

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Acquiring Special Case Schemata In Explanation-Based Learning

Permalink

<https://escholarship.org/uc/item/15s810dd>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 9(0)

Authors

Shavlik, Jude W.

DeJong, Gerald F.

Ross, Brain H.

Publication Date

1987

Peer reviewed

ACQUIRING SPECIAL CASE SCHEMATA IN EXPLANATION-BASED LEARNING *

Jude W. Shavlik[†]
Gerald F. DeJong
Artificial Intelligence Research Group
Coordinated Science Laboratory

Brian H. Ross
Department of Psychology

University of Illinois at Urbana-Champaign

ABSTRACT

Much of expertise in problem-solving situations involves rapidly choosing a tightly-constrained schema that is appropriate to the current problem. The paradigm of *explanation-based learning* is being applied to investigate how an intelligent system can acquire these "appropriately general" schemata. While the motivations for producing these specialized schemata are computational, results reported in the psychological literature are corroborated by a fully-implemented computer model. Acquiring these *special case* schemata involves combining schemata from two different classes. One class contains domain-independent problem-solving schemata, while the other class consists of domain-specific knowledge. By analyzing solutions to sample problems, new domain knowledge is produced that often is not easily usable by the problem-solving schemata. Special case schemata result from constraining these general schemata so that a known problem-solving technique is guaranteed to work. This significantly reduces the amount of planning that the problem solver would otherwise need to perform elaborating the general schema in a new problem-solving situation. The model and an application of it in the domain of classical physics are presented.

INTRODUCTION

We are investigating the role of specialized knowledge in schema-based problem solvers. This research illustrates the importance of special-purpose knowledge and demonstrates how the knowledge can be acquired. Furthermore, this specialized level of knowledge is computationally motivated and corroborates a number of findings in the psychological literature.

In a schema-based approach to problem solving, a few general schemata are brought to bear on a problem. Very little searching is performed, therefore the system can only solve those new problems that can easily be made to fit into an existing general schema. We are examining how a computer system can learn new problem-solving schemata for itself. The paradigm we adopt is *explanation-based learning* [DeJong81, DeJong86, Mitchell86], for which there is already some psychological evidence [Ahn87]. In this type of learning, a specific problem solution is generalized into a form that can be later used to solve conceptually similar problems. The generalization process is driven by the *explanation* of why the example solution worked. Extensive knowledge about the domain at hand allows the explanation to be developed and then extended. The resulting schema is quite general. Its generality is limited in part by characteristics of the observed example, but primarily by the system's domain model.

One would expect that such problem-solving schemata should be as general as possible so that they might each cover the broadest class of problems. Indeed, our previous research has been primarily aimed at the acquisition of such maximally general schemata. Recently, for computational reasons, we have adopted an intermediate level of generalization for various schemata. The class of intermediate generality schemata improves the performance of the problem solver by supplying "appropriately general" schemata instead of forcing the system to rely on its maximally general schemata. This results in much improved efficiency at a relatively minor cost in

* This research was partially supported by the National Science Foundation under grant NSF IST 85-11542.

[†] University of Illinois Cognitive Science/Artificial Intelligence Fellow.

generality.

Automatically acquiring schemata of the intermediate level of generality requires that the system's schemata be organized into two classes: a) schemata that represent knowledge of the domain of application and b) schemata that represent general problem-solving knowledge, which apply across many application domains. New schemata are learned by our implemented system as described in [Shavlik87a]. As well as storing the new schema in its general form, the system also stores special cases. These special cases are the result of composing the new, general schema with a small number of problem-solving schemata. A successful composition results in a specialization which is guaranteed to work using the composed problem-solving technique. This frees the problem solver from performing the planning that would otherwise be required to elaborate the general schema to fit the current problem-solving episode. The system can, of course, always resort to its collection of maximally general schema.

An example (discussed further in a later section) that illustrates this idea involves momentum conservation, a fundamental concept in physics. The explanation-based generalization of a sample collision problem leads to a physics formula that describes how external forces change a system's momentum. This general schema is broadly-applicable, but ascertaining that it will lead to the solution of a given problem requires a good deal of work. The constructed special case states that when there are no external forces, momentum is conserved.

