Experience, Introspection, and Expertise:
Learning to Refine the Case-Based Reasoning Process*

David B. Leake
Computer Science Department
Indiana University
Bloomington, IN 47405
leake@cs.indiana.edu
812-855-9756
August 5, 1995

Journal of Experimental and Theoretical Artificial Intelligence, special issue *The Psychology of Expertise: Human and Artificial.* In press.

Running head: "Experience, Introspection, and Expertise."

*This work was supported by the National Science Foundation under Grant No. IRI-9409348.

Experience, Introspection, and Expertise:
Learning to Refine the Case-Based Reasoning Process

Abstract

The case-based reasoning paradigm models how reuse of stored experiences contributes to expertise. In a case-based problem-solver, new problems are solved by retrieving stored information about previous problem-solving episodes and adapting it to suggest solutions to the new problems. The results are then themselves added to the reasoner's memory in new cases for future use. Despite this emphasis on learning from experience, however, experience generally plays a minimal role in models of how the case-based reasoning process is itself performed. Case-based reasoning systems generally do not refine the methods they use to retrieve or adapt prior cases, instead relying on static pre-defined procedures. The thesis of this article is that learning from experience can play a key role in building expertise by refining the case-based reasoning process itself. To support that view and to illustrate the practicality of learning to refine case-based reasoning, this article presents ongoing research into using introspective reasoning about the case-based reasoning process to increase expertise at retrieving and adapting stored cases.

1 Introduction

Artificial intelligence models of expert reasoning often base their reasoning on fixed sets of rules representing fundamental domain principles. Such models are aimed at reflecting a fixed body of knowledge that an expert possesses after having acquired expertise; their design assumes that their rule libraries are complete and correct. However, this assumption may be difficult to realize in practice. Eliciting the needed rule libraries from experts can be difficult, and even a perfect rule library may suddenly become obsolete due to changes in the task or domain.

One way to address the problem of maintaining adequate domain knowledge is to replace the focus on acquiring fixed and final "expert" knowledge with a focus on the process by which expertise is acquired and refined during problem-solving. Including the development of expertise as an organic part of computer models of expertise is appealing both to help alleviate the knowledge acquisition burden of having to reconstruct an expert's rules and to help illuminate the development of human expertise.

One widely-investigated artificial intelligence account of the role of experience in expertise is presented by the case-based reasoning (CBR) paradigm (e.g., Kolodner (1993), Riesbeck & Schank (1989)). In CBR models, expert performance is viewed as arising largely from reuse of appropriate cases from a rich library of previous problem-solving episodes.[1] Proponents of the case-based reasoning model observe that because case-based reasoning bases problem-solving on entire prior situations rather than general rules, it obviates the need to extract general rules from problem-solving experiences, which can simplify knowledge acquisition in poorly-understood task domains. In addition, the case-based reasoning process naturally addresses the need to update problem-solving knowledge in response to changing circumstances, in that novel experiences are stored as new cases to guide future problem-solving.

Yet despite the benefits of reasoning from specific experiences, the availability of cases does not entirely account for high levels of problem-solving performance. People with comparable levels of experience may exhibit widely divergent levels of skill, sometimes reaching

[1]This paper focuses on case-based reasoning for problem solving. We will not discuss another important side of case-based research, case-based reasoning for interpretation and classification (e.g., Ashley (1990), Ashley & Rissland (1987), Bareiss (1989), Branting & Porter (1991)).

plateaus of performance that fall short of mastery of their tasks (e.g., Bereiter & Scardamalia (1993), Ericsson & Smith (1991), Lesgold, Rubinson, Feltovitch, Glaser, Klopfer, & Wang (1988)). Thus expertise depends not only on experiences, but also on factors affecting how that experience is used. For models of case-based reasoning, a key factor affecting the benefit of a given library of experiences is the reasoner's expertise at the case-based reasoning process itself: its ability to retrieve the right cases and to adapt them appropriately to fit new situations.

This article examines the requirements for successful case-based reasoning and presents a method for acquiring expertise at the case-based reasoning process. The approach involves learning from introspective reasoning about the reasoner's needs for information to adapt a case, about the organization of the cases in memory, and about the desired behavior of the case-based reasoning process itself. The article illustrates the approach with descriptions of computational models being developed to use introspective learning as the basis for addressing two classic problems of case-based reasoning: how to adapt cases to fit new situations and how to retrieve appropriate cases from memory.

2 Overview

The first section of this article outlines the relationship between our approach and other perspectives on expertise. The next section highlights main points of the case-based reasoning paradigm with a simple "first-pass" view of how expertise arises from acquiring a library of problem-solving cases. With that foundation in place, the following section takes a more critical view, identifying the assumptions on which the first-pass model depends and the related issues that must be addressed by case-based reasoning models of expertise. In response to those issues, the next section suggests a "second-pass" account of case-based reasoning in which introspective reasoning about the case-based reasoning process enables a case-based reasoner to develop expertise at applying its cases. The remainder of the article discusses ongoing research on two computational models that use introspection about their reasoning processes to refine their use of case-based reasoning.

3 Perspective

Several different criteria enter into notions of what it means to be an "expert" at a particular task. One is a social criterion, depending on whether society is willing to bestow the status of "expert" on a particular person (e.g., Agnew, Ford, & Hayes (1994), Fuller (1994), Sternberg & Frensch (1992)). Another is a performance criterion, based on factors such as that person's problem-solving speed and the quality of the solutions attained. Yet another is a person's level of experience; for example, "novice/expert" comparisons implicitly contrast inexperience with expertise. These three factors obviously tend to be interrelated—society anoints as "experts" those it views as performing with particular skill, and acquiring that skill may require long periods of practice—but they reflect different considerations and are not equivalent. For example, neither social recognition nor experience assures expert-level performance (e.g., Camerer & Johnson (1991)).

This article focuses on the relationship between two of the previous criteria for expertise: performance and experience. Its basic perspective comes from research on case-based reasoning, in which experience (in the form of stored prior problem-solving cases) is used as the starting point for solving new problems. In most treatments of case-based reasoning, case acquisition is the primary means by which problem-solving ability increases. The new contribution of this article is to address how experience can also lead to refinement of the mechanisms for applying stored cases. In this way, it addresses the question of how *expertise at doing case-based reasoning* is developed from introspective reasoning about the case-based reasoning process.

Because our approach focuses on how expertise is refined, it is in sympathy with recent perspectives from psychology and education that advocate focusing on the development of expertise rather than treating expertise as a single fixed state to be achieved. For example, Campbell, Brown, and DiBello (1992) argue for developmental studies of human expertise and observe that traditional studies, which pigeonhole complex differences in levels of skill into a novice/expert dichotomy, can blur important distinctions in performance. Likewise, Bereiter and Scardamalia (1993) argue that the ongoing refinement of abilities is an integral part of the notion of being an expert (even if the expert has already achieved a high level of performance). Whether or not one accepts their view in general, it is clear that for many tasks, the continual refinement of skill plays an important role in expertise. (For example, a scientific researcher who does not continually extend current skills will soon cease to be an expert). This article presents a model of the role of experience and introspective reasoning in the ongoing process of refining expertise.

