Cognition as Case-Based Reasoning

David B. Leake

Computer Science Department, Indiana University

Lindley Hall 215, Bloomington, IN 47405

leake@cs.indiana.edu

September 27, 1996

To appear in Bechtel, W., and Graham, G., eds, *The Blackwell Companion to Cognitive Science*.

Introduction

In some views of cognitive science, episodic memory is simply one of many pieces of the puzzle of cognition: memory, reasoning and learning can be studied largely independently from each other. The case-based reasoning (CBR) stance takes the contrasting position that human reasoning, memory, and learning are inextricably bound together. In this view, reasoning is primarily based on remembering and reapplying the lessons of prior episodes. A case-based problem-solver, for example, solves new problems by retrieving traces of relevant prior problems from memory, establishing correspondences between those problems and the new situation, and adapting the prior solutions to fit the problem at hand. Learning is a natural side-effect of this process, as traces of new episodes are stored for future use.

The case-based model of reasoning and learning provides an account of how people can reason effectively even in complex situations and using imperfect information. For example, consider the problem faced by a teacher deciding how to help a troubled student. The decision must be made despite the lack of guaranteed rules for successful intervention, and despite imperfect information on the student's state of mind and the myriad outside influences on the student. Nevertheless, if the teacher is experienced, he or she can make a decision based on prior cases of students with similar problems. By remembering what did—and did not—work with them, and why, the teacher can reuse successful strategies and avoid repeating past mistakes. The teacher can also learn, by remembering the current situation, decision, and results, how to handle similar problems in the future.

Observations of human reasoning from experience have prompted interest in the CBR process both from artificial intelligence practitioners aiming to develop computer systems with intelligent behavior, and from cognitive scientists aiming to understand human reasoning and learning. Understanding human case-based reasoning requires addressing questions about multiple research areas of cognitive science. For example, case-based reasoning is fundamentally analogical: it solves new problems and interprets new situations by analogy to prior episodes (see ANALOGY in this volume). Likewise, case-based reasoning is intimately connected to areas such as memory and similarity. However, the case-based reasoning stance involves investigating these areas from a new perspective. First, it examines each process as part of an integrated whole, in context of the others. Second, it calls for studying them in the pragmatic context of large-scale real-world tasks, such as planning and explanation. This requires analyzing not only the reasoning processes involved, but also the structure, content, and organization of the knowledge required for those processes to succeed for real-world tasks.

The case-based reasoning stance points to three central research questions for cognitive science. The first question concerns *knowledge access,* or *indexing*—how memory is organized and accessed to retrieve relevant information. The second concerns *knowledge adaptation*—how relevant episodes are modified to apply to a new situation. The third concerns *learning*—how lessons of new experiences are extracted, represented and organized.

Each of the three questions addresses how people perform complex cognition for real-world tasks, using their real-world knowledge. This makes them very difficult to answer with laboratory experiments. However, case-based reasoning research is beginning to provide answers to each one, through a methodology of developing computer models whose designs are strongly influenced by observations and studies of human reasoning and learning processes, but whose development is largely guided by examining functional requirements for the processes and knowledge involved.

For example, one way that case-based reasoning research attacks the indexing problem is to analyze the types of reasoning tasks that cases can serve, the circumstances under which cases are likely to be useful, and the indices required for them to be retrieved at the right times. The resulting theories of indexing are tested and refined by being implemented in computer models whose behavior can be examined. Empirical results from the computer models provide information about the theories' functional sufficiency and about gaps in the theories that must be addressed. Computer models of human case-based reasoning also suggest testable hypotheses about how people learn and how to structure education to aid human learning. For example, empirical evaluations of the learning process of Michael Redmond's CELIA (Redmond, 1992)—a computer model of how problem-solving changes with experience as cases and other types of knowledge are acquired—predict the relative role of cases and domain models in learning, the relative speed with which each body of knowledge grows, and the value of presenting a range of examples early on in the learning process.

Taking case-based reasoning seriously as a human reasoning process suggests that education should be aimed towards facilitating students' natural learning through case-based reasoning, as well as towards teaching students to perform case-based reasoning as effectively as possible. A strong new direction in case-based reasoning research focuses on developing innovative curricula and classroom activities built on the lessons of case-based reasoning research.

