Fifteenth Annual Conference of the Cognitive Science Society, pp. 17–26, Boulder, CO, June 1993.

Creative Conceptual Change

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

Abstract

Creative conceptual change involves (a) the construction of new concepts and of coherent belief systems, or theories, relating these concepts, and (b) the modification and extrapolation of existing concepts and theories in novel situations. I discuss these and other types of conceptual change, and present computational models of constructive and extrapolative processes in creative conceptual change. The models have been implemented as computer programs in two very different task domains, autonomous robotic navigation and fictional story understanding.

Introduction

Much research in conceptual change has focussed on developmental conceptual change in children, and scientific conceptual change in expert adults. Keil (1989), for example, is concerned with the nature of children's concepts, their differences from concepts that adults have, and how children's concepts change through cognitive development. Such conceptual change is qualitative; not only do children learn new concepts, the nature of the concepts themselves changes through development. The study of scientific conceptual change is concerned with how new conceptual structures come to replace existing conceptual structures through scientific revolutions (Kuhn, 1962) or through longer-term enterprise (Gruber, 1989). Nersessian (1991) argues that "the problem-solving strategies scientists have invented and the representational practices they have developed over the course of the history of science are very sophisticated and refined outgrowths of ordinary reasoning and representational processes."

The conceptual change that I am concerned with here is the everyday kind. It involves everyday reasoning by reasoning systems, human or machine, in situations that allow (or require) creativity and learning. Conceptual change requires two kinds of creative processes: the construction of new concepts from input information, and the extrapolation of existing concepts in novel and unfamiliar situations. The first kind of process involves reformulating low-level information, such as sensorimotor data, into higher-level abstractions. For example, a reasoner in a strange environment may improve its ability to act in that environment by learning about the effects of its actions in that environment (for example, learning to control a car on the highway). The actions themselves may be new and unfamiliar; a reasoner may need to learn about its own actions and the interactions of these actions with the environment (for example, learning to drive a car in the first place). The reasoner may also need to learn about the structure of the environment itself (for example, learning the layout of the roads in a city). All of these scenarios require creative conceptual change of a particular kind: the construction of conceptual representations to represent causal and predictive relationships between sensory inputs, motor actions, and the environment. I will call this *constructive conceptual change* since it involves the construction of new concepts from sensorimotor experience. Although this process is not usually thought of as "creative," I will argue that the process is in fact so because it results in representations that are novel, useful, and qualitatively different from those that the reasoner initially starts out with.

Another kind of process involved in creative conceptual change is that commonly associated with fictional and imaginative scenarios. Reading a science fiction story, for example, requires a temporary suspension of disbelief and the extension or adaptation of existing concepts to create a conceptual model of the described situation (which may be very different from the reasoner's real-world experience). I will call this *extrapolative conceptual change* since it involves extrapolation from existing concepts to create new ones. In addition to guiding the reasoner in the current situation, the new concepts (or systems of concepts) may be useful in other contexts as well. As I will argue, the mechanisms and knowledge involved in such reasoning are not unique to understanding fiction; they are really no different from the mechanisms and knowledge involved in reasoning in nonfictional or real-world situations. Although models of creativity and conceptual change have traditionally been developed separately from models from everyday reasoning, the constructive and extrapolative processes discussed here are not viewed as being extraordinary or special; they, and the creative conceptual change that they result in, are an integral part of everyday reasoning.

Both constructive and extrapolative conceptual change have much in common with each other, as well as with developmental and scientific conceptual change. Keil (1989) argues that systematic belief systems, or "theories," are important in developmental conceptual change, and that causal relations are essential and more useful in such theories than other sorts of relations (see also Neisser, 1987). Causal belief systems are critical in extrapolative conceptual change as well since they guide and constrain the creative adaptations performed by the reasoner. Keil views concepts as partial theories in that they embody explanations or mental models of the relations between their constituents, of their origins, and of their relations to other clusters of features (see also Johnson-Laird, 1983; Murphy & Medin, 1985). Similarly, the representations constructed through extrapolative and constructive conceptual change also embody such explanations (albeit not always "correct" ones). Analogy and mental modelling play a crucial role in theories of scientific conceptual change (e.g., Nersessian, 1991), and in extrapolative conceptual change as well. All these types of conceptual change rely both on inductive and analytical reasoning processes, though sometimes to different extents. Typically, analytical processes are used when appropriate theories are available to support analysis (such as in experts), and inductive processes are used when such theories are not available (such as in novices). In addition to the creation of individual concepts and their gradual evolution through experience, conceptual change may also involve the reorganization of an entire system of concepts.

The decomposition of the processes of conceptual change into constructive and extrapolative is a functional one. Rather than discuss conceptual change in children and adults, in laypersons and scientists, or in physics and mathematics, I will focus on the underlying *functions* of conceptual change (the construction and evolution of concepts), on the *mechanisms* that achieve these functions, and on the *knowledge* that these mechanisms rely on. Such a decomposition is methodologically useful because it allows us to study the types of knowledge and processes that underlie conceptual change and their commonalities across different performance tasks, domains, and levels of expertise of the reasoners. In this paper, I will discuss computational models of constructive and extrapolative conceptual change, focussing in particular on two computer programs that instantiate the models in two very different "everyday" task domains. The computer programs aid in the development and evaluation of the models, and provide an experimental framework for further exploration of theoretical ideas. I will conclude with a discussion of a framework for the integration of these (and other) methods of conceptual change into a single "multistrategy" system.

Case studies in creative conceptual change

The computer programs presented here serve as case studies of constructive and extrapolative processes in conceptual change. The first program, called SINS (Self-Improving Navigation System) is an autonomous robotic navigation system that learns to navigate in an obstacle-ridden world (Ram & Santamaria, 1993). Autonomous robotic navigation is the task of finding a path along which a robot can physically move through a given environment and then executing the actions to carry out the movement in a real or simulated world. The ability to adapt to changes in the environment, and to learn from experiences, is crucial to adequate performance and survivability in the real world. SINS uses fast robotic control augmented with multiple learning methods that allow the system to adapt to novel environments and to learn from its experiences. The core of the system is a constructive conceptual change mechanism that autonomously and progressively constructs representational structures that encapsulate the system's experiences. These structures comprise a higher-level representation of the system's perceptual and sensorimotor interactions with its environment, and are used to aid the navigation task in two ways: they allow the system to dynamically select the appropriate robotic control behaviors in different situations, and they also allow the system to adapt selected behaviors to the immediate demands of the environment.

