Goal-Driven Learning: Fundamental Issues (A Symposium Report)*

David Leake
Computer Science Department
Indiana University
Bloomington, Indiana 47405
(812) 855-9756
E-mail: leake@moose.cs.indiana.edu

Ashwin Ram
College of Computing
Georgia Institute of Technology
Atlanta, Georgia 30332-0280
(404) 853-9372
E-mail: ashwin@cc.gatech.edu

June 7, 1994

Abstract

In Artificial Intelligence, Psychology, and Education, a growing body of research supports the view that learning is a goal-directed process. Psychological experiments show that people with different goals process information differently; studies in education show that goals have strong effects on what students learn; and functional arguments from machine learning support the necessity of goal-based focusing of learner effort. At the Fourteenth Annual Conference of the Cognitive Science Society, a symposium brought together researchers in AI, psychology, and education to discuss goal-driven learning. This article presents the fundamental points illuminated by the symposium, placing them in the context of open questions and current research directions in goal-driven learning.

*Appears in *AI Magazine*, 14(4):67-72, 1993

1 Introduction

Learning is a central area of study for researchers interested in human cognition as well as those interested in machine intelligence. Its study has benefited greatly from the multiple perspectives provided by disciplines such as psychology, artificial intelligence, and education. In artificial intelligence, machine learning research has developed a rich repertoire of learning mechanisms. Less attention, however, has been given to understanding the issues involved in applying those methods—the issues of when learning should occur, what should be learned, and which learning strategies are appropriate in a given context. Standard machine learning systems address the question of when to learn by attempting to learn in response to every input; they address the question of what to learn by learning a user-supplied target concept (either explicit in the input provided to the system, or implicit in the user's choice of training examples); and they address the question of how to learn by applying a single fixed learning method. Although such systems provide a useful testbed for examining individual learning mechanisms, they are inadequate for use as real-world learners. The problem is that real-world situations offer countless opportunities for learning and each of those opportunities licenses the learning of infinitely-many concepts—few of which will actually be useful. Consequently, an indiscriminate learning system will expend enormous processing effort to learn things that may provide little or no benefit. In order to learn effectively in such situations, a learning system needs ways to focus the learning process.

The need for focusing concept formation is widely recognized in AI, and standard focusing methods have emerged. Inductive learning systems depend on built-in biases to constrain what they learn (e.g., (Michalski, 1983)); explanation-based learning systems depend on pre-specified target concepts and appropriate operationality criteria (e.g., (Mitchell, Keller, & Kedar-Cabelli, 1986)). In most systems, these focusing criteria are fixed (cf. (Utgoff, 1986)). However, as circumstances change, the need for learning changes as well. Because inappropriate learning may actually degrade system performance (e.g., (Minton, 1988)), effective performance depends on assuring that what is learned actually furthers system goals.

Goal-driven learning takes system goals as a starting point in the learning process. The idea of goal-driven learning is that because the value of learning depends on how well it satisfies system goals, system goals should direct decisions of when and what to learn. In this way, goal-driven learning follows the spirit of research on failure-driven learning systems, in which learning is motivated by deficiencies in system performance (*e.g.*, (Sussman, 1975; Riesbeck, 1981; Schank, 1982; Collins & Birnbaum, 1988; Hammond, 1989; Ram & Cox, 1993; Schank & Leake, 1989)). Likewise, it is in the spirit of explanation-based learning research on how to form useful target concepts (Kedar-Cabelli, 1987) and on judging the utility of learning (e.g., (Keller, 1987; Minton, 1988)). Goal-driven learning, however, takes a broader view, examining the relationships between the many possible motivations for learning and the many strategies to achieve it.

The effectiveness of goal-driven learning depends on being able to make good decisions about when and what to learn, and on selecting the best strategies for achieving the desired learning. Unlike the passive and static process used in many learning systems, goal-driven learning is itself a planful process in which selection of target concepts and learning strategies is guided by desires and needs for knowledge (Hunter, 1990).