4 Expertise and CBR

When case-based reasoning is applied to problem-solving, the problem-solving process is based on the lessons suggested by a memory of specific problem-solving episodes.[2] When confronted with a new problem, a case-based problem solver generates a solution by retrieving a previous solution for a similar problem and adapting it to fit the new circumstances. Thus the problem-solving process depends on retrieving an appropriate prior case, performing analogical reasoning to associate old and new solutions, and adapting the solution suggested by the prior case in order to fit new constraints.

Also integral to case-based reasoning is learning from problem-solving episodes. After a solution is generated, lessons from the current problem are added to the case library to be available for future use. If a successful solution was generated, that solution is stored for future re-use; if problems occurred, information on how to anticipate those problems is stored to allow those problems to be remembered in similar future situations as warnings of errors to avoid. As cases are added to the case-based reasoning system's memory, both prior solutions and warnings of errors become available for a wider range of problems. This increases the likelihood that appropriate prior lessons will be available in new situations and the likely similarity of the retrieved case to the new situation, decreasing the amount of adaptation that must be done and consequently increasing the speed of problem-solving.

As an illustration of how case-based reasoning applies to a particular task, consider the problem of generating explanations in order to understand surprising events in news

[2] In most case-based reasoning systems the initial set of episodes is simply provided to the reasoning system. In general it may be built by storing results of reasoning from scratch (e.g., Goel, Callantine, Donnellan, & de Silva Garza (1993), Koton (1988), Veloso (1994)), gathered from observations of the performance of other actors, or presented to the system in other ways. For example, a case-based reasoning system that creates recipes can start with the recipes provided by a cookbook (Hammond, 1989).

stories. Unlike explanation generation methods that build explanations from scratch (e.g., Hobbs, Stickel, Appelt, & Martin (1993)), the case-based approach to explanation generation builds new explanations by adapting explanations of similar previous events to fit the new situations (Leake, 1992b; Schank & Leake, 1989; Schank, Riesbeck, & Kass, 1994). This phenomenon of "explanation based on remindings" is often observed in human explainers who must explain despite partial information. For example, when newspapers first reported the attack on Olympic figure skater Nancy Kerrigan, many candidate explanations could be generated. However, some readers were reminded of a previous attack on tennis star Monica Seles and used that reminding as the basis of their explanations for the Kerrigan attack. Seles was stabbed by an obsessed fan of one of her rivals, Steffi Graf; the fan hoped to assure Graf's victory by preventing Seles from competing. That reminding suggested considering the explanation that Kerrigan's attacker, Shane Stant, was an obsessed fan who wanted one of Kerrigan's rivals to win. That explanation was only partially applicable—in fact, Stant was hired to do the crime rather than being the instigator. However, it provided a starting point by suggesting focusing on foul play to aid a competitor.

Eventually, it was determined that the husband of Kerrigan's rival Tonya Harding hired Stant to attack Kerrigan, to prevent her from competing, so the suggestion provided by the reminding was in fact correct. Although there is no guarantee that a given reminding will actually be useful in a new situation (and we discuss later the factors that affect the relevance of retrieved cases), experiments by Read and Cesa (1991) show that human explainers favor explanations based on remindings of prior explanations for similar anomalies.

4.1 Motivations for research on the CBR model

Research on the case-based reasoning model of expertise is motivated by two types of factors: cognitive considerations for modeling human reasoning and functional motivations based on the desire to achieve improved problem-solving performance.

Psychological support as a cognitive model: Although a number of psychologically-inspired models of problem-solving focus on the process of forming and applying generalized rules or procedures (e.g., Anderson (1983), Laird, Rosenbloom, & Newell (1990)), there is considerable evidence that the retrieval and reapplication of specific problem-solving experiences also plays an important role in human problem-solving. People reason from prior experiences both in the early phases of learning a domain and after they have achieved expertise (e.g., Campbell et al. (1992), Faries & Schlossberg (1994), LeFevre (1988), Pirolli & Anderson (1985), Ross (1989)). For example, Lancaster and Kolodner (1987) show that both novice and expert mechanics use specific experiences to help to generate hypotheses about problems and to help select appropriate tests; Klein and Calderwood (1988) show that expert decision-makers in complex and changing situations use analogs to suggest starting points for problem-solving and to help evaluate candidate solutions.

Case-based reasoning also appears to play an important role in human reasoning for tasks such as real estate appraisal (Burstein, 1994) and labor mediation (Sycara, 1987), in which it is difficult to enumerate and weigh all the factors that may be relevant, and in which changing circumstances preclude reasoning from a fixed library of rules. For example, the decision to build a prison in a particular neighborhood would have an important effect on property values that might not have been anticipated in a predefined rule base, but that would automatically be reflected in an appraisal based on the recent selling prices of similar houses in the area.

Functional support: Case-based reasoning offers three potential functional benefits compared to problem-solving models based on static rule libraries. First, by storing prior solutions rather than re-deriving them from scratch, it can increase problem-solving efficiency over time, as additional cases are stored and become accessible to be used in similar new situations. For example, comparisons by Koton (1989) showed that the case-based diagnostic system CASEY performed two to three orders of magnitude faster than the rule-based system from which it was derived.

Second, in poorly-understood or changing domains, CBR helps overcome the lack of perfect domain knowledge by augmenting the reasoner's domain theory with records of specific experiences. Instead of neutrally giving equal consideration to all solutions licensed by the reasoner's domain theory (which may be inconsistent or incorrect), the case-based reasoning process focuses on solutions that have proven successful in similar prior situations. In a regular world, these may be more likely to apply than solutions based entirely on uncertain *a priori* knowledge.

The third functional motivation for case-based reasoning is to simplify the knowledge acquisition process. Unlike rule acquisition, which requires analyzing the interactions between all individual factors in a situation, in case-based reasoning an entire episode can be treated as a unit from which to reason. Reports from developers of AI systems corroborate the benefits of using cases as the primary unit of domain knowledge to acquire (e.g., Kolodner (1993, pp. 93-94)).

Despite the appeal of these arguments for case-based reasoning, however, there has been little examination of the factors affecting whether these desired benefits are actually realized in practice. In the following section we address this issue as an introduction to our investigation of introspective reasoning as a means to increase expertise at case-based reasoning.

5 Requirements for successful case-based reasoning

The performance of a case-based reasoning system depends on three types of factors: the reasoner's experience at relevant problems; the reasoner's additional reasoning capabilities, using methods other than case-based reasoning; and the reasoner's knowledge of how to apply the cases in its memory.

