Proceedings of the 1994 AAAI Spring Symposium on Goal-Driven Learning, pp. 1-11

A Framework for Goal-Driven Learning[1]

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

David Leake
Computer Science Department
Indiana University
Bloomington, Indiana 47405-4101
E-mail: leake@cs.indiana.edu

Why goals?

In cognitive science, artificial intelligence, psychology, and education, a growing body of research supports the view that learning is largely a goal-directed process. Experimental studies show that people with different goals process information differently; work in machine learning presents functional arguments for goal-based focusing of learner effort. Recent work in these fields has focussed on issues of how learning goals arise, how they affect learner decisions of when and what to learn, and how they guide the learning process. It is increasingly evident that investigation of goal-driven learning can benefit from bringing these perspectives together in a multidisciplinary effort (Leake & Ram, 1993).

The central idea underlying goal-driven learning is that because the value of learning depends on how well the learning contributes to achieving the learner's goals, the learning process should be guided by reasoning about the information that is needed to serve those goals. The effectiveness of goal-driven learning depends on being able to make good decisions about when and what to learn, on selecting appropriate strategies for achieving the desired learning, and on guiding the application of the chosen strategies. Research into such topics includes the development of computational models for goal-driven learning, the testing of those models through psychological experiments and empirical experiments with computer programs, the justification of the models through functional arguments about the role and utility of goals in learning, and the use of models of goal-driven learning in guiding the design of educational environments. The common themes in these research efforts are the investigation of types of learning goals, the origins of learning goals, and the role of goals in the learning process.

Research on goal-driven learning in artificial intelligence has been motivated largely by computational arguments. The problem of combinatorial explosion of inferences is well known; in any realistic task domain, time and resource constraints prohibit consideration of all but a few of the possible inferential paths. Consequently, any reasoner, human or machine, must focus its attention and resources on pursuing those inferential paths that are likely to be most useful. Similarly, in any realistic situation, there are several different types of learning that a reasoner might perform, several kinds of new knowledge that a reasoner might acquire, and several kinds of reformulation or reorganization of existing knowledge that a reasoner might carry out. Again, due to time and resource constraints it is only practical to perform a few of these operations. Consequently, the reasoner must focus its attention and resources on executing the learning operations that are likely to be most useful. Because the utility of an inference or a piece of knowledge can best be evaluated relative to a particular task or goal, goal-based considerations must guide reasoning and learning.

In addition to these computational arguments for goal-driven learning, research in goal-driven learning has a cognitive basis in psychological research. This research has established much evidence for the influence of goals on human learning, and for the use of active, strategic, and goal-driven processes in many kinds of learning that humans perform. However, many questions remain concerning the kinds of goals that people pursue, the conditions under which those goals influence learning, and the kinds of learning that are influenced by those goals.

Research in cognitive science combines the cognitive perspective of psychology with the computational perspective of artificial intelligence, developing computational models of human learning that are evaluated using computational metrics as well as by comparison with human performance. Research in education has also been concerned with psychological data about human learning, but from a pragmatic perspective. This research has attempted to use empirical evidence to guide the design of instructional and educational scenarios so as to facilitate learning, taking as its starting point the evidence for facilitation of certain kinds of learning by particular kinds of goals. These scenarios have also been used as the basis for further psychological experimentation to validate the underlying theories. In this paper, we describe a framework for goal-driven learning and its relationship to prior and current theories from each of these perspectives.

An everyday example

Goal-driven learning is triggered when a reasoner needs to learn in order to improve its performance at some task. A goal-driven learner determines *what* to learn by reasoning about the information it needs, and determines *how* to learn by reasoning about the relative

[1] This is an extended abstract of a chapter to appear in A. Ram and D. Leake (eds.), *Goal-Driven Learning*, MIT Press/Bradford Books.

merit of alternative learning strategies in the current circumstances. For example, for a first-time stereo buyer, the goal of getting good buy on a stereo may give rise to at least two learning goals: a goal to learn the best sources for sound equipment and to a goal to learn how to judge the merits of competing equipment. Each of these learning goals may trigger learning subgoals. In order to learn the best place to buy sound equipment, the buyer may first have to learn general criteria for what constitutes a good store for buying sound equipment, and then specifics about prices, service, etc. to classify different stores. In order to learn how to judge particular equipment, the buyer will have to learn about the classes of alternatives available and about specific equipment within those classes. Thus some learning goals involve gathering information in the external world, while others involve reformulating or changing information that is already known, by operations such as forming generalizations or reorganizing memory.

In order to perform the desired learning, the stereo buyer must select strategies for accomplishing each of its learning goals. For example, the buyer may choose between learning strategies including asking others' opinions, reading magazine articles, forming inductive or explanation-based generalizations from demonstrations of equipment, or even disassembling equipment to determine the quality of its electronic components. Learning strategy selection depends on factors such as the buyer's prior knowledge, the buyer's resources (e.g., how much time the buyer can spend on the shopping process), opportunities (e.g., happening to meet an expert on sound equipment at a party), and the buyer's own abilities (e.g., whether the buyer has the expertise to judge the quality of equipment by disassembling it).

This example illustrates the value of goal-driven learning in focusing learner effort, and also suggests the range of roles that goals can play in influencing learning. Goals determine how much effort to allocate to performance tasks (e.g., the task of buying a stereo), indirectly influencing the resources available for the learning that will be performed as part of that task. Goals also determine the focus of attention when new information is received as input (e.g., focusing attention on announcements of stereo sales). They determine what should be learned (e.g., determining that it is worthwhile to generalize about relationships between store types and prices). They give criteria for evaluating the results of learning and deciding what learned information to store (in this example, the value of learning is its usefulness for guiding the shopping decision). Table 1 summarizes these and other possible roles of goals in learning.

Guiding the performance task by:

- Determining the resources made available to the performance task
- Guiding the control or search procedure used in the performance task
- Guiding retrieval of plans, problem solutions, and other types of knowledge
- Focusing attention on certain aspects of the input
- Guiding the evaluation of the outcome of the performance task

Guiding the learning task, by:

- Specifying the target of learning (desired output of a learning algorithm)
- Selecting the learning algorithms to be used
- Constraining the learning process (for example, influencing the policies under which the learning algorithms operate)
- Focussing the search for information needed to carry out the learning
- Determining when learning should be attempted
- Aiding evaluation of results of learning with respect to the desired output

Guiding storage, by:

- Selecting what to store
- Determining how learned knowledge is indexed

Table 1: Ways in which goals can influence learning.

Towards a planful model of learning

As the previous example illustrates, a goal-driven learner makes decisions about what, how, and when to learn in order to further its goals. Consequently, its learning can be considered a "planful" process (e.g., Hunter, 1990; Leake, to appear; Michalski & Ram, to appear; Pryor & Collins, 1992; Quilici, this volume; Ram & Cox, 1994; Ram, Cox, & Narayanan, to appear; Ram & Hunter, 1992; Redmond, 1992; Schank & Abelson, 1977; Xia and Yeung, 1988). This learning process is analogous to models of problem solving in which the reasoner uses task goals to formulate action plans for achieving these goals (e.g., Newell & Simon, 1972; Greeno & Simon, 1988; VanLehn, 1989). Learning actions or schemas are selected, combined, and invoked appropriately on the basis of existing learning goals and available environmental opportunities for learning. Learning is a behavior explicitly carried out to seek information, driven by needs arising from the reasoner's performance on a task that learning is intended to facilitate, and mediated by the formulation and manipulation of explicit learning goals.