The second case study is based on a computer program called ISAAC (Integrated Story Analysis And Comprehension), which is a natural language understanding system that reads short stories from the science fiction genre (Moorman & Ram, 1993). Such stories require creative understanding, in which the reader must learn enough about an alien world in a short text in order to accept it as the background for the story, and simultaneously must understand the story itself. ISAAC implements a process of extrapolative conceptual change which is based on the creative extrapolation, modification, or extension of existing concepts and theories to invent new ones. The extrapolation is constrained by the content of the story, by the system's existing concepts and theories, and by the requirements of the reading and understanding task.

As the case studies will reveal, there is much in common between these two systems despite their superficial differences. Both systems use multiple types of knowledge, and multiple types of reasoning processes. Both rely on multiple sources of constraints on these processes, including theories, knowledge and knowledge organization, and actual experience. Creative conceptual change in both systems is a process of gradual evolution of concepts to create better approximations of the observed world. Both systems learn autonomously through experience. The new concepts contribute significantly to the systems' abilities to carry out their respective tasks, and may be very different from those that the systems initially started out with.

The differences between the systems are also of interest. SINS relies directly on its experiences in the real world, whereas ISAAC's real world is that of natural language texts which vicariously describe fictional world experiences of fictional characters. ISAAC integrates its processes using explicit arbitration and control; thus, conceptual change in ISAAC is guided by the particular needs and goals of the program. SINS, in contrast, learns "automatically" through its task performance, and thus is better characterized as having an implicit orientation or goal to learn (Barsalou, discussed in Leake & Ram, 1993).

The two systems are discussed in more detail below.

Constructive conceptual change

Many machine learning and conceptual change systems have traditionally been used in problem domains that can be adequately described using discrete, symbolic representations. However, an important type of conceptual change is that which occurs in continuous problem domains. In order to actually perform a task in the real world, for example, an agent (human or robot) must be able to accept perceptual or sensory inputs from the environment, select an appropriate action based on its goals, the input, and the task at hand, and then carry out that action through appropriate motor control commands to its effectors. Perception and action are inherently continuous in three ways: they require representations of continuous information, they require continuous performance (for example, driving a car), and they require continuous adaptation and learning.

For example, consider the problem of spatial representation and exploration in a real-world environment. An agent learning about its physical environment through exploration might build a cognitive map representing topological and metrical information about the space around it. Several studies have suggested that cognitive maps are organized into layers (e.g., Lynch, 1960; Piaget & Inhelder, 1967; Siegel & White, 1975). The cognitive map contains information about space, locations, connectivity, and distance, learned gradually through interaction with and exploration of the environment. These studies have motivated computational models of robot map-learning as well. For example, Kuipers & Byun (1991) describe a simulated robot, NX, that learns a hierarchy of types of spatial knowledge organized into sensorimotor, control, procedural, topological, and metrical knowledge. At the lowest level, the robot has access to raw sensory data from the environment. The robot's representation of the space surrounding it undergoes a series of conceptual changes as sensorimotor data (which is continuous and numerical) is reformulated and abstracted into successively higher-level descriptions (which are discrete and symbolic). This is an example of what I am calling constructive conceptual change in this paper.

The SINS system discussed here also learns from continuous sensorimotor information, but addresses a somewhat different problem in constructive conceptual change: that of learning the appropriate concepts for dynamic and adaptive control of action. In addition to learning about the environment around it, an agent must also learn about the interactions of its behaviors with the environment. It must learn what effects its actions have and when different actions are appropriate. This problem is different from the map-learning problem because it involves constructing representations, not just of the environment, but of the agent's interactions with the environment. Often, action and learning are incremental of necessity because the agent's knowledge is limited and because the environment is unpredictable; the agent can at best execute the most promising short-term actions available to it and then re-evaluate its progress. An agent navigating in an unfamiliar environment, for example, may not know where obstacles lie until it actually encounters them. As the problems encountered become more varied and difficult, it becomes necessary to use available knowledge in an incremental manner to act, and to rely on continuous feedback from the environment to adapt actions and learn from experiences. The problem solving and learning process must operate continuously; there is no time to "stop and think," nor a logical point in the process at which to do so. Through this on-going process, the agent must construct higher-level conceptual representations that constitute its "understanding" of the world and of its interactions with the world.

SINS addresses this problem by constructing conceptual structures that encapsulate continuous sensorimotor experience. These structures are modified continuously even as they are used to guide action. Through experience, these structures evolve into stable perception-action models and result in improved performance on a wide range of input environments.

Technical details: The SINS system

Autonomous robotic navigation is defined as the task of finding a path along which a robot can move safely from a source point to a destination point in a obstacle-ridden terrain (path planning) and executing the actions to carry out the movement in a real or simulated world (plan execution). Several methods have been proposed for this task, ranging from high-level planning methods to reactive methods.

High-level planning methods use extensive world knowledge and inferences about the environment they interact with (e.g., Fikes, Hart & Nilsson, 1972; Sacerdoti, 1975). Knowledge about available actions and their consequences is used to formulate a detailed plan before the actions are actually executed in the world. These methods can successfully perform the path-finding required by the navigation task, but only if an accurate and complete representation of the world, and of available actions and their effects, is available to the agent. Situated or reactive control methods have been proposed as an alternative to high-level planning methods (e.g., Arkin, 1989; Brooks, 1986; Kaelbling, 1986; Payton, 1986). In these methods, no

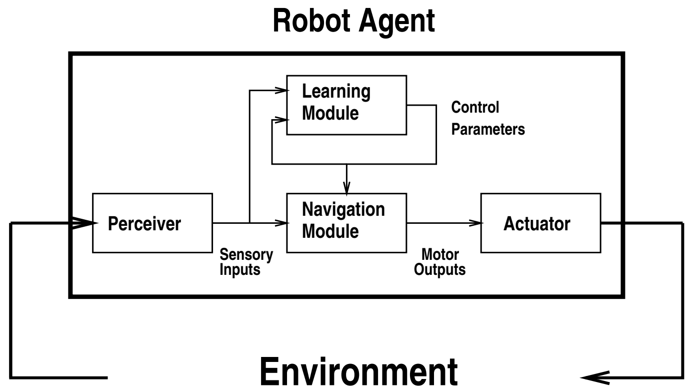

Figure 1: Architecture of the self-improving robot navigation system.

planning is performed; instead, a simple sensory representation of the environment is used to select the next action that should be performed. Actions are represented as simple behaviors, which can be selected and executed rapidly, often in real-time. These methods can cope with unknown and dynamic environmental configurations, but only those that lie within the scope of predetermined behaviors.