Recent research provides growing support for goal-driven approaches to learning, both on cognitive and on functional grounds. In psychology, learner goals have been shown to have strong effects on the human learning process (e.g., (Barsalou, 1991; Zukier, 1986)); in education, learner goals have been shown to have a strong effect on student performance (e.g., (Ng & Bereiter, 1991; Scardamalia & Bereiter, 1991)); and in AI, a growing body of recent research presents functional justifications for making decisions about the usefulness of potential learning and for guiding learning according to learner goals (e.g. (desJardins, 1992; Hunter, 1990; Krulwich, Birnbaum, & Collins, 1992; Leake, 1992; Ram & Cox, 1993; Ram & Hunter, 1992; Ram, 1991)).

At the Fourteenth Annual Conference of the Cognitive Science Society, a symposium was organized to bring together researchers addressing goal-driven learning from diverse perspectives. The symposium provided a forum to present recent results and new directions in goal-driven learning and to examine the fundamental issues in goal-driven learning—how learning goals arise, how they affect learner decisions of when and what to learn, and how they guide the learning process. The symposium was jointly organized by **David Leake**, of the Computer Science Department of Indiana University, and **Ashwin Ram**, of the College of Computing of the Georgia Institute of Technology; the symposium organizers participated in a panel with **Lawrence Barsalou**, of the Psychology Department of the University of Chicago; **Ryszard Michalski**, of the Computer Science Department of George Mason University; **Evelyn Ng**, of the Faculty of Education of Simon Fraser University; and **Paul Thagard**, of the Philosophy Department of the University of Waterloo.

Prior to the symposium, a questionnaire was circulated to the panelists to clarify positions on goal-driven learning and to identify key issues. The symposium itself included the presentation of position papers by each panelist followed by open discussion among the panelists and with the audience. This overview discusses some of the fundamental issues and perspectives involved in the work discussed at the workshop, highlighting the contributions and challenges of goal-driven learning compared to traditional methods.

2 What is goal-driven learning?

One of the purposes of the symposium was to clarify the nature of goal-driven learning. We cannot simply describe goal-driven learning as "learning that reflects system goals"; this definition includes all learning systems, because any learning system may be viewed as having a built-in "goal" to perform a particular type of learning. Thus built-in "attitudes," general desires to learn, or "goals" for self-improvement are equivalent to the implicit focuses of traditional machine learning systems. To clarify the difference between goal-driven learning and traditional approaches, Thagard proposed distinguishing between *goal-relevant* learning, in which learning is relevant to goals in a weak sense, and *goal-directed* learning, in which the content of what is learned is driven by the general goals of the learner. All learning in real systems is goal-relevant in some sense, but not necessarily goal-directed.

To the symposium participants, a crucial aspect of goal-driven learning was that it be dynamically focused by current information needs. This sharpens the definition of goal-driven learning, specifying that a goal-driven learner chooses its own target concepts in response to information needs from goals *outside* the learning process. In this view, non-goal-driven learning is done by a system as routine processing, regardless of changes in its overarching goals; goal-driven learning is learning tailored towards providing the specific information (or type of information) currently needed to further overarching goals.

In order to flexibly adjust learning strategies towards satisfying system goals, a learning system must be able to reason about the information it needs. Consequently, Ram proposed that in truly goal-driven learning systems, the goals must be explicitly represented and the learning system must be able to reason about them to guide its search for information (Ram, 1991; Ram & Cox, 1993; Ram & Hunter, 1992). In this view, goal-driven learning is a process in which the learner forms and executes plans for learning needed concepts, guided by its knowledge both of what it needs to learn and of the ways that learning strategies can be applied.

This description of goal-driven learning highlights key differences between goal-driven and non-goal-driven learning systems. Traditional learning systems take a fixed "target concept" as an input (e.g., (Mitchell, Keller, & Kedar-Cabelli, 1986)), or rely on the user to select appropriate examples to define the appropriate concept (e.g., (Winston, 1975)); they also apply fixed learning strategies regardless of circumstances. Goal-driven learning systems decide when to learn, formulate their own target concepts, and decide how best to carry out their learning.