Experience at relevant problems: Because case-based reasoning solves new problems by applying reasoning from prior problems, it is obvious that the effectiveness of the reasoning process will depend on the relevance of the lessons from prior problems to the new situation. In a completely novel situation or a situation to which no prior experience applies, little or no benefit may accrue from using prior cases. For example, a library of planning cases acquired for cooking stews may not be very helpful as the starting point for planning how to cook a cake. (However, it has also been argued that even in novel situations to which no prior case applies directly, the use of case-based reasoning can have benefits for creative problem-solving; see Schank (1986) and Schank & Leake (1989).)

Additional reasoning capabilities: Case-based reasoning is one of many reasoning strategies. It is not necessarily appropriate for all tasks, and consequently can be applied most effectively in conjunction with other reasoning methods. For example, in domains for which a complete domain theory is available and for which optimal solutions are essential, reasoning from first principles may be appropriate. The integration of reasoning from general knowledge and from specific cases (e.g., Ashley & Rissland (1987), Branting & Porter (1991), Hinrichs (1992), Redmond (1992)) can also play an important role in the reasoning process. For example, Lancaster and Kolodner (1988) show that expert mechanics use both remindings of specific problems and reasoning from abstract domain models when they diagnose automobile problems.

Knowledge of how to guide case application: The knowledge required for case-based reasoning goes beyond cases alone: CBR systems depend both on their cases and on knowledge of how to apply those cases. In order for a case-based reasoning system to function effectively, it must retrieve appropriate cases and adapt those cases effectively.

In general, there is no guarantee that either human or artificial case-based reasoners will apply stored cases effectively. For example, psychological experiments show that people are not necessarily reminded of the most relevant prior cases and may fail to notice important similarities between old and new cases (Gentner, Ratterman, & Forbus, 1993; Gick & Holyoak, 1980; VanLehn, 1989). However, there is evidence that remindings based on goal-relevant features do occur in task-driven reasoning (Seifert, 1988), and that even novice programmers can retrieve prior problem cases based on structural features of the problem being addressed, rather than being misled by superficial similarities (Faries & Schlossberg, 1994).

Nor are human case application abilities static. For example, developmental studies show a shift in the criteria that children use when determining relevant features during adaptation of previous stories to new circumstances. Experiments by Gentner and Toupin (1986) gave children the task of adapting previously-encountered stories to fit new characters. In the experiments, children first acted out stories with toy animals as the characters. They were then presented with the same beginnings of the stories, with different toys representing the characters, and asked to act out the remainder. In some trials, corresponding characters in the first and second stories had similar appearances (e.g., a chipmunk and a squirrel might play corresponding roles). In others, the characters in corresponding roles had lower surface similarity (e.g., as if the role initially filled by a squirrel were filled by an elephant).

In adapting the stories to use the new characters, both older children (8-10 years old) and younger children (5–7 years old) were influenced by surface similarities; both sets of children did better at mapping when similar animals played similar roles. However, considerations related to systematicity (e.g., that one character was a friend of another) aided the older children in making the correct mappings between characters as they adapted the old stories to the new characters. Gentner and Toupin observed that older children would sometimes make mapping errors, having an animal act the same way that a similar-looking animal had in the first story, and then correct themselves to focus on structural features (e.g., by noting that the character was greedy and consequently should be the one to do greedy things, regardless of whether it shared surface features with the greedy animal in the initial story).

Adult experts are better than non-experts at recognizing important similarities when applying old experiences to new problems; for example, a study by Novick (1988) showed that when given math problems for which the same solution procedure is appropriate but which have very different surface features, experts show a strong tendency to recognize the relevance of the solution procedure despite the surface dissimilarities, while non-experts do not. A survey of psychological literature on development of analogical reasoning by Goswami (1991) stresses the importance of knowledge in the development of analogical reasoning skills, which raises the question of how recognition of important features might be learned as humans or machines acquire expertise. Learning relevant features is the focus of the introspective reasoning process that we describe in the following sections.

6 Using introspection to refine case-based reasoning

As discussed in the previous sections, the quality of case-based reasoning depends on expertise at retrieving the right cases from memory and adapting those cases to fit new situations. In computational models of case-based reasoning, the standard approach to providing this expertise is to attempt to "build in" the requisite case application knowledge for a particular domain. As Simoudis, Ford, and Canas (1992) point out, that task involves the same types of knowledge acquisition problems that have proven a serious impediment to developing rule-based expert systems. Given that one of the pragmatic motivations for modeling learning from experience is to alleviate knowledge acquisition problems, a natural question is whether learning from experience could be applied to refining expertise at the case-based reasoning process itself.

In order to decide how to refine its reasoning process, a CBR system must be able to reason about that process: to reason introspectively about the motivations for its reasoning, the requirements that must be satisfied in order for the results of the CBR process to satisfy the reasoner's needs, and the way in which the CBR process is expected to perform.

A rich literature addresses issues such as reflection, introspection, and metacognition (see for example Piaget (1976), Campbell & Bickhard (1986)), and human experts appear to have greater awareness of their own problem-solving process than less expert performers (e.g., Chi, Bassok, Lewis, Reimann, & Glaser (1989)). However, introspective reasoning has received little attention in studies of the case-based reasoning process. In the following sections we discuss models of how introspective reasoning can be used to refine two aspects of the case-based reasoning process. The first section discusses how reasoning about the CBR system's memory organization can contribute to learning how to find the information needed to adapt cases to new situations. The second section discusses how, after a case has been applied to a situation, introspective reasoning about prior processing can be used to detect and repair sub-optimal case retrieval criteria. Initial computer implementations of both models have been developed and are now being extended and refined.

6.1 Using introspection about memory search to increase expertise at case adaptation

A case-based reasoning system's flexibility comes from its ability to adapt prior cases to fit new situations. Unfortunately, pre-defining the needed knowledge is a hard problem (Allemang, 1993; Kolodner, 1991; Leake, 1994b). The difficulty in hand-coding case adaptation knowledge is so acute that many CBR applications do not even include case adaptation, leaving adaptation to be performed manually by the users of the systems (e.g., Bayles & Das

(1994), Blevis, Burke, Glasgow, & Duncan (1991), Domeshek & Kolodner (1992), Hennessey & Hinkle (1991), Simoudis & Miller (1991), Slator & Riesbeck (1991)). This makes the systems into "expert memories" without the capability to apply the contents of those memories themselves. We are addressing the problem of defining adaptation criteria by modeling how a case-based reasoning system can, starting from limited case adaptation knowledge, learn from experience to refine its case adaptation process.