The motivation for the goal-driven approach is to control processing in a rich world. Simply put, knowledge that is valid in principle need not necessarily be useful (Mitchell & Keller, 1983); thus, it is desirable to avoid the effort involved in learning knowledge that does not contribute to the reasoner's overall purpose. More specifically, some of the motivations for goal-based approaches include (see also Cox & Ram, this volume):

- **Alleviating problems of computational complexity:** The ability of a reasoner to make decisions about its reasoning and learning processes helps to alleviate problems caused by the computational complexity of reasoning in an open world, by enabling the reasoner to focus its efforts towards processing that serves its goals (Cox, 1993; Hunter, 1990; Leake, 1992; Leake, this volume; Ram & Hunter, 1992). An analysis of the utility of learning can help in determining the target of learning (desJardins, 1992), in guiding learning processes (Gratch, DeJong, & Chien, this volume; Provost, this volume), and also in deciding whether to learn at all (Markovitch & Scott, 1993; Minton, 1990).
- **Facilitating the use of opportunities to learn:** If a reasoner does not have sufficient resources at the time it realizes it has a need to learn, or if the requisite knowledge is not available at that time, the reasoner can suspend its learning goals in memory so that they can be retrieved and pursued at a later time (Hunter, 1990; Hammond, Converse, Marks, & Seifert, 1993; Ram, 1991, 1993; Ram, Cox, & Narayanan, to appear; Ram & Hunter, 1992).
- **Improving the global effectiveness of learning:** Taking goal priorities and goal dependencies into account when deciding what to learn and how to coordinate multiple learning strategies improves the effectiveness of learning in a system with multiple goals. Learning strategies, represented as methods for achieving learning goals, can be chained, composed, and optimized, resulting in learning plans that are created dynamically and pursued in a flexible manner (Cox, 1993; Cox & Ram, this volume; Gratch, DeJong, & Chien, this volume; Hadzikadic & Yun, 1988; Hunter, 1990; Michalski, 1993; Michalski & Ram, to appear; Ram & Hunter, 1992; Redmond, 1992; Stroulia & Goel, this volume).
- **Increasing the flexibility of learning:** In situations involving multiple reasoning failures, multiple active and suspended learning goals, multiple applicable learning strategies, and limited resources, direct mapping from specific types of failures to individual learning strategies is impossible, and an active, planful approach becomes necessary. For a given failure, there may be more than one algorithm which needs to be applied for successful learning and, conversely, a given algorithm may apply to many different types of failures (Cox, 1993; Cox & Ram, this volume; Ram, Cox, & Narayanan, to appear; Krulwich, Birnbaum, & Collins, 1992). A planful model of learning allows decoupling of many-to-many relationships, leading to more flexible behavior (Cox, 1993; Cox & Ram, this volume).
- **Improving management of interactions between learning processes:** Explicit formulation of learning goals facilitates detection of dependency relationships, so that goal violations can be avoided (Cox, 1993; Cox & Ram, this volume). When multiple items are learned from a single episode, the changes resulting from one learning algorithm may affect the knowledge structures used by another algorithm. Such dependencies destroy any implicit assumption of independence built into a given learning algorithm that is used in isolation. For example, one learning algorithm may split a concept definition into separate schemas or otherwise modify the definition. Therefore, an indexing algorithm that uses the attributes of concepts to create indices must necessarily follow the execution of any algorithm that changes the conceptual definition.

Psychological evidence also supports the existence of goal-based influences on human focus of attention, inference, and learning (e.g., Barsalou, 1991; Faries & Reiser, 1988; Hoffman, Mischel, & Mazze, 1981; Ng & Bereiter, 1991; Seifert 1988; Srull & Wyer, 1986; Wisniewski & Medin, 1991; Zukier, 1986; see also discussion by Hunter, 1990). These ideas are related to the "goal satisfaction principle" of Hayes-Roth and Lesser (1976), which states that more processing should be given to knowledge sources whose responses are most likely to satisfy processing goals, and to the "relevance principle" of Sperber and Wilson (1986), which states that humans pay attention only to information that seems relevant to them. Those principles make sense because cognitive processes are geared to achieving a large cognitive effect for a small effort. To achieve this, the understander must focus its attention on what seems to it to be the most relevant information available.

Goals can facilitate learning even when they are not generated internally by the reasoner; for example, Steinbart (1992) shows that asking users questions (i.e., "creating" knowledge goals in people) can help them learn from a computer-assisted training program, and Patalano, Seifert, and Hammond (1993) show that presenting users with a goal and a plan to achieve it can facilitate later detection of relevant features of a situation. There is also much research on the origins of goals; for example, Graesser, Person, and Huber (1992) discuss several types of questions, or goals to seek information, and the cognitive mechanisms that generate them.

A significant body of psychological research points to the influence of "metacognition"—cognition by a person concerning that person's own cognitive processes—in human performance (e.g., Forrest-Pressley, MacKinnon & Waller, 1985; Weinert, 1987; Wellman, 1985, 1992). Gavelek and Raphael (1985) discuss a form of metacognition, called metacomprehension, which addresses the abilities of individuals to adjust their cognitive activity in order to promote more effective comprehension, in particular, the manner in which questions generated by sources external to the learner (i.e., from the teacher or text), as well as those questions generated by the learners themselves, serve to promote their comprehension of text. White and Gunstone (1989) argue that resolution of conflicting beliefs and permanent conceptual change requires "metalearning"—control over one's learning. For example, they discuss a study by Gauld (1986) that shows that students who learn new scientific beliefs often revert to their original beliefs over time because they have merely accepted the new knowledge without any real commitment to it. They argue that deep reflection on one's beliefs is a key part of the awareness and control over one's learning, and suggest methods for promoting metalearning in science classrooms.

The goal-driven learning framework does not imply that all processing is explicitly goal-driven. A reasoner that was completely goal-driven would only notice what it was looking for already; it would not be able to respond to and learn from unexpected input. Instead, it is reasonable to assume that there would be some automatic, bottom-up, or non-goal-driven processing during reasoning and learning, which would support strategic, top-down, or goal-driven processes such as those discussed here (e.g., Barsalou, to appear; Kintsch, 1988; Leake, 1992; McKoon & Ratcliff, 1992; Ram, 1991). It is clear, of course, that humans cannot exert explicit metacontrol over all their learning processes, and the level of control that can be exerted, as well as how it is exerted, remain open questions. It is also possible (though, in our opinion, unlikely) that it may turn out not to be efficient to use this framework as a technological basis for the design of computer programs that learn. Nevertheless, the framework presented here may be used to take an intentional stance (Dennett, 1987) towards a reasoner for the purposes of building a computational model of learning. In such a stance, the competence of the reasoner can be modeled using goals, learning decisions, learning actions, and so forth as the basic theoretical constructs in task-level and algorithm-level descriptions of the reasoner. That stance can be taken without any commitment to existence of these constructs at the implementational level of, say, neural representations and processes in the human brain, or to the degree of conscious self-awareness of these processes in human thought.

A framework for goal-driven learning

In order to form a unified view of the diverse research results on goal-driven learning, we propose a general framework for describing the goal-driven learning process. While no single piece of research to date has investigated this framework as a whole or exactly as stated, the framework serves to provide an integrative structure into which individual research efforts fit as pieces of the puzzle of goal-driven learning. The key idea behind our framework is to model learning as an *active* (explicitly goal-driven) and *strategic* (rational and deliberative) process in which a *reasoner*, human or machine, explicitly identifies its goals in learning and attempts to learn by determining and pursuing appropriate learning actions via explicit reasoning about its goals, its abilities, and environmental opportunities.