This chapter provides a brief overview of the case-based reasoning stance towards cognitive science, of its development, and of some of the research directions it has engendered. Janet Kolodner (1994) provides an extensive discussion of case-based reasoning as a cognitive model and presents important computer models of human case-based reasoning. Kolodner's (1993) textbook provides a thorough examination of case-based reasoning both as a cognitive model and as an artificial intelligence technology inspired by observations of human reasoning . Case studies of selected computer models for multiple tasks are presented in volumes edited by Christopher Riesbeck and Roger Schank (1989). A volume edited by David Leake (1996) provides an overview of the principles of case-based reasoning research, a tutorial, and case studies addressing key issues.

Reasoning from Remindings

Remindings are commonplace in everyday human reasoning. In children and adults, and in novices and experts, across a broad spectrum of cognitive tasks, memories of previous episodes provide useful information to guide understanding, problem-solving and prediction.

Previous episodes may be used in a straightforward way for routine processing of simple situations. For example, once when taking my two-year-old on a walk, I asked him where he would like us to turn at an intersection. He picked the direction we had turned the day before to go to the supermarket, and explained his decision by his memory of the previous trip: "I have a memory: Buy donut."

Memories of previous episodes can also guide the reasoning of experts performing difficult tasks. Radiologists, for example, use records of previous treatment cases as a basis for choosing X-ray treatments; lawyers argue and judges reach verdicts by comparing current circumstances to those of precedent cases (see COGNITIVE SCIENCE IN THE REAL WORLD: LAW); politicians shape their strategies according to memories of their prior successes and failures; governmental policy-makers weigh new plans in terms of the experiences of the past. When people are confronted with complex situations, similar previous episodes provide a basis for new reasoning. When the stock market crashed in 1987, the *New York Times'* analysis of the crash was framed by the question, "Is 1987 like 1929?" When Americans debate the wisdom of intervening in a foreign country, the debate is often framed by the question, "Will this be another Vietnam?"

Reasoning from prior episodes has a number of functional benefits compared to traditional rule-based models of reasoning. First, it helps for generating effective solutions in situations whose causal structure is not completely understood. Rule-based models assume that conclusions are drawn by chaining together generalized rules, and that learning involves deriving and storing new generalized rules for future use. However, real-world events are often too complex and too imperfectly-understood to immediately distill them into generalized rules. The case-based reasoning process avoids this problem by deriving conclusions only when they are needed, directly from the episodes themselves, in context of the new situation. The Vietnam war provides an example of the benefit of the case-based approach. The Vietnam war is a potential source of countless useful lessons for the future, but attempting to anticipate all of them *a priori,* or to determine the precise circumstances in which a given lesson should be applied in the future, would be an impossible task. Decisions about whether Vietnam is a relevant precedent for a new situation, about which of its lessons apply to that situation, and about how they apply, are best done in the context of the new situation.

Second, even when general rules are available, cases can usefully augment generalized knowledge. Rules intended to describe real-world phenomena are often unreliable, because they cannot enumerate the infinite set of factors that might affect their validity in a particular situation. Consequently, even when there exist rules from which a solution could be derived, it may be valuable to favor solutions that have proven successful in similar prior situations. Likewise, cases can capture and warn of exceptions to the rules.

Reasoning from cases rather than rules can also lead to more efficient processing. Consider, for example, the task of filling out an income tax form. Although it is possible to categorize types of income using the general IRS rules, it is much easier to simply reuse the categories selected and applied the year before. Empirical tests comparing artificial intelligence systems using case-based reasoning to rule-based systems show that case-based reasoning can provide significant speedups.

Another benefit of case-based reasoning is that cases retain sufficient detail to allow them to be used for novel purposes, even if those purposes were not anticipated at storage time. This provides a basis for creative problem-solving. For example, the chief engineer at Honda reported that one of their technological innovations was based on applying a specific observation in a new context:

> We could not think of a way to smoothly change the camshaft profile [for the VTEC-E engine]. It was a severe problem which we finally solved by inserting and removing a pin from the valve rocker gear. The person who developed this apparently got the idea by taking a serious look at a chicken skewer in a yakitori bar. (*Automobile Magazine,* January 1996)

Learning from Experiences

After a case-based reasoner—human or machine—has applied a prior case to a new situation, it learns by storing the results in memory as a new case. If the results were successful, remembering that case in the future can provide suggestions and shortcuts for future problem-solving. If the results were unsuccessful, remembering the case can warn about potential problems. In addition, explanations of current problems can help to avoid future difficulties, by enabling re-indexing of erroneously retrieved cases to make sure they are not mistakenly retrieved in the future. Thus in case-based reasoning, every processing episode provides an opportunity for learning.