6.1.1 The case adaptation problem

Case adaptation involves performing operations such as adding, deleting, and substituting components of a retrieved solution, in order to generate a new solution that applies to the current problem. The types of structural transformations involved in modifying an encoded solution can generally be described in terms of a very small set of basic operations that are then combined as needed to perform complex adaptations (e.g., Carbonell (1983), Hammond (1989), Hinrichs (1992), Kass (1990), Kolodner (1993)). However, applying those transformations to particular cases can require retrieving a wide range of supplementary domain-specific knowledge. For example, if part of the evidence that applied to a previous explanation of a crime is implausible in the current situation, a possible transformation is to *substitute evidence* (Koton, 1988) to replace the implausible evidence with evidence that is more believable. In order to do that replacement, it is necessary to find new supporting evidence, which may be quite difficult.

The problem is illustrated by the previous example in which case-based reasoning was applied to explaining the attack on Nancy Kerrigan. In that example, the reminding of the attack on Monica Seles suggested an explanation—an attack instigated and carried out by a crazed fan of an opponent. That explanation only partially fits the Kerrigan attack, because Shane Stant, Kerrigan's attacker, was hired to perform the crime; he did not instigate that attack. Consequently, applying the retrieved explanation depends on adapting it to reflect that Stant was hired by someone else. That adaptation requires adding the new instigator's role to the causal chain leading to the attack. Despite the differences, the explanation for the attack on Seles still suggests a motive to consider when searching for an instigator: That the attack was instigated by a fan who hired Shant to carry out the attack.

In order to complete the explanation it would be necessary to identify a fan who might have carried out the attack. In general there is no guarantee that relevant information will be in memory (and for this example, it is very unlikely that a newspaper reader would know enough about those involved to even form a reasonable conjecture). Even if the appropriate information were in memory, however, it could be difficult to retrieve it. No pre-existing category in an average reader's memory is likely to group "fans of competitors in Kerrigan's upcoming competitions." Consequently, it is necessary to find the information through an indirect search method, perhaps first searching for possible future competitors of Kerrigan (which is a memory search problem in its own right and might require using additional strategies such as trying to retrieve episodes of skating competitions in which Kerrigan was involved, to see who competed there, or trying to retrieve information about the Olympic skating team). Next, it is necessary to find ways to focus on the competitors' fans who are sufficiently deranged or obsessive to resort to violence. The Kerrigan example shows that finding the specific information needed to apply an adaptation rule can be a difficult part of the case adaptation process.

We can view the problem of finding the information needed to adapt cases as a problem in *operationalizing* (Mostow, 1983) abstract structural transformations by gathering the domain-specific information needed to apply them Leake (1993). The need to search memory for information to operationalize abstract adaptation rules arises not only during case adaptation for explanation, as in the previous example, but for any task being solved by case-based reasoning. For example, in case-based planning, one problem that can arise in reapplying a previous plan is that a side-effect of a step in a retrieved plan has bad effects in the current context. An abstract adaptation rule to repair that problem is *add a step to remove harmful side-effect* (Hammond, 1989). In that abstract form, the rule is not operational; applying it requires searching memory for the plan step to add, and determining that step may be a hard problem. If the case-based planning system is attempting to build a plan for X-ray treatment, and the X-ray dose needed to destroy a tumor will result in an excessive radiation dose to healthy tissue, considerable domain expertise may be needed to decide what steps should be added to the treatment plan.

A common way of addressing the problem of gathering the information needed for case adaptation is for the developer of the CBR system to define domain-specific adaptation rules with the needed information built in. For example, more specific versions of *add a step to remove harmful side-effect* can be tailored to adaptation problems for X-ray treatment plans, resulting in rules such as *add the step "rotate radiation sources" to remove harmful side-effect "excess radiation"*, a rule used in the ROENTGEN system (Berger & Hammond, 1991). In general, a case-based reasoner whose adaptation rules are tailored to its task can adapt cases more effectively. However, developing domain-specific adaptation rules requires knowledge of the domain and of the types of adaptation problems that the reasoner is likely to encounter. Neither type of knowledge may be available *a priori*, which makes knowledge acquisition a particularly difficult problem for the case adaptation process.

6.1.2 Acquiring adaptation expertise by introspective learning

We address the problem of acquiring expertise at case adaptation by focusing on how to learn to find the domain-specific information needed to adapt cases. Given a novel adaptation problem, our model uses introspective reasoning to generate a new memory search plan to find needed information (Leake, 1995c). The search plan is then stored as a case in memory. As similar adaptation problems are encountered, the search plan is re-applied by case-based reasoning to deal with the new problem. Thus our model addresses both the formulation of original memory search plans and their reuse (Leake, 1994a, 1995b; Leake, Kinley, & Wilson, 1995).

The following sections discuss the major phases of this process: How the need to adapt a case results in a "knowledge goal" to retrieve a particular type of information from memory; how that knowledge goal is the starting point for a planning process to search for needed information; and how the results are stored and reused to facilitate future case adaptation.

Generating knowledge goals: In order to search for the information needed to adapt a case, a reasoner must be able to describe the goal of its search. This section discusses how our model identifies the type of information it needs for an adaptation problem in order to represent it as an explicit *knowledge goal* (Leake & Ram, 1993; Ram, 1987)—a goal to acquire a particular type of knowledge—to guide the later information search process. The only input information required is the standard type of input that case-based reasoning systems provide to their adaptation components: a description of the problem that necessitates adapting the case. (We use a vocabulary of problem types based on Leake (1992b).) The problem description is used as an index to attempt to retrieve stored cases for similar adaptation problems. If one has been encountered previously, the same adaptation method used in the previous episode is re-applied, as described later in this section. Otherwise, the model uses simple heuristics to select an abstract transformation to repair the problem (e.g., if part of the explanation does not apply to the new situation, one heuristic is to try to apply *substitute component* to find an alternative playing a similar role in the explanation). Once a transformation is selected, the transformation determines information that must be found in order to apply the transformation. For example, when substituting a component in an explanation, memory search must find a component that plays the same causal role that the original component played in the explanation.

To provide a concrete example, we return to the Kerrigan episode. The explanation for the attack on Monica Seles only partially fits the Kerrigan incident, because Shane Stant does not fit the role of instigator. The need to identify a substitute instigator corresponds to a knowledge goal: the goal to find someone excessively devoted to the victory of a competitor to Kerrigan.

Reasoning introspectively about how to search memory to satisfy knowledge goals: Our model's memory search process builds on Kolodner's (1984) approach to memory search as a deliberative process. In her model, queries to memory that cannot be answered directly are transformed into new queries, according to elaboration rules, in order to eventually obtain a query that can be used successfully as a retrieval index. In a similar spirit but using a more general mechanism, we model the memory search process for finding the information needed to perform case adaptation as a process of *knowledge planning* (Hunter, 1990). In knowledge planning, information search is conducted by a planning process in which "mental" operations are selected based on explicit reasoning about needs for information and how to satisfy them.

