Goal-Driven Learning in Multistrategy Reasoning and Learning Systems*

Ashwin Ram, College of Computing

Michael T. Cox, College of Computing

S. Narayanan, School of Industrial & Systems Engineering

Georgia Institute of Technology

Atlanta, Georgia 30332-0280

1 Introduction

This chapter presents a computational model of introspective multistrategy learning, which is a deliberative or strategic learning process in which a reasoner introspects about its own performance to decide what to learn and how to learn it. The reasoner introspects about its own performance on a reasoning task, assigns credit or blame for its performance, identifies what it needs to learn to improve its performance, formulates learning goals to acquire the required knowledge, and pursues its learning goals using multiple learning strategies. Our theory models the following characteristics of goal-driven learning: (i) that *learning is active*, and strategic, goal-driven processes underlie much of the learning that occurs during the performance of analytical tasks in complex, real-world domains; (ii) that *learning is experiential* and occurs incrementally through the performance of a reasoning task; (iii) that *learning is opportunistic*, and learning goals that are not immediately satisfiable are remembered so that the reasoner can recognize and use later opportunities to pursue them; (iv) that *learning is diverse* and involves multiple different strategies for acquiring new

*To appear in A. Ram and D. Leake (eds.), *Goal-Driven Learning*, MIT Press/Bradford Books, Cambridge, MA, forthcoming.

knowledge, modifying existing knowledge, and reorganizing the knowledge base; and finally (v) that *learning is introspective* and involves reflecting upon one's own performance, monitoring the state of one's own knowledge, and analyzing one's own reasoning processes.

This chapter also describes two computer systems that implement our theory. The systems are active and goal-driven, starting out with an incomplete understanding of a novel domain and learning through experience using multiple learning strategies. Their learning goals are functional to the purpose of the systems, and are identified during the pursuit of the performance task. The systems reason about the best way to perform a task, introspectively analyze their own successes and failures in performing their tasks, reason about what they need to learn, select appropriate learning strategies to acquire that information, and invoke the learning algorithms which then cause them to acquire new knowledge, modify existing knowledge, or reorganize knowledge in memory. Our theory is motivated by cognitive as well as computational considerations, and provides a framework for the development of integrated, multistrategy learning systems for real-world tasks.

2 The nature of goal-driven learning

In our model, goal-driven learning is viewed as an active, experiential, opportunistic, multistrategy, and introspective process. Although some of these properties have individually been the basis of other approaches to learning, our model integrates all these properties into a single uniform framework and is thus capable of effective learning across a variety of learning situations in a variety of task domains. Let us discuss these properties at greater length.

Active learning: Some traditional approaches in machine learning have assumed that the knowledge to be learned has already been identified by an external agent (e.g., the "target concept" in explanation-based generalization [Mitchell, Keller, and Kedar-Cabelli , 1986]). In some approaches, the learning process has no target or goal at all; the program has no sense of what it is trying to learn or why it is trying to learn it. Recently, several researchers have argued that the identification and pursuit of what might be called the "learning goals" or "knowledge goals" of the reasoner is an important aspect of the learning problem (e.g., [Collins, Birnbaum, Krulwich, and Freed , 1993; desJardins, 1992; Hunter, 1990; Laird, Rosenbloom, and Newell , 1986; Michalski, 1993; Mitchell, Utgoff, and Banerji , 1983; Quilici, to appear; Ram, 1989; Ram, 1991; Ram and Cox, 1994; Ram and Hunter, 1992; Redmond, 1992]). We argue that active, goal-based learning is important for functional or computational reasons as well as for cognitive reasons. Thus formulating learning goals, asking questions, focussing attention, and pursuing learning actions are essential components of our learning model. Since the need to learn often arises from a reasoning failure, credit or blame assignment also plays a central role in learning [Hammond, 1989; Minsky, 1985; Ram and Cox, 1994; Schank, 1982; Stroulia, Shankar, Goel, and Penberthy , 1992; Sussman, 1975; Weintraub, 1991].

Experiential learning: Learning is an incremental process of theory formation, theory revision, and conceptual change, which occurs as the reasoner accumulates experience in some task domain. Through this experience, the reasoner learns to avoid the mistakes made during the performance task. Learning and reasoning, therefore, are tightly coupled into a single integrated framework. Introspective analysis of a reasoning experience in a given situation allows the reasoner to use an appropriate learning strategy (or, as in our models, multiple learning strategies) to learn exactly that which it needs to know to improve its ability to process similar situations in the future. This is essentially a case-based or experience-based approach, which relies on the assumption that it is worth learning about one's experiences since one is likely to have similar experiences in the future (see, e.g., [Hammond, 1989; Kolodner, 1993; Kolodner and Simpson, 1984; Ram, 1993; Schank, 1982]).

Opportunistic learning: A corollary of the active and experiential nature of goal-driven learning is that learning is opportunistic. Often, a desired piece of knowledge will not be immediately available in the input, and so the corresponding learning goal will not be immediately satisfiable. In such cases, the reasoner must be able to suspend its learning goals, and reactivate them later when an appropriate opportunity arises. Because learning goals are indexed in memory, it is quite likely that an understander will find information relevant to goals other than the ones that are currently "active." In other words, learning goals can be satisfied opportunistically during the course of understanding [Birnbaum, 1986; Dehn, 1989; Hammond, Converse, Marks, and Seifert , 1993; Ram, 1989; Ram, 1991], leading to opportunistic learning of information previously identified as being useful to obtain (e.g., [Ram, 1993; Ram and Hunter, 1992]). In order for this to happen, the reasoner must be able to remember what it needs to learn, and recognize opportunities to learn the desired knowledge.