The learning done by case-based reasoners combines aspects of both inductive and theory-driven learning. Adding cases to a case library is a form of inductive learning, but one that differs markedly from models of inductive learning (such as neural network models) that abstract away from specific examples to learn only generalizations. Models of case-based reasoning also differ in taking seriously the role of prior knowledge in learning. Selection and generalization of indices, for example, are often based on explanations of why particular features of a situation are relevant in a particular task context. Nevertheless, large chunks of knowledge from specific experiences temper the domain theory and supplant general knowledge. As in the chunking process of SOAR, and the knowledge compilation process of ACT* (see PRODUCTION SYSTEMS in this volume), the CBR learning process increases the availability of specific knowledge, and later processing favors that specific knowledge over initial rules. Case-based reasoning, by learning and reusing specific cases, reapplies the most specific knowledge possible.

The Case-Based Reasoning Process

Cases record specific, operational knowledge tied to a larger context. They encapsulate both the circumstances of that context, and the outcomes or lessons learned. For example, a case might represent an anomaly and its explanation, a plan and its outcome, or a legal argument and its results. Cases can be used for two types of purposes: *interpretation* and *problem-solving* (see PROBLEM-SOLVING in this volume). Problem-solving CBR adapts prior solutions to solve new problems. For example, an architect designing a building may base his or her design on memories of successful designs for similar buildings, adapting them to fit new requirements and to avoid pitfalls encountered in the past.

Interpretive CBR uses prior cases as the basis for forming a judgment about or a classification of a new situation, by comparing and contrasting a new example with cases that have already been classified. This is in the same spirit as exemplar-based models of categorization (see CONCEPTS AND CATEGORIZATION). Exemplar-based models posit that category information is stored in the form of ungeneralized instances of a category, and that each new example is categorized by assigning it to the category of the "closest" prior exemplar. For example, it is natural to predict whether someone will enjoy a movie by remembering the reaction to the most similar movie he or she saw before. Interpretive CBR may also involve sophisticated reasoning about the relevance of particular cases. For example, when cases are used to make legal decisions, a reasoner must make arguments comparing and contrasting a potentially large set of precedent cases to the new situation (Ashley & Rissland, 1987).

Case-based reasoning is made up of a number of subprocesses, with specific subprocesses depending on whether case-based reasoning is being used for problem-solving or interpretation. Both problem-solving and interpretive case-based reasoning begin with a process called **situation assessment,** which elaborates the problem situation and characterizes it in a form compatible with the indices used for the case library. Situation assessment is needed because it is impossible to anticipate all the contexts in which a case may be useful. As a result, the indices used when storing a case may not reflect all the situations under which that case should be applied. Consequently, in novel situations, it may be necessary to redescribe the current situation in order to generate indices that will lead to useful retrievals.

An example of the redescription process is provided by SWALE, a model of case-based reasoning for generating creative explanations of anomalous events (see Schank, Riesbeck, & Kass, 1994). One of the anomalies the system explains is the sudden death of a star racehorse named Swale. To explain the death, the system attempts to retrieve prior explanations indexed under deaths of racehorses, but finds none in its memory; it then generalizes the index to deaths of animals, and succeeds in retrieving prior cases with candidate explanations. However, none of those explanations applies to Swale, so it must generate other indices to find additional explanations. To do so, the system recharacterizes Swale's death in other ways (e.g., as destruction of valuable property, as the death of a young performer, and as death despite peak physical condition). Using the index "death despite peak physical condition," it retrieves a case involving the death of the runner Jim Fixx, who died when exertion overtaxed a hidden hereditary heart defect. This case suggests a plausible explanation for Swale's death.

Either during or after retrieval, analogical **mapping processes** establish correspondences between the current situation and the retrieved case, and **similarity assessment processes** identify relevant similarities and differences. (Some CBR systems reason from multiple cases, but this discussion simplifies the process by assuming that only a single case has been retrieved.) In problem-solving CBR, problematic differences are repaired by **adaptation processes** that fit the retrieved case to the new situation; in interpretive CBR, problematic differences may be accounted for by explaining why those differences are unimportant in the current context. **Evaluation processes** identify problems in the result, guiding incremental adaptation or further justification. The results of the reasoning process, combined with feedback about the success of the solution, are the raw material for **learning processes** that install new episodes in memory for future use and refine indices. This learning forms provides the basis for better future reasoning.