In our model of reasoning about memory search, generating a plan depends on using knowledge of the interrelationships between concepts in memory. Given information about the meanings of memory links and the type of information sought by memory search, a planning process can reformulate queries and follow sequences of links to extract the information it needs (Leake, 1995c). This is illustrated by the reasoning required to find a fan who might have instigated the attack on Kerrigan. First, to find a fan of one of Kerrigan's potential competitors, it is necessary to identify who those competitors are. That requires a memory search process to locate potential competitors (e.g., by retrieving information about prior skating competitions in which Kerrigan competed and noting who competed against her there, which might involve first finding general information about the schedules of figure-skaters, etc.) After identifying the competitors, it is necessary to search for their fans, who may not be explicitly represented in memory, but for whom candidates may be generated by considering people devoted to them, such as spouses, lovers, siblings, or close friends. Finally, after candidates have been identified, evaluation of the candidates is needed to decide whether they are plausible attackers.

In order to decide which memory links to traverse during its search for information, a reasoner must have self-knowledge about the relationships that its memory links represent. This contrasts with most memory models, in which relationships are labeled by naming the links in memory but the names have no meaning to the memory search system itself. For example, to reason about whether a spouse is appropriate to consider when searching for fans, the memory must represent not only that a link named "spouse" exists, but the meaning that the link reflects: the memory searcher must have metaknowledge of the relationships underlying its memory organization.

Being able to treat memory search as a planning process and to reformulate queries increases the flexibility of the memory search process but also increases potential processing cost. In a rich memory, there will be numerous possible paths to any particular piece of information, and numerous possible wrong turns during memory search. Consequently, building up a memory search chain from scratch each time new information is needed is likely to be prohibitively expensive, especially for long search paths. The next section addresses the question of how a reasoner can learn from its memory search process, acquiring expertise at memory search and case adaptation.

Learning memory search strategies: Many sophisticated memory search schemes have been developed in CBR research, but they are normally driven by opaque procedures, rather than being accessible to explicit reasoning and learning. The planful memory search process described in the previous section makes it possible to reason about the memory search procedures as plans. A key benefit is that memory search plans can be subject to the same learning methods that have been applied to planning in other contexts (e.g., Segre (1988), Hammond (1989), Birnbaum, Collins, Freed, & Krulwich (1990)).

A question is which of these learning methods to apply. Initially, it appears that explanation-based generalization (Mitchell, Keller, & Kedar-Cabelli, 1986; DeJong & Mooney, 1986) is the method of choice, because it has been widely used for forming operational generalizations of specific plans. However, explanation-based generalization requires a perfect domain theory to form valid generalizations. The memory search rules described previously are just heuristics for finding concepts with particular relationships; whether they succeed in a given instance depends not only on the rules themselves but on the idiosyncratic contents of memory. Consequently, what is needed is a learning process that can start from those unreliable general rules and form more specific and reliable rules reflecting the specific contents of memory. This suggests recursively applying case-based reasoning to the memory search task within case-based reasoning.

By recursively applying case-based reasoning to a CBR system's own memory search process, it is possible to build up a library of cases reflecting how particular heuristics apply to the idiosyncratic contents of a particular memory. For example, in a memory reflecting an athlete's knowledge it might be reasonable to search for competitions as part of a knowledge structure for a competitive season, while in a memory reflecting the knowledge of an apathetic stadium employee it might be necessary to search memory under another category, such as "overtime days." Thus the memory search cases that are learned reflect specifics of the system's knowledge, task, and memory organization.

The blame assignment problem: We note that the success of the memory search process depends both on using the right search strategies and on whether memory actually contains the needed information. When there is insufficient information in memory, or when the wrong memory search strategies are used, the search process will fail to find the needed information. Determining what was the source of the failure is important for guiding remedial learning. For example, if the failure was caused by insufficient knowledge, augmenting the contents of memory by consulting reference sources, etc., is appropriate; no new search strategies need to be learned. (In fact, the series of questions that a detective might ask during interviews when investigating a crime can be viewed as a way of applying the memory search strategies of the detective to a memory richer in relevant knowledge than the detective's own.) However, when the failure results from a deficient search strategy not finding available information, new search strategies are needed.

Distinguishing between a search failure resulting from lack of information and one resulting from a bad choice of search strategy is the "blame assignment" problem for memory search. In general, it is difficult to resolve: How can a reasoner know that its search strategy is deficient (i.e., that its memory actually contained the needed information but memory search failed to find it)? A method we are investigating for addressing the blame assignment problem is to try as many search plans as are possible within a limit on the number of memory links traversed, until the desired information is found or the resource limit is exceeded. If the information is found by one memory search plan but was not found by previous plans, the previous plans were flawed; learning can reflect that the final plan was superior and that the previous plans failed. If the information is not found by any plan, the problem could be either the search plans used or a lack of information in memory, and no immediate learning can be done unless external feedback is available to identify the source of the problem. However, if failed search plans are stored in memory, indexed by the information they sought, learning can be done opportunistically if the needed information is later found to be in memory.

A summary of knowledge sources: In our model, introspective reasoning about case adaptation depends on four types of knowledge. The first is knowledge of types of *abstract transformations,* that can be used to make structural changes when adapting a case (e.g., the transformation that substitutes a new component for an existing component in the case). The second type of knowledge is knowledge about the reasoner's own memory organization, encoded as *memory search rules.* Memory search rules reflect information about the types of memory links to follow to gather particular types of information and how queries can be transformed into other queries. For example, a memory search rule would represent that a way to identify the abstractions of a concept is to follow the "IS-A" links connected to that concept. Thus memory search rules form the building blocks for generating plans to search memory. Both abstract transformations and memory search rules are provided to the system.

When memory search rules have been combined in a plan to find a particular type of information, the plan as a whole is stored in a third structure, a *memory search case.* Memory search cases store knowledge about specific episodes of memory search. They include information about the knowledge goal that motivated the memory search, the memory search plan that was used, and information about its success or failure. By retrieving memory search plans for similar previous searches, the reasoner can re-use successful reasoning and avoid unsuccessful paths when generating memory search plans for novel adaptation problems.

The fourth type of knowledge, *adaptation strategies* (Kass, 1990), stores knowledge about the entire episode of case adaptation. An adaptation strategy includes the adaptation problem being addressed, the abstract transformation used for it, and the memory search case used to find the information needed for that transformation. This allows the reasoner to re-use the information from an entire previous case adaptation when a similar problem arises. We view the availability of appropriate adaptation strategies (and, when none are available, appropriate memory search cases) as fundamental to expertise at case adaptation.

As we continue development of the model we are refining this framework for the content and organization of memory search knowledge. Because it has been shown that in some cases, the learning of control knowledge may actually degrade processing speed (Minton, 1985), another question to be examined is the factors affecting utility of learning as memory search cases are acquired.