Although the motivation for this intermediate level of generalization was computational, the use of this level helps to reconcile our approach with a variety of psychological evidence showing that problem solvers use highly specific schemata [Chase73, Hinsley77, Schoenfeld82, Sweller85]. Much of expertise consists of rapidly choosing a tightly-constrained schema appropriate to the current problem. However, the difference between the knowledge of an expert and a novice cannot be explained on the basis of number of schemata alone. The scope and organization of these schemata have been shown in psychological experiments to be qualitatively different [Chi81, Larkin80, Schoenfeld82]. In representing a problem, novices make great use of the specific objects mentioned in the problem statement, while experts first categorize according to the techniques appropriate for solving the problem.

We have found that the intermediate-level schemata generated by our system are similar in scope of applicability to those that human experts appear to possess. For example, the conservation of momentum problem results in a special case schema characterized by the absence of external forces and the specification of a *before* and *after* situation. These features are those cited by experts as the relevant cues for the principle of conservation of momentum (see table 12 of [Chi81]). It should be noted that it was not our explicit intent to model this psychological data. Rather, computational efficiency considerations led to a system that produced results matching this empirical data.

OVERVIEW OF THE MODEL

Our model involves combining two types of schemata: domain-independent problem-solving schemata (e.g., a schema for utilizing a conserved quantity to solve a problem) and schemata that represent domain-specific knowledge (e.g., Newton's laws). People with mature problem-solving backgrounds possess the first type of schema and are told schemata of the second type when introduced to a new domain. Through study they acquire a large collection of schemata that combine aspects of both types, thereby increasing their performance in the domain. Combining general problem-solving techniques with domain specific knowledge produces schemata that, when applied, lead to the rapid solution of new problems.

Figure 1 contains an overview of our model. We assume a known problem-solving schema is used to understand a solution to a specific problem. The explanation-based analysis of the solution may lead to the construction of a new broadly-applicable schema. The generalization process often produces a new schema that, in its fullest form, is not usable by the originally applied problem-solving schema. Constraining the general result so that this problem-solving schema does apply produces a special case. In the special case schema, the constrained schema, its constraints, and the

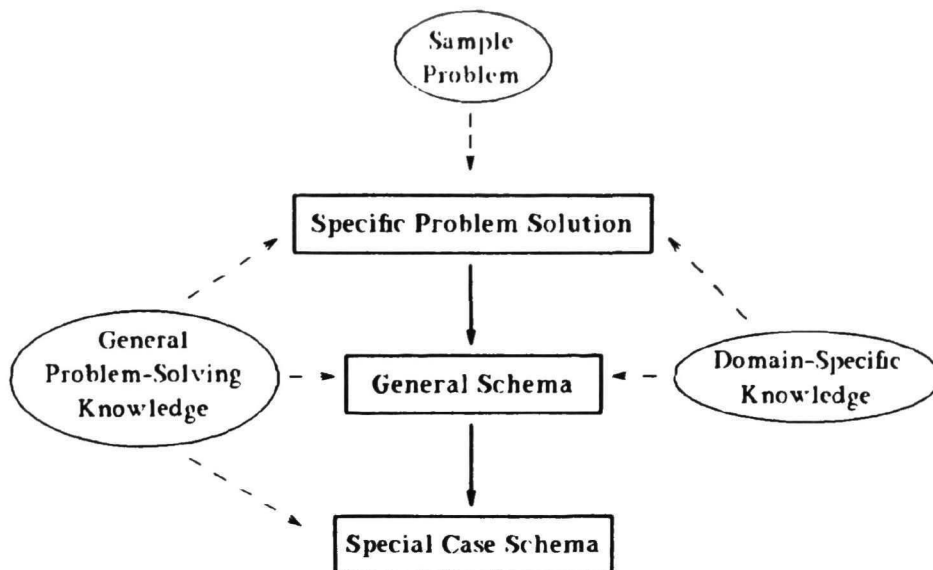


Figure 1. Overview of the Model

original problem-solving schema are packaged together to produce a specialized problem-solving strategy.

Figure 2 shows the relation we propose between a general schema and its special cases. Although not shown in the figure, there can be special cases of the special cases. Retrieval cues directly index the special cases.¹ If the indexed special case is not applicable, the general concept is then accessed. Besides being constructed when the general case is acquired, a new special case may be created whenever the general case is used to solve a later problem.

The next section presents examples of the construction of general and specific case schemata in the domain of classical physics.