Multistrategy learning: There are several things one might learn from any experience, and several different ways of learning these. Once the reasoner has identified what to learn, it still needs to identify the method best suited for performing the desired learning. In many cases, a combination of learning strategies is necessary. For example, if the reasoner is presented with a novel explanation for a problem, it needs to be able both to acquire such an explanation in a general way (explanation generalization) and to remember it again in future situations in which it is likely to be applicable (index learning). Furthermore, a single learning strategy may be applicable in a number of different reasoning situations. For example, the reasoner may need to learn a new index to an explanation, both when the explanation is newly acquired and when the explanation is already known but incorrectly indexed in memory. Identifying appropriate learning strategies is called the "strategy selection problem", and is particularly important in multistrategy learning systems (e.g., [Cox and Ram, 1992b; Hunter, 1990; Reich, 1993; Ram and Cox, 1994]).

Introspective learning: There are several fundamental problems to be solved before we can build intelligent systems capable of general multistrategy learning, including: determining the cause of a reasoning failure (*blame assignment*), deciding what to learn (*learning goal formulation*), and selecting the best learning strategies to pursue these learning goals (*strategy selection*). Although previous research has led to algorithms for learning in particular situations, no general theory of learning exists which allows the system to determine its own learning goals and to learn using multiple learning strategies. We claim that a reasoning system that can do this in a general manner must be able to reflect or introspect about its own internals. Determining the reasons why a failure occurred often involves not just analyzing a plan of action that the system created or understanding the events in the world that resulted from it, but also understanding the reasoning process by which the plan was created and the knowledge which was used in that reasoning process.

For example, consider a situation in which an agent runs out of gas on a vacation. There are several possible causes of this failure, some attributable to factors in the external world (such as a hole in the car's fuel lines) and some to factors in the "internal" world: Did the agent not formulate the goal to fill up the car with gas? Or, did the agent, having formulated that goal, forget to achieve it before setting out on the trip? Did the agent not know that cars needed to be refueled? While such as analysis may not always help in recovering from the failure in the present situation, it is essential in learning to avoid such a failure in the future. As this simple example illustrates, in order to decide what to learn (and hence how to learn it), the agent must analyze its internal cognitive processes and knowledge.

Our approach to *introspective multistrategy learning* is as follows. First, we identify classes of learning situations based on an analysis of the types of reasoning failures that might occur. The taxonomy of reasoning failures follows from the functional architecture of the reasoning system. For each type of reasoning failure, we identify how the conclusions were drawn (a description of a chain of reasoning that led up to those conclusions), why these conclusions were drawn (a description of the bases for the processing decisions underlying that chain of reasoning), why the conclusions were faulty or inappropriate (an explanation of why the drawn conclusions were incorrect), what the correct conclusions ought to have been (a description of the desired conclusions), and how the reasoner should have drawn them (a description of a chain of reasoning that would lead to the desired conclusions).[1] Finally, for each type of reasoning failure, we identify what needs to be learned to avoid such a failure, and associated learning strategies that can learn the desired knowledge.

This information is represented explicitly in the system using a meta-model describing the reasoning process itself. In addition to the world model that describes its domain, the reasoning system has access to meta-models describing its reasoning processes, the knowledge that this reasoning is based on, and the indices used to organize and retrieve this knowledge. A meta-model is used to represent the system's reasoning during a performance task, the decisions it took while performing the reasoning, and the results of the reasoning. If a difficulty or failure is encountered, the system introspectively examines the record of its own reasoning processes to determine where the problem lies, and uses this introspective understanding to improve itself using the appropriate learning strategies [Cox and Ram, 1992b; Ram and Cox, 1994].

3 An architecture for introspective multistrategy learning

Figure 1 illustrates a general architecture for an integrated multistrategy reasoning and learning system. The reasoner receives input from the outside world, and focuses its attention on the part of the input which is relevant or interesting to it. It then uses one or more reasoning strategies, selected using a set of heuristics, to process the input appropriately. The strategies rely on knowledge that is indexed in the system's memory. The reasoner also records a trace of its reasoning process, which is introspectively examined for the purposes of blame assignment, deciding what to learn, and selecting the appropriate learning strategy(s).[2]

The key representational entity in our learning theory is a *meta-explanation pattern* (Meta-XP), which is a causal, introspective explanation structure that explains how and why an agent reasons, and that helps the system in the learning task. There are two broad classes of Meta-XPs. *Trace Meta-XPs* record a declarative trace of the reasoning performed by a system, along with causal links that explain the decisions taken. The trace holds explicit information concerning the manner in which knowledge gaps are identified, the reasons why particular hypotheses are generated, the strategies chosen for verifying candidate hypotheses, and the basis for choosing particular reasoning methods for each of these. Trace Meta-XPs are similar to "reasoning traces" [Carbonell, 1986; Minton, 1988; Veloso and Carbonell, 1993] or "justification structures" [Collins, Birnbaum, Krulwich, and Freed , 1993; deKleer, Doyle, Steele, and Sussman , 1977; Doyle, 1979], with the difference that Trace Meta-XPs represent, in addition to the subgoal structure of the problem and justifications for operator selection decisions, information about the structure of the (possibly multistrategy) reasoning process that generated a solution. For example, at the highest level of granularity, a node in a Trace Meta-XP might represent the choice of a reasoning method such as case-based reasoning, and at a more detailed level a node might represent the process of retrieving a case from memory. These structures could, therefore, be viewed as representing the "mental operators" underlying the reasoning process.