In a complex and dynamic environment, an agent needs to develop a combination of the above abilities: a fast and accurate perception process, the ability to reliably map sensory inputs to higher-level representations of the world, the ability to reliably predict the effects of its actions, and the ability to respond immediately to unexpected situations. Furthermore, to ensure adequate performance and survivability in the real world, the agent's ability to perform these functions must adapt to changes in the environment and improve through experience. In the SINS system, we have focussed on the problem of constructing representations of the agent's interactions with its environment. These representations model the environment and the effects of the agent's actions in that environment, and provide a basis for selecting appropriate actions in a possibly unfamiliar environment.

SINS uses schema-based reactive control for fast performance (Arkin, 1989), augmented with multistrategy learning methods that allow the system to adapt to novel environments and to learn from its experiences (see figure 1). The system autonomously and progressively constructs representational structures that encapsulate its experiences into "cases" that are then used to aid the navigation task in two ways: they allow the system to dynamically select the appropriate robotic control behaviors in different situations, and they also allow the system to adapt selected behaviors to the immediate demands of the environment (see figure 2).

The system's cases are automatically constructed using a hybrid case-based and reinforcement learning method without extensive high-level reasoning. The learning and navigation modules function in an integrated manner. The learning module is always trying to find a better model of the interaction of the system with its environment so that it can tune the navigation module to perform its function better. The navigation module provides feedback to the learning component so it can build a better model of this interaction. The behavior of the system is the result of an equilibrium point established by the learning module, which is trying to refine the model, and the

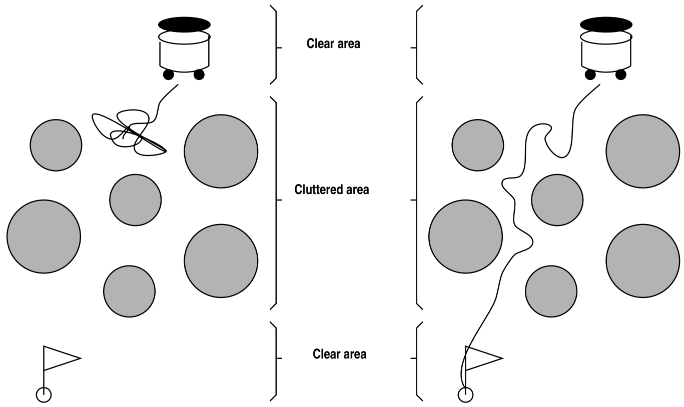

Figure 2: Typical navigational behaviors of the autonomous robotic system. The figure on the left shows the non-learning system with high obstacle avoidance and low goal attraction. On the right, the learning system has lowered obstacle avoidance and increased goal attraction, allowing it to "squeeze" through the obstacles and then take a relatively direct path to the goal.

environment, which is complex and dynamic in nature. This equilibrium may shift and need to be re-established if the environment changes drastically; however, the model is generic enough at any point to be able to deal with a very wide range of environments.

The learning methods are based on a combination of ideas from case-based reasoning and learning, which deals with the issue of using past experiences to deal with and learn from novel situations (e.g., Hammond, 1989), and from reinforcement learning, which deals with the issue of updating the content of system's knowledge based on feedback from the environment (e.g., Sutton, 1992). Each case in SINS represents an observed regularity between a particular environmental configuration and the effects of different actions, and prescribes the values of the control parameters that are most appropriate (as far as the system can determine based on its previous experience) for that environment.

The learning module performs the following tasks in a cyclic manner: (1) **perceive** and represent the current environment; (2) **retrieve** a case which represents an environment most similar to the current environment; (3) **adapt** the motor control parameters in use by the navigation module based on the recommendations of the case; and (4) **learn** new associations and/or adapt existing associations represented in the case to reflect any new information gained through the use of the case in the new situation to enhance the reliability of their predictions.

Since learning is not supervised by an outside expert, one of the issues to be addressed is how the system can determine whether the current experience should be used to modify and improve an existing case, or whether a new case should be created. In SINS, this is done through an inductive procedure that uses information about prior applications of the case. When a case is retrieved and applied to the current situation, a "relative similarity measure" is used to quantify how similar the current environment configuration is to the environment configuration encoded by the case, relative to how similar the environment has been in previous utilizations of the case. Intuitively, if a case matches the current situation better than previous situations it was used in, it is likely that the situation involves the very regularities that the case is beginning to capture; thus, it is worthwhile modifying the case in the direction of the current situation. Alternatively, if the match is not quite as good, the case should not be modified because that will take it away from the regularity it has been converging towards. Finally, if the current situation is a very bad fit to the case, it makes more sense to create a new case to represent what is probably a new class of situations.

A case in SINS represents a set of associations between sensory inputs and control parameters. Sensory inputs provide information about the configuration of the environment, and control parameters specify how to adapt the motor outputs of the navigation module in the environments to which the case is applicable. Each type of information is represented as a vector of analog values. Each analog value corresponds to a quantitative variable (a sensory input or a control parameter) at a specific time, and a vector of such values represents the trend or recent history of the corresponding variable. This representation has three essential properties. First, the representation is capable of capturing a wide range of possible associations between of sensory inputs and schema parameters. Second, it permits continuous progressive refinement of the associations. Finally, the representation captures trends or patterns of input and output values over time. This allows the system to detect patterns over larger time windows rather than having to make a decision based only on instantaneous values of perceptual inputs.