An example processed by the computer model CHEF illustrates the major processes used in problem-solving CBR. CHEF was developed to investigate planning as a case-based reasoning process (Hammond, 1989a). CHEF built recipes—cooking plans—by recalling prior recipes and adapting them to fit current needs. The system remembered its failures, so that they could be avoided in the future; its successes, so that they could be reused; and the repairs it made to failed recipes, so that they could be reapplied. One of the planning problems given to CHEF involved building a recipe for the stir-fry dish "beef and broccoli." To generate a new recipe, CHEF begins by retrieving from memory the recipe whose goals most closely match the set of goals for the current task (in this example, the goals are to make a stir-fry dish, to include beef, and to include broccoli). During similarity assessment, differences between the goals of old and new dishes are weighted by a metric for comparing the importance of particular types of differences, to select the "most similar" prior recipe—the one expected to provide the best starting point. For example, it is harder to adapt a stir-fry dish into a dessert than to adapt a beef dish into a pork dish, so differences involving the type of dish are weighted more heavily than goals involving different meats. Initially, the recipe in CHEF's memory most similar to beef and broccoli is one for beef and green beans. CHEF generates the desired recipe with a simple adaptation, substituting broccoli for green beans and adding a new step required by broccoli but not green beans: chopping the broccoli into pieces. Based on its knowledge of the recipe taken as a starting point, CHEF predicts the characteristics of the new dish, such as tender beef, a savory taste, and crisp broccoli.

When CHEF tests the new recipe in its simulated world, one of its expectations fails: the broccoli is soggy. CHEF diagnoses the problem as involving a bad interaction between two concurrent steps in the plan. Because the beef and broccoli were both stir-fried simultaneously, liquid released by the beef during stir-frying caused the broccoli to be cooked in a pan that was wet rather than dry. As a result, the broccoli steamed rather than fried. (This was not a problem in the recipe taken as the starting point, because green beens are more robust than broccoli.) Based on its explanation, CHEF repairs the recipe by applying the strategy *split and reform*—changing the recipe to stir-fry each ingredient separately and combine their results. This produces a successful recipe which CHEF stores for future use. In addition, based on its explanation of the problem, CHEF refines its indices. First, it forms a generalization about when to anticipate the problem in the future: whenever meat is stir-fried with a crisp vegetable. From this generalization CHEF learns that when it needs to generate a stir-fry recipe involving any meat and crisp vegetable, it should try to retrieve recipes designed to deal with the liquid problem.

When CHEF is later called upon to generate a recipe for chicken and snow peas, it anticipates the potential problem of inadvertently steaming the snow peas, and searches memory for a recipe for a similar dish that addresses the problem. As a result, CHEF retrieves the beef and broccoli recipe as its starting point, in preference to other recipes that (before its learning) would have appeared to have more in common with the requested one (e.g., chicken and green beans).

Development of the Case-Based Reasoning Stance

Many contributions to development of the case-based reasoning stance came from studies of the role of memory in story understanding. In the 1970's, considerable research was directed towards schema-based models of understanding. According to such models, the understanding process is based on reapplication of existing structured knowledge describing normative expectations about a situation. For example, Roger Schank and Robert Abelson (1977) proposed *scripts* as a type of event schema. Each script describes a stereotyped sequence of events that take place in a particular context, such as the standard events in going to a restaurant (wait for a table, sit down, order, etc.). After a script has been identified as relevant to a story, the script is used to guide inferencing about the reasons for events stated in the story and to fill in unspecified details.

Psychological experiments by Gordon Bower, John Black and Terrence Turner (1979) supported the role of scriptal knowledge in human understanding, but also pointed to an important discrepancy between the predictions of script-based computer models and human memory. In script-based computer models, all information about a particular script was localized in a single package, independent of other scripts. If this were taken as a model of human memory, it would suggest that people who process a story with a script will store it in a localized form as well. To examine how script-based information was stored, Bower, Black, and Turner had subjects read stories involving scripts that were distinct, but related to each other, such as the scripts for visiting a doctor and for visiting a dentist. The stories subjects read did not include all the events in each script; parallel script actions were selected to be omitted in one story but included in the other. For example, subjects might be asked to read a story about a doctor visit that explicitly included the event "John arrived at the doctor's office," and to read a story about a dentist visit that did not make an explicit mention of arriving at the dentist's office. After an intervening task, the subjects were asked to recall the stories that they had read. It was found that subjects were more likely to report script actions that had been left unstated if they had read stories which stated the parallel actions. Such cross-schema influences are not accounted for by models in which memory is organized by schemas that are independent of each other.