Figure 1 should be placed near here.

We also propose a new kind of meta-explanation structure to represent classes of reasoning errors along with the types of learning needed in those situations. For example, a common type of reasoning failure is one where a reasoner mistakenly applies an inappropriate knowledge structure (e.g., a rule or a case) in a novel situation because it did not have the right piece of knowledge for that situation. In such a situation, the reasoner must learn the new knowledge, and also revise the indices (or applicability conditions) of the incorrect knowledge structure to prevent it from being applied in similar situations in the future. In contrast, a situation in which the reasoner does have the appropriate, but incorrectly indexed, knowledge structure, results in the need to learn a new index to an existing knowledge structure. Analysis of the reasoning trace to identify the nature of the reasoning failure is an essential component of a multistrategy learning system that can automatically identify and correct its own shortcomings.

Instead of simply representing a trace of the reasoning process, therefore, we also represent the knowledge required to analyze these traces in order to determine what to learn and how to learn it. If the reasoner encounters a reasoning failure, it uses *Introspective Meta-XPs* to examine the declarative reasoning chain. Introspective Meta-XPs are structures used both to explain a reasoning failure and to learn from it. They associate a failure type with learning goals and the appropriate set of learning strategies for pursuing those goals, and point to likely sources of the failure within the Trace Meta-XP. Thus an Introspective Meta-XP performs three functions: it aids in blame assignment (determining which knowledge structures are missing, incorrect or inappropriately applied); it aids in the formulation of appropriate learning goals to pursue; and it aids in the selection of appropriate learning algorithms to recover and learn from the reasoning error. Such meta-explanations augment a system's ability to introspectively reason about its own knowledge, about gaps within this knowledge, and about the reasoning processes which attempt to fill these gaps. The use of explicit Meta-XP structures allows direct inspection of the reasons by which learning goals are posted and processed, thus enabling a system to improve its ability to reason and learn.

4 A taxonomy of reasoning failures

Based on the above architecture, we can characterize the types of reasoning failures that a reasoner might encounter. The term "reasoning failure" includes not simply performance errors, but also expectation failures [Schank, 1986], anomalous situations which the reasoner failed to predict, and other types of reasoning failures as well. By using failures to guide learning, the reasoner can focus on the problems that actually arise rather than searching through the more complex space of problems that can theoretically arise [Hammond, Converse, Marks, and Seifert , 1993]. In addition, since not all potential learning is actually useful to perform, failures encountered during the performance of actual tasks can focus learning effort on knowledge that is likely to be useful in performing the tasks that the reasoner is engaged in. In our model of goal-driven learning, reasoning failures are used to focus the reasoner's attention on its own inadequacies and hence to trigger and guide the goal-driven learning process. Unlike successful processing where there may or may not be anything to learn, failure situations are guaranteed to provide a potential for learning, otherwise the failure would not have occurred [Hammond, 1989; Minsky, 1985; Schank, 1982; Sussman, 1975]. In our model, an unexpected success also counts as a reasoning "failure" because the reasoner was unable to correctly predict the outcome of the task.

If reasoning is modelled as goal-directed processing of an input using some knowledge, there is only a limited number of classes of factors that may be responsible for the success or failure of the reasoning process. A failure could stem from the goal, the process, the input, or the domain knowledge. Furthermore, if both knowledge and sets of reasoning strategies are organized in order to facilitate access to them, so that appropriate knowledge and strategies can be retrieved and brought to bear on a given situation, the organization of knowledge and reasoning strategies may be responsible for a failure as well. Finally, failures may arise from the generation, selection, and opportunistic triggering of goals as well.

Table 1 should be placed near here.

These dimensions of reasoning failures, shown in table 1, identify classes of factors that bear on the blame assignment problem and, hence, on learning. The top row lists failure categories for each column; these follow from the assumptions that reasoning is goal-directed processing of input using some knowledge, and that knowledge, process, and goals all have an organizational component in memory which affects retrieval or selection. The bottom row is a general characterization of these categories, in terms of the introspective multistrategy learning architecture, that relate them to more traditional distinctions (as between action and perception). From left to right across the top, we see that reasoning failures may be attributed to some piece of background knowledge about the domain, to the indices used to organize and access domain knowledge, to reasoning strategies or methods used to pursue goals, to the heuristics used to choose such strategies, to the reasoner's goals and the manner in which they are generated, to the opportunistic selection of suspended goals (or tasks) to pursue, or, finally, to the input. Each column in the table represents one of these dimensions, each of which could be missing or incorrect (the failure cases) or correct (the successful case). Finally, as suggested by the subtable in the input column, if a reasoner interacts with other agents in the world, blame for a failure may be attributed to the goals, strategies, input and knowledge of other agents.

Learning is triggered when a reasoning failure is detected. For example, consider the domain knowledge column in table 1. A reasoner may fail to explain a novel situation due to incomplete domain knowledge. Without the necessary knowledge, the reasoner cannot produce an explanation and thus "draws a blank." Alternatively, it might detect an anomaly by noticing a contradiction between the input and its domain knowledge. Such an anomaly corresponds to either incorrect domain knowledge or to erroneous or noisy input.