3 Effects of Goal-Driven Learning on Human and Machine Learning

The intuitive appeal of goal-driven learning is clear—to focus learning according to what the learner needs to know. However, Barsalou pointed out that because almost any cognitive processing can be construed as serving some goal, it is vacuous to make a blanket statement that human learning is goal-driven. In his view, the study of human goal-driven learning must analyze specific goals, the processing that achieves them, the learning produced during processing, and positive effects of that learning on subsequent goal achievement. Thus the study of human goal-driven learning must focus on the processing that involves learning as a side-effect. He discussed two classes of goals, explicit problem solving goals and implicit orientation goals for maintaining a world model. He argued that learning taking place during explicit problem solving tends to be goal-specific, in that the information stored centers on achievement of the particular goal being achieved. This differs from learning based on general goal orientations, which results in learning that may later serve a wide range of explicit goals.

Data collected by Ng substantiate the influence of particular classes of goals in shaping human learning (Ng & Bereiter, 1991). Based on her studies of student learning, Ng distinguished goal-driven learning (the learning that results from setting and pursuing goals beyond the completion of an assignment or learning procedure) from learning that is not motivated by learning goals (e.g., learning that occurs as an incidental outcome of other processing, or that occurs when the student attempts to complete an assignment but not to learn the concepts that the assignment is intended to teach). Although some students carry out learning tasks in order to advance their knowledge, others perform learning tasks primarily to maintain favorable appearances, to avoid criticism or to gain praise (e.g., (Dweck, 1985)).

Ng collected protocols from students learning BASIC and divided the students into three classes, according to the relative frequencies of the statements they made relating to the following types of goals: (1) goals simply to complete the current task successfully, (2) goals to learn what the assignment was intended to teach, or (3) goals to build knowledge relevant to an outside agenda for learning. Students with explicit knowledge-building goals were better able to deal constructively with problems, raised more questions, and tagged unsolved problems for future investigation.

Other panelists argued that goal-driven learning promises to have a profound effect on machine learning as well. Leake, Michalski and Ram pointed out that the necessity to function effectively despite limited processing resources makes it important for programs to be able to decide what it would be useful to know; Ram discussed the issue of deciding whether to pursue the related knowledge goal immediately, or to suspend it until a better opportunity to satisfy it arises in the future. Ram and Michalski also highlighted the need to reason about the ways in which learning goals can be achieved, that is, the different learning strategies that can be used. Thagard pointed out that this reasoning must consider not only how to achieve given goals, but how to reconcile conflicts between those goals; in this way goal-driven learning both furthers goal achievement and clarifies what the system's goals are.

4 Issues in Goal-Driven Learning

Building goal-driven learning systems depends on addressing concrete issues affecting the goal-driven learning process. The symposium panelists identified five fundamental questions that must be addressed by theories of goal-driven learning:

1. What are the types of learning goals?
2. How do learning goals arise?
3. How do learning goals affect the learning process?
4. How do different types of learning goals relate to each other?
5. How are learning goals represented?

4.1 What are the types of learning goals?

Developing goal-driven learning mechanisms depends on identifying the goals that can drive learning. Panelists proposed four ways of taxonomizing learning goals:

By overarching tasks: Most of the panelists taxonomized learning goals by the tasks that give rise to them. Barsalou distinguished two types of goals, those serving explicit problem solving and those reflecting implicit orientations. Leake described a taxonomy of task-based information requirements relevant to the generation of explanations. Ng categorized the goals underlying student effort on class assignments into three types: *task completion* goals, which are achieved by simply completing the assignment satisfactorily, *instructional* goals, which reflect what the program of instruction is intended to teach, and *knowledge-building* goals, which relate to the student's own purposes or agenda for learning (e.g., to use the knowledge that the exercise is designed to impart in order to further some other goal).