Such models also failed to account for cross-contextual learning. When people learn about how to pay for a visit to a doctor's office, what they learn appears to be naturally transferred to other contexts, such as paying in a dentist's office. However, models based on fully independent schemas do not support this type of learning, nor do they account for how spontaneous remindings can occur across contexts (e.g., being reminded, when someone takes out a credit card to pay in a dentist's office, of a problem when a doctor's office refused to accept one). Thus three phenomena needed to be explained: cross-schema confusions during retrieval, cross-contextual learning, and cross-contextual remindings.

Roger Schank (1982) responded by proposing Dynamic Memory theory. Dynamic Memory theory postulated a new type of knowledge structure, Memory Organization Packages (MOPs). Like scripts, MOPs characterize event sequences; unlike scripts, MOPs are hierarchical. They are composed of other MOPs and lower-level components, called scenes, which can be shared with other MOPs. For example, the act of payment after a professional office visit tends to be standard regardless of the particular profession involved. Although M-doctor-visit and M-dentist-visit are distinct MOPs, both share a single PAY scene. Consequently, what is learned about paying in one context, changing the PAY scene in memory, is automatically reflected in all MOPs that share the PAY scene. In addition, each processing structure serves as a memory structure. When episodes are retrieved, they must be reconstructed from the information stored under the MOPs and scenes used to process their constituent parts.

Thus in Dynamic Memory theory, processing structures, memory structures, and episodic memory are inextricably bound together. Spontaneous remindings occur during understanding because prior episodes—the cases of a case-based reasoner—are stored in memory under the structures used to process them, and are accessible to provide guidance when those processing structures are used in the future. One of the first cognitive models investigating this integrated view of memory was Janet Kolodner's CYRUS, developed in the early 1980's. CYRUS examined the relationship between encoding, storage, and retrieval, and implemented a model of the retrieval process based on constraints implied by human retrieval. For example, because people have difficulty enumerating long lists of items (e.g., enumerating the 50 states), the CYRUS computer model does not allow the members of categories in its memory to be enumerated. Instead, it models retrieval as a two-step cycle which first formulates a sufficiently-detailed description of the desired item to distinguish it from other similar items, and then discovers whether that item is in fact in memory. This research developed criteria for the types of schemas useful for storing episodes, how schemas change over time, and how schemas organize specific episodes.

Likewise, Christopher Riesbeck developed the ALFRED system to investigate the ramifications of failure-driven reminding for incremental learning. ALFRED used rules that reflect everyday knowledge to understand newspaper articles in the domain of economics. If one of the system's rules of thumb failed, an *exception episode* describing the failure and the recovery procedure was stored in memory, indexed under the faulty rule. If the rule failed again, the exception episode was remembered and the rule was revised, using the exception episode to provide data about the types of situations under which the rule failed to apply.

By the mid-1980's, computer models of case-based reasoning were being investigated in multiple task domains, for tasks such as planning (e.g., in CHEF), diagnosis, interpretation during legal reasoning, criminal sentencing, story understanding, creative explanation of anomalous events, and design. This research resulted in functionally-motivated models of memory organization, retrieval, similarity assessment, adaptation and storage.

At the same time, studies of human learning and problem-solving supported the role of remindings in human reasoning in many tasks. For example, it was shown that people learning to use a word processor reason about how to perform new tasks by remembering how they carried out old ones; that both novice and expert mechanics use prior experiences to help generate hypotheses about problems and to help select appropriate diagnostic tests; that expert decision-makers use analogs to suggest starting points for problem-solving and to help evaluate candidate solutions; and that human explainers favor explanations based on prior experiences as they explain anomalies. These results are surveyed in (Kolodner, 1993, 1994).

Some Research Directions

Memory organization: How human memory is organized, how memory organization is refined, and how retrieval reflects changing task contexts are central questions for case-based reasoning. One source of clues to human memory is to consider the remindings that occur in people and ask what types of features are shared by the old and new situations. Analyzing why a particular episode is retrieved in a given situation—and, equally important, why other episodes in memory that may share features with the new situation are *not* retrieved—provides information about the content of the indices involved.

