

Encoding Planning Knowledge for Recognition, Construction, and Learning*

Charles Dolan
Hughes Aircraft AI Center
and
Artificial Intelligence Lab
Computer Science Department
UCLA

Michael Dyer
Artificial Intelligence Lab
Computer Science Department
UCLA

Abstract

This paper discusses a method for representing thematic level structures, i.e. abstract plan/goal combinations. We make the case that the processes for both recognition and construction of plans use the same memory structures. In particular, we are looking at the knowledge structures for recognizing and avoiding bad planning. The learning procedure we describe starts by observing the bad planning behavior of narrative characters and combines old descriptions of planning errors to create new abstract structures. The learning method discussed is a one-trial, schema acquisition method, which is similar to DeJong's [DeJong 1983]. The method used involves taking schemas for planning situations that are found in an actual narrative situation, and using causal reasoning to construct a new schema which better characterizes the situation. This work is part of the MORRIS project at UCLA [Dyer 1983a]. The planning situations are represented using Thematic Abstraction Units (TAUs) [Dyer 1983b].

1 Introduction

In the real world, tasks cannot always be accomplished by using simple subgoal partitioning and recursive problem analysis. Both real world agents and narrative characters often must apply plans that require cooperation from other agents, adjust plans that conflict with an agent's concurrent goals, and manage plans that contribute simultaneously to more than one goal. A classification of real world plans is found in [Schank and Abelson 1977], and a taxonomy of goal/plan interactions can be found in [Wilensky 1978].

There are two reasons for looking at poor planning behavior. (1) Knowledge structures that encode poor planning can be used to critique plans and point-out possible weak

* This work was supported in part by grants from the Hughes Aircraft AI Center, the UCLA Center for Experimental Computer Science, and the Keck Foundation.

points. In this way they are similar in function to planning critics and meta-plans [Sussman 1973, Sacerdoti 1977, and Wilensky 1983]; however, they are much more varied than the fixed number of very general critics or meta-plans other researchers propose, and as this paper shows, they are learnable. (2) It is important to represent planning errors, not only to critique plans, but also so that in counter-planning situations [Carbonell, 1979], a planner can try to trick another agent into making a planning error. Any situation which is bad for a planner is generally a good situation into which to force an opponent.

Dyer [1983, 1981] showed how a class of planning errors could be represented by Thematic Abstraction Units (TAUs), and how these planning errors might be recognized in stories. This paper will present a representation for planning errors that also facilitates the combination of planning descriptions and use of thematic structures in planning.

The combination method requires an example narrative situation that contains a new planning error. The input example is conceptually analyzed to discover whether known planning errors are present in the example. In addition to the representations for TAUs, goals, and plans mentioned above, the examples here also rely on Schank's Conceptual Dependency theory [Schank 1972] for the output of the conceptual analysis. Planning errors in the story are examined to find out if they can be combined into new structures, and to establish the links between the component structures that make up the newly discovered planning error.

TAUs are used in planning by treating an agent's plans as if they were a story, and using the TAU recognition process to point out potential planning errors. Specific TAUs are indexed at memory nodes for plans, goals, and actions which are likely to be contained in those TAUs. The memory organization that supports this indexing is a dynamic memory similar to [Schank 1982], [Kolodner 1980], and [Lebowitz 1980]. Depending on the planning error detected, a particular set of heuristics is applied to try and fix the plan.

2 An Example Planning Situation

Here we will see a planning situation that contains three planning errors whose abstract descriptions are already known to the system. From this situation we will see how we can generate two specializations of planning errors, and one novel planning error construct.

The Fox and the Crow

The Crow was sitting in the tree with a piece of cheese in her mouth. The Fox walked up to the bottom of the tree and said to the Crow, "Crow, what a beautiful voice you have; please sing for me." The Crow was very flattered and began to sing. When she did, the cheese dropped out of her mouth. The Fox grabbed the cheese and ran away laughing.