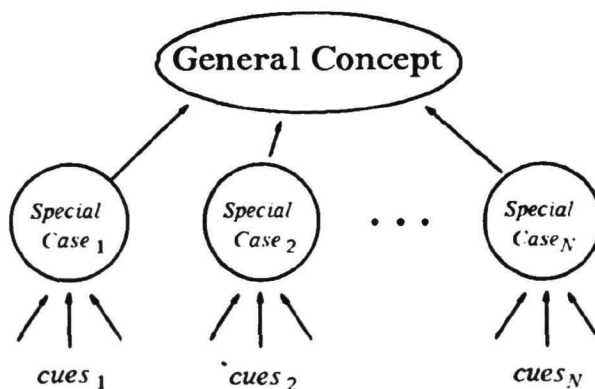


Figure 2. Inter-Schema Organization

¹ The method of using these special case cues to select the appropriate schema is not addressed in this paper. Possible indexing techniques include approaches based on *discrimination nets* [Feigenbaum63, Kolodner84, Schank82] and approaches based on *spreading activation* [Anderson83, Quillian68].

AN APPLICATION OF THE MODEL IN CLASSICAL PHYSICS

We have implemented a psychologically-plausible model of the process by which a mathematically-sophisticated student becomes a better problem solver in a new domain. In particular, we have been investigating the transition from novice to expert problem solver in the field of classical physics. Our model, implemented in a computer system named **Physics 101**, assumes the student has an understanding of mathematics through introductory calculus.

There are three main components of our model. The first is a model of how operators are chosen during problem solving [Shavlik86b]. The second explores the processes by which one can understand and generalize solutions to novel problems [Shavlik85]. The third, the topic of this paper, addresses the process of storing learned results so that they can be used to improve subsequent problem solving.

The next sections discuss two sample problems analyzed by **Physics 101**. One involves momentum conservation, and the other, energy conservation. The general schema produced in these two cases, the resulting special cases, and their selection cues are presented.

The schema that makes use of conserved quantities during problem solving is contained in table 1. (Terms beginning with a question mark are universally instantiated variables.) To apply this schema, a formula that is constant with respect to some variable is needed. This formula is instantiated at two different points. If the values of all but one variable at these two points are known, simple algebra can be used to easily find the unknown.

Example 1 - Momentum Conservation

One of the problems presented to **Physics 101** involves a collision among three balls. In this one-dimensional problem (shown in figure 3), there are three balls moving in free space, without the influence of any external forces. Nothing is specified about the forces between the balls. Besides their mutual gravitational attraction, there could, for example, be a long-range electrical interaction and a very complicated interaction during the collision. In the initial state (**state A**) the first ball is moving toward the two stationary ones. Some time later (**state B**) the second and third balls are recoiling from the resulting collision. The task in this problem is to determine the velocity of the first ball after the collision.

Table 1. Conserved Quantity Schema

Preconditions

```
(AND (IsaFormula ?formula )
      (ConstantWithRespectTo ?formula ?x )
      (SpecificPointOf ?x1 ?x )
      (SpecificPointOf ?x2 ?x )
      (≠ ?x1 ?x2)
      (= ?leftHandSide (InstantiatedAt ?formula ?x1))
      (= ?rightHandSide (InstantiatedAt ?formula ?x2))
      (= ?equation (CreateEquation ?leftHandSide ?rightHandSide ))
      (AllButOneValueKnown ?equation ) )
```