The case-based reasoning literature presents many examples of remindings based on abstract, thematic features that are shared by episodes in very different contexts. For example, Roger Schank (1982) described being reminded of "Romeo and Juliet" when seeing "West Side Story." Accounting for such remindings depends on identifying the fundamental similarity that exists between the two stories, despite their many differences. Schank proposed that these types of similarity can be described by the type of goal being pursued—here, the mutual goal of the members of each couple to act together—and the conditions under which the goal is being pursued—here, outside opposition. To account for such remindings, Schank proposed that both stories are organized by a single memory structure that captures their abstract similarity, one organizing instances of "mutual goal; outside opposition." Schank called the abstract memory structures *thematic organization packages* (TOPs).

Human cross-contextual remindings support that thematic indices play a role in human memory organization. However, an important research question concerns the circumstances under which thematic remindings occur. Experiments in the 1980's on analogical reasoning showed that people may not be reminded of the most appropriate cases during problem-solving, and suggested the importance of "surface features" in guiding remindings. On the other hand, a series of more recent studies demonstrates the ability of human subjects to achieve remindings based on important features and to avoid being mislead by superficial similarities. For a sampling of perspectives framing the issues involved and relationship between human analogical reasoning and CBR, see (Hammond, 1989b, pp. 125-152).

At the same time that psychological studies continue to examine human retrieval, computer models are being used to examine the functional requirements for memory processes and the content of indices in order for useful remindings to occur. This is a particularly active area of current case-based reasoning research.

Similarity assessment and knowledge adaptation: Modeling the similarity assessment and knowledge adaptation processes is a continuing challenge for case-based reasoning research. Deciding whether a previous case should be considered similar to a new situation, for example, depends on addressing the problem of identifying relevant features in particular task contexts. Similarity assessment is closely tied to the knowledge adaptation problem: Similarity assessment identifies the differences between old and new situations that need to be repaired by case adaptation.

Models of the case adaptation process are often rule-based, relying on fixed sets of rules, raising questions about how adaptation rules are acquired and how the cost of case adaptation is controlled. A number of early models of case-based reasoning proposed applying case-based reasoning to the adaptation process itself, relying on prior examples of adaptations to guide the adaptation process rather than reasoning from scratch. Later models, such as

CHEF, stressed the role of causal explanations of problems in determining how cases should be adapted. The mid-1990's saw a resurgence of research on case adaptation, including research by David Leake and colleagues on methods in which general adaptation principles are gradually augmented with learned domain-specific cases to guide adaptation and memory search. Attention also began to be directed towards the role of metareasoning in case-based reasoning to refine the case-based reasoning process. This work is surveyed in (Leake, 1996).

The integration of multiple knowledge sources and reasoning processes: A natural question is how CBR relates to other forms of reasoning. Case-based reasoning can be viewed as one component of an integrated problem-solving architecture (e.g., see Veloso in Leake, 1996), or within an architecture for which CBR is the primary reasoning process, supported by other reasoning mechanisms (e.g., Hinrichs, 1992). Another question is how multiple cases and other knowledge sources are brought to bear on complex problems. Solving complex problems depends on the ability to reason from multiple remindings and generalizations, throughout the reasoning process.

CBR and creativity: A common misconception about case-based reasoning is that it only applies if new problems are very similar to those solved in the past. However, the case-based reasoning architecture provides a framework for understanding creative reasoning: creativity can enter into case-based reasoning through flexible retrieval processes, that result in novel starting points for solving new problems; through mapping processes that form novel correspondences; and through flexible case adaptation to generate novel solutions. These types of processes have been studied as a basis for modeling creativity for tasks such as explanation (Schank, Riesbeck, & Kass, 1994), and design (e.g., see Wills and Kolodner in Leake, 1996).

CBR and education: Case studies have long been the basis of legal and medical education, and the "learning by doing" framework that case-based reasoning suggests is at the heart of problem-based learning methods in education. However, the lessons from computational models of CBR go farther, providing new concrete suggestions about what makes a good problem, the range of problems that students should solve, and the kinds of resources that should be made available to student learners (Kolodner, Hmelo, & Narayanan, 1996). Case-based reasoning highlights, for example, the importance of providing problems that present difficulties for students, enabling students to identify and repair gaps in their knowledge. It also suggests supporting students' learning by presenting them with cases chosen to help point out potential issues to address, to guide their learning; to suggest potential solutions or parts of solutions to consider; and to help in anticipating the effects of their decisions.