Sets of sensory inputs and control parameters are associated by grouping their vectors together into a single case. This grouping induces (albeit implicitly) a set of concepts that can be used to describe a control strategy or an environmental regularity. For example, if SINS is getting deeper into a crowded area, the values of the sensory inputs responsible for object detection will increase over time. A useful strategy in such a situation might be to back out and go around the obstacles. However, such a strategy cannot be expressed in purely perceptual terms; it requires the concepts of crowdedness, retreat, and so on, which are qualitatively different from the sensorimotor information that is initially available to the system.

Since learning and adaptation are based on a relative similarity measure, the overall effect of this process is to cause the cases to converge on stable associations between environment configurations and control parameters. Stable associations represent regularities in the world that have been identified by the system through its experience, and provide the predictive power necessary to navigate in future situations. The assumption behind this method is that the interaction between the system and the environment can be characterized by a finite set of causal patterns or associations between the sensory inputs and the actions performed by the system. The method allows the system to learn these causal patterns and to use them to modify its actions by updating its motor control parameters as appropriate.

One disadvantage of the analog representations is that they are not easy to interpret, making it difficult for a human observer to characterize the regularities and concepts that SINS actually learns in a given environment. To evaluate the method, we have developed a three-dimensional interactive visualization of a robot navigating through a simulated obstacle-ridden world, and used it to test the SINS system through extensive empirical simulations on a wide variety of environments using several different performance metrics. The system is very robust and can perform successfully in (and learn from) novel environments without any user intervention or supervisory input, yet it compares favorably with traditional reactive methods in terms of speed and performance (Ram & Santamaria, 1993). Furthermore, the system designers do not need to foresee and represent all the possibilities that might occur since the system develops its own "understanding" of the world and its actions.

SINS carries out a constructive conceptual change process in which new conceptual representations of regularities in system-environment sensorimotor interactions are created through experience. The process results in a qualitative shift in the system's internal "theory" of perception and action, and results in new concepts that are creative by virtue of being both original and useful (Koestler, 1964; Turner, 1991). As one might expect, the creation of new concepts in SINS (and in other systems such as NX) is an incremental process and involves, in addition to the abstraction of low-level inputs into higher-level representations, the modification of such representations in response to future experiences. In this sense, constructive conceptual change involves some degree of extrapolation as well. However, since this extrapolation does not require the kinds of creative leaps as those needed in the ISAAC system, the latter provides a better case study of extrapolative conceptual change and is discussed next.

Extrapolative conceptual change

In developing the SINS system, we were interested in the problem of constructing conceptual representations from continuous sensorimotor experience. Another type of conceptual change, however, is that which occurs when conceptual representations are used to understand a new and unfamiliar domain. The more different the domain, the more radical the change. In the ISAAC system, we are focussing on the construction of new concepts (and associated theories) through creative theory-guided transfer of existing concepts to a new domain. This process is largely analytical and involves analogical and metaphorical reasoning. There are two central issues here: what are the processes by which existing theories are extrapolated, and what is the nature of the constraints on these processes?

ISAAC explores these ideas in the domain of reading short stories from the science fiction literature. Consider the following short story, *Men Are Different* by Alan Bloch (1963).

> I'm an archaeologist, and Men are my business. Just the same, I wonder if we'll ever find out about Men—I mean *really* find out what made Man different from us Robots—by digging around on the dead planets. You see, I lived with a Man once, and I know it isn't as simple as they told us back in school.
>
> We have a few records, of course, and Robots like me are filling in some of the gaps, but I think now that we aren't really getting anywhere. We know, or at least the historians say we know, that Men came from a planet called Earth. We know, too, that they rode out bravely from star to star; and wherever they stopped, they left colonies—Men, Robots, and sometimes both—against their return. But they never came back.
>
> Those were the shining days of the world. But are we so old now? Men had a bright flame—the old word is "divine," I think—that flung them far across the night skies, and we have lost the strands of the web they wove.
>
> Our scientists tell us that Men were very much like us—and the skeleton of a Man is, to be sure, almost the same as the skeleton of a Robot, except that it's made of some calcium compound instead of titanium. Just the same, there are other differences.
>
> It was on my last field trip, to one of the inner planets, that I met the Man. He must have been the last Man in this system, and he'd forgotten how to talk—he'd been alone so long. I planned to bring him back with me. Something happened to him, though.
>
> One day, for no reason at all, he complained of the heat. I checked his temperature and decided that his thermostat circuits were shot. I had a kit of field spares with me, and he was obviously out of order, so I went to work. I pushed the needle into his neck to operate the cut-off switch, and he stopped moving, just like a Robot. But when I opened him up he wasn't the same inside. And when I put him back together I couldn't get him running again. Then he sort of weathered away—and by the time I was ready to come home, about a year later, there was nothing left of him but bones. Yes, Men are indeed different.

In order to understand this story, the reader must infer that the narrator is a robot, that robots are the dominant lifeform in the future, that humans have practically died out, that robots are capable of making logical errors such as the ones that the narrator made, and so on. The reader must construct an appropriate model of this world, and interpret the story with respect to this model even as the model evolves. The reader must also be willing to suspend disbelief to understand concepts which do not fit into a standard world view.

In ISAAC, new theories (and associated concepts) are constructed through extrapolation and modification of existing theories and concepts. The extrapolation is constrained by the actual content of the story, by the system's existing theories and concepts, and by the cognitive constraints on the reading and understanding mechanisms that are responsible for processing the story. No reader, machine or human, could have the time, memory, and other resources to read every single word in a story in-depth and to consider all the ramifications of each word. The reader's environment (the story), knowledge (existing concepts), goals and tasks (e.g., Ram & Hunter, 1992), and cognitive resources available to the processing machinery (e.g., Just & Carpenter, 1992) interact to constrain the possible extrapolation to a more manageable level.

The story understanding processes in ISAAC are not unique to science fiction stories, of course. Understanding any fictional story requires similar kinds of processing. The same is true of nonfictional stories as well as unfamiliar real-world scenarios, although the types and degree of conceptual modifications required may be different.