Schema Body

```
(SolveForSingleUnknown ?equation )
```

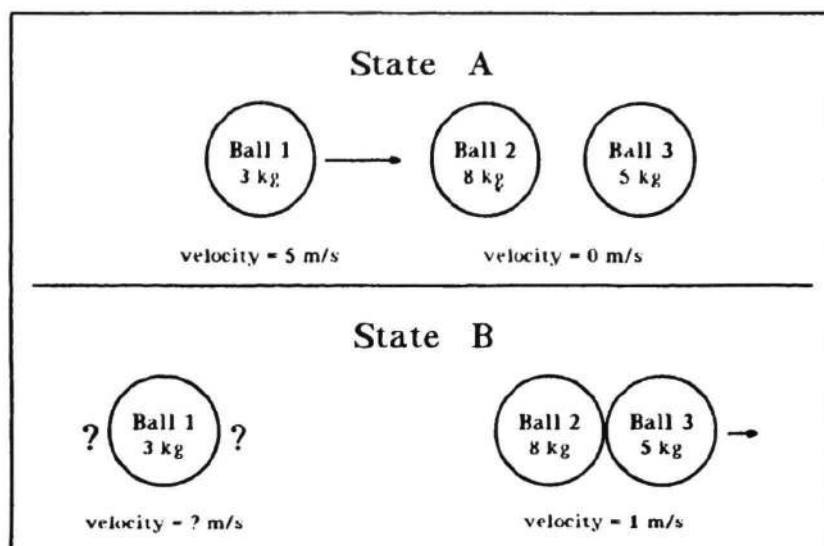


Figure 3. A Three-Body Collision Problem

A teacher's solution to figure 3's problem is analyzed by **Physics 101**. The teacher's solution uses the concept of momentum ($mass \times velocity$) conservation to solve the problem. Since this is a conservation law, the time between the two states need not be known.² (Another important attribute of momentum conservation is that the properties of the inter-object forces need not be provided.) In verifying the provided solution, the system applies a problem-solving schema in which a constant function is equated at two different points (table 1). In accordance with the explanation-based learning approach, the system's justification of the provided solution is generalized as far as possible while maintaining the veracity of the solution technique. This results in the general schema presented in table 2. (See [Shavlik86a] or [Shavlik87a] for more details on the construction of this schema.)

The explanation-based approach results in a formula that applies to situations significantly different from the sample problem. In addition to not being restricted to problems containing exactly three objects, the newly-acquired formula is not restricted to situations where the external forces are all zero. Instead, an understanding of how the external forces effect momentum is obtained. This process also determines that there is no constraint that restricts this formula to the x -direction. It applies equally well to the y - and z -components of velocity. Hence, the acquired formula is a vector law. The mathematical operations used in the specific solution require, for the solution strategy to be valid, that all objects have non-zero mass and that these masses are constant over time. Finally, the generalization algorithm determines that the inter-object forces need not be known, since they are algebraically cancelled during the derivation of the momentum law.

Notice that the result in table 2 is *not* a conservation law. It describes how the momentum of a system evolves over time. Although this new formula applies to a large class of problems, recognizing its applicability is not easy. The external forces on the system must be summed and a possibly complicated differential equation needs to be solved. Applying this law requires more than counting the number of unknowns in the equation, determining there is only one, and then using simple algebra to find its value.

In order for the originally used problem-solving schema (table 1) to be applicable to this new formula, it must be the case that momentum be constant with respect to time and easily calculable at two different times. This means that the derivative of momentum be zero, which leads to the

² When the inter-state time is unknown, simply solving the equations of motion resulting from Newton's laws is not possible.

Table 2. The General Momentum Law

Formula

$$\frac{d}{dt} \sum_{i=1}^n mass_i velocity_{i,\gamma_c}(t) = \sum_{i=1}^N force_{external,i,\gamma_c}(t)$$

Preconditions

- (AND (IsaComponent ?c)
- $\forall i \in 1, \dots, n$ (NOT (ZeroValued $mass_i$))
- $\forall i \in 1, \dots, n$ (IndependentOf $mass_i$ t))

Eliminated Terms

$$\forall i \forall j \neq i force_{i,j,\gamma_c}(t)$$

requirement that the external forces sum to zero. When this occurs, the momentum of a system can be equated at *any* two distinct states. The special case schema for momentum conservation is contained in table 2sc. Since this is a conservation schema, the time at which each state occurs need not be provided in a problem for this schema to apply.

Table 2sc. The Special-Case Momentum Law

Formula

$$\sum_{i=1}^n mass_i velocity_{i,\gamma_c}(?t_1) = \sum_{i=1}^n mass_i velocity_{i,\gamma_c}(?t_2)$$

Preconditions

- (AND (IsaComponent ?c) (Time ?t₁) (Time ?t₂) ($\neq ?t_1 ?t_2$))
- $\forall i \in 1, \dots, n$ (NOT (ZeroValued $mass_i$))
- $\forall i \in 1, \dots, n$ (IndependentOf $mass_i$ t))

Eliminated Terms

$$\forall i \forall j \neq i force_{i,j,\gamma_c}(t), ?t_1, ?t_2$$

Special Case Conditions

$$\forall i \in 1, \dots, n force_{external,i,\gamma_c} = 0$$




Figure 4. A Falling Brick

Example 2 - Energy Conservation

A second problem (figure 4) presented to **Physics 101** involves a brick falling under the influence of gravity. Again, information at two different states is presented. The mass of the brick, its initial velocity, and its height in the two states are provided. The goal is to find its velocity in the second state. The teacher's solution to this problem uses energy conservation. The kinetic energy ($\frac{1}{2} mass \times velocity^2$) plus the potential energy ($mass \times g \times height$) in the two states is equated. The general law **Physics 101's** produces by analyzing the sample solution is presented in table 3.