By knowledge gap or failure necessitating learning: Ram proposed that learning goals can also be characterized by the type of reasoning failure from which they arise, as when forgetting gives rise to a goal to learn a new index (Ram & Cox, 1993); Leake examined the types of learning goals that arise from failures during understanding—anomalies—that must be resolved (Leake, 1992).

By the learning that results: Ram and Michalski pointed out that taxonomies may also be based on the type of learning that results. For example, knowledge acquisition goals seek to learn by acquiring specific types of new knowledge, and knowledge organization goals seek to learn by reorganizing or reindexing existing knowledge.

By the learning activity: Michalski suggested that learning goals can also be characterized by the particular learning activities that they involve, such as generalizing given knowledge, discovering regularities, or placing that knowledge in more operational form.

4.2 How do learning goals arise?

Given the ramifications of learning goals on processing, an important question is how particular learning goals arise. Both Leake and Ram discussed the role of failure in triggering learning, Leake concentrating on learning in response to anomalies during understanding, and Ram developing a more general taxonomy of failures that result from multiple types of tasks. In their view, this goal activation process can be characterized by the sequence:

Failure $\Longrightarrow$ Analysis of failure $\Longrightarrow$ Repair strategy $\Longrightarrow$ Knowledge required $\Longrightarrow$ Missing knowledge $\Longrightarrow$ Learning goal

Other panelists commented on other influences on the genesis of learning goals. Thagard pointed out that new goals and goal priorities may arise from reasoning about the relationships of existing goals. In his view, the decision-making task itself may involve balancing many goals. This process may shed new light on which goals apply or may change the learner's goals and goal priorities, changing the course of learning. Michalski distinguished between learning goals that are hard-wired and that arise from other influences in human learners. Ng pointed out that in human learners, culture may have significant effects on background learning goals; consequently, developing effective instructional materials depends on understanding and influencing background learning goals.

4.3 How do learning goals affect the learning process?

Focusing on storage of information during processing, Barsalou observed that most psychological theories assume that the storage of information is done unintentionally; a problem solver attempting to solve a problem simply stores a trace of its processing without attention to its future relevance. However, Ng's previously-mentioned studies show that for a different class of task, learning goals have a strong effect on the learning performance of human learners. A future question is to identify the limits of goal-driven processing in human learners.

From an AI perspective, other panelists focused on how connections are made between needs for information and relevant learning strategies. Ram pointed out that goals can affect many parts of the learning process such as focusing attention, controlling what is learned, and determining how to select and combine learning strategies (Ram & Cox, 1993). He stressed the importance of having an explicit representation of learning goals, to support the goal-driven learning process by allowing the reasoner to notice and avail itself of unexpected opportunities to learn something previously identified as important or interesting. He also advanced that the goal-driven learner's decision process must involve ways to reason about the relative priorities of pending goals, to select and combine learning strategies, and to suspend and opportunistically trigger learning goals when circumstances make them appropriate.

In the context of explanation, Leake identified broad classes of tasks, such as prediction, prevention and repair tasks, and connected them to requirements for information; those connections allow an explainer to judge whether candidate explanations provide the information needed for useful learning. In his model, requirements for filling system knowledge gaps also direct explanation generation, by guiding retrieval and revision of explanations during case-based explanation construction (Leake, 1992). In the context of analogical mapping, Thagard pointed out that goals, semantic constraints and syntactic constraints all affect both analogical mapping (Holyoak & Thagard, 1989) and retrieval of potential analogues (Thagard, Holyoak, Nelson, & Gochfeld, 1990).

Michalski described the Inferential Theory of Learning, a theoretical framework in which learning depends on input information, prior knowledge, and learning goals. In the theory, current learner knowledge is transformed into desired knowledge according to a set of transmutations for searching knowledge space, such as generalization, discrimination, and re-formulation of concepts (Michalski, 1993). As the bridge between learning goals and learning strategies, Michalski presented a taxonomy associating inference types with knowledge transmutations that serve those types of inferences.