The process for applying each knowledge source is summarized in figure 1. (This high-level description omits details concerning issues such as how to deal with failures of the memory search plans to find the needed information.) Through this process, a case-based reasoning system simultaneously adapts a case and learns information to facilitate future case adaptation. The model is being developed and extended with Andrew Kinley and David Wilson of Indiana University; the implementation of the model is discussed in (Leake, 1994a, 1995b; Leake et al., 1995).

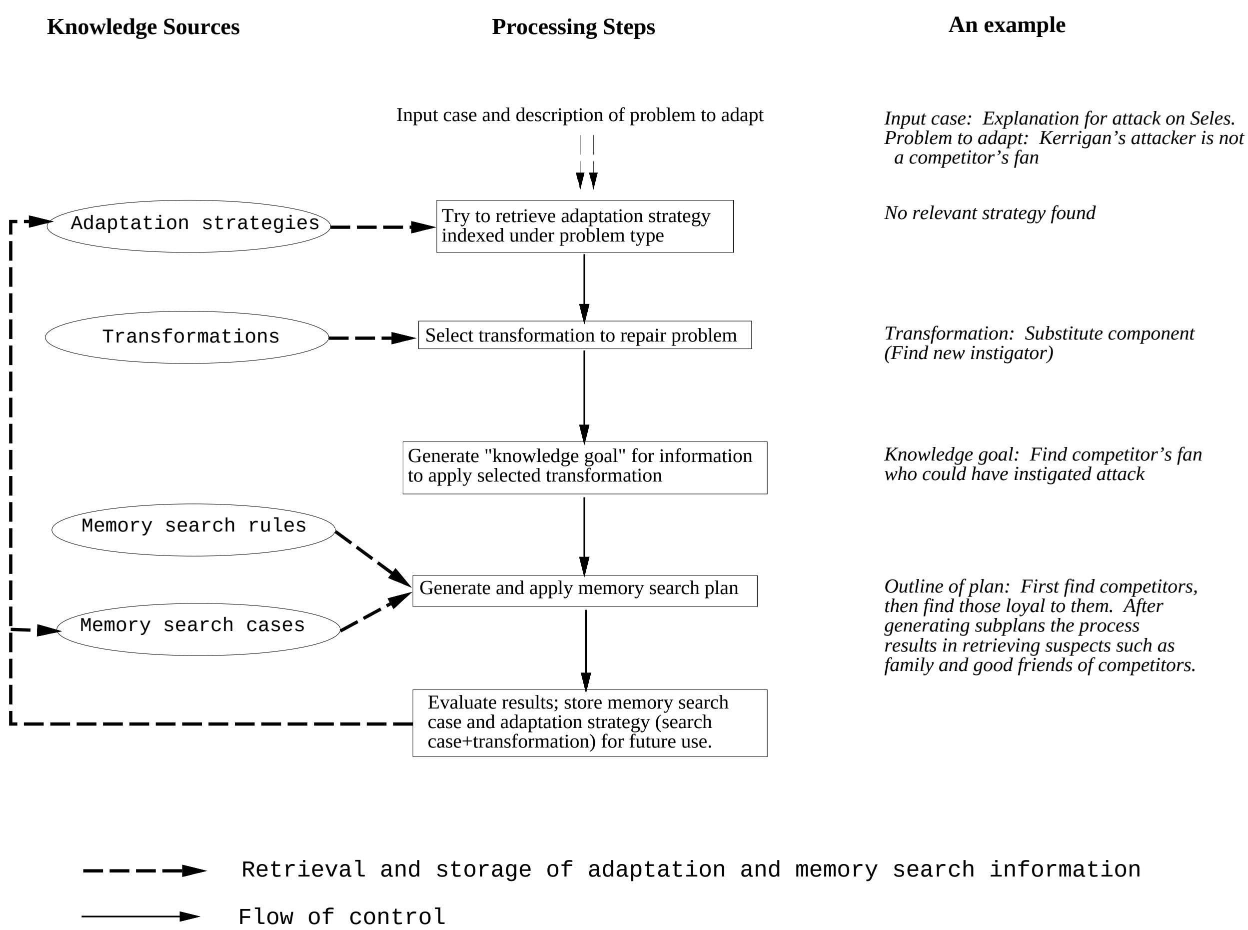

Figure 1: The basic process for learning adaptation strategies.

6.2 Acquiring retrieval expertise by introspective reasoning

In a second project, with Susan Fox at Indiana University, we are investigating how introspective analysis of the effectiveness of the case-based problem-solving process can guide refinement of retrieval criteria. As a benchmark for this self-analysis, we endow a case-based reasoning system with a model of the desired behavior of its own case-based reasoning process and the processes underlying that behavior. When its performance falls short of the desired behavior, that failure triggers learning aimed at developing performance closer to the idealized model. Our primary focus is not on the learning triggered by failures in the *outcome* of processing, which has been extensively studied in the case-based reasoning literature, but instead on how a reasoner can learn from sub-optimal reasoning performance even when it results in wasted reasoning effort rather than bad final results.

Using a model of the CBR process itself to detect processing failures: *Model-based reasoning* is a widely-studied method for diagnosing device failures (for an overview, see Davis (1988)). In that work, faults are identified by comparing a model of how the device is expected to perform to the device's actual behavior. Birnbaum, Collins, Freed and Krulwich (1991) point out that model-based reasoning can also be used in a different way: Rather than having the model reflect the expected behavior of a system, they suggest using a model of the *ideal* behavior of a reasoning system—behavior that may be beyond expectations for actual performance—as a benchmark against which to compare actual system performance. Discrepancies identify points for improvement.

Although Birnbaum et al. applied the method to the task of refining a rule-based planning system, they proposed its application to self-improving case-based reasoning systems as well (Birnbaum et al., 1991). The remainder of this section discusses the application of that method to refining indexing criteria in a case-based path planning system (Fox & Leake, 1994, 1995b, 1995c, 1995a). By this process, the system acquires expertise at case retrieval.

Identifying hidden retrieval problems by model-based reasoning: In most CBR systems, the success of reasoning is judged entirely by whether the system's solution results in a successful problem-solving outcome. When attempts to apply the system's solution result in a failure, the failure reveals that the solution was flawed and prompts the system to learn in order to avoid similar future failures (e.g., Hammond (1989)). Such learning is important, but another important issue in acquiring expertise at case-based reasoning is how to detect and learn from what we call "hidden failures"—deficiencies in the reasoning process that make the process more costly but that do not necessarily cause erroneous problem-solving results. The reason that these deficiencies may remain hidden is that each component of a case-based reasoning system can compensate to some extent for flaws in the others. For example, a correct solution can result even if the retrieved case is not the best precedent in memory, provided that case adaptation can still fit the retrieved case to the new situation. If so, the result of retrieving the wrong case is not an execution failure but instead that case adaptation is unnecessarily costly, because the retrieved case was not the one most similar to the new situation, even if the eventual solution is nevertheless correct.