The knowledge selection column in table 1 characterizes reasoning failures that can occur if the reasoner, while possessing the relevant knowledge, does not have it organized sufficiently or labelled correctly so as to allow retrieval of the knowledge at the appropriate time. A missing index may lead to a failure to retrieve a relevant piece of knowledge, leading to a kind of "forgetting" [Cox and Ram, 1992a]. Similarly, an erroneous index may bring inappropriate knowledge to bear, leading to inappropriate inferences or conclusions.

The processing strategy and goals columns can be thought of in much the same way as the domain knowledge column; in addition to being either missing or incorrect, they too have a organizational component and are thus subject to retrieval problems. As described earlier, opportunistic reasoners suspend goals that cannot currently be achieved, indexing them in memory where they can be retrieved at a later time. The reasoner may fail to formulate the appropriate reasoning goals to pursue, or it may select the wrong goal to pursue next. For example, a reasoning failure may occur because a goal was labelled with an index that did not match the cues present in the input at a later retrieval time.

A multistrategy reasoning system must have some method of selecting appropriate reasoning strategies to apply in a given situation. Such a method relies on heuristic knowledge of the applicability conditions of these strategies. To facilitate introspection, reasoning strategies are represented declaratively using knowledge structures that the reasoner can inspect; applicability heuristics, then, may be viewed as "indices" to these knowledge structures. A reasoning failure might occur if the reasoner does not have the appropriate reasoning strategy with which to process the input (represented in the processing strategy column of table 1), or if its heuristics do not correctly identify the best strategy to use even it is is available (represented in the strategy selection column). This may lead the reasoner to come to an incorrect conclusion, or to come to a correct conclusion less efficiently than it might have.

The most complicated column is the one representing the input to the reasoner, or, more generally, the interaction between the reasoner and the external world. Errors of reasoning may occur due to incorrect (noisy) input, and also due to missing input. Furthermore, as suggested by the subtable in that column, a reasoner that interacts with other agents in the world may need to assign responsibility for a reasoning failure to the goals, strategies, input and knowledge of other agents. Thus, for example, noise in the input may actually be due to intentional deception motivated by conflicting goals of an external agent. Although this makes blame assignment exceedingly difficult, it cannot be ignored in open world situations. The input column is interesting in that it models the interaction between the input, the perception of the input, and the interpretation of the perception of the input.

These and other types of reasoning failures are detected through an analysis of the Trace Meta-XP that represents the reasoning trace (see table 2). Reasoning failures are associated with learning goals which specify what the reasoner needs to learn, which, in turn, are associated with learning strategies. An Introspective Meta-XP, then, could be viewed as a "reasoning pattern," representing a trace of a typical reasoning process, the failures encountered during the reasoning process, and the learning necessary in that situation. Examples and further details may be found in [Cox, 1993; Cox and Ram, 1992b; Ram and Cox, 1994].

5 Two case studies

To ensure generality of the theory, we have performed two case studies in developing reasoning systems in two very different task domains. One system uses “shallow” knowledge to troubleshoot in an electronics assembly plant, and the other uses “deep” causal knowledge to understand natural language stories. Although space limitations preclude detailed discussion of the two systems, they are discussed briefly below. Although each system currently focuses on a few types of reasoning failures, we are continuing to extend these systems to deal with all the types of reasoning failures discussed above.

5.1 Meta-TS

The first task is that of knowledge-based diagnostic problem solving. Meta-TS[3] is an introspective multistrategy learning system that learns to troubleshoot in a real-world domain of electronic assembly manufacturing. The system is based on observations of human troubleshooting operators and protocol analysis of the data gathered in the test area of NCR’s electronics assembly manufacturing plant located near Atlanta. The computational model integrates several strategies for learning the associative and heuristic knowledge used by the operators in the problem solving task. The system takes as input the in-circuit test ticket information and the printed circuit board information available to the human troubleshooter.[4] The output of the system is a diagnosis and a set of repair actions. The system also learns from each troubleshooting episode in order to improve its own performance at the task.

Since a particular problem solving episode may involve several pieces of knowledge (potentially of different types) and real-world interactions (such as tests on the printed circuit board), the troubleshooter, whether human or machine, must be able to examine the reasons for successes and failures during problem solving in order to improve its performance on future problems. To accomplish this, the system must be able to examine its own problem solving processes. After each troubleshooting episode, Meta-TS uses declarative representations of the knowledge and methods used for problem solving in order to facilitate critical self-examination. A trace of the problem solving process is constructed during the troubleshooting episode, and introspectively analyzed during the learning phase to determine what the system might learn from that episode.

The analysis also helps the system select the learning strategy appropriate for that type of learning. The introspective multistrategy learning architecture provided a framework for the construction of a computational model of this learning task.