The general energy conservation law applies whenever the total force on an object is known. Notice, though, that a rather complicated vector integral involving the scalar (dot) product of two vectors needs to be computed if this general law is to be used. To use this formula, it is not sufficient to possess knowledge of the values of variables at two different times. A problem solver must also know how the net force depends on position for a continuum of times. In the specific problem there is a *constant* net force (gravity). When the force is constant the problem is greatly simplified. Integrating a constant force leads to a potential energy determined by that constant force multiplied by the object's position. The position only needs to be known at the two distinct times, and not for *all* intervening times. The special case schema for energy conservation is contained in table 3sc. Again, since this is a conservation schema, the time at which each state occurs need not be known.

SIMILARITY-BASED APPROACHES TO LEARNING SPECIAL CASES

A common induction scheme is to posit that learners compare particular instances of a concept (such as specific problems of a problem type) and abstract out those aspects that are common to both problems [Anderson83, Michalski83, Mitchell78, Posner68]. The fact that problem solvers

Table 3. The General Energy Law

Formula

$$\frac{d}{dt} [\frac{1}{2} mass_{\gamma_i} velocity^2_{\gamma_i}(t) - \int \overline{force}_{net, \gamma_i}(t) \cdot \overline{dposition}_{\gamma_i}] = 0 \frac{kg \ m^2}{s^3}$$

Preconditions

(AND (Object ?i) (IndependentOf $mass_{\gamma_i}$ t) (NOT (ZeroValued $mass_{\gamma_i}$)))

Table 3sc. The Special-Case Energy Law

Formula

$$\begin{aligned} & \frac{1}{2} \text{mass}_{\gamma_i} \text{velocity}^2_{\gamma_i} (?t_1) + \text{mass}_{\gamma_i} g \text{position}_{\gamma_i, \gamma_c} (?t_1) \\ &= \frac{1}{2} \text{mass}_{\gamma_i} \text{velocity}^2_{\gamma_i} (?t_2) + \text{mass}_{\gamma_i} g \text{position}_{\gamma_i, \gamma_c} (?t_2) \end{aligned}$$

Preconditions

(AND (Object ?i) (Component ?c) (Time ?t₁) (Time ?t₂) (\neq ?t₁ ?t₂)
(IndependentOf *mass* _{γ_i} t) (NOT (ZeroValued *mass* _{γ_i})))

Eliminated Terms

?t₁, ?t₂

Special Case Conditions

$$\overline{\text{force}}_{\text{net}, \gamma_i}(t) = \text{mass}_{\gamma_i} g \hat{c}$$

use highly specific schemata supports such a view, since these schemata would arise whenever two problems from an intermediate level problem type are compared. Although we believe that similarity-based generalization is an important means of learning, especially for novices [Gentner87, Ross84, Ross87], the research described in this paper shows that many of these highly specific schemata can arise from an explanation-based approach. Even some strong proponents of example comparison learning have begun to incorporate some explanation-based ideas in order to account for how much is learned from one example [Anderson87].

Because explanation-based learning requires extensive domain knowledge, it clearly is not appropriate for all learning in a new domain. However, it may be useful even in early learning if the new domain relies heavily upon a domain for which the novice does have substantial knowledge. Because mathematics underlies many other domains, a novice with some mathematical sophistication may be able to make use of explanation-based techniques without extensive knowledge of the new domain.

CONCLUSION

Much of expertise in problem-solving situations involves rapidly choosing a tightly-constrained schema that is appropriate to the current problem. We are applying the paradigm of explanation-based learning to investigate how an intelligent system can acquire these "appropriately general" schemata. While our motivations for producing these specialized schemata are computational, results reported in the psychological literature are corroborated by our fully-implemented computer model.