In this framework, learning is motivated by the *performance tasks* that the reasoner is attempting to perform in the world. The performance tasks give rise to *task goals*, as well as subgoals of those goals, and subtasks to achieve them. As the tasks and subtasks are performed, the reasoner formulates explicit *learning goals* to perform types of learning which, if successful, would improve its ability to carry out those performance tasks or subtasks. The learning goals, in turn, guide the learning behavior of the reasoner, leading it to focus attention, allocate resources, and select appropriate *learning algorithms* or *learning strategies* when opportunities to learn arise. In our previous example, the top-level task goal would be to get a good buy on a stereo, which would spawn subtasks such as going to a store and purchasing the stereo. These subtasks give rise to learning goals to learn information needed to select the store and the stereo to buy. Some of those learning goals may seek to gather information about the external world, while others may seek to create generalizations, test hypotheses, reorganize memory, or otherwise change existing knowledge. Those learning goals prompt the choice of learning strategies such as "shopping around," looking at reviews in magazines, and so forth.

The goal-driven learning process involves not only learning about the world, but also learning to improve the reasoner's own reasoning process. In order to identify the learning that needs to occur, the reasoner needs to be able to analyze its reasoning process in addition to the knowledge that the reasoner invokes during the reasoning process. To facilitate this, the reasoner maintains a *reasoning trace* of its internal decision-making. The reasoning trace provides the basis for *introspective reasoning* or *meta-reasoning* to guide learning and improve its reasoning performance.

More concretely, goal-driven learning can be modeled as a two-step process. The first step involves the generation of learning goals based on the performance tasks and task goals of the reasoner. This step can be thought of as the process of deciding what to learn, and results in the formulation of learning goals that specify the desired learning that is to occur as well as the origin of the need for this learning. The second step involves the pursuit of learning goals based on the reasoner's needs, its resources, and on environmental factors that determine the timeliness of pursuing certain learning actions in a given situation. This step can be thought of as the process of deciding how and when to learn and carrying out the learning. When the learning actually occurs,

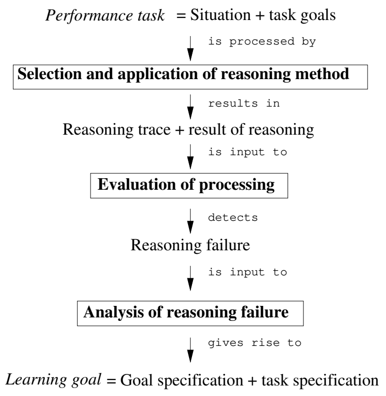

Figure 1: Generation of learning goals

this step results in the satisfaction of one or more of the reasoner's learning goals.

Step 1: Generating learning goals: Figure 1 describes the process by which learning goals are generated. The reasoner is assumed to be pursuing a performance task that can be characterized in terms of the current situation and task goals specifying the desired result of the task. In the stereo example, the situation might be that the shopper lives in New York, knows nothing about stereos, and has $500 to spend; the task goal would be to buy a stereo that was a good value for the $500 price range.

Given the performance task, the reasoner performs reasoning in support of that task and maintains a trace reflecting its reasoning process. The reasoning trace records the goal-subgoal decompositional structure of the task goals, the choice of methods for achieving them and other decisions taken, the factors influencing those decisions, and descriptions of other reasoning actions (e.g., attempts to retrieve information) and their outcomes (Carbonell, 1986; Ram & Cox, 1994). For example, forming an executable plan to get a good buy on a stereo requires knowing which stereo to buy and where to buy it. If the reasoner does not know, a reasoning failure occurs because current knowledge is insufficient to make a decision.

At a suitable point in processing, the reasoning trace and its results are evaluated in light of the reasoner's task goals. If any problems arose during processing, learning is needed to enable the reasoner to avoid similar problems in the future. In being driven by deficiencies in the reasoner's knowledge, the process for generating learning goals is in the spirit of impasse-driven or failure-driven learning (e.g., Chien, 1989; Collins & Birnbaum, 1988; Hammond, 1989; Kocabas, this volume; Krulwich, this volume; Laird, Newell, & Rosenbloom, 1986; Mooney & Ourston, 1993; Mostow & Bhatnagar, 1987; Newell, 1990; Owens, 1991; Park & Wilkins, 1990; Ram & Cox, 1994; Riesbeck, 1981; Schank, 1982; Schank & Leake, 1989; Sussman, 1975; VanLehn, 1991a). There are several kinds of failures that may be involved, for example, expectation failures, retrieval failures, and knowledge application failures. In our framework, an unexpected success is also treated as an expectation "failure". Ram, Cox, and Narayanan (to appear) present a taxonomy of possible types of failures and discuss their relationship to goal-driven learning.

Even if no failure has yet occurred, anticipation of a reasoning failure may trigger learning. For example, a reasoner may realize that it cannot perform a task and decide to perform the necessary learning before even attempting the task. In our framework, all these motivations for learning—reasoning failures, difficulties, impasses, suboptimalities, surprises, and other types of processing problems or anticipated processing problems—will be collectively and simply referred to as *failures.*

Different kinds of failures give rise to different kinds of learning goals. For example, a reasoner may need to acquire additional knowledge if its reasoning reached an impasse due to missing knowledge, as in the case of a novice stereo buyer who has no knowledge of which brand of stereo to buy. If the reasoner possessed sufficient knowledge but did not retrieve it at an appropriate time, it may need to reorganize memory (Ram & Cox, 1994; Ram, Cox, & Narayanan, to appear). A reasoner may need to modify the underlying representational vocabulary if its vocabulary is found to be inadequate (e.g., Schlimmer, 1987; Wrobel, 1988). In some situations, a reasoner might also need to add to its repertoire of reasoning strategies (e.g., Leake, 1993).

When a reasoning failure is detected, the reasoning trace is analyzed, in a process called credit/blame assignment, to find the source of the failure (Freed & Collins, this volume; Hammond, 1989; Krulwich, this volume; Minsky, 1963; Ram & Cox, 1994; Weintraub, 1991). Blame assignment may be thought of as a process of model-based diagnosis of the reasoner itself (Birnbaum, Collins, Freed & Krulwich, 1990; Stroulia, Shankar, Goel, & Penberthy, 1992). If the failure is attributed to faulty knowledge, learning is needed to improve the reasoner's performance, and a learning goal is generated to repair that knowledge. In our framework, the learning goal is characterized in terms of two pieces of information: The desired learning—*what* learning is needed—and a description of the task that motivates learning—*why* learning is needed. The additional information about why learning is needed is important to allow the reasoner to carry out its tasks in an opportunistic manner, with learning goals (and the tasks that they support) being suspended until circumstances are favorable to their pursuit (Ram, 1991, 1993; Ram & Hunter, 1992).