Technical details: The ISAAC system

The ISAAC system consists of six "supertasks," each of which is made up of several subtasks that interact with each other. The tasks are based on research in psycholinguistics (e.g., Holbrook, Eiselt & Mahesh, 1992; van Dijk & Kintsch, 1983), reading comprehension (e.g., Black & Seifert, 1981; Graesser, Golding, & Long, 1991), story understanding (e.g., Birnbaum, 1986; Ram, 1991; Rumelhart, 1977), episodic memory (e.g., Kolodner, 1984; Schank, 1982), analogy (e.g., Falkenhainer, 1987; Gentner, 1989), creativity (e.g., Gruber, 1989; Schank

& Leake, 1990), and metacognition (e.g., Gavelek & Raphael, 1985; Schneider, 1985; Weinert, 1987; Wellman, 1985). The supertasks and their functions are summarized below.

Language understanding is responsible for low-level text understanding, including lexical retrieval, syntactic parsing, pronoun reference, punctuation analysis, and tense analysis.

Story structure understanding focusses on details of the text which relate to story structure, including character identification (protagonist, antagonist), setting identification (time, location), plot description, and genre identification.

Episodic understanding carries out the event representation (agent, action, state, object, location), agent modelling (agents' goals, knowledge, and beliefs), and action modelling tasks that are central to understanding fictional, narrative or real-world episodes.

Explanation and reasoning is responsible for high-level reasoning and learning tasks, including those supporting specific language understanding tasks such as unknown word definition, and general tasks such as belief management, inference, creative analogy, interest management, and learning.

Memory management carries out memory storage and retrieval, including spontaneous reminding and case construction.

Metacontrol is responsible for integration of the other supertasks, and for focus of attention, time management, and suspension of disbelief. Since it is unreasonable to assume that the system would have complete metacognitive access to all its internal processes (Nisbett & Wilson, 1977), metacontrol and metareasoning operate on supertasks and do not access the individual tasks directly. The supertasks in turn control the individual tasks that they are responsible for.

We chose science fiction stories as the domain for ISAAC because it is a particularly good one to study what one might call "creative understanding." People can comprehend stories which have no basis in fact, and which may require invention of concepts and theories which are radically different from those in the real world. The process of understanding the un-understandable involves the extrapolative type of creative conceptual change. A central requirement is the willingness of the reader to suspend his or her disbelief of the material being presented or the assumptions being made about the fictional world (Corrigan, 1979). Consider the ambiguous title of a Larry Niven (1973) story, *Flight of the Horse*. This phrase could refer to a fleeing horse, a horse on an airplane flight, or to a flying horse. If a story understanding system relied on a belief in the validity of world knowledge, it would disambiguate the phrase to eliminate the latter meaning since it "knows" horses cannot fly. This may be incorrect if the story was about a flying horse (or a pegasus), which is perfectly reasonable in a science fiction or mythological story. As I argued earlier, these considerations are not unique to science fiction stories; even factual stories (such as newspaper stories) in domains that are not completely understood may require the system to consider the possibility that its current understanding of the domain is incomplete or incorrect (e.g., Ram, 1993).

To understand concepts which do not fit into a standard world view, the system attempts to modify existing concepts (Schank, 1986). This usually involves extending or adapting not just a single concept, but systems of concepts—that is, theories. This modification can occur in several ways. Definitional constraints may be relaxed to produce concepts with alternative constraints. For example, relaxing the definitional constraint that a horse's primary mode of locomotion is its legs may result in a "horse" with wings—a pegasus. Another option is to add new constraints or features to existing concepts, or to combine two concepts together. Suitcases, for example, do not normally have a mode of locomotion; adding one may result in an independently mobile suitcase, much like the one depicted in Terry Pratchett's (1983) story, *The Colour of Magic*. Creativity may also result from relaxed constraints on memory search processes, such as in the "imaginative memory" of Turner's (1991) MINSTREL system.

	Physical	Mental	Social	Emotional	Temporal
Agents	person	consciousness	boss	Ares	entropy
Actions	walking	thinking	selling	loving	getting closer to March
Objects	rock	idea	teacher-student relationship	hatred	second
States	young	lack of knowledge	public dishoner	being angry	early

Figure 3: Knowledge representation grid.

A problem with such concept manipulation is that it is difficult to specify principled constraints on this process. Could a toaster be a good mode of horse locomotion? Up to a certain limit, constraint manipulation will result in concepts which could be called creative, after which the resulting concepts may be too bizarre to be useful. However, utility and interestingness are not inherent in particular concepts, but can only be evaluated with respect to the reasoner's knowledge, the organization of this knowledge, the reasoner's goals, the task at hand, the environment in which the reasoner is carrying out its tasks (in the case of ISAAC, the story), and general processing heuristics (Pinto, Shrager & Berthenthal, 1992; Ram, 1990).

In ISAAC, the knowledge organization scheme provides a structure for the conceptual change process. ISAAC's knowledge base is organized into a semantic network, which is indexed through a multidimensional grid (see figure 3). The rows of the grid represent "thematic roles" for adaptation; for concepts representing events, these include action, agent, state, and object. The columns of the grid represent "conceptual domains," such as physical, mental, social, emotional, and temporal. For example, a transfer is a generic action. Different types of transfers can be represented as physical (e.g., the PTRANS primitive of Schank, 1972), mental (e.g., MTRANS), and social (e.g., ATRANS). The grid also allows the system to represent emotional and temporal transfers (see also Domeshek, 1992).

Concept extrapolation is accomplished by moving around the grid, leading to creative and metaphorical interpretations of known concepts. Each type of movement incurs a cost to the system, depending on the degree to which the concept has been altered. Movement within a single cell is the easiest type to perform, movement along a single row or a single column is more difficult, and adaptations requiring movement across both rows and columns are the most difficult. Although the details of the grid are still under development, the point is that the system tries to perform the least amount of adaptation necessary, guided by the grid, such that the resulting concepts can explain and provide a structure for the input.