4.4 How do different types of learning goals relate to each other?

Both Michalski and Ram advanced models that treat learning as a planful process. In their models, explicit reasoning about information needs and information requirements determines subgoals for learning activity. Thagard focused on the need to reason about the relationship between learning goals to determine which goals are relevant in a given situation: agents often simultaneously have a number of interrelated goals that must be balanced and reasoned about to determine which goals to keep, which to abandon, and the goal orderings that apply. He described current research on the coherence theory of decision, and on the DECO system which uses connectionist algorithms for parallel constraint satisfaction to both make decisions about actions and adjust goal priorities.

4.5 How are learning goals represented?

In order for a system to reason about its information needs, it must be able to represent what those needs are. Ram proposed representations including the desired knowledge (possibly partially-specified) and the reason that the knowledge is sought. Leake focussed on the representation of the knowledge required to resolve anomalies (which depends on a vocabulary of anomaly characterization structures to describe the information needed to resolve an anomaly), and on dimensions for representing the types of information that must be provided by explanations for different tasks. Michalski presented the goal-dependency network, a representation including the general goal being served, subordinate goals, attributes relevant to the goals, and the relationships connecting them. He also showed how a learning system using such a goal-dependency network creates very different conceptual classifications of input data depending on the top-level goal and associated subordinate goals.

5 Properties of Goal-Driven Learners

The properties of goal-driven learning described above suggest some basic design principles for goal-driven learning systems. First, because the systems reason about current goals to make decisions about which information to seek and how to seek it, learning in such systems is an *active* process, unlike the passive process of systems that accept the information they are given and apply to it a fixed learning procedure. Likewise, this view requires that goal-driven learners reason about how to acquire needed information, making learning a *planful* process (Hunter, 1990; Ram & Hunter, 1992).

Finally, in order to generate learning goals, goal-driven learning systems must be *introspective:* they must be able to notice gaps in their knowledge and to reason about the information needed to fill those gaps. Developing introspective systems requires that the system have a representation of its own processes, in order to detect deviations that show that learning is needed (e.g. (Ram & Cox, 1993; Krulwich et al., 1992)). Experimental results in the metacognition literature also suggest that introspective or metacognitive reasoning can facilitate human learning (e.g., (Schneider, 1985; Weinert, 1987)).

6 The Relationship between Goal-Driven and Non-Goal-Driven Learning

Goal-driven learning offers significant advantages over non-goal-driven methods, but it plays a complementary role to those methods; a number of panelists argued that both are needed for successful performance. Goal-driven learning is important because in complex domains, learners are faced with an overwhelmingly large set of alternatives that could be learned; this led to functional arguments in favor of goal-driven learning, advanced by Barsalou, Leake, Michalski and Ram, focusing on the importance of constraining the field of possible generalizations, thus restricting the effort of a learning system and allowing it to be applied more effectively (Leake, 1992; Michalski, 1993; Ram & Hunter, 1992).

However, non-goal-driven learning is needed because it is impossible to anticipate all future needs for information; learning exclusively in service of current goals might not take advantage of opportunities for low-cost learning. Thus controlled non-goal-driven learning can also be beneficial, provided that it can be done at sufficiently low cost. Along these lines, Barsalou observed that one form of human learning—storage of information—seems to be influenced by goals only indirectly, in that it maintains a trace of processing that may be goal-directed; both Barsalou and Leake stressed the importance of maintaining an accurate world model to support future goal-based activity, even though, in general, the task of maintaining the model is only indirectly related to goals that are active at the time of learning.

The tradeoff between goal-driven and non-goal-driven learning was also addressed by Thagard, who observed that if goals are the sole influence in deciding what to learn, to the exclusion of other semantic or pragmatic constraints, the objective accuracy of learning may be compromised. (An example of this sort of distortion is a student who is eager to avoid blame for a bad test score, and attributes the score to the teacher's unfairness, even though the true explanation is that the student had not studied hard enough. The explanation may serve the student's goals, but what is learned is nevertheless false.) The issue here is the balance between learning what is most useful to the learner's current tasks, and maintaining an accurate view of the world. Thagard discussed a model of analogical learning in which such influences are balanced.