Step 2: Pursuing learning goals: In the goal-driven view of learning, learning goals are treated analogously to task goals in the world. Just as task goals are achieved through a planning process using available methods for reasoning and action, learning goals are achieved through a *knowledge planning* process using available learning methods or strategies (Hunter, 1990; Quilici, this volume; Ram & Hunter, 1992; Redmond,

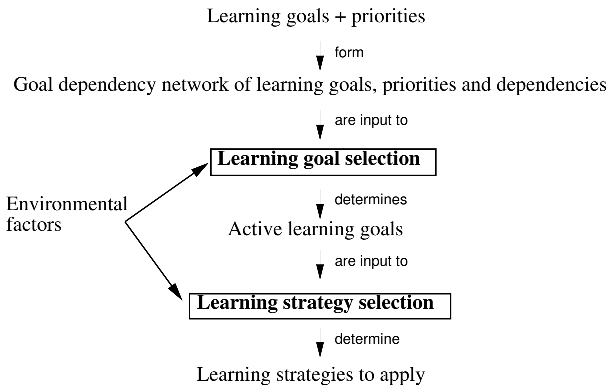

Figure 2: Pursuing learning goals using appropriate learning strategies

1992). In the knowledge planning process, explicit reasoning is done about learning goals, their relative priorities, and strategies by which they can be achieved. These learning goals, also called *knowledge goals* (Ram, 1987, 1990; Ram & Hunter, 1992), can be represented in a goal dependency network (Michalski, 1993; Michalski & Ram, to appear), which is used to select and combine learning actions into learning strategies that are appropriate for current learning goals and for the learning opportunities provided by the current environment.

Individual learning actions may include performing knowledge acquisition (e.g., asking a friend to recommend a stereo) knowledge reorganization (e.g., grouping stores by the size of their stereo departments), knowledge reformulation or transmutation (e.g., forming new generalizations from stored episodes concerning others' experiences with particular sound equipment), and so on. Their application is guided by the learning goals of the reasoner (Gratch, DeJong, & Chien, this volume; Hunter, 1990; Michalski & Ram, to appear; Pryor & Collins, 1992; Ram & Cox, 1994; Ram & Hunter, 1992). Figure 2 sketches the second step of the goal-driven learning process. This step begins with reasoning about the relationships and relative priorities of learning goals in order to form a goal dependency network. Based on the information contained in the goal dependency network and on environmental factors affecting the appropriateness of different goals, the reasoner selects the learning goals to pursue. Learning strategies for achieving those goals in the current environment are then selected and applied.

Perspective on the framework: The model of learning embodied in the above steps contrasts with the approach to learning taken in traditional machine learning systems in artificial intelligence. Typically, in those systems, learning is primarily a passive, data-driven process of applying a single learning algorithm (or a predetermined combination of a few learning algorithms) to training examples presented to the system. Goal-driven learning, in contrast, is an active and strategic process driven by reasoning about information needs, alternative learning strategies, and opportunities in the environment. In our framework, the process of determining what to learn is an integral part of the computational model of learning, as is the process of deciding (on a dynamic basis) how and when to learn it.

Our view of goal-driven learning implies a tightly coupled relationship between learning and the "rest of reasoning." This view is consistent with recent models of intelligence that are framed as *integrated intelligent architectures,* in which the knowledge and reasoning tasks underlying learning and performance are integrated into a complete interacting system (see Laird, 1991, and VanLehn, 1991b, for examples of this work). A common theme in this research, and one that is compatible with the goal-driven learning framework proposed here, is the explicit representation of task goals, reasoning goals, and learning goals, and their role in a multistrategy reasoning process that integrates learning with performance tasks such as problem solving or comprehension (e.g., see Ram, Cox, & Narayanan, to appear).

Types of goals

In order to understand how goals can relate to one another and to learning, it is useful to consider the classes of goals that influence learning. As Barsalou (to appear) observes, in some sense any reasoner executing a built-in procedure can be viewed as having a "goal" to perform that type of processing, so that any learner could be considered trivially "goal-driven". To distinguish between built-in behaviors and behaviors that are more explicitly goal-driven, Barsalou differentiates between implicit background orientations and explicit problem solving or task goals. Explicit task goals are the goals that guide a problem solving process in which a person intends to achieve a set of goals, assesses what must be performed to achieve them, and executes the needed actions. In contrast, an implicit background orientation is a behavior that is performed without explicit reasoning about when and how it should be pursued. For example, one such implicit orientation is the orientation to constantly maintain a world model that adequately represents the reasoner's environment (e.g., Barsalou, to appear; Leake, 1989, 1992), although in some formalisms this is expressed in terms of an explicit goal (see, e.g., Van de Velde, 1988).

Explicit goals are traditionally expressed as specifications of a target or desired outcome of a problem solving or learning task (e.g., Fikes, Hart, & Nilsson, 1972; Newell & Simon, 1972). However, Ram and Hunter (1992; Hunter, 1990; Ram & Cox, 1994) argue that capturing the introspective nature of the goal-driven learning process requires a richer characterization in which a goal is not merely a specification of a target. They argue that a target specification or an orientation is a goal only if the reasoner can actively plan to accomplish the goal, can make decisions about it, and can even decide to suspend it or not to pursue it.

In order for the reasoner to make such decisions, goals must be explicitly represented, and the reasoner must be able to reflect on its goals, how to achieve them, and their relative priorities and interdependencies. Ram and Hunter (1992) discuss a representation of learning goals in terms of the desired knowledge to be learned as well as the reason that the knowledge is needed. Additional representational issues concern the kinds of decision-making relationships that goals can enter into (Thagard & Millgram, to appear) and the intergoal relationships and interdependencies in which goals can play a role (Cox & Ram, this volume; Michalski & Ram, to appear; Schank and Abelson, 1977; Slade, 1993; Wilensky 1983).

In our framework, the "goals" in goal-driven learning research can be divided into three classes: *task goals*, *learning goals*, and *policies*. Broadly, task goals determine *why* the reasoner is learning in the first place, learning goals specify *what* the reasoner needs to learn, and policies influence *how* learning occurs. Task goals, exemplified by early planning programs (e.g., Fikes, Hart, & Nilsson, 1972; Sacerdoti, 1977) are specifications of desired outcomes from a performance task in the external world, which are explicitly pursued through planful reasoning processes or, in some recent models, goal-directed reactive processes (e.g., Earl & Firby, this volume; Freed & Collins, this volume; Maes, 1990). Because task goals characterize a desired state of affairs, they can also be used to describe the need for information that a planner requires to achieve that state of affairs (e.g., Leake, 1991, Ram & Leake, 1991), to understand interactions between task goals (e.g., Freed & Collins, this volume), and to influence or bias learning strategies (e.g., Martin, this volume). In some models, task goals (and resulting learning goals) are decomposed into subgoals or task structures to facilitate planning and learning (e.g., Karlsson, this volume; Stroulia & Goel, this volume).

Other computational models explicitly describe goals for learning, rather than implicitly characterizing it in terms of the external task. These learning goals differ from task goals in that, while they too specify a desired state, the specified state is an internal or mental state—a state of knowledge or belief that the learner is attempting to achieve. Task goals are satisfied through problem solving in the external (usually physical) world, while learning goals are satisfied through a learning process that, in the goal-driven learning framework, is viewed as problem solving in the "informational" world. These learning goals have been characterized in different ways, including as knowledge goals, knowledge acquisition goals, knowledge-building goals, questions, learning goals, and delta-knowledge goals (e.g., Cox & Ram, this volume; desJardins, 1992; Hunter, 1990; Michalski, 1993; Ng & Bereiter, 1991; Oehlmann, Sleeman, & Edwards, 1992; Quilici, this volume; Ram, 1987, 1990, 1991; Ram & Cox, 1994; Ram & Hunter, 1992; Schank & Abelson, 1977). In our framework, learning goals, in addition to specifying the desired outcome of learning, specify the reason that the desired learning is required (e.g., "task specifications" (Ram, 1991; Ram & Hunter, 1992) specify the suspended reasoning task that is awaiting the knowledge to be learned).