A major issue in explanation-based learning concerns the *operationality/generalizability* trade-off [DeJong86, Mitchell86, Segre87]. A schema whose relevance is easy to determine may only be useful in an overly-narrow range of problems. Conversely, a broadly-applicable schema may require extensive work before a problem solver can recognize its appropriateness. Other approaches to selecting the proper level of generality involve pruning easily-reconstructable portions of the explanation structure. Our approach to this problem is to produce as general a schema as possible

from the analysis of a specific solution, and then construct a *special case* of this general schema. In constructing the general schema, the original explanation structure is often substantially altered during generalization [Shavlik87a, Shavlik87b]. Augmentation of the explanation is needed in order to generalize such things as the number of entities in a concept or the number of times some action is performed. A special case is produced by constraining a general schema in such a way that its relevance is easily checked. This results in additional features that a situation must possess if the special case is to apply.

Acquiring these special case schemata involves combining schemata from two different classes. One class contains domain-independent problem-solving schemata, while the other class consists of domain-specific knowledge. In our model, learning by analyzing sample problem solutions produces broadly-applicable schemata that, often, are not usable by the originally applied problem-solving schemata. Special case schemata result from constraining these general schemata so that the originally used problem-solving techniques are guaranteed to work. This significantly reduces the amount of planning that the problem solver would otherwise need to perform elaborating the general schema to match a new problem-solving episode.

Besides improving a problem solver's efficiency, special cases also indicate good assumptions to make. For instance, if you do not know what the external forces are, assume they are zero. Physics problems often require one to assume things like "there is no friction", "the string is massless", "the gravity of the moon can be ignored", etc. Problem descriptions given to students contain cues such as these, and students must learn how to take advantage of them. Facts in the initial problem statement suggest possible problem-solving strategies, while any additional requirements of the special case situations indicate good assumptions to make (provided they do not contradict anything else that is known).

This paper demonstrates that these highly-specific schemata can arise in an explanation-based fashion. Explanation-based learning requires extensive knowledge, and seems particularly suited for modelling learning by experts. Although all learning cannot be of this type, explanation-based learning can prove useful even in early learning in a new domain. This can occur if the new domain relies heavily upon another domain in which the novice learner has substantial abilities.

REFERENCES

- [Ahn87] W. Ahn, R. J. Mooney, W. F. Brewer and G. F. DeJong, "Schema Acquisition from One Example: Psychological Evidence for Explanation-Based Learning," Technical Report, AI Research Group, Coordinated Science Laboratory, University of Illinois, Urbana, IL, February 1987.
- [Anderson83] J. R. Anderson, *The Architecture of Cognition*, Harvard University Press, Cambridge, MA, 1983.
- [Anderson87] J. R. Anderson and R. Thompson, "Use of Analogy in a Production System Architecture," in *Similarity and Analogical Reasoning*, S. Vosniadou and A. Ortony (ed.), 1987.
- [Chase73] W. G. Chase and H. A. Simon, "Perception in Chess," *Cognitive Psychology* 4, (1973), pp. 55-81.
- [Chi81] M. T. Chi, P. J. Feltovich and R. Glaser, "Categorization and Representation of Physics Problems by Experts and Novices," *Cognitive Science* 5, 2 (1981), pp. 121-152.
- [DeJong81] G. F. DeJong, "Generalizations Based on Explanations," *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, Vancouver, B.C., Canada, August 1981, pp. 67-70.
- [DeJong86] G. F. DeJong and R. J. Mooney, "Explanation-Based Learning: An Alternative View," *Machine Learning* 1, 2 (April 1986), pp. 145-176.
- [Feigenbaum63] E. A. Feigenbaum, "The Simulation of Natural Learning Behavior," in *Computers and Thought*, E. A. Feigenbaum and J. Feldman (ed.), McGraw-Hill, New York, NY, 1963.