To allow learning in response to flaws in the problem-solving *process,* as opposed to flaws only in the *outcome* of problem-solving, we augment traditional failure-driven learning from bad outcomes with learning based on analysis of the solution process itself. A key problem is how to perform that analysis: how to determine whether the problem-solving process was flawed even though it generated a correct solution. In general, the knowledge required to judge the problem-solving process will not be available while the problem is being solved. (If that knowledge *were* available, the reasoner could simply use it during initial problem-solving to avoid following an incorrect problem-solving path.) For example, when a CBR system attempts to retrieve a case to apply to a new situation, case retrieval must usually be based on partial knowledge of the relevant features of the problem, because the features that are relevant to retrieval may not be apparent until the problem has been solved (e.g., Kolodner (1993, pp. 371-372) and Leake (1992a, 1995a)).

However, after the problem has been solved, additional information is available, and that information can be used to determine features that should be considered during future case retrieval. Consequently, our approach uses "hindsight" to identify incorrect problem-solving paths: the information provided by the successful solution is used to illuminate how the solution should have been generated.

We are applying this approach to determining whether retrieval criteria were sufficient to retrieve the right case or if new types of features should be considered during future case retrieval. The process is as follows. After problem-solving, the execution of the case-based reasoning process is compared to assertions about its desired performance; deviations prompt learning. For example, one of the assertions about desired performance is that the retrieved case should be the case "closest" to the new situation (i.e., that it will be the case easiest to adapt to the new situation). To verify this assertion, our method compares the solution to other cases in memory, to decide whether the retrieved case was really the best case to retrieve (i.e., the case most similar to the solution). If not, it adjusts its retrieval criteria to reflect the features that should have been considered in order to retrieve the right case. In this way, it is possible to learn new similarity criteria that enable case retrieval to focus on cases that are likely to be easiest to adapt.

An example: As an illustration of this process, consider the task of generating plans for traveling within a city. A natural criterion for retrieving relevant prior plans is to retrieve the prior plans whose starting and ending points are geographically closest to the desired origin and destination points. Intuitively, it seems obvious that this should result in retrieving the plan for the path that is "closest"—easiest to adapt—to the desired path.

However, sometimes this criterion will lead a planner astray. For example, suppose that the city is divided by a river that can only be traversed at a single bridge, and the goal requires reaching a point across the river from the point of origin. For this task, a plan retrieved by the criteria in the previous paragraph might have very little in common with the final solution. Figure 2 illustrates this situation for two possible stored plans, plan A and plan B. Plan A is the plan favored by the retrieval strategy we described. However, straightforward re-use of plan A would result in arriving at a point close to the destination but on the wrong side of the river, requiring backtracking to a point near the starting point to cross the bridge and then additional adaptation to cross the bridge and reach the destination. Plan B, which is not most similar according to the retrieval criteria, better matches the solution to the problem and is a much more useful starting point towards a plan for the goal.

If the planner does not initially know that the side of the river is an important feature to consider when planning paths, it will have no way to avoid retrieving the wrong case. Once the path to the destination has been generated, however, that solution can be analyzed to

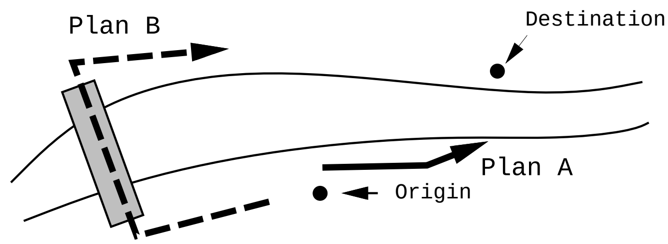

Figure 2: An illustration of the need to learn new features to guide case retrieval: How the location of bridges affects the relevance of path plans.

judge which case in memory, *given knowledge of the solution*, would have been the best starting point for generating that solution. Based on that analysis, the CBR system can explain the features that were significant—in this case, the fact that the origin and destination were on opposite sides of the river—and can use them to refine future retrieval criteria.

We performed initial tests of this type of learning in a path planning system with very limited adaptation capabilities. In tests involving a library of 20-25 cases and a randomly-selected set of origin and goal points for which to plan, the learning of new indexing criteria resulted in the system narrowing the set of prior cases that it considered relevant to a given problem (the average number of cases considered for a problem decreased by 33%), and the cases that were retrieved required less adaptation than the cases retrieved initially or when learning was restricted to acquiring new cases without learning new indexing criteria. Without introspective learning of new retrieval criteria, 25% of the plans retrieved could not be adapted by its limited adaptation component; with learning of retrieval criteria to improve the appropriateness of the cases retrieved, all plans in the experiments could be successfully adapted. For details on the experiments and results, see (Fox & Leake, 1994). More extensive experimental tests are described in (Fox & Leake, 1995a, 1995c).

6.3 Related Work

Some previous case-based reasoning systems have the capability to learn knowledge useful for guiding case adaptation. For example, the program CHEF (Hammond, 1989), which does case-based planning for the task of generating recipes, has a static library of domainindependent plan repair strategies but augments that library with learned *ingredient critics* that suggest adaptations appropriate to particular ingredients. Likewise, PERSUADER (Sycara, 1988) uses a combination of heuristics and case-based reasoning to guide adaptation, searching memory for similar prior adaptations to apply. In these systems, however, the adaptation information learned is quite domain and task specific—the cases learned must be applied to future problems in the same domain and with very specific similarities. The introspective reasoning model we are developing increases the flexibility of the learning process: The search strategies generated in response to current constraints can also be re-applied to a wider range of new situations.

In a similar spirit to our method is research by Veloso and Carbonell (1993) on storing and replaying the reasoning used to derive solutions for problems. However, that research does not address the knowledge planning problems specific to the case adaptation task. Also related, with respect to the memory search process, is the approach to flexible memory search in a heuristic search framework described in Rissland, Skalak, and Friedman (1994), but that research does not address the learning issues needed to acquire expertise in the search process.

It is also useful to contrast our approach to other methods for failure-driven refinement of reasoning criteria. For example, systems such as CHEF refine their indexing criteria when bad outcomes result from applying an incorrect case, and other projects apply introspective monitoring to detect reasoning failures in order to respond to them (Cox, 1994; Ram & Cox, 1994). The focus of that work differs from our approach to refining retrieval, however, in that their approaches trigger learning only in response to failures in the outcomes of processing steps. Our method of index refinement enables learning to occur not only in response to bad outcomes, but—and perhaps even more importantly for refining established expertise—in response to processing errors from which the reasoner successfully recovered.

7 Conclusion

Case-based reasoning research generally focuses on the role of case acquisition in the development of expertise. The benefits accrued from new cases, however, depend on the case-based reasoner's level of expertise at using its cases—on the quality of its strategies for retrieving and adapting the cases in its memory.

