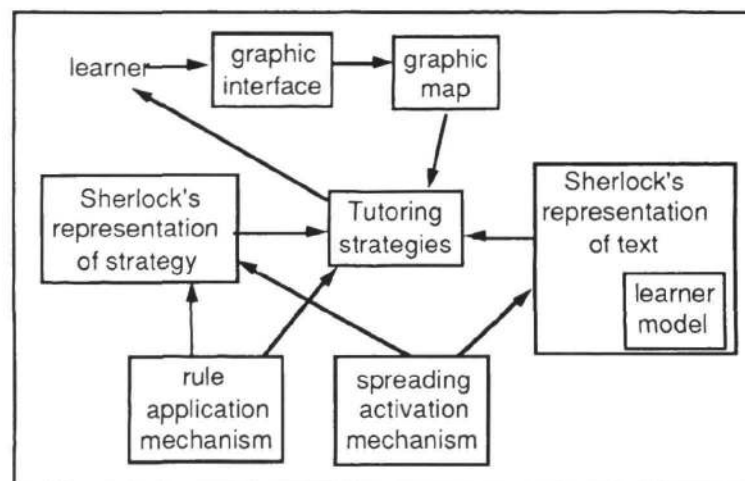


Figure 10: Overview of Sherlock's Components

The function of each of these components is briefly described below.

1. The **graphic interface** interprets the learner's clicks on the mouse buttons to build the graphic map.
2. The learner builds the **graphic map** using the icons provided by Sherlock. The icons are linked together using the six gm-links described above.
3. **Sherlock's representation of text** is a localist spreading activation network made up of nodes for concepts and links for the relationship between concepts. The semantic network is used to represent Sherlock's understanding of a text's content and background. The representation of the text is hand-coded into the network.
4. **Sherlock's representation of strategy** consists of rules or plans for building a graphic map. Each of the rules is of the form:

IF	the following things are true about the relationship between two icons
THEN	make this kind of gm-link between them

These rules are hand-coded into Sherlock and are based on strategies described by subjects during initial pilot studies.

5. The **learner model** represents what Sherlock currently believes to be the learner's understanding of the text, graphic elements, and mapping skills. It is built by Sherlock by modifying elements in the hand-coded representations of the text content and graphic mapping strategy.

The learner model is not fully implemented at this time. The only aspect of the learner model which Sherlock currently represents is the learner's interpretation of the screen icons.
6. The **tutoring strategies** are rules or plans in a form similar to that of the strategies for building a graphic map. The rules in the tutoring strategies are hand-coded to represent Sherlock's pedagogical knowledge.
7. The **spreading activation mechanism** operates on the semantic network to generate inferences. It is based on a mechanism developed by Michael Gasser (1988) for classifying concepts.
8. The **rule application mechanism** operates on the rules in Sherlock's representation of strategy and the tutoring strategies to determine what action Sherlock should take at any given time.

Sherlock's Tutoring Strategy

Sherlock determines the relationship between the two icons that the learner just linked by using spreading activation on a semantic network representation of the text. Sherlock then uses a production system representation of graphic mapping strategy to determine whether the learner's link is appropriate. If there is no plan that would justify the gm-link that the learner just made, Sherlock asks the learner to indicate the reasons for making the gm-link.

Sherlock uses the learner's answer to separately evaluate the learner's plan and the facts that the learner believes. Sherlock classifies the learner's plan as an instance of one its graphic mapping

FEIFER

rules using the representations of these rules in the semantic network. If Sherlock cannot recognize the plan the learner used, it will inform the learner of this. If Sherlock does recognize the plan, but it is a bad plan, it will inform the learner that he has used a bad plan. If Sherlock recognizes the plan and it is a good plan, it will not give any feedback on the plan.

To evaluate the learner's factual beliefs, Sherlock compares the learner's answer to its own factual beliefs. If there is a good³ match between the learner's fact beliefs and Sherlock's, no further action is taken on facts. If there is not a good match, Sherlock will consider the possibility that the learner is using an alternative interpretation for the icons. If any alternative interpretations for the icons exist, that have not already been tried, Sherlock will use one of these interpretations and start again from the top.

If all possible alternative interpretations have been tried, Sherlock will pick the interpretation that led to fact beliefs which were closest to the learner's fact beliefs. Using this interpretation Sherlock will determine two things:

1. Are there any false facts that the learner believes to be true, that led to the bad gm-link that the learner made? If so, Sherlock will bring them to the learner's attention.
2. Are there any true facts that the learner does not believe, which would have led to a better gm-link? If so, Sherlock will bring them to the learner's attention.

EVALUATION

To evaluate the accuracy of Sherlock's diagnosis a study was conducted in the summer of 1988. Four subjects were video taped during a two-hour session using Sherlock. Subjects were instructed to think aloud while building their map.

As a first step, I inferred the subjects' beliefs looking at their actions and words. Inferences were made without looking at Sherlock's feedback or the transcript of Sherlock's processing. There is no claim that the beliefs inferred actually reflect the subjects' beliefs; only that they represent one human tutor's best guess as to what the subjects believe. Sherlock's diagnosis is referred to as correct if it agrees with these inferred beliefs.