Meta-TS uses Introspective Meta-XPs to introspect into its troubleshooting processes (represented using Trace Meta-XPs), to decide what to learn, and to select an appropriate learning strategy. In this task domain, we found that operators rely predominantly on “shallow” reasoning using heuristic and context-sensitive associative knowledge during problem solving; furthermore, they often make use of a human expert. Thus, in the Meta-TS system we have focussed primarily on the supervised and unsupervised acquisition, modification, and deletion of associative knowledge through the analysis of reasoning traces which, however, do not contain detailed domain knowledge. For example, if the system arrives at a correct diagnosis using heuristic knowledge alone, it compiles that knowledge into an association using a learning method similar to knowledge compilation [Anderson, 1989] in order to reduce the number of intermediate steps when the system encounters a similar problem in the future. An interactive knowledge acquisition strategy is used to acquire a new association from an expert troubleshooter (the operator’s supervisor) when the system arrives at an incorrect solution and the supervisor is available. Other strategies recommend postponement of a learning goal when expert input is unavailable, or deletion of an association that is rendered obsolete due to changes in the manufacturing process. Furthermore, given the architectural framework of the system, it is relatively straightforward to incorporate additional learning strategies. Meta-TS has been evaluated both qualitatively and quantitatively using several test problems from the data gathered at the NCR plant. The results show that Meta-TS can be justified both as a plausible model of human learning and as a computational framework for the design of machine learning systems. Further details may be found in [Ram, Narayanan, and Cox , 1993].

5.2 Meta-AQUA

The second task is that of causal and motivational analysis of conceptual input in order to infer coherence-creating structures, or explanations, that tie the input together. The performance task in Meta-AQUA[5] is to understand input stories in the domain of drug smuggling, i.e., to construct explanations of the actions observed in the story that causally relate the actions to the goals, plans and beliefs of the actors and planners of the actions. Reasoning failures arise when the system's background knowledge is incomplete or incorrect, or when the appropriate background knowledge is present but indexed incorrectly in memory. In order to learn from the failure and to avoid repeating the same mistake in the future, the system needs to identify the cause of the failure and then, depending on the cause, apply a given learning strategy.

Unlike Meta-TS, the knowledge used by Meta-AQUA's reasoning module is deep, explanatory knowledge about physical causality and human motivations. Trace Meta-XPs are instantiated to represent the system's explanation process, and the system's use of this knowledge for the explanation process. Learning strategies are selected using Introspective Meta-XPs. In general, Introspective Meta-XPs are built out of reasoning chains involving successful predictions, expectation failures, retrieval failures and incorporation failures. For example, a common type of failure arises when the reasoner finds an explanation that it thinks is appropriate, but the correct explanation turns out to be a different, novel explanation that the reasoner did not know about. This situation, or "reasoning pattern," is represented by a composite Meta-XP that consists of two basic Meta-XPs: `XP-Novel-Situation` and `XP-Mis-Indexed-Structure`. `XP-Novel-Situation` directs an explanation-based generalization algorithm to be applied to the node representing the novel explanation. The new explanation is then indexed in memory using an index learning algorithm. `XP-Mis-Indexed-Structure` directs the indexing algorithm to the old, incorrectly applied explanation. It recommends that the explanation be re-indexed so that it is not retrieved in similar situations in the future. A "knowledge planner" controls the learning process and manages potential interactions between learning goals [Cox and Ram, 1994].

Further details of the Meta-AQUA system can be found in [Cox, 1993; Cox and Ram, 1992b; Ram and Cox, 1994], along with a taxonomy of reasoning failures and representations of Trace and Introspective Meta-XPs. The process of goal-driven learning in the Meta-AQUA and Meta-TS systems is based on the algorithm presented in table 2.

Table 2 should be placed near here.

6 Conclusions

The focus of our research is on the integration of different kinds of knowledge and reasoning processes into goal-driven, real-world systems that can learn through experience. In particular, we are interested in modelling the kind of active, goal-driven learning processes that underlie deliberative learning during the performance of complex reasoning tasks. Our model of introspective multistrategy learning makes several claims about the nature of learning, reasoning and introspection that are supported by research in psychology and metacognition. In psychology, goals have been shown to have strong effects on the human learning process (e.g., [Barsalou, 1991]); in social cognition, goals have been shown to focus attention and inferencing (e.g., Zukier's [1986] review); and in education, learning goals have been shown to have a strong effect on student performance (e.g., [Ng and Bereiter, 1991; Scardamalia and Bereiter, 1991]). The meta-explanations in our approach are similar to self-explanations [Chi and VanLehn, 1991; Pirolli and Bielaczyc, 1989]. Research in self-explanation shows that formulation of self-explanations while understanding input examples significantly improves the subjects' ability to learn from the examples. One difference between the two approaches is that self-explanations are explanations about events and objects in the world, whereas our meta-explanations are explanations about events and objects in the reasoner's "mind". Pirolli and Recker [in press] make a similar distinction between learners that perform self-explanation and those that engage in self-reflection. Experimental results in the metacognition literature suggests that introspective reasoning of the kind we propose here can facilitate reasoning and learning (see, e.g., [Schneider, 1985; Weinert, 1987; Wellman, 1992]).

Our approach can be justified on functional and computational grounds as well. Computational considerations about the complexity of reasoning in open world situations point to the need for making decisions about the utility of learning, and for guiding learning based on the goals of the reasoner. Explicit models of learning goals facilitate the use of unexpected 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. A goal-driven model of learning provides additional computational benefits in a multistrategy learning system. The reasoner can consider goal priorities and goal dependencies to decide what to learn and how to coordinate multiple learning strategies to achieve this learning. Learning strategies, represented as methods for achieving learning goals, can be chained, composed, and optimized, resulting in learning plans [Cox and Ram, 1994; Hunter, 1990; Redmond, 1992; Ram and Hunter, 1992] that are created dynamically and pursued in a flexible manner. 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.