For example, many temporal metaphors can be represented as analogies between the physical and temporal columns of the grid (Lakoff & Johnson, 1980). In a sentence such as "Time has passed her by," for example, a temporal event is described in physical terms, and an abstract object (time) is described as the agent of the physical action. Similarly, in the second paragraph of *Men Are Different*, "we aren't really getting anywhere" is a metaphorical use of knowledge of physical actions to describe a mental action. Such a metaphor requires a larger creative leap than an adaptation within the physical column alone, such as in Schank's (1986) example in which an analogy is drawn between a jogger and a racehorse. Continuing with the earlier horse locomotion examples, a horse with wings involves an adaptation in which a known mode of locomotion (wings) is substituted for another one (legs), and is less bizarre than an independently mobile suitcase with wings in which an inanimate object is viewed as an animate agent with an invented (but plausible) mode of locomotion where none existed previously. As before, however, utility and interestingness are not absolute; a suitcase with wings (perhaps airplane wings rather than bird wings) might make sense in the right context.

In *Men Are Different*, robots, which in the real world are physical objects used as tools in manufacturing, are conceptualized as independent volitional agents. The reader must adopt this view to build an appropriate story model. Interestingly, the irony in this story derives from the fact that the robot in the story performs what one might view as the reverse inference, conceptualizing the man as a physical object to be repaired in a manner that one might use to repair a physical robotic device. It is important to note that the invented concepts are "real" within the context of the story, in contrast to the "bright flame of Men" which is metaphorical even within the fictional world. Similarly, a sentence such as "Winter is rapidly approaching" uses a spatial metaphor to describe a temporal event, whereas time travel may in fact be a "real" concept in a story. Understanding this concept involves adapting knowledge about actions, states, and causality from the physical column of the grid to the temporal. Such adaptation is the heart of the extrapolative conceptual change process. Once the new concepts and theories are built, they can be used to understand the story within the framework of these concepts; in turn, this may result in further modification of the concepts.

In addition to aiding in the story comprehension process, the new concepts and theories can also provide a basis for future problem solving in the real world (e.g., Koestler, 1964). For example, reading about a fictitious device may prompt the reader to develop a similar device in the real world, or may help the reader understand a similar device when it is actually encountered at some later point. Motorola's MicroTAC hand-held personal cellular phone, for instance, has a strong resemblance to the hand-held personal communicators used in the *Star Trek* television series. Goodman, Waterman & Alterman's (1991) SPATR system uses a similar case-based reasoning process to understand novel devices (such as an Airphone) and natural language instructions for using these devices based on hierarchical spatial models of known devices (such as an ATM). Reading about a creative problem solving episode may also allow the reader to replay the observed solution process on a real-world problem in a manner similar to Carbonell's (1986) derivational analogy.

Stories that are not creative can also be understood and used in such ways, of course. The mechanisms of conceptual change discussed here are an integral part of ordinary reasoning. Creative understanding in ISAAC is not implemented through a separate "creativity" process, but rather through normal processes of reasoning and learning (Gruber, 1989). Similarly, conceptual change in SINS also occurs through the normal processes of perception and control of action. Everyday reasoning is robust, adaptive, and creative; no special process need be postulated to model or explain these capabilities.

Discussion

On the surface, the models of constructive and extrapolative conceptual change presented above appear very different. The SINS model is inherently experiential and can be characterized as constructive induction of representations from sensorimotor input, whereas the ISAAC model is based on vicarious experience and can be characterized as theory-guided transfer of concepts to a new domain. The former is mostly inductive, whereas the latter is mostly analytical. In ISAAC, multiple processes are integrated through (some degree of) explicit arbitration and control; in SINS, the processes are automatic and the integrative control mechanisms are implicit.

It is an open question how these models which are, in some sense, at opposite ends of the spectrum of creative conceptual change might be unified into a single framework. Quine (1977) suggests that early concepts may be more perceptual, being defined inductively using an "innate similarity notion or spacing of qualities," and later concepts may become more "scientifically sophisticated," conceptual, and theory-embedded (see also Keil, 1989). Quine was interested in the issue of development of natural kinds, but perhaps a similar idea could be used to integrate perceptual and conceptual change in an "adult" reasoning system.

To facilitate integration, it is useful to look at commonalities between the models. Although the SINS model is closer to actual perceptual features in real world and the ISAAC model is closer to theories and mental models, both are based in real experience (whether personal or vicarious), and are constrained by the interaction between the system and the environment. Both are creative processes, and result not just in learning but in conceptual change as well. In SINS, raw sensorimotor information is encapsulated into predictive perception-action models, and in ISAAC, existing theories are modified to provide a belief structure for new and unfamiliar concepts. Both require inductive and analytical processes (although to different degrees), and both combine multiple methods of learning, concept formation, and conceptual change. Both are based on multiple types of knowledge. In both, existing knowledge provides constraints on reasoning and learning processes. Both types of creative conceptual change model a gradual evolution of concepts to better approximate the observed world and in both, evolving concepts are used in the performance task even as they are modified. These points also highlight many of similarities between the models of constructive and extrapolative conceptual change presented here and other models of conceptual change, including models of developmental and scientific conceptual change.

One framework for integration of these (and other) methods of conceptual change is through a multistrategy learning model, in which various learning methods are combined into a unified framework. Recent attention to such models is evident in machine learning (e.g., Carbonell, Knoblock & Minton, 1991; Michalski & Tecuci, 1993) and cognitive psychology (e.g., Anderson, 1983; Wisniewski & Medin, 1991). Multistrategy approaches provide the flexibility and power required in practical, real-world domains.

There are several methods of integrating multiple learning algorithms into a single system (see Michalski & Tecuci, 1993). One such framework is that used in the Meta-AQUA and Meta-TS systems (Ram & Cox, 1993; Ram, Cox & Narayanan, 1992). In this model, the reasoning system actively selects and combines learning methods based on an analysis of its learning goals which are represented explicitly in the system. Some learning goals may be low-level and always active, such as in SINS and NX. These systems can be described as performing "goal relevant" learning, in that learning is relevant to the overall goals of the system (Thagard) but the system only has an implicit goal to learn (Barsalou, both discussed in Leake & Ram, 1993). Other learning goals may be selected based on a higher-level analysis of utility of knowledge and relevance to the system's tasks, such as in ISAAC, Meta-AQUA, IVY (Hunter, 1990), and PAGODA (desJardins, 1992). These systems are better described as "goal directed" since goals are explicitly represented and used to drive the selection and execution of reasoning and learning strategies (Leake & Ram, 1993; Ram & Cox, 1993; Ram & Hunter, 1992).