Note that this story can be looked at in two ways: (1) as an instance of bad planning on the part of the Crow and (2) as an instance of good planning on the part of the Fox. Here we will show how to recognize a planning error from its representation, how to learn new planning errors, and how to apply planning information to avoid goal failures. TAUs can also be turned around and used to find situations where other agents will make mistakes.

The first of the three planning errors we will discuss is the most basic. When the Crow sings, she does not realize that she is already using her mouth to hold the cheese. This planning error is characterized at an abstract level by TAU-CONF-ENABL (confused enablement). The full representation for TAU-CONF-ENABL is given in Figure 1 below.

Our representation of TAUs always defines them from the point of view of the poor planner. The slots 'mistake', 'consequence', and 'failed-goal' of the TAU memory structure provide pointers into a large conceptual graph of **PLAN/GOAL** interactions, which defines the meaning of the TAU.

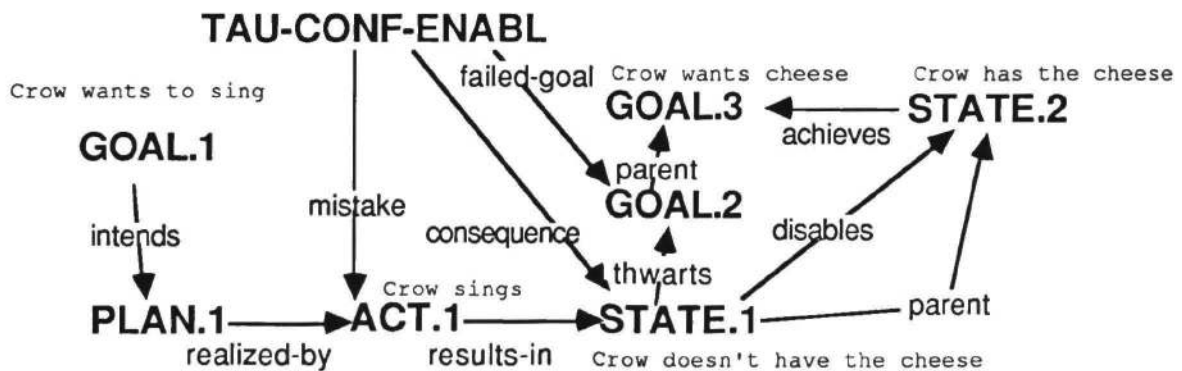


Figure 1

The 'parent' links indicate conceptual structures which are derived from other structures over time. Hence, **GOAL.2** is the same as **GOAL.3** except that its status is 'FAILED'. The rest of the links are derived from Dyer's I-links [1983a].

The abstract situation this structure characterizes is one where an agent has a goal, **GOAL.3**, which has failed, and where the goal was to preserve possession of some object. The cause of the goal failure is an act, **ACT.1**, which attempted to accomplish another goal, **GOAL.1**.

The processes of recognizing and indexing TAUs are covered more fully in [Dyer 1983] and [Dolan 1984]. Dolan [1984] gives a detailed explanation of the comprehension process and memory model that allows the recognition of TAUs in this format.

*The situation is taken from Aesop's fables; the version below is taken from Bewick's collection [Bewick 1973].

As we mentioned above, **The Fox and the Crow** instantiates two other TAUs: (1) TAU-VANITY is the planning error of allowing personal vanity to dictate plan choice; (2) TAU-ULTERIOR is the planning error of not considering another agent's possible motives before acting.

These TAUs can be combined to form new planning heuristics. There are two key problems in TAU acquisition:

1. How does a program know which TAUs to select and examine for combination attempts?
2. Once selected, how are TAUs actually combined to form new planning and indexing structures?