Finally, several computational models reflect other types of influences and constraints on learning that are goal-related. Although these are not "goals" in the sense of driving the learning process in an explicit manner, they may play an important role in influencing that process. Such influences include goal concepts, target concepts, purposes, operationality criteria, bias, policies, quality metrics, and utility metrics (desJardins, 1992; Gordon & Perlis, 1989; Gratch, DeJong, & Chien, this volume; Kedar-Cabelli, 1987; Keller, 1988; Laird, Rosenbloom, & Newell, 1986; Leake, 1991, 1992; Markovitch & Scott, 1993; Martin, this volume; Michalski, 1983; Minton, 1990; Mitchell, 1982; Mitchell, Keller, & Kedar-Cabelli, 1986; Perez, this volume; Provost, this volume; Utgoff, 1986). Policies and constraints are not learning goals in the sense that the learner does not actively seek to satisfy them; instead, they influence the learning processes that the learner uses to achieve its learning goals. In particular, they describe the policies under which the learning task should operate in order to better achieve the overarching learning goals, and describe relevant constraints on the processes that carry out the learning task. Note that a learner might formulate explicit learning goals to learn these criteria. For example, a learner might formulate an explicit goal to learn appropriate biases for a given type of learning situation, and pursue an explicit learning agenda to learn such biases.

The underlying commonality among these constructs is that each reflects an intention to influence learning according to needs that are external to the learning process itself. However, quite different focuses are apparent in the formulations described in the previous sections. Consequently, developing a general theory of goal-driven learning depends on analyzing the relationships of these constructs and their role in reasoning and learning.

To relate the previous perspectives, we refer to the general class of *goals* to describe theoretical constructs that refer to mental entities that are explicitly represented and actively pursued through a planful reasoning process.[2] *Task goals* refer to goals which specify desired effects in the world external to the reasoner. *Learning goals* or *knowledge goals* refer to goals which specify desired effects within the reasoner such as acquiring new knowledge or augmenting, reorganizing, or reformulating existing knowledge. Learning goals describe not only the desired processing outcome, but how the desired knowledge will be used when it is acquired. *Reasoning goals* refer to more general internal goals to form conclusions or inferences through learning or other reasoning processes. *Target concepts* specify a desired concept to be learned, but not necessarily learned through a goal-driven learning process; and general *policies* or *orientations* influence learning without being explicitly represented or available for manipulation by the reasoner's reasoning or learning process, including *constraints* on the formulation of hypotheses such as biases, operationality criteria, and utility metrics. Table 2 summarizes these distinctions, and

[2] Note that this definition does not imply that goals or goal-driven processing must necessarily be conscious, nor that the reasoner must necessarily be able to report externally about this processing.

	Explicitly represented?	Range of effects (internal to reasoner or in external world)	Influences selection of solution algorithm?	Solution process	Effect on solution generation
Goals	Yes	Either	Yes	Planning	Guidance
Task goals	Yes	External	Yes	Planning actions in external world	Guidance
Reasoning goals/ learning goals	Yes	Internal	Yes	Knowledge planning	Guidance
Policies	Sometimes	Internal	Sometimes	Unspecified	Constraint
Target concepts	Yes	Internal	No	Unspecified	Guidance
Operationality criteria	Yes	Internal	No	Unspecified	Constraint

Table 2: Types of goals and policies.

Table 1 summarizes the different roles that such constructs play in learning.

Note, however, that these classes of goals can overlap and influence each other. Task goals have been used to guide learning and performance in several systems, and can also be used to formulate learning goals to acquire information necessary for a given task (Ram & Leake, 1991) or to come to a better understanding of the task itself (Freed & Collins, this volume). In conjunction with knowledge or theories, they can guide learning processes (Barsalou, 1991; Ng & Bereiter, 1991; Wisniewski & Medin, 1991). Likewise, although target concepts are generally provided to a learning system as input by a human user, in some models target concepts are generated from aspects of the performance task in a manner similar to the generation of learning goals. For example, Kedar-Cabelli (1987) discusses a method for generating target concepts from standard constraints on artifacts to be used in particular plans. Keller (1987) also sketches a process for generating learning goals from higher-level performance objectives. Similarly, policies (such as bias, which is usually formulated as a passive, background constraint on learning) may be actively monitored and modified by the reasoner to guide the learning task (Gordon & Perlis, 1989; Martin, this volume; Provost, this volume; Provost & Buchanan, 1992; Utgoff, 1986).

Several models include learning goals as an explicit part of their formulation of the learning process. Learning goals have been used to guide resource allocation, information search, hypothesis evaluation, and other aspects of learning; to select and combine learning strategies; to guide and to learn about the reasoning process itself; and to model active learning in educational contexts.

Pragmatic implications of goal-driven learning

Goal-driven learning can provide considerable power in intelligent systems, whether those systems are viewed as computational models of human intelligence, or purely as artificial intelligence systems. In learning systems, goals can be used to focus learning and to avoid unrestricted search and inferencing. They can also be used to guide the information-seeking process and to make decisions about what, when, and how to learn.

Applying a planful model of learning promises to be fruitful for many applications, including perception (Pryor & Collins, 1992), intelligent information retrieval (Ram & Hunter, 1991), learning through apprenticeship (Redmond, 1992), knowledge acquisition (Quilici, this volume), information search during explanation (Leake, this volume), robotics (Earl & Firby, this volume; Karlsson, this volume), medical diagnosis (Hunter, 1990), natural language understanding (Cox & Ram, this volume; Ram, 1991), manufacturing (Perez, this volume; Ram, Narayanan, & Cox, 1993), and scientific discovery (Kocabas, this volume).

In addition, goals can be used as a theoretical device to build computational models of strategic and active reasoning and learning processes, and such models have practical ramifications for the design of instructional material. Ng and Bereiter (1991) show that different kinds of goals facilitate different kinds of reasoning and result in different kinds of learning. Such results suggest principles for the design of computer-based tools for education (Scardamalia & Bereiter, 1991). For example, van Berkum, Hijne, de Jong, van Joolingen, and Njoo (1991) use goal-driven learning both as a theoretical framework for decomposing the education problem and as a guide toward designing simulation-based instructional software. Schank proposes that because of the importance of goals in motivating and guiding learning, instruction should be conducted using a particular type of simulation environment—a *goal-based scenario*—to exploit the role of learning goals (Schank, Fano, Jona, & Bell, to appear). In goal-based scenarios, students play roles that are connected to their goals, and whose successful completion requires acquisition of the skills to be taught. In that way, goal-based scenarios provide a framework for students to perform goal-driven learning to acquire the skills to be taught.

Summary

In goal-driven learning, decisions of when to learn, what to learn, and how to learn are determined by explicit reasoning about needs for information. Although many aspects of goal-driven learning have been investigated in diverse fields, that research has been conducted in a piecemeal fashion, largely segregated by field. Even when multiple studies have been conducted in a single field, as is the case for artificial intelligence, each study has tended to concentrate on a few aspects of the problem without placing those aspects within a unifying framework and examining their larger implications.