Although I do not want to suggest that humans have perfect or conscious metacognitive knowledge of, and control over, their learning processes, such a model could be used to take an intentional stance (Dennett, 1987) towards a computational theory of multistrategy reasoning, both as a description of human reasoning processes and as a basis for the design of creative AI systems. Meta-TS, for example, implements a computational model of human operators learning to troubleshoot physical devices (Narayanan & Ram, 1992). The model is based on observations of human troubleshooting operators and protocol analysis of the data gathered in the test area of an operational electronics assembly manufacturing plant. In Meta-TS, multiple learning methods for knowledge compilation (Anderson, 1989), interactive transfer of expertise (Davis, 1979), postponement (Ram, 1991, 1993), and forgetting are integrated through metacognitive analysis. Experimental results in metacognition also suggest that such analysis can facilitate reasoning and learning (e.g., Alexander, 1992; Carr, 1992; Schneider, 1985; Weinert, 1987). An open issue in the design of such models is the integration of "automatic" learning strategies, such as those used in SINS, that are goal-relevant rather than explicitly goal-directed.

In conclusion, creative conceptual change is an everyday process involving multiple integrated mechanisms that are constrained by existing knowledge and by the task at hand. This process is situated in, and therefore also constrained by, the real world, and results in original, useful and qualitatively different representations of systems of beliefs. The process involves the on-going construction and extrapolation of concepts and theories in the context of, in service of, and in response to a real-world performance task. The constructive and extrapolative processes are modelled computationally through specification of functions (tasks), mechanisms and knowledge; these models are then instantiated as computer programs and evaluated empirically. In this paper, I have used robotic navigation in dynamic environments and comprehension of actual science fiction short stories as the task domains in which to present two case studies of creative conceptual change. These case studies highlight the issues involved in conceptual change, provide a basis for the development and evaluation of models that address these issues, and raise several new issues for future research.

Acknowledgements

Juan Carlos Santamaria and Kenneth Moorman contributed to the research and program implementations described here. I would also like to thank Mike Cox and Eric Domeshek for their comments on an earlier draft of this paper.

References

Alexander, J. (1992). Metacognition and Giftedness. Paper presented at *The SouthEast Cognitive Science Conference*, abstracted in Technical Report #1, Cognitive Science Program, Georgia Institute of Technology, Atlanta, GA.

Anderson, J.R. (1983). *The Architecture of Cognition*. Harvard University Press, Cambridge, MA.

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

Arkin, R.C. (1989). Motor Schema-Based Mobile Robot Navigation. *The International Journal of Robotics Research*, 8(4):92–112.

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

Black, J.B. & Seifert, C.M. (1981). The Psychological Study of Story Understanding. Technical Report #18, Yale University, New Haven, CT.

Brooks, R. (1986). A Robust Layered Control System for a Mobile Robot. *IEEE Journal of Robotics and Automation*, RA-2(1):14–23.

Bloch, A. (1963). Men Are Different. In I. Asimov & G. Conklin, editors, *50 Short Science Fiction Tales*, Macmillan Publishing Company, New York.

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*, Morgam Kaufman Publishers, San Mateo, CA.

Carbonell, J.G., Knoblock, C.A., & Minton, S. (1991). PRODIGY: An Integrated Architecture for Planning and Learning. In K. Van Lehn, editor, *Architectures for Intelligence*, pages 241–278, Lawrence Erlbaum Associates, Hillsdale, NJ.

Carr, M. (1992). Metacognitive Knowledge as a Predictor of Decomposition Strategy Use. Paper presented at *The SouthEast Cognitive Science Conference*, abstracted in Technical Report #1, Cognitive Science Program, Georgia Institute of Technology, Atlanta, GA.

Corrigan, R.W. (1979). *The World of the Theatre*. Scott, Foresman and Company, Glenview, IL.

Davis, R. (1979). Interactive Transfer of Expertise: Acquisition of New Inference Rules. *Artificial Intelligence*, 12:121–157.

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.

Domeshek, E.A. (1992). *Do the Right Thing: A Component Theory for Indexing Stories as Social Advice*. Ph.D. thesis, Yale University, Department of Computer Science, New Haven, CT.

Falkenhainer, B. (1987). Scientific Theory Formation through Analogical Inference. In *Proceedings of the Fourth Inter-*

national Workshop on Machine Learning, Morgan Kaufman Publishers, Los Altos, CA.

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

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.

Gentner, D. (1989). Mechanisms of Analogical Learning. In S. Vosniadou & A. Ortony, editors, *Similarity and Analogical Reasoning*, Cambridge University Press, London.

Goodman, M., Waterman, S., & Alterman, R. (1991). Interactive Reasoning about Spatial Concepts. *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*, pages 734–738, Lawrence Erlbaum Associates, Hillsdale, NJ.

Graesser, A., Golding, J.M., & Long, D.L. (1991). Narrative Representation and Comprehension. In R. Barr, M.L. Kamil, P. Mosenthal & P.D. Pearson, editors, *Handbook of Reading Research*, volume 2, chapter 8, Longman Publishing Group, White Plains, NY.

Gruber, H.E. (1989). The Evolving Systems Approach to Creative Work. In D.B. Wallace & H.E. Gruber, editors, *Creative People at Work*, pages 3–24, Oxford University Press, New York.

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

Holbrook, J.K., Eiselt, K.P., & Mahesh, K. (1992). A Unified Process Model of Syntactic and Semantic Error Recovery in Sentence Understanding. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, pages 195–200, Lawrence Erlbaum Publishers, Hillsdale, NJ.

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

Johnson-Laird, P.N. (1983). *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*. Harvard University Press, Cambridge, MA.

Just, M.A. & Carpenter, P.A. (1992). A Capacity Theory of Comprehension: Individual Differences in Working Memory. *Psychological Review*, 99(1):122–149.