Both 1. and 2. are non-trivial. A sophisticated planner will have numerous stories indexed by multiple TAUs in memory. Attempting to combine TAUs arbitrarily would lead to combinatoric problems. Fortunately, memorable stories (such as Aesop's fables) are designed to give novel planning advice through illustrating planning errors. Thus, TAU selection can be governed by the following strategy:

```
+-----+
| WHENEVER two TAUs share concepts in an observed |
| planning situation,                               |
| TRY to combine them to form a novel planning    |
| construct                                         |
+-----+
```

This heuristic can only be applied *after* reading a story and thus the heuristic serves as a form of learning by example. The *comprehension of the story* thus directs the learning algorithm to the planning errors to combines and indicates which concepts are shared.

There are two ways to combine TAUs based on the way they share concepts:

- (1) specialization,
- (2) combination (chunking).

Recent work in specialization learning includes [DeJong 1983], [Lebowitz 1980], and [Kolodner 1980]. All three researchers formulate methods for creating new planning knowledge through specialization, but do not have a method extensible to chunking. Most research in learning by chunking has been in domains where there is no counterplanning [Laird 1984, Rosenbloom et al 1986].

A detailed discussion of the learning process for TAUs can be found in [Dolan and Dyer 1985]. Here we will simply outline the procedure and highlight the features of the representation that facilitate learning.

3 Creating New TAUs through Specialization

TAU-ULTERIOR represents the situation in which ACTOR1 tells ACTOR2 information that motivates ACTOR2 to perform an act which results in the disablement of one of ACTOR2's goals, while at the same time achieving a goal of ACTOR1's which ACTOR2 did not foresee. In the representation for TAU-ULTERIOR in Figure 2 below, outlined concepts are those which have ACTOR1 in the 'actor' slot.

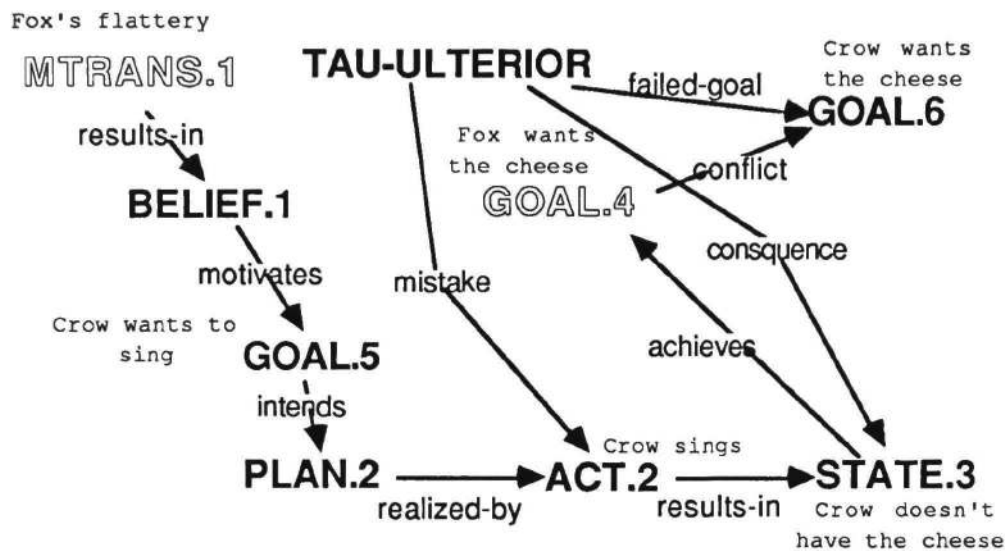


Figure 2

From the description of TAU-ULTERIOR we can see that in some sense it contains TAU-CONF-ENABL because all of the distinct concepts, those not linked by 'parent' links from TAU-CONF-ENABL, are present in TAU-ULTERIOR. The only concepts that TAU-ULTERIOR has that TAU-CONF-ENABL does not are **MTRANS.1**, **BELIEF.1**, and **GOAL.4**. TAU-CONF-ENABL brings two relationships to this specialization: the relationship between **GOAL.2** and **GOAL.3** is that a goal which was achieved in the past has now failed; the relationship between **STATE.2** and **STATE.1** is that a current state disables a previous state. These relationships denote time dependencies not present in TAU-ULTERIOR.