Our theory relies on a declarative representation of meta-models for reasoning and learning to achieve goal-driven learning in multistrategy reasoning and learning systems. There are several advantages of maintaining such structures in memory. Because these structures represent reasoning processes explicitly, the system can directly inspect the reasons underlying a given processing decision it has taken, evaluate the progress towards a goal, and compare its reasoning to past instances of reasoning in similar contexts. Thus, these traces can also be used in credit/blame assignment, to analyze why reasoning errors occurred, and to facilitate learning from these errors. Furthermore, because both the reasoning process and the knowledge base are represented using the same type of declarative representations, processes which identify and correct gaps in a knowledge base can also be applied to the reasoning process itself. Finally, these knowledge structures provide a principled basis for integrating multiple reasoning and learning strategies, and a unified framework which makes it relatively straightforward to incorporate additional types of failures and additional learning strategies.

Knowledge about reasoning and learning processes is often encoded procedurally in current artificial intelligence systems; that is, this knowledge is embodied in the behavior of the programmed system rather than explicit knowledge structures and is therefore beyond the system's scope of analysis. As far back as McCarthy [1958/1985], the argument that declarative representation is a precursor to successfully changing a system's knowledge has been a compelling one. This argument applies not just to a system's knowledge of an external domain, but also to knowledge about the operation of the system itself. For example, Provost and Buchanan [1992] argue that for a system to successfully choose an appropriate bias for learning, the bias and the policy that the bias supports should be declarative rather than implicit. Likewise, we claim that knowledge of a system's own reasoning abilities and episodes of reasoning experience is simply another piece of knowledge that should be represented declaratively. Using such representations, a system can better modify and understand its reasoning and can manipulate it in consistent manner.

Our systems can reason introspectively about their own reasoning process and hence determine both what they need to learn and what learning strategies they should use. The approach is novel because it allows systems to reason about themselves and make decisions that would normally be hard-coded into their programs by the designer, adding considerably to the power of such systems. This ability is central to a general theory of goal-driven learning in multistrategy reasoning and learning systems. To realize this ability, we have developed algorithms for learning and introspection, as well as representational methods using which a system can represent and reason about meta-models describing itself.

Acknowledgements

This research was supported by the National Science Foundation under grant IRI-9009710 and by the Georgia Institute of Technology. We would like to thank David Leake and several anonymous reviewers for their detailed comments on earlier drafts of this chapter.

References

[Anderson, 1989] J. Anderson. A Theory of the Origins of Human Knowledge. *Artificial Intelligence*, 30:313–351, 1989.

[Barsalou, 1991] L. W. Barsalou. 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, 1991.

[Birnbaum, 1986] L. Birnbaum. *Integrated Processing in Planning and Understanding*. Ph.D. thesis, Yale University, Department of Computer Science, New Haven, CT, December 1986. Research Report #489.

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

[Chi and VanLehn, 1991] M. T. H. Chi and K. A. VanLehn. The Content of Physics Self-Explanations. *The Journal of the Learning Sciences*, 1(1):69–105, 1991.

[Collins, Birnbaum, Krulwich, and Freed , 1993] G. Collins, L. Birnbaum, B. Krulwich, and M. Freed. The Role of Self-Models in Learning to Plan. In A. Meyrowitz and S. Chipman, editors, *Foundations of Knowledge Acquisition: Machine Learning*, pages 117–143. Kluwer Press, 1993.

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

[Cox and Ram, 1992a] M. T. Cox and A. Ram. An Explicit Representation of Forgetting. In *Proceedings of the Sixth International Conference on Systems Research, Informatics and Cybernetics*, Baden-Baden, W. Germany, 1992.

[Cox and Ram, 1992b] M. T. Cox and A. Ram. Multistrategy Learning with Introspective Meta-Explanations. In *Machine Learning: Proceedings of the Ninth International Conference*, Aberdeen, Scotland, July 1992. Morgan Kaufman, San Mateo, CA.

[Cox and Ram, 1994] M. T. Cox and A. Ram. Choosing Learning Strategies to Achieve Learning Goals. In *Proceedings of the AAAI Spring Symposium on Goal-Driven Learning*. AAAI Press, Menlo Park, CA, 1994.

[Dehn, 1989] N. Dehn. *Computer Story Writing: The Role of Reconstructive and Dynamic Memory*. Ph.D. thesis, Yale University, Department of Computer Science, New Haven, CT, 1989. Research Report #792.

[deKleer, Doyle, Steele, and Sussman , 1977] J. deKleer, J. Doyle, G. L. Steele, and G. J. Sussman. Explicit Control of Reasoning. *SIGPLAN Notices*, 12(8), 1977.

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

[Doyle, 1979] J. Doyle. A Truth Maintenance System. *Artificial Intelligence*, 12:231–272, 1979.

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

[Hammond, 1989] K. J. Hammond. *Case-Based Planning: Viewing Planning as a Memory Task*. Perspectives in Artificial Intelligence. Academic Press, Boston, MA, 1989.

[Hunter, 1990] L. E. Hunter. Planning to Learn. In *Proceedings of the Twelvth Annual Conference of the Cognitive Science Society*, pages 261–268, Boston, MA, July 1990.

[Kolodner and Simpson, 1984] J. L. Kolodner and R. L. Simpson. A Case for Case-Based Reasoning. In *Proceedings of the Sixth Annual Conference of the Cognitive Science Society*, Boulder, CO, 1984. Institute of Cognitive Science and University of Colorado, Boulder.