We have presented a unifying picture of existing goal-driven learning research in terms of a new framework for modeling goal-driven learning, in terms of the types of goals that may guide learning, and in terms of the ways those goals can influence learning. The framework presented here is not suggested as a final theory of goal-driven learning, but rather a device for understanding the relationships of different results relevant to goal-driven learning and for suggesting issues that must be addressed with further investigation through a coordinated multidisciplinary research effort. The individual models and perspectives of the following papers illuminate specific aspects of the framework and the issues that remain to be addressed in future research.

Acknowledgements

This research was supported in part by the National Science Foundation under grant IRI-9009710, by the Georgia Institute of Technology, and by Indiana University. We would like to thank Betty Stanton for her enthusiasm and support for our work on the volume *Goal-Driven Learning* in which the full version of this chapter will appear, and to thank the MIT Press anonymous reviewers for their detailed comments on earlier versions of that chapter. We would also like to thank Michael Cox, Eric Domeshek, Susan Fox, Anthony Francis, Janet Kolodner, Kenny Moorman, Paul Panaro, and Raja Sooriamurthi for their many helpful comments on earlier drafts of that chapter.

References

Barsalou, L. (1991). Deriving Categories to Achieve Goals. In G.H. Bower, editor, *The Psychology of Learning and Motivation: Advances in Research and Theory*, Volume 27, Academic Press, New York, NY.

Barsalou (to appear). Storage Side Effects: Studying Processing to Understand Learning. In A. Ram & D.B. Leake, editors, *Goal-Driven Learning*, MIT Press/Bradford Books, Cambridge, MA.

Birnbaum, L., Collins, G., Freed, M., & Krulwich, B. (1990). Model-Based Diagnosis of Planning Failures. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, pages 318–323, Boston, MA.

Carbonell, J.G. (1986). Derivational Analogy: A Theory of Reconstructive Problem Solving and Expertise Acquisition. In R.S. Michalski, J.G. Carbonell, & T.M. Mitchell, editors, *Machine Learning II: An Artificial Intelligence Approach*, Morgan Kaufman Publishers, San Mateo, CA.

Chien, S. (1989). Using and Refining Simplifications: Explanation-based Learning of Plans in Intractable Domains. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 590-595, Detroit, MI.

Collins, G. & Birnbaum, L. (1988). An Explanation-Based Approach to the Transfer of Planning Knowledge Across Domains. In *Proceedings of the AAAI Spring Symposium on Explanation-based Learning*, pages 107–111, Stanford, CA.

Cox, M.T. (1993). *Introspective Multistrategy Learning.* Technical Report GIT-CS-93/02, Cognitive Science Program, College of Computing, Georgia Institute of Technology, Atlanta, GA.

Cox, M.T. & Ram, A. (this volume). Choosing Learning Strategies to Achieve Learning Goals. In M. desJardins & A. Ram, editors, *Proceedings of the AAAI Spring Symposium on Goal-Driven Learning*, Stanford, CA, 1994.

Dennett, D. (1987). *The Intentional Stance.* Bradford Books/MIT Press, Boston, MA.

desJardins, M. (1992). Goal-Directed Learning: A Decision-Theoretic Model for Deciding What to Learn Next. In *Proceedings of the ML-92 Workshop on Machine Discovery*, pages 147–151, Ninth International Machine Learning Conference, University of Aberdeen, Scotland.

Earl, C. & Firby, R.J. (this volume). An Integrated Action and Learning System. In M. desJardins & A. Ram, editors, *Proceedings of the AAAI Spring Symposium on Goal-Driven Learning*, Stanford, CA, 1994.

Faries, J. & Reiser, B. (1988). Access and Use of Previous Solutions in a Problem Solving Situation. In *Proceedings of the Tenth Annual Conference of the Cognitive Science Society*, pages 433–439, Montreal.

Fikes, R.E., Hart, P.E., & Nilsson, N.J. (1972). Learning and Executing Generalized Robot Plans. *Artificial Intelligence*, 3:251–288.

Forrest-Pressley, D.L., MacKinnon, G.E., & Waller, T.G. (1985), editors. *Metacognition, Cognition, and Human Performance, Volume 1 (Theoretical Perspectives)*, Academic Press, New York.

Freed, M. & Collins, G. (this volume). Learning to Cope with Task Interactions. In M. desJardins & A. Ram, editors, *Proceedings of the AAAI Spring Symposium on Goal-Driven Learning*, Stanford, CA, 1994.

Gauld, C. (1986). Models, meters and memory. *Research in Science Education*, 16:49–54.

Gavelek, J.R. & Raphael, R.E. (1985). Metacognition, Instruction and the Role of Questioning Activities. In D.L. Forrest-Pressley, G.E. MacKinnon, & T.G. Waller, editors, *Metacognition, Cognition, and Human Performance, Volume 2 (Instructional Practices)*, pages 103–136, Academic Press, New York.

Gordon, D. & Perlis, D. (1989). Explicitly Biased Generalization. *Computational Intelligence*, 5(2):67-81.

Graesser, A.C., Person, N., & Huber, J. (1992). Mechanisms that Generate Questions. In T.W. Lauer, E. Peacock, and A.C. Graesser, editors, *Questions and Information Systems*, pages 167–187, Lawrence Erlbaum Associates, Hillsdale, NJ.

Gratch, J., DeJong, G., & Chien, S.A. (this volume). Deciding When and How to Learn. In M. desJardins & A.

Ram, editors, *Proceedings of the AAAI Spring Symposium on Goal-Driven Learning*, Stanford, CA, 1994.

Greeno, J.G. & Simon, H.A. (1988). Problem Solving and Reasoning. In R.C. Atkinson, H. Hernstein, G. Lindzey, & R.D. Luce, editors, *Stevens' Handbook of Experimental Psychology*, Wiley, New York.

Hadzikadic, M. & Yun, D.Y.Y. (1988). Concept Formation by Goal-Driven, Context-Dependent Classification. In Z.W. Ras & L. Saitta, editors, *Methodologies for Intelligent Systems*, 3:322–332, 1988.

Hammond, K.J. (1989). *Case-Based Planning: Viewing Planning as a Memory Task.* Academic Press, Boston, MA.

Hammond, K., Converse, T., Marks, M., & Seifert, C.M. (1993). Opportunism and Learning. *Machine Learning*, 10(3):279–309.

Hayes-Roth, F. & Lesser, V. (1976). Focus of attention in a distributed logic speech understanding system. In *Proceedings of the IEEE International Conference on ASSP*, Philadelphia, PA.

Hoffman, C., Mischel, W., & Mazze, K. (1981). The Role of Purpose in the Organization of Information about Behavior: Trait-based versus Goal-Based Categories in Human Cognition. *Journal of Personality and Social Psychology*, 39:211–255.

Hunter, L.E. (1990). Planning to Learn. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, pages 261–276, Boston, MA.

Karlsson, J. (this volume). Task Decomposition in Reinforcement Learning. In M. desJardins & A. Ram, editors, *Proceedings of the AAAI Spring Symposium on Goal-Driven Learning*, Stanford, CA, 1994.

Keller, R. (1987). The Role of Explicit Contextual Knowledge in Learning Concepts to Improve Performance. Ph.D. Dissertation, Technical Report ML-TR-7, Department of Computer Science, Stanford University.

Keller, R. (1988). Defining Operationality for Explanation-Based Learning. *Artificial Intelligence*, 35:227–241.

Kedar-Cabelli, S. (1987). Formulating Concepts According to Purpose. In *Proceedings of the Sixth Annual National Conference on Artificial Intelligence*, pages 477-481, Seattle, WA.