All of the distinct components from TAU-CONF-ENABL are shared with TAU-ULTERIOR, but TAU-CONF-ENABL has additional constraints. We say that TAU-ULTERIOR contains TAU-CONF-ENABL. The containment relation allows us to create a new TAU, a specialization of TAU-ULTERIOR. The new TAU is formed by taking the extra constraints from TAU-CONF-ENABL ("extra" meaning those not already in TAU-ULTERIOR) and conjoining them with those in TAU-ULTERIOR. This TAU, which we will call TAU-MISDIRECTED-ENABL is used to characterize situations where ACTOR1 says something to ACTOR2 to deflect ACTOR2's attention away from an enablement condition on a current goal. TAU-MISDIRECTED-ENABL, shown in

Figure 3 below, is planning advice specific to possession goals and ulterior motives.

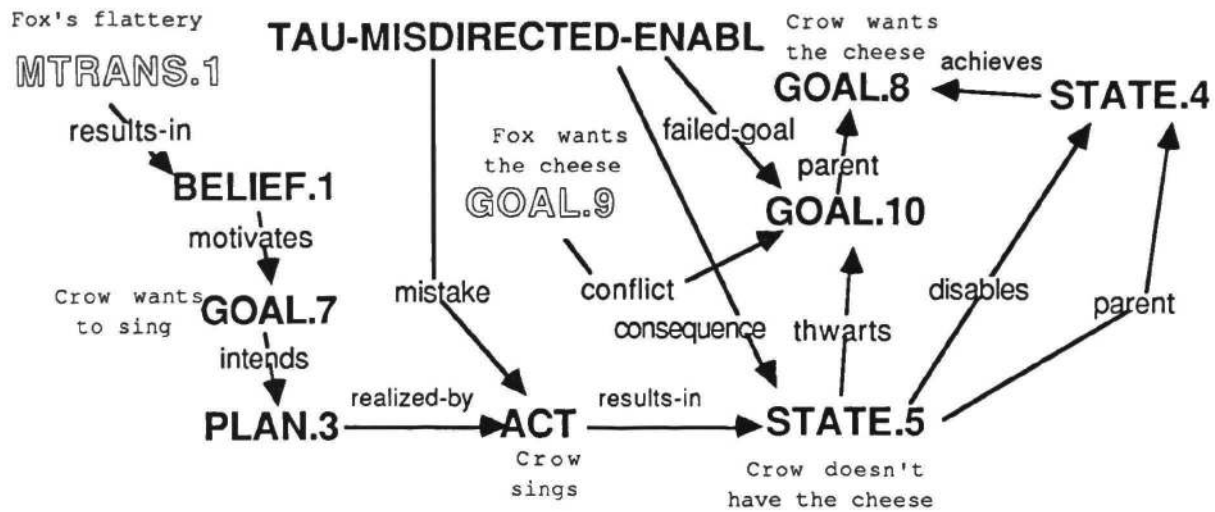


Figure 3

This TAU is also used to process stories such as the following one,

Joey's Waffle

Joey was sitting at the breakfast table about to eat his waffle. His sister Mary said, "Look Joey, Halley's comet!". When Joey looked up Mary grabbed the waffle and started eating it.

Here the goal that failed for Joey was to maintain possession of the waffle. The enablement condition on that goal was to attend his eyes to the waffle and guard it. The condition was removed when he looked up.

In some sense, when one TAU contains another as with TAU-ULTERIOR and TAU-CONF-ENABL, the contained TAU mediates the application of the containing TAU. An example of how another TAU can mediate for TAU-ULTERIOR is,

Bogus Job Advice

Fred was a post-doc working at the college. He told his boss, Henry, that there was a faculty position open at the university across town. Henry decided to give his notice at the college and apply for the job. When Henry left, Fred took his position at the college. As it turned out, Henry was not quite right for the new position and didn't get the job.