A total of 70 links were made by the subjects. Of these, 11 were aborted before Sherlock's final analysis. These 11 links were aborted because the subject decided that he or she was doing something wrong. Included in the analysis are the 59 links that the subjects allowed Sherlock to analyze.

Of the 59 completed links, Sherlock determined that 33 (55.93%) were appropriate. There were actually 32 appropriate links. I classified a link as appropriate if it reflected a correct understanding of the text and the graphic mapping strategy. Of the 26 links that Sherlock labelled as wrong, 25 were actually wrong. Thus Sherlock's determination of whether a link was right or wrong was correct for 96.61% of the links (Table 1).

³"Good" is currently defined as meaning that the learner and Sherlock agree on 80% of the facts.

FEIFER

For the 25 links which Sherlock correctly determined were wrong, the diagnosis matched only 56% of the time. This percentage of match means that Sherlock was able to correctly determine that a link was right or wrong and provide a correct diagnosis for a total of 77.97% of the links.

	# of links	% of total
good link correctly identified	32	54.24
good link incorrectly seen as bad	1	1.69
bad link correctly identified		
reason found	14	23.73
reason not found	11	18.64
bad link incorrectly seen as good	1	1.69
Total	59	100

Table 1: Overall Accuracy of Links

CONCLUSION

Accuracy is important in a tutoring system; wrong feedback can have negative impact on learning. Telling a learner he is wrong when he is right can lead the learner to a loss of self confidence. Telling a learner he is right when he is wrong can reinforce incorrect beliefs. Sherlock was able to correctly determine if a link was right or wrong over 96% of the time. The question is: Is this accurate enough? One way to answer this is to compare Sherlock to other forms of instruction.

A human tutor will certainly be more accurate than any machine-based tutor. In terms of automated tutoring, the highest accuracy can be achieved with multiple-choice CAI. If the questions and distractor choices in a multiple-choice format are carefully written, close to 100% accuracy can be reached. The problem is that multiple-choice responses are the least indicative of what a learner understands. Open-ended responses can be much more useful for ascertaining what a learner understands. But if even single-word responses are allowed in CAI the accuracy dramatically drops because it is difficult to anticipate every potentially correct response.

Sherlock offers a compromise in that there is a finite range of responses possible. That range, however, is quite large. With 25 icons there are 3600 possible learner actions⁴. It would be possible, but very difficult, to record in advance which of the 3600 possible actions is correct. Add the qualifier that certain responses are only acceptable if the learner is using particular interpretations, and it becomes impossible to completely anticipate correct responses. Considering the range of responses allowed, Sherlock's performance is certainly comparable to any automated tutoring approach that currently exists.

⁴Each learner action is a possible combination of from-icon, gm-link type, and to-icon. Thus there are 25 (number of possible from-icons) * 6 (number of gm-link types) * 24 (number of possible to-icons) possible combinations.

FEIFER

Sherlock was less successful at determining why an action was wrong. Comparing this ability to other forms of tutoring is more difficult. In automated tutoring, only intelligent tutoring systems make any claim for diagnosing the cause of a misunderstanding. And in intelligent tutoring systems research, evaluations have only been done on a system's ability to identify errors. No system has been evaluated for its accuracy in identifying the cause of an error.

Sherlock has demonstrated that it is possible to understand unanticipated responses and diagnose the cause of misunderstanding without degrading the system's ability to make a bottom-line determination of whether a learner's response is correct.

REFERENCES

- Anderson, T.H. (1979). Study skills and learning strategies. In H.F. O'Neil and C.D. Spielberger (Eds.). *Cognitive and affective learning strategies*. New York Academic Press.
- Dansereau, D.F. (1978). The development of a learning strategies curriculum. In H.F. O'Neil, Jr., (Ed.), *Learning Strategies*. New York: Academic Press.
- Dansereau, D.F., Collins, K.W., McDonald, B.A., Holley, C.D., Garland, J.C., Diekhoff, G.M., & Evans, S.H. (1979a). Development and evaluation of a learning strategy training program. *Journal of Educational Psychology*, 71, 64-73.
- Dansereau, D.F., McDonald, B.A., Collins, K.W., Garland, J.C., Holley, C.D., Diekhoff, G.M., & Evans, S.H. (1979b). Evaluation of a learning strategy system. In H.F. O'Neil, Jr., & C.D. Spielberger (Eds.), *Cognitive and affective learning strategies*. New York: Academic Press.
- Feifer, R.G., Dyer, M.G., & Baker, E.L. (1988). Learning procedural and declarative knowledge. *Proceedings of the Intelligent Tutoring Systems Conference - 88*, Montreal, Canada.
- Feifer, R.G. (1989). *A intelligent tutoring system for graphic mapping strategies*. Doctoral dissertation. (Technical Report UCLA-AI-89-04). University of California, Los Angeles, Computer Science Department, Artificial Intelligence Laboratory.
- Gasser, M.E. (1988). *A connectionist model of sentence generation in a first and second language*. Doctoral dissertation. (Technical Report UCLA-AI-88-13). University of California, Los Angeles, Computer Science Department, Artificial Intelligence Laboratory..
- Novak, J.D., Gowin, B., & Johnsen, G.T. (1983). The use of concept mapping and knowledge vee mapping with junior high school science students. *Science Education*, 67, 625-645.