[Kolodner, 1993] J. L. Kolodner. *Case-Based Reasoning*. Morgan Kaufmann, San Mateo, CA, 1993. In press.

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

[McCarthy, 1958/1985] J. McCarthy. Programs with Common Sense. Reprinted in R. Brachman and H. Levesque, editors, *Readings in Knowledge Representation*, pages 299–308. Morgan Kaufman, San Mateo, CA, 1985.

[Michalski, 1993] R. S. Michalski. Inferential Theory of Learning: Developing Foundations for Multistrategy Learning. In R. S. Michalski and G. Tecuci, editors, *Machine Learning: A Multistrategy Approach, Volume IV*. Morgan Kaufman Publishers, Inc., 1993. To appear.

[Minsky, 1985] M. Minsky. *The Society of Mind*. Simon and Schuster, New York, NY, 1985.

[Minton, 1988] S. Minton. *Learning effective search control knowledge: An explanation-based approach*. Ph.D. thesis, Carnegie-Mellon University, Computer Science Department, Pittsburgh, PA, 1988. Technical Report CMU-CS-88-133.

[Mitchell, Utgoff, and Banerji , 1983] T. M. Mitchell, P. E. Utgoff, and R. Banerji. Learning by Experimentation: Acquiring and Refining Problem-Solving Heuristics. In R. S. Michalski, J. G.

Carbonell, and T. M. Mitchell, editors, *Machine Learning: An Artificial Intelligence Approach*, chapter 6, pages 163–189. Morgan Kaufman, Inc., Los ALtos, CA, 1983.

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

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

[Pirolli and Bielaczyc, 1989] P. Pirolli and K. Bielaczyc. Empirical Analyses of Self-Explanation and Transfer in Learning to Program. In *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society*, pages 450–457, Ann Arbor, MI, 1989.

[Pirolli and Recker, in press] P. Pirolli and M. Recker. Learning Strategies and Transfer in the Domain of Programming. *Cognition and Instruction*, in press.

[Provost and Buchanan, 1992] F. J. Provost and B. G. Buchanan. Inductive Policy. In *Proceedings of the Tenth National Conference on Artificial Intelligence*, pages 255–161, San Jose, CA, 1992.

[Quilici, to appear] A. Quilici. Toward Automatic Acquisition of an Advisory System's Knowledge Base. *Applied Intelligence*, to appear.

[Ram and Cox, 1994] A. Ram and M. T. Cox. Introspective Reasoning using Meta-Explanations for Multistrategy Learning. In R. S. Michalski and G. Tecuci, editors, *Machine Learning: A Multistrategy Approach, Volume IV*. Morgan Kaufman Publishers, Inc., 1994.

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

[Ram, Narayanan, and Cox , 1993] A. Ram, S. Narayanan, and M. T. Cox. *Multistrategy learning of diagnostic knowledge for a real-world troubleshooting task*. Technical Report GIT-CC-93/67, Georgia Institute of Technology, College of Computing, Atlanta, GA, 1993.

[Ram, 1989] A. Ram. *Question-driven understanding: An integrated theory of story understanding, memory and learning*. Ph.D. thesis, Yale University, Department of Computer Science, New Haven, CT, May 1989. Research Report #710.

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

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

[Redmond, 1992] M. Redmond. *Learning by Observing and Understanding Expert Problem Solving*. Ph.D. thesis, Georgia Institute of Technology, College of Computing, Atlanta, Georgia, 1992.

[Reich, 1993] Y. Reich. Macro and Micro Perspectives of Multistrategy Learning. In R. S. Michalski and G. Tecuci, editors, *Machine Learning: A Multistrategy Approach, Volume IV*. Morgan Kaufman Publishers, Inc., 1993. To appear.

[Scardamalia and Bereiter, 1991] M. Scardamalia and C. Bereiter. 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):37–68, 1991.

[Schank, 1982] R. C. Schank. *Dynamic Memory: A Theory of Learning in Computers and People*. Cambridge University Press, New York, NY, 1982.

[Schank, 1986] R. C. Schank. *Explanation Patterns: Understanding Mechanically and Creatively*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1986.

[Schneider, 1985] W. Schneider. Developmental Trends in the Metamemory-Memory Behavior Relationship: An Integrative Review. In D. L. Forrest-Pressley, G. E. MacKinnon, and T. G. Waller, editors, *Metacognition, Cognition, and Human Performance, Volume 1*. Academic Press, New York, NY, 1985.

[Stroulia, Shankar, Goel, and Penberthy , 1992] E. Stroulia, M. Shankar, A. Goel, and L. Penberthy. 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, June 1992.

[Sussman, 1975] G. J. Sussman. *A Computer Model Of Skill Acquisition*, volume 1 of *Artificial Intelligence Series*. American Elsevier, New York, 1975.

[Veloso and Carbonell, 1993] M. Veloso and J. Carbonell. Derivational Analogy: Automating Case Acquisition, Storage and Utilization. *Machine Learning*, 1993. To appear.

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

[Weintraub, 1991] M. A. Weintraub. *An Explanation-Based Approach to Assigning Credit*. Ph.D. thesis, The Ohio State University, Columbus, OH, 1991.