In this case, TAU-ULTERIOR is mediated by the TAU that encodes the adage, "Look before you leap." The way the two other TAUs mediate TAU-ULTERIOR is show in Figure 4 below. We see that this organization of TAUs points out the differences and similarities of the two stories on a thematic level.

4 Creating New TAUs through Combination

As we saw above, we can get a non-trivial specialization of a TAU by discovering containment in a particular situation. In general TAU combination, however, neither TAU contains the other. In these cases we must examine the relationships or constraints among the concepts *not* shared between the two TAUs. For an example, consider TAU-VANITY. This TAU represents a situation in which ACTOR2 believes he has a special skill and is thus motivated to have a goal (of "showing off" in the Fox and Crow story) that will interfere with pre-existing goals.



Figure 4

TAU-ULTERIOR and TAU-VANITY, as they are instantiated by **The Fox and the Crow**, share a number of concepts. The concepts that they do not share are **MTRANS.1** and **GOAL.4** from TAU-ULTERIOR and **GOAL.14** from TAU-VANITY.

The key to combining TAUs is finding the relationship between concepts that are not common to two TAUs being combined. Although **BELIEF.1** and **BELIEF.2** are instantiated by the same concept, they do not belong to identical parts of the graph and this fact provides the point at which we attach the two TAUs. The TAU that we get by combining TAU-VANITY and TAU-ULTERIOR is TAU-SUCKERED. A detailed example of this process is given in [Dolan and Dyer, 1985]. This TAU embodies the planning failure of allowing someone else to take advantage of one's dormant goals by providing one of the missing enablement conditions on that goal. In the case of **The Fox and the Crow** the dormant goal is the Crow's goal to show off. The missing enablement condition is a receptive audience. The Fox provides that audience and so tricks the Crow into defeating her standing goal of keeping the cheese.

TAU-SUCKERED is a new *counterplanning technique* which can be used by a planning program in situations where the appropriate constraints have been met. The

representation for TAU-SUCKERED is shown in Figure 5.

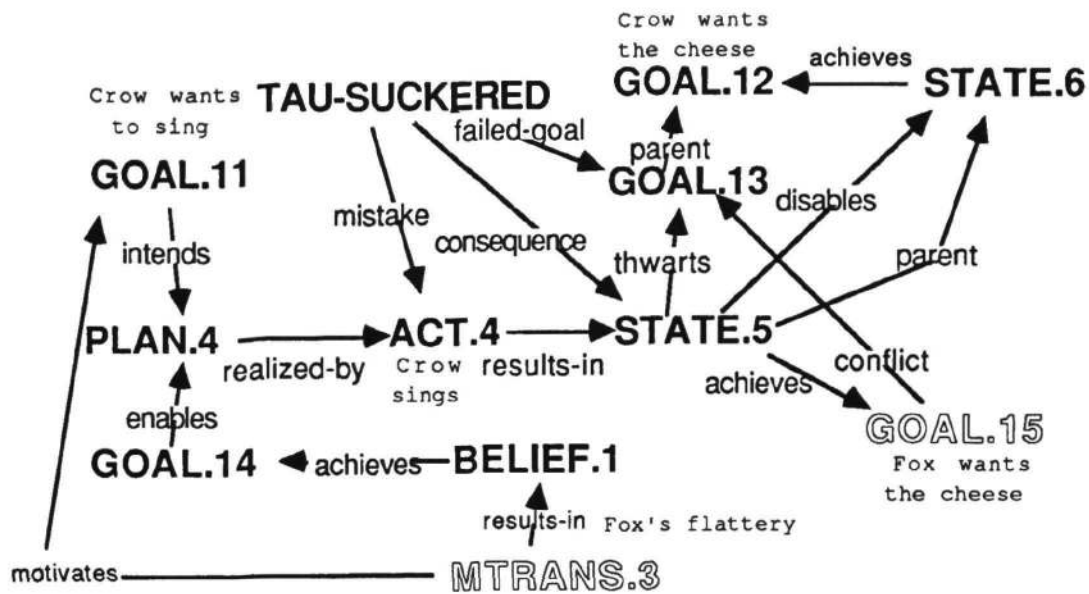



Figure 5

It is intuitively obvious to people TAU-SUCKERED represents the proper lesson to be learned by a reader who knows about vanity and physical enablements, but who has never seen this kind of trickery before.