- [Gentner87] D. Gentner, "Mechanisms of Analogical Learning," in *Similarity and Analogical Reasoning*, S. Vosniadou and A. Ortony (ed.), 1987.
- [Hinsley77] D. A. Hinsley, J. R. Hayes and H. A. Simon, "From Words to Equations: Meaning and Representation on Algebra Word Problems," in *Cognitive Processes in Comprehension*, P. A. Carpenter and M. A. Just (ed.), Lawrence Erlbaum and Associates, Hillsdale, NJ, 1977, pp. 89-105.
- [Kolodner84] J. L. Kolodner, *Retrieval and Organizational Strategies in Conceptual Memory*, Lawrence Erlbaum and Associates, Hillsdale, NJ, 1984.
- [Larkin80] J. H. Larkin, J. McDermott, D. P. Simon and H. A. Simon, "Models of Competence in Solving Physics Problems," *Cognitive Science* 4, 4 (1980), pp. 317-345.
- [Michalski83] R. S. Michalski, "A Theory and Methodology of Inductive Learning," in *Machine Learning: An Artificial Intelligence Approach*, R. S. Michalski, J. G. Carbonell, T. M. Mitchell (ed.), Tioga Publishing Company, Palo Alto, CA, 1983, pp. 83-134.
- [Mitchell78] T. M. Mitchell, "Version Spaces: An Approach to Concept Learning," Ph.D. Thesis, Stanford University, Palo Alto, CA, 1978. (Also appears as Technical Report STAN-CS-78-711, Stanford University)
- [Mitchell86] T. M. Mitchell, R. Keller and S. Kedar-Cabelli, "Explanation-Based Generalization: A Unifying View," *Machine Learning* 1, 1 (January 1986), pp. 47-80.
- [Posner68] M. J. Posner and S. W. Keele, "On the Genesis of Abstract Ideas," *Journal of Experimental Psychology* 77, 3 (July 1968), pp. 353-363.
- [Quillian68] M. R. Quillian, "Semantic Memory," in *Semantic Information Processing*, M. L. Minsky (ed.), MIT Press, Cambridge, MA, 1968.
- [Ross84] B. H. Ross, "Reminders and their Effects in Learning a Cognitive Skill," *Cognitive Psychology* 16, (1984), pp. 371-416.
- [Ross87] B. H. Ross, "Reminders in Learning and Instruction," in *Similarity and Analogical Reasoning*, S. Vosniadou and A. Ortony (ed.), 1987.
- [Schank82] R. C. Schank, *Dynamic Memory*, Cambridge University Press, Cambridge, England, 1982.
- [Schoenfeld82] A. H. Schoenfeld and D. Herrmann, "Problem Perception and Knowledge Structure in Expert and Novice Mathematical Problem Solvers," *Journal of Experimental Psychology: Learning, Memory, and Cognition* 8, 5 (1982), pp. 484-494.
- [Segre87] A. M. Segre, "Explanation-Based Learning of Generalized Robot Assembly Tasks," Ph.D. Thesis, Department of Electrical and Computer Engineering, University of Illinois, Urbana, IL, January 1987. (Also appears as UIIU-ENG-87-2208, AI Research Group, Coordinated Science Laboratory, University of Illinois at Urbana-Champaign.)
- [Shavlik85] J. W. Shavlik and G. F. DeJong, "Building a Computer Model of Learning Classical Mechanics," *Proceedings of the Seventh Annual Conference of the Cognitive Science Society*, Irvine, CA, August 1985, pp. 351-355.
- [Shavlik86a] J. W. Shavlik and G. F. DeJong, "Computer Understanding and Generalization of Symbolic Mathematical Calculations: A Case Study in Physics Problem Solving," *Proceedings of the 1986 Symposium on Symbolic and Algebraic Computation*, Waterloo, Ontario, Canada, July 1986, pp. 148-153.
- [Shavlik86b] J. W. Shavlik and G. F. DeJong, "A Model of Attention Focussing During Problem Solving," *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*, Amherst, MA, August 1986, pp. 817-822.
- [Shavlik87a] J. W. Shavlik and G. F. DeJong, "Analyzing Variable Cancellations to Generalize Symbolic Mathematical Calculations," *Proceedings of the Third IEEE Conference on Artificial Intelligence Applications*, Orlando, FL, February 1987.
- [Shavlik87b] J. W. Shavlik and G. F. DeJong, "BAGGER: An EBL System that Extends and Generalizes Explanations," Technical Report, AI Research Group, Coordinated Science Laboratory, University of Illinois, Urbana, IL, February 1987.
- [Sweller85] J. Sweller and G. A. Cooper, "The Use of Worked Examples as a Substitute for Problem Solving in Learning Algebra," *Cognition and Instruction* 2, 1 (1985), pp. 59-89.