In addition, the case-based reasoning stance—that human reasoning is largely case-based—suggests that an important aspect of education is developing students' abilities as case-based reasoners. For example, when students reason from prior cases, they may be misled if they blindly apply the first case they retrieve: it is crucial for students to learn to evaluate the relevance of prior cases and the solutions those cases suggest.

Research groups led by Roger Schank and by Janet Kolodner are applying principles from the case-based reasoning cognitive model to develop a new generation of curricula, learning activities, and computer systems to support the learning process. (See, for example, Schank in (Leake, 1996).) The resulting instructional systems provide students with opportunities for learning by doing in rich task environments, driven by compelling goals, and also support students' learning by presenting students with useful information about others' experiences, in the form of relevant cases at appropriate times.

Summary

The case-based reasoning stance treats experiencing, remembering, and learning as inextricably bound together. The key questions it proposes for cognitive science concern the structure, content, and organization of knowledge in memory, the knowledge and processes required to retrieve relevant information, how to adapt retrieved information to fit new situations, and how to extract, represent and organize the lessons of new experiences.

Case-based reasoning emphasizes the role of concrete, operational knowledge. Rather than focusing on how basic knowledge can be composed to generate new solutions, case-based reasoning focuses on how large structures—cases—can be modified to fit new situations, and views the learning of new cases as an integral part of the reasoning process.

A series of computer models, inspired by observations of human behavior, has built up understanding of the functional requirements for processes involved in case-based reasoning and the knowledge they require. These models provide computational definitions of central cognitive science concepts such as similarity and relevance. In addition, examining case-based reasoning as a cognitive architecture suggests concrete answers to classic questions ranging from the sources of creativity to how to aid students' learning. At the same time, CBR research is identifying further questions to be answered by a combination of computer modeling and psychological experimentation. Thus the case-based reasoning stance suggests not only central issues for cognitive science, but also how to address them with a research methodology built on the interplay between studies of human behavior and analysis of functional constraints on the processes and knowledge needed to account for that behavior.

Acknowledgments

This work was supported in part by the National Science Foundation under Grant No. IRI-9409348. I would like to thank Rob Goldstone, Janet Kolodner, Ashwin Ram, and the Editors of this volume for their very helpful comments on a draft of this chapter.

David B. Leake

References

Ashley, K. & Rissland, E. (1987). Compare and contrast, a test of expertise. In *Proceedings of the Sixth Annual National Conference on Artificial Intelligence*, pp. 273–284 San Mateo, CA. AAAI, Morgan Kaufmann.

Bower, G., Black, J., & Turner, T. (1979). Scripts in memory for text. *Cognitive Psychology, 11*, 177–220.

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

Hammond, K. (Ed.). (1989b). *Proceedings of the DARPA Case-Based Reasoning Workshop*, San Mateo. Morgan Kaufmann.

Hinrichs, T. (1992). *Problem Solving in Open Worlds: A Case Study in Design.* Lawrence Erlbaum Associates, Hillsdale, NJ.

Kolodner, J., Hmelo, C., & Narayanan, N. (1996). Problem-based learning meets case-based reasoning. In *Proceedings of the Second International Conference on the Learning Sciences* Evanston, IL.

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

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

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

Schank, R., Riesbeck, C., & Kass, A. (Eds.). (1994). *Inside Case-Based Explanation.* Lawrence Erlbaum Associates, Hillsdale, NJ.

Recommended Readings

Kolodner, J. (1993). *Case-Based Reasoning.* Morgan Kaufmann, San Mateo, CA.

Kolodner, J. (1994). From natural language understanding to case-based reasoning and beyond: a perspective on the cognitive model that ties it all together. In Schank, R. & Langer, E. (Eds.), *Beliefs, Reasoning, and Decision Making: Psycho-Logic in Honor of Bob Abelson*, pp. 55–110. Lawrence Erlbaum, Hillsdale, NJ.

Leake, D. (Ed.). (1996). *Case-Based Reasoning: Experiences, Lessons, and Future Directions.* AAAI Press, Menlo Park, CA.

Riesbeck, C. & Schank, R. (1989). *Inside Case-Based Reasoning.* Lawrence Erlbaum Associates, Hillsdale, NJ.

Word count (counted by unix detex and wc, including references): 6076 words.