5 Summary of TAU Learning

In both cases of TAU learning we have examined here, the representation played a central role in guiding the learning. In specialization, the presence of the 'parent' links allowed us to add extra constraints onto TAU-ULTERIOR. In the case of combination, the mapping between abstract concepts in the graph and concepts in the story showed us where to connect the two TAUs together.

The problem of learning a new TAU from two old ones, either by specialization or by combination, can be thought of as a graph union problem. The elements of the 'binding-spec' are nodes and the 'constraints' are arcs. The union of the graphs represents the intersection of the situations that conform to the constraints. In matching abstract graphs, however, we have a problem. If we want to take the union of two graphs below:



We have to make sure that the constraints 'a' and 'b' don't conflict with each other, making their union an inconsistent graph. That problem is solved here because we

have a situation, found in the story, which has all the constraints among the components of the story. Therefore we know that the new structure is a *possible* combination. This is why *learning from examples* is important because *the examples provide a model of what is actually possible*.

6 Using TAUs in Planning

Using TAUs in planning requires two phases: (1) discovering a potential planning error, and (2) taking corrective action. The first phase is accomplished in the same manner as TAU recognition in the task of comprehending stories. As the planner examines the state of the world and formulates various plans, the concepts are "played back" as if they were occurring in a story. When a TAU trigger of TAU-SUCKERED, such as MTRANS.3, the use of flattery, is encountered we have a potential planning error. Possible outcomes of the situation are examined to see if they meet the constraints of the TAU.

In these cases, where some outcome of the current situation might result in a goal failure, corrective action needs to be determined. In informal protocols taken from **The Fox and the Crow**, people seem to think that the Crow could have done three things: (1) put the cheese down and sang, (2) eaten the cheese and sang, or (3) refused to sing. Along with the structure present in TAU-CONF-ENABL, three general heuristics yield these two solutions. The first is,

```
+-----+
| IF a new goal might cause a standing goal to fail |
| TRY removing a possible resource contention         |
| between plans for the two goals                   |
+-----+
```

The standing goal is to maintain control over the cheese. The new goal is the goal to sing. The common resource is the Crow's mouth. Since there is no way to achieve the new goal without using the Crow's mouth, the system would choose to find another way to support the cheese.

The second solution, that of eating the cheese on the spot, is actually an instance of removing the need for the standing goal. In that case the guiding heuristic is,

```
+-----+
| IF a new goal might cause a standing goal to fail |
| TRY relieving the need for the standing goal       |
+-----+
```

Here the system would see that the support for having the goal of controlling the cheese was to satisfy a hunger goal. If the hunger goal is satisfied, then the need to control the cheese goes away.

The function of TAUs such as TAU-SUCKERED is to point out goals that need to be

dropped. The goal conflict shown in TAU-SUCKERED points to the goal we need to examine. The heuristic,

```
+-----+
| IF a potential goal conflict is found           |
| LOOK for a current plan which would thwart the |
| conflicting goal and reconsider the need support|
| for the current plan                           |
+-----+
```

directs the program to look at 'thwarting' constraints for the ACT that could defeat the goal. On the other hand, the heuristic,

```
+-----+
| IF a potential goal conflict is found           |
| LOOK for a current plan that would achieve the |
| conflicting goal for the other actor, and reconsider |
| the support for the current plan               |
+-----+
```

directs the program to look at 'achievement' constraints for the ACT which could defeat the goal. In the case of TAU-SUCKERED, both these heuristics give the same result.

Our hypothesis is that the number of these heuristics is small, but that a program with a large number of TAUs could correct most of its planning errors using a few heuristics, the knowledge contained in TAUs, and well chosen stories as input.