Kaelbling, L. (1986). *An Architecture for Intelligent Reactive Systems*. Technical Note #400, SRI International.

Keil, F.C. (1989). *Concepts, Kinds, and Cognitive Development*. MIT Press, Cambridge, MA.

Koestler, A. (1964). *The Act of Creation*. MacMillan Publishers, New York.

Kolodner, J.L. (1984). *Retrieval and Organizational Strategies in Conceptual Memory: A Computer Model*. Lawrence Erlbaum Associates, Hillsdale, NJ.

Kuhn, T.S. (1962). The Structure of Scientific Revolutions, University of Chicago Press, Chicago, IL.

Kuipers, B.J. & Byun, Y.-T. (1991). A Robot Exploration and Mapping Strategy Based on a Semantic Hierarchy of Spatial Representations. *Robotics and Autonomous Systems*, 8(1-2):47–63.

Lakoff, G. & Johnson, M. (1980). *Metaphors We Live By*. University of Chicago Press, Chicago, IL.

Leake, D. & Ram, A. (1993). *Goal-Driven Learning: Fundamental Issues and Symposium Report*, Technical Report #85, Indiana University, Cognitive Science Program, Bloomington, IN.

Lynch, K. (1960). *The Image of the City*. MIT Press, Cambridge, MA.

Michalski, R.S. & Tecuci, G. (1993), editors. *Machine Learning: A Multistrategy Approach, Volume IV*, Morgan Kaufman Publishers, San Mateo, CA. (In press.)

Moorman, K. & Ram, A. (1993). A New Perspective on Story Understanding. In *Proceedings of the Thirty-First Southeast ACM Conference*, Birmingham, AL.

Murphy, G.L. & Medin, D.L. (1985). The Role of Theories in Conceptual Coherence. *Psychological Review*, 92:289–316.

Narayanan, S. & Ram, A. (1992). Learning to Troubleshoot in Electronics Assembly Manufacturing. In *Proceedings of the ML-92 Workshop on Integrated Learning in Real-World Domains*, Ninth International Machine Learning Conference, University of Aberdeen, Scotland.

Neisser, U. (1987), editor. *Concepts and Conceptual Development: Ecological and Intellectual Factors in Categorization*. Cambridge University Press, 1987.

Nersessian, N.J. (1991). How Do Scientists Think? Capturing the Dynamics of Conceptual Change in Science. In R.N. Giere, editor, *Minnesota Studies in the Philosophy of Science, Volume XV (Cognitive Models of Science)*, University of Minnesota Press, Minneapolis, MN.

Nisbett, R.E. & Wilson, T.D. (1977). Telling More Than We Can Know: Verbal Reports on Mental Processes. *Psychological Review*, 84(3).

Niven, L. (1973). *The Flight of the Horse*. Ballantine Books, New York.

Payton, D. (1986). An Architecture for Reflexive Autonomous Vehicle Control. In *Proceedings of the 1986 IEEE Conference on Robotics and Automation*, pages 1838–1845, IEEE.

Piaget, J. & Inhelder, B. (1967). *The Child's Conception of Space*. Norton Publishers, New York.

Pinto, J., Shrager, J. & Berthenthal, B.I. (1992). Developmental Changes in Infants' Perceptual Processing of Biomechanical Motions. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, pages 60–65, Lawrence Erlbaum Associates, Hillsdale, NJ.

Pratchett, T. (1983). *The Colour of Magic*. Penguin Books, New York.

Quine, W.V.O. (1977). Natural Kinds. In S.P. Schwarts, editor, *Naming, Necessity, and Natural Kinds*, Cornell University Press, Ithaca, NY.

Ram, A. (1990). Knowledge Goals: A Theory of Interestingness. In *Proceedings of the Twelvth Annual Conference of the Cognitive Science Society*, pages 206–214, Lawrence Erlbaum Associates, Hillsdale, NJ.

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. (In press.)

Ram, A. & Cox, M.T. (1993). 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. (In press.)

Ram, A., Cox, M.T., & Narayanan, S. (1992). An Architecture for Integrated Introspective Learning. In *Proceedings of the ML-92 Workshop on Computational Architectures for Supporting*

Machine Learning and Knowledge Acquisition, Ninth International Machine Learning Conference, University of Aberdeen, Scotland.

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. & Santamaria, J.C. (1993). A Multistrategy Case-Based and Reinforcement Learning Approach to Self-Improving Reactive Control Systems for Autonomous Robotic Navigation. In *Proceedings of the Second International Workshop on Multistrategy Learning*, Center for Artificial Intelligence, George Mason University, Fairfax, VA. (To appear.)

Rumelhart, D.E. (1977). Understanding and Summarizing Brief Stories. In D.L. Berge and J. Samuels, editors, *Basic Processes in Reading and Comprehension*, Lawrence Erlbaum Associates, Hillsdale, NJ.

Sacerdoti, E.D. (1975). *A Structure for Plans and Behavior*. Technical Note #109, SRI International. Summarized in P.R. Cohen & E.A. Feigenbaum, *Handbook of AI, Volume III*, pages 541–550.

Schank, R.C. (1972). Conceptual Dependency: A Theory of Natural Language Understanding. *Cognitive Psychology*, 3(4):552–631.

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

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

Schank, R.C. & Leake, D.B. (1990). Creativity and Learning in a Case-Based Explainer. In J.G. Carbonell, editor, *Machine Learning: Paradigms and Methods*, MIT Press, Cambridge, MA.

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

Siegel, A.W. & White, S.H. (1975). The Development of Spatial Representations of Large-Scale Environments. In H.W. Reese, editor, *Advances in Child Development and Behavior*, Academic Press, New York.

Sutton, R.S. (1992), editor. *Machine Learning*, 8(3/4), special issue on Reinforcement Learning.

Turner, S.R. (1991). A Case-Based Model of Creativity. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*, pages 933-937, Lawrence Erlbaum Associates, Hillsdale, NJ.

van Dijk, T.A. & Kintsch, W. (1983). *Strategies of Discourse Comprehension*. Academic Press, New York.

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.

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.

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.