7 Conclusions

Goal-driven learning systems learn in response to explicit goals for knowledge. Goal-driven learning allows flexibility of processing that is otherwise impossible in learning systems: a goal-driven learning system's choices of what to learn, when to learn, and which learning strategies to employ can be tailored towards achieving effective learning. In addition to these functional supports for the goal-driven learning process, psychological experiments support its validity as a cognitive model.

The symposium on goal-driven learning revealed the common ground and differences between disparate research efforts on goal-driven learning. It identified fundamental questions about the goal-driven learning process, and also pinpointed a number of important avenues for future research. Major issues include developing appropriate representations for learning goals, developing principles for resolving contradictions among competing goals, and developing theoretical principles and practical mechanisms for reflecting goals in the learning process.

The symposium is only a first step towards developing a paradigm whose ramifications on learning systems are likely to be considerable. Viewing learning as a goal-driven process suggests a new generation of active, planful and introspective learning systems that can learn effectively in complex situations by reasoning about when, what and how to learn.

8 Acknowledgments

The authors would like to thank Ryszard Michalski and Michael Cox for their helpful comments on a draft of this article.

References

Barsalou, L., 1991. Deriving Categories to Achieve Goals. In *The Psychology of Learning and Motivation: Advances in Research and Theory*, ed. G. Bower, 1–64, New York: Academic Press. 27.

Collins, G. and Birnbaum, L., 1988. An Explanation-Based Approach to the Transfer of Planning Knowledge Across Domains. In *Proceedings of the 1988 AAAI Spring Symposium on Explanation-based Learning,* 107–111. Stanford, CA: American Association for Artificial Intelligence.

desJardins, M., 1992. Goal-Directed Learning: A Decision-Theoretic Model for Deciding What to Learn Next. In *Proceedings of the Machine Discovery Workshop, Ninth International Machine Learning Conference*, 147–151. San Mateo: Morgan Kaufmann.

Dweck, C., 1985. Intrinsic motivation, perceived control, and self-evaluation maintenance: An achievement goal analysis. In *Research on Motivation in Education: The Classroom Milieu*, eds. C. Ames and R. Ames, 289–305. New York: Academic Press.

Hammond, K., 1989. *Case-Based Planning: Viewing Planning as a Memory Task*. San Diego: Academic Press.

Holyoak, K. & Thagard, P., 1989. Analogical Mapping by Constraint Satisfaction. *Cognitive Science*, *13*(3): 295–355.

Hunter, L., 1990. Planning to Learn. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, 261–268. Hillsdale, NJ: Lawrence Erlbaum.

Kedar-Cabelli, S., 1987. Formulating Concepts According to Purpose. In *Proceedings of the Sixth Annual National Conference on Artificial Intelligence*, 477–481. Menlo Park, CA: Morgan Kaufmann.

Keller, R., 1987. The Role of Explicit Contextual Knowledge in Learning Concepts to Improve Performance. Ph.D. diss, Dept. of Computer Science, Rutgers University.

Krulwich, B., Birnbaum, L., & Collins, G., 1992. Learning Several Lessons from one Experience. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, 242–247. Hillsdale, NJ: Lawrence Erlbaum.

Leake, D., 1992. *Evaluating Explanations: A Content Theory*. Hillsdale, NJ: Lawrence Erlbaum.

Michalski, R., 1983. A Theory and Methodology of Inductive Learning. In *Machine Learning: An Artificial Intelligence Approach*, eds. R. Michalski, J. Carbonell, and T. Mitchell, 83-134. Cambridge, MA: Tioga Publishing Company.

Michalski, R., 1993. Inferential Theory of Learning: Developing Foundations for Multistrategy Learning. In *Machine Learning: A Multistrategy Approach*, eds. R. Michalski and G. Tecuci. San Mateo, CA: Morgan Kaufmann. Forthcoming.