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

[Zukier, 1986] H. Zukier. The Paradigmatic and Narrative Modes in Goal-Guided Inference. In R. Sorrentino and E. Higgins, editors, *Handbook of Motivation and Cognition: Foundations of Social Behavior*, pages 465–502. Guilford Press, Guilford, CT, 1986.

Notes

[1]In general, a complete explanation of a reasoning failure would have all these components. In any given situation, however, the reasoner might only be able to construct a partial explanation, which would then determine what the reasoner can learn from that experience.

[2]Although learning is, in our view, a type of "reasoning", it is functionally distinct and is thus shown as a separate functional module in figure 1.

[3]TS stands for TroubleShooter.

[4]A typical printed circuit board manufacturing line has two major test and repair areas, the "in-circuit test" area and the "functional test" area. The in-circuit test area uses an automated tester on which a populated printed circuit board is placed for testing. The tester performs several test procedures on the board and outputs a ticket reading indicating the results of the test. If the board fails to meet the desired specifications, the tester may also provide some additional symptomatic information which can be used by the human operator in the troubleshooting process.

[5]AQUA stands for Asking Questions and Understanding Answers (see [Ram, 1991; Ram, 1993]).

Figures

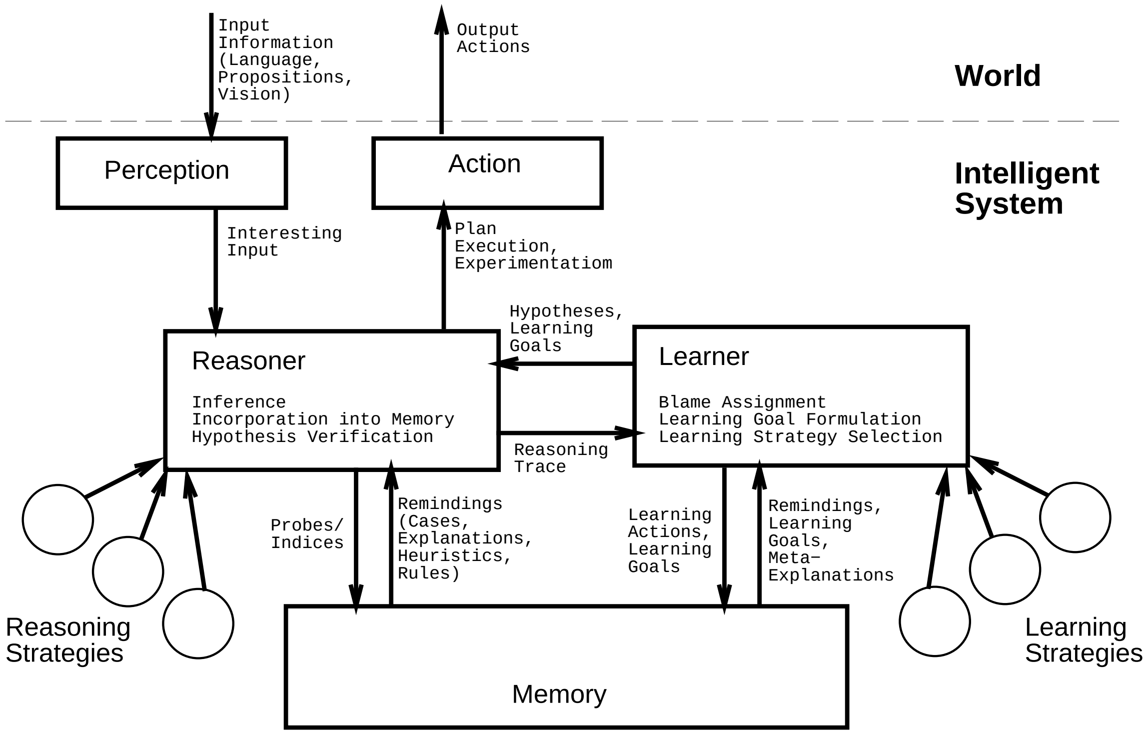

Figure 1: The Introspective Multistrategy Learning functional architecture.

Step 0: Perform and record reasoning in Trace Meta-XP (TMXP).

Step 1: Analyze TMXP to identify reasoning failures.

Step 2: If analysis reveals a failure, then learn:

Step 2A: Blame Assignment

- Compute index based on characterization of reasoning failure
- Retrieve Introspective Meta-XP (IMXP)
- Apply IMXP to reasoning trace represented in TMXP
- If IMXP application was successful, then:
  - Evaluate introspective explanation provided by IMXP
  - If IMXP not fully verified, perform introspective questioning and go to step 0

  else go to step 0

Step 2B: Learning Goal Formulation

- Create learning goals based on introspective analysis
- Compute tentative goal priorities

Step 2C: Learning Plan Formulation

- Choose learning algorithm(s)
- Expand subgoals
- Create learning plan
- Compute data dependencies
- Order learning actions

Step 2D: Learning Plan Execution

- Apply learning algorithm(s)

Step 3: If learning occurs, evaluate results of learning

Table 2: Conceptual algorithm for introspective multistrategy learning in a goal-driven learning system. Note that step 2C is not necessarily performed immediately after 2B; in some cases, it

may be performed at a later time. The reasoner may, for example, not have enough information to satisfy a learning goal at the time it is posted, or it may decide to suspend the goal because it has a lower priority than other reasoning goals that are currently being pursued. Suspended learning goals may be triggered opportunistically and/or pursued strategically at a later time.