Kocabas, S. (this volume). Goal-Directed Discovery and Explanation in Particle Physics. In M. desJardins & A. Ram, editors, *Proceedings of the AAAI Spring Symposium on Goal-Driven Learning*, Stanford, CA, 1994.

Kintsch, W. (1988). The Role of Knowledge in Discourse Comprehension: A Construction-Integration Model. *Psychological Review*, 95(2):163–182.

Krulwich, B. (this volume). A Component-Model Approach to Determining What to Learn. In M. desJardins & A. Ram, editors, *Proceedings of the AAAI Spring Symposium on Goal-Driven Learning*, Stanford, CA, 1994.

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*, pages 242–247, Bloomington, IN.

Laird, J.E., Rosenbloom, P.S., & Newell, A. (1986). Chunking in Soar: The Anatomy of a General Learning Mechanism. *Machine Learning*, 1:11–46.

Laird, J.E. (1991), editor. Special Section on Integrated Cognitive Architectures. *SIGART Bulletin*, 2(4):12–184.

Leake, D.B. (1989). Anomaly Detection Strategies for Schema-Based Story Understanding. In *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society,* pages 490–497, Ann Arbor, MI.

Leake, D.B. (1991). Goal-Based Explanation Evaluation. *Cognitive Science*, 15:509–545.

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

Leake, D.B. (1993). Learning Adaptation Strategies by Introspective Reasoning about Memory Search. In D.B. Leake, editor, *Proceedings of the AAAI-93 Workshop on Case-Based Reasoning*, pages 57–63, AAAI Press, Menlo Park, CA.

Leake, D.B. & Ram, A. (1993). *Goal-Driven Learning: Fundamental Issues (A Symposium Report). AI Magazine*, 14(4):67–72.

Leake, D.B. (this volume). Issues in Goal-Driven Explanation. In M. desJardins & A. Ram, editors, *Proceedings of the AAAI Spring Symposium on Goal-Driven Learning*, Stanford, CA, 1994.

Leake, D.B. (to appear). Goal-Driven Integration of Explanation and Action. In A. Ram & D.B. Leake, editors, *Goal-Driven Learning*, MIT Press/Bradford Books, Cambridge, MA.

Maes, P. (1990). Situated agents can have goals. *Robotics and Autonomous Systems*, 6:49–70.

Markovitch, S. & Scott, P.D. (1993). Information Filtering: Selection Mechanisms in Learning Systems. *Machine Learning*, 10:113–151.

Martin, J.D. (this volume). DC1: Supervised and Unsupervised Clustering. In M. desJardins & A. Ram, editors, *Proceedings of the AAAI Spring Symposium on Goal-Driven Learning*, Stanford, CA, 1994.

Michalski, R.S. (1983). A Theory and Methodology of Inductive Learning. *Artificial Intelligence*, 20:111-161.

Michalski, R.S. (1993). Inferential Theory of Learning as a Conceptual Basis for Multistrategy Learning. *Machine Learning*, 11(2/3):111–151, 1993.

Michalski, R.S. & Ram, A. (to appear). Machine Learning as Goal-Driven Inference. In A. Ram & D.B. Leake, editors, *Goal-Driven Learning*, MIT Press/Bradford Books, Cambridge, MA.

Minsky, M. (1963). Steps towards Artificial Intelligence. In E.A. Feigenbaum & J. Feldman, editors, *Computers and Thought*, pages 406–450, McGraw-Hill, New York, NY.

Minton, S. (1990). Quantitative Results Concerning the Utility of Explanation-Based Learning. *Artificial Intelligence*, 42:363-391.

Mitchell, T.M. (1982). Generalization as Search. *Artificial Intelligence*, 18:203–226.

Mitchell, T.M. & Keller, R. (1983). Goal Directed Learning. In *Proceedings of the International Machine Learning Workshop*, pages 117–118, Monticello, IL.

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

McKoon, G. & Ratcliff, R. (1992). Inference During Reading. *Psychological Review*, 99: 440-466.

Mooney, R.J. & Ourston, D. (1993). A Multistrategy Approach to Theory Refinement. In R.S. Michalski & G. Tecuci, editors, *Machine Learning: A Multistrategy Approach, Volume IV*, Morgan Kaufman Publishers, San Mateo, CA.

Mostow, J. & Bhatnagar, N. (1987). FAILSAFE—A Floor Planner that Uses EBG to Learn from its Failures. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, pages 249–255, Milan, Italy.

Newell, A. (1990). *Unified Theories of Cognition.* Harvard University Press. Cambridge, MA.

Newell, A. & Simon, H. (1972). *Human Problem Solving.* Prentice-Hall, Englewood Cliffs, NJ.

Ng, E. & Bereiter, C. (1991). Three Levels of Goal Orientation in Learning. *The Journal of the Learning Sciences*, 1(3&4):243–271.

Oehlmann, R., Sleeman, D., & Edwards, P. (1992.) Self-Questioning and Experimentation in an Exploratory Discovery System. In *Proceedings of the ML-92 Workshop on Machine Discovery*, pages 41–50, Ninth International Machine Learning Conference, University of Aberdeen, Scotland.

Owens, C. (1991). A Functional Taxonomy of Abstract Plan Failures. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*, pages 167–172, Chicago, IL.

Patalano, A., Seifert, C., & Hammond, K. (1993). Predictive Encoding: Planning for Opportunities. *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society*, pages 800–805, Boulder, CO.

Park, Y. & Wilkins, D. (1990). Establishing the Coherence of an Explanation to Improve Refinement of an Incomplete Knowledge Base. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, pages 511-516, Boston, MA.

Perez, M.A. (this volume). The Goal is to Produce Better Plans. In M. desJardins & A. Ram, editors, *Proceedings of the AAAI Spring Symposium on Goal-Driven Learning*, Stanford, CA, 1994.

Provost (this volume). Goal-Directed Inductive Learning: Trading off Accuracy for Reduced Error Cost. In M. desJardins & A. Ram, editors, *Proceedings of the AAAI Spring Symposium on Goal-Driven Learning*, Stanford, CA, 1994.

Provost, F.J. & Buchanan, B.G. (1992). Inductive Policy. In *Proceedings of the Tenth National Conference on Artificial Intelligence*, pages 255–261, San Jose, CA.

Pryor, L. & Collins, G. (1992). Planning to Perceive. In *Proceedings of the 1992 AAAI Spring Symposium on Selective Perception*, Stanford, CA.

Quilici, A. (this volume). Modeling the Goal-Driven Learning of Novice UNIX Users. In M. desJardins & A. Ram, editors, *Proceedings of the AAAI Spring Symposium on Goal-Driven Learning*, Stanford, CA, 1994.

Ram, A. (1987). AQUA: Asking Questions and Understanding Answers. In *Proceedings of the Sixth Annual National Conference on Artificial Intelligence*, pages 312–316, Seattle, WA.

Ram, A. (1990). Knowledge Goals: A Theory of Interestingness. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, pages 206–214, Cambridge, MA.

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

Ram, A. (1993). Indexing, Elaboration and Refinement: Incremental Learning of Explanatory Cases. *Machine Learning*, 10:201–248.

Ram, A. & Cox, M.T. (1994). Introspective Reasoning using Meta-Explanations for Multistrategy Learning. In R.S. Michalski & G. Tecuci, editors, *Machine Learning: A Multistrategy Approach, Volume IV*, Morgan Kaufman Publishers, San Mateo, CA.