Minton, S. 1988. *Learning Search Control Knowledge: An Explanation-Based Approach.* Boston: Kluwer Academic Publishers.

Mitchell, T., Keller, R., & Kedar-Cabelli, S., 1986. Explanation-Based Generalization: A Unifying View. *Machine Learning*, *1*(1): 47–80.

Ng, E. & Bereiter, C., 1991. Three levels of goal orientation in learning, *Journal of the Learning Sciences*, *1*(3 and 4): 243–271.

Ram, A., 1991. A Theory of Questions and Question Asking. *The Journal of the Learning Sciences*, *1*(3 and 4): 273–318.

Ram, A. & Cox, M., 1993. Introspective Reasoning using Meta-Explanations for Multistrategy Learning. In *Machine Learning: A Multistrategy Approach*, eds. R. Michalski and G. Tecuci. San Mateo, CA: Morgan Kaufmann. Forthcoming.

Ram, A. & Hunter, L., 1992. The Use of Explicit Goals for Knowledge to Guide Inference and Learning. *Applied Intelligence*, *2*(1): 47–73.

Riesbeck, C., 1981. Failure-driven Reminding for Incremental Learning. In *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, 115–120 Vancouver, B.C: International Joint Conferences on Artificial Intelligence.

Scardamalia, M. & Bereiter, C., 1991. Higher Levels of Agency for Children in Knowledge Building: A Challenge for the Design of New Knowledge Media. *The Journal of the Learning Sciences*, *1*(1): 37–86.

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

Schank, R. & Leake, D., 1989. Creativity and Learning in a Case-Based Explainer. *Artificial Intelligence*, *40*(1-3): 353–385.

Schneider, W., 1985. Developmental Trends in the Metamemory-Memory Behavior Relationship: An Integrative Review. In *Metacognition, Cognition, and Human Performance*, eds. D. Forrest-Pressley, G. MacKinnon, and T. Waller, 57-109. New York: Academic Press. 1.

Sussman, G., 1975. *A Computer Model of Skill Acquisition.* New York: American Elsevier.

Thagard, P., Holyoak, K. J., Nelson, G., & Gochfeld, D., 1990. Analog Retrieval by Constraint Satisfaction. *Artificial Intelligence*, *46*(3), 259–310.

Utgoff, P. E., 1986. Shift of Bias for Inductive Concept Learning. In *Machine Learning: An Artificial Intelligence Approach: Volume II*, eds R. Michalski, J. Carbonell, and T. Mitchell, 107–148. San Mateo: Morgan Kaufmann.

Weinert, F., 1987. Introduction and Overview: Metacognition and Motivation as Determinants of Effective Learning and Understanding. In *Metacognition, Motivation, and Understanding,* eds. Weinert, F. & Kluwe, R., 1–16. New York: Lawrence Erlbaum.

Winston, P., 1975. Learning Structural Descriptions from Examples. In *The Psychology of Computer Vision*, ed. Winston, P., 157-209. New York: McGraw Hill.

Zukier, H., 1986. The Paradigmatic and Narrative Modes in Goal-Guided Inference. In *Handbook of Motivation and Cognition: Foundations of Social Behavior*, eds. Sorrentino, R. & Higgins, E. T., chap. 16, 465–502. New York: Guilford Press.

David Leake is an Assistant Professor of Computer Science and a member of the Cognitive Science faculty at Indiana University. He received his Ph.D. from Yale University in 1990. His primary research interests are case-based reasoning and explanation, which are the focus of his book *Evaluating Explanations: A Content Theory* (Erlbaum, 1992).

Ashwin Ram is an Assistant Professor in the College of Computing of the Georgia Institute of Technology, and an Adjunct Professor in the School of Psychology. He received his Ph.D. from Yale University in 1989. His research interests lie in the areas of machine learning, natural language understanding, explanation, and cognitive science.

Ashwin Ram and David Leake are editors of the forthcoming book *Goal-Driven Learning* (MIT Press).