6 Progress and Future Work

The CRAM is under development as part of this research. Currently, CRAM is able to understand stories that are parsed by a phrasal parser using a conceptual memory implemented in a frame representation system. CRAM finds the planning errors in each story and characterizes them in terms of one or more TAUs. These TAUs are then used to index the story in memory for later retrieval. CRAM is also able to create new planning error descriptions based on the methods outlined in this paper.

The next step in the on-going development of the theory presented here is to give CRAM enough real-world knowledge so that it can understand a large number of stories, and to observe its performance as it learns more planning errors.

In addition, CRAM will be able to give advice to correct the character's planning errors.

7 Conclusions

The approach presented here allows both specialization-based learning and chunking-based learning of planning errors in multiple planning agent domains. The structures learned can be used both for critiquing plans and also for generating counterplanning advice.

References

- Bewick, T., Illustrator, *Treasury of Aesop's Fables*, Avenel Books, New York, 1973.
- Carbonell, J. G., *Subjective Understanding: a Computer Model of Belief Systems*, Yale Ph.D Dissertation, Technical Report 150, 1979.
- DeJong, G., "Acquiring Schema Through Understanding and Generalizing Plans", in *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, 1983.
- Dolan, C. P., *Memory Based Processing for Cross Contextual Reminding: Reminding and Analogy Using Thematic Structures*, UCLA Computer Science Department, M.S. Thesis, 1984.
- Dolan, C. P. and Dyer, M. G., "Learning Planning Heuristics through Observation", in *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, 1985.
- Dyer, M. G., *In-Depth Understanding: A Computer Model of Integrated Processing for Narrative Comprehension*, The MIT Press, Cambridge, Mass., 1983.
- Dyer, M. G., "Understanding Stories through Morals and Reminders", in *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, 1983.
- Dyer, M. G., "The Role of TAUs in Narrative Processing", in *The Proceedings of the Third Annual Meeting of the Cognitive Science Society*, 1981.
- Franklin, B., *Poor Richard: The Almanacks for the Years 1733-1758*, Heritage Press, New York, 1964.
- Kolodner, J. L., *Retrieval and Organization Strategies in Conceptual Memory: A Computer Model*, Technical Report 187, Yale University, Department of Computer Science, Ph.D. Dissertation, 1980.
- Laird, J. E., Rosenbloom, P. S., and Newell, A., "Towards Chunking as a General Learning Mechanism", in *Proceeding of the National Conference on Artificial Intelligence*, 1984.
- Lebowitz, M., *Generalization and Memory in an Integrated Understanding System*, Technical Report 186, Yale University, Department of Computer Science, Ph.D. Dissertation, 1980.
- Meehan, J., *The Metanovel: Writing Stories by Computer*, Technical Report 74, Yale University, Computer Science Department, Ph.D. Dissertation, 1979.

Rosenbloom, P., et al. "Current Research on Learning in SOAR, in *Machine Learning: A Guide to Current Research*, Mltchell, Carbonell, and Michalski (Eds), Kluwer Academic Publishers, 1986.

Sacerdoti, E., *A Structure for Plans and Behavior*, Elssvier North-Holland, 1977.

Schank, R. C., "Conceptual Dependency: A Theory of Natural Language Understanding", *Cognitive Psychology*, Vol 3, No 4, 1972.

Schank, R. C. and Abelson, R. P., *Scripts, Plans, Goals, and Understanding: An Inquiry into Human Knowledge Structures*, Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1977.

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

Sussman G. J., *A Computational Model of Skill Aquisition*, AI Technical Report 297, AI Laboratory, MIT, Ph.D. Dissertation, 1973.

Wilensky, R., *Understanding Goal-Based Stories*, Technical Report 140, Yale University, Department of Computer Science, 1978, Ph.D. Dissertation.

Wilensky, R., *Plannning and Understanding*, Addison-Wesley, 1983.