Ram, A., Cox, M.T., & Narayanan, S. (to appear). Goal-Driven Learning in Multistrategy Reasoning and Learning Systems. In A. Ram & D.B. Leake, editors, *Goal-Driven Learning*, MIT Press/Bradford Books, Cambridge, MA.

Ram, A. & Hunter, L. (1991). A Goal-Based Approach to Intelligent Information Retrieval. In L. Birnbaum & G. Collins, editors, *Machine Learning: Proceedings of the Eighth International Workshop*, Chicago, IL.

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

Ram, A. & Leake, D.B. (1991). Evaluation of Explanatory Hypotheses. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*, pages 867–871, Chicago, IL.

Ram, A., Narayanan, S., & Cox, M.T. (1993). Learning to troubleshoot: Multistrategy learning of diagnostic knowledge for a real-world problem solving task. Technical Report GIT-CC-93/67, College of Computing, Georgia Institute of Technology, Atlanta, Georgia 30332-0280, 1993.

Redmond, M. (1992). *Learning by Observing and Understanding Expert Problem Solving.* Ph.D. Dissertation, Technical report GIT-CC-92/43, College of Computing, Georgia Institute of Technology, Atlanta, GA.

Riesbeck, C. (1981). Failure-driven Reminding for Incremental Learning. In *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, pages 115–120, Vancouver, British Columbia.

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

Schank, R.C. & Abelson, R. (1977). *Scripts, Plans, Goals and Understanding.* Lawrence Erlbaum Associates, Hillsdale, NJ.

Schank, R.C., Fano, A., Jona, K., & Bell, B. (to appear). The Design of Goal-Based Scenarios. *The Journal of the Learning Sciences*, in press.

Schank, R.C. & Leake, D.B. (1989). Creativity and Learning in a Case-Based Explainer. *Artificial Intelligence*, 40(1–3):353–385. Also appears in J.G. Carbonell, editor, *Machine Learning: Paradigms and Methods*, MIT Press, Cambridge, MA, 1990.

Schlimmer, J.C. (1987). Incremental Adjustment of Representations for Learning. In *Proceedings of the Fourth International Workshop on Machine Learning*, pages 79–90, Irvine, CA.

Seifert, C. (1988). Goals in Reminding. In J.L. Kolodner, editor, *Proceedings of the Case-Based Reasoning Workshop*, pages 190–208, Morgan Kaufmann Publishers, San Mateo, CA.

Slade, S. (1993). *Goal-Based Decision Making: An Interpersonal Model.* Lawrence Erlbaum Associates, Hillsdale, NJ, in press.

Sperber, D. & Wilson, D. (1986). *Relevance: Communication and Cognition.* Language and Thought Series, Harvard University Press, Cambridge, MA.

Srull, T. & Wyer, R. (1986). The Role of Chronic and Temporary Goals in Social Information Processing. In R. Sorrentino and E. Higgins, editors, *Handbook of Motivation and Cognition: Foundations of Social Behavior*, pages 503–549. Guilford Press, Guilford, CT.

Steinbart, P.J. (1992). The Role of Questioning in Learning from Computer-Based Decision Aids. In T.W. Lauer, E. Peacock, and A.C. Graesser, editors, *Questions and Information Systems*, pages 273–285, Lawrence Erlbaum Associates, Hillsdale, NJ.

Stroulia, E. & Goel, A.K. (this volume). Task Structures: A Language for Learning Goals. In M. desJardins & A. Ram, editors, *Proceedings of the AAAI Spring Symposium on Goal-Driven Learning*, Stanford, CA, 1994.

Stroulia, E., Shankar, M., Goel, A.K., & Penberthy, L. (1992). A Model-Based Approach to Blame Assignment in Design. In J.S. Gero, editor, *Proceedings of the Second International Conference on AI in Design*, pages 519–537.

Sussman, G. (1975). *A Computer Model of Skill Acquisition.* Artificial Intelligence Series, Volume 1, American Elsevier, New York, NY.

Thagard, P. & Millgram, E. (to appear). Inference to the Best Plan: A Coherence Theory of Decision. In A. Ram & D.B. Leake, editors, *Goal-Driven Learning*, MIT Press/Bradford Books, Cambridge, MA.

Utgoff, P.E. (1986). Shift of Bias for Inductive Concept Learning. In R.S. Michalski, J.G. Carbonell, & T.M. Mitchell, editors, *Machine Learning II: An Artificial Intelligence Approach*, Morgan Kaufman Publishers, San Mateo, CA.

van Berkum, J.J.A., Hijne, H., de Jong, T., van Joolingen, W.R., & Njoo, M. (1991). Aspects of Computer Simulations in an Instructional Context. *Education and Computing*, 6:231–239, 1991.

Van de Velde, W. (1988). Quality of Learning. In *Proceedings of the European Conference on Artificial Intelligence*, pages 408–413, Munich, Germany.

VanLehn, K. (1989). Problem Solving and Cognitive Skill Acquisition. In M.I. Posner, editor, *Foundations of Cognitive Science*, MIT Press, Cambridge, MA.

VanLehn, K. (1991a). Rule Acquisition Events in the Discovery of Problem Solving Strategies. *Cognitive Science*, 15(1):1–47.

VanLehn, K. (1991b), editor. *Architectures for Intelligence: The Twenty-Second Carnegie Mellon Symposium on Cognition*, Lawrence Erlbaum Associates, Hillsdale, NJ.

Weinert, F.E. (1987). Introduction and Overview: Metacognition and Motivation as Determinants of Effective Learning and Understanding. In F.E. Weinert & R.H. Kluwe, editors, *Metacognition, Motivation, and Understanding*, Lawrence Erlbaum Associates, Hillsdale, NJ.

Weintraub, M. (1991). *An Explanation-Based Approach to Assigning Credit,* Ph.D. Dissertation, Computer Science Department, The Ohio State University, Columbus, OH.

Wellman, H.M. (1985). The Origins of Metacognition. In D.L. Forrest-Pressley, G.E. MacKinnon, & T.G. Waller, editors, *Metacognition, Cognition, and Human Performance*, Volume 1, Academic Press, New York.

Wellman, H.M. (1992). *The Child's Theory of Mind.* MIT Press/Bradford Books, Cambridge, MA.

White, R.T. & Gunstone, R.F. (1989). Metalearning and Conceptual Change. *International Journal of Science Education*, 11:577-586.

Wilensky, R. (1983). *Planning and Understanding.* Addison-Wesley, Reading, MA.

Wisniewski, E.J. & Medin, D.L. (1991). Harpoons and Long Sticks: The Interaction of Theory and Similarity in Rule Induction. In D. Fisher & M.J. Pazzani, editors, *Concept Formation: Knowledge and Experience in Unsupervised Learning*, Morgan Kaufman Publishers, San Mateo, CA.

Wrobel, S. (1988). Automatic Representation Adjustment in an Observational Discovery System. In D. Sleeman, editor, *Proceedings of the Third European Working Session on Learning*, pages 253–261, Pitman Publishing, London.

Xia, X. & Yeung, D.-Y. (1988). A Learning Model for the Selection of Problem Solving Strategies in Continuous Physical Systems. In *Proceedings of the Fourth IEEE Conference on Artificial Intelligence Applications*, pages 200–206, San Diego, CA.

Zukier, H. (1986). The Paradigmatic and Narrative Modes in Goal-Guided Inference. In R. Sorrentino and E.T. Higgins, editors, *Handbook of Motivation and Cognition: Foundations of Social Behavior*, Guilford Press, New York, NY.