Modeling the role of introspective reasoning and learning in case-based reasoning is a step towards accounting for how case-based reasoners, whether human or machine, can become more expert at applying their experiences to new problems. This article describes ongoing research on developing models of how introspective reasoning about the case-based reasoning process can contribute to expertise at case retrieval and adaptation. In this research, self-knowledge concerning the tasks, memory organization, and reasoning processes of the case-based reasoning system guide learning about how to be an effective case-based reasoner. As problems are solved, the solution process itself becomes an object of learning to allow more effective use of prior cases in the future. In this way, learning from experience plays a key role in acquiring expertise at applying prior experience to new situations.

References

Agnew, N., Ford, K., & Hayes, P. (1994). Expertise in context: personally constructed, socially selected, and reality-relevant?. *International Journal of Expert Systems*, *7*(1), 65–88.

Allemang, D. (1993). Review of the first European workshop on case based reasoning EWCBR-93. *Case-Based Reasoning Newsletter*, *2*(3). Electronic newsletter published by the special interest group AK-CBR of the German Society for Computer Science.

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

Ashley, K. (1990). *Modeling legal argument: reasoning with cases and hypotheticals*. MIT Press, Cambridge.

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 Palo Alto. AAAI, Morgan Kaufmann, Inc.

Bareiss, R. (1989). *Exemplar-Based Knowledge Acquisition: A Unified Approach to Concept Representation, Classification, and Learning.* Academic Press, Inc., San Diego.

Bayles, S. & Das, B. (1994). Using artificial intelligence to support traffic flow management problem resolution. In *Proceedings of the AAAI-94 Workshop on Case-Based Reasoning*, pp. 133–137 Seattle, WA. AAAI.

Bereiter, C. & Scardamalia, M. (1993). *Surpassing Ourselves: An Inquiry into the Nature and Implications of Expertise.* Open Court, Chicago.

Berger, J. & Hammond, K. (1991). ROENTGEN: a memory-based approach to radiation therapy treatment. In Bareiss, R. (Ed.), *Proceedings of the Case-Based Reasoning Workshop*, pp. 203–214 San Mateo. DARPA, Morgan Kaufmann, Inc.

Birnbaum, L., Collins, G., Brand, M., Freed, M., Krulwich, B., & Pryor, L. (1991). A model-based approach to the construction of adaptive case-based planning systems. In Bareiss, R. (Ed.), *Proceedings of the Case-Based Reasoning Workshop*, pp. 215–224 San Mateo. DARPA, Morgan Kaufmann, Inc.

Birnbaum, L., Collins, G., Freed, M., & Krulwich, B. (1990). Model-based diagnosis of planning failures. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, pp. 318–323 Boston, MA. AAAI.

Blevis, E., Burke, R., Glasgow, J., & Duncan, N. (1991). The life analysis & depreciation integrated exemplar system (ladies). *International Journal of Expert Systems*, *4*(2), 141–155.

Branting, K. & Porter, B. (1991). Rules and precedents as complementary warrants. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, pp. 3–9 Anaheim, CA. AAAI.

Burstein, M. (1994). Case age: selecting the best exemplars for plausible reasoning using distance in time or space. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, pp. 106–111 Atlanta, GA. Cognitive Science Society.

Camerer, C. & Johnson, E. (1991). The process-performance paradox in expert judgement. In Ericsson, K. & Smith, J. (Eds.), *Toward a General Theory of Expertise*, chap. 9. Cambridge, New York.

Campbell, R. & Bickhard, M. (1986). *Knowing levels and developmental stages*. S. Karger, Basel, Switzerland.

Campbell, R., Brown, N., & DiBello, L. (1992). The programmer's burden: developing expertise in programming. In Hoffman, R. (Ed.), *The Psychology of Expertise: Cognitive Research and Empirical AI*, chap. 16, pp. 269–294. Springer Verlag, New York.

Carbonell, J. (1983). Learning by analogy: formulating and generalizing plans from past experience. In Michalski, R., Carbonell, J., & Mitchell, T. (Eds.), *Machine Learning: An Artificial Intelligence Approach*. Tioga Publishing Company, Cambridge, MA.

Chi, M., Bassok, M., Lewis, M., Reimann, P., & Glaser, R. (1989). Self-explanations: how students study and use examples in learning to solve problems. *Cognitive Science*, *13*, 145–182.

Cox, M. (1994). Machines that forget: learning from retrieval failure of mis-indexed explanations. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, pp. 225–230 Atlanta, GA.

Davis, R. (1988). Model-based reasoning: troubleshooting. In Shrobe, H. (Ed.), *Exploring Artificial Intelligence: Survey Talks from the National Conferences on Artificial Intelligence*. Morgan Kaufmann, Palo Alto.

DeJong, G. & Mooney, R. (1986). Explanation-based learning: an alternative view. *Machine Learning*, *1*(1), 145–176.

Domeshek, E. & Kolodner, J. (1992). Towards a case-based aid for conceptual design. *International Journal of Expert Systems*, *4*(2), 201–220.

Ericsson, K. & Smith, J. (1991). Prospects and limits of the empirical study of expertise: an introduction. In Ericsson, K. & Smith, J. (Eds.), *Toward a General Theory of Expertise*, chap. 1. Cambridge, New York.

Faries, J. & Schlossberg, K. (1994). The effect of similarity on memory for prior problems. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, pp. 278–282 Atlanta, GA. Cognitive Science Society.

Fox, S. & Leake, D. (1994). Using introspective reasoning to guide index refinement in case-based reasoning. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, pp. 324–329 Atlanta, GA.

Fox, S. & Leake, D. (1995a). Learning to refine indexing by introspective reasoning. In *Proceedings of First International Conference on Case-Based Reasoning* Sesimbra, Portugal. In press.

Fox, S. & Leake, D. (1995b). Planning for repairing reasoning failures. In *Proceedings of the 1995 AAAI Spring Symposium on Representing Mental States and Mechanisms* Stanford, CA. AAAI. Available from ftp.cs.indiana.edu:/pub/leake/p-95-02.ps.Z.

Fox, S. & Leake, D. (1995c). Using introspective reasoning to refine indexing. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence* Montreal. IJCAI. In press.

Fuller, S. (1994). The constitutively social character of expertise. *International Journal of Expert Systems*, *7*(1), 51–64.

Gentner, D., Ratterman, M., & Forbus, K. (1993). The roles of similarity in transfer: separating retrievability from inferential soundness. *Cognitive Psychology*, *25*, 524–575.

Gentner, D. & Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, *10*(3), 277–300.

Gick, M. L. & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology*, *12*, 306–355.

Goel, A., Callantine, T., Donnellan, M., & de Silva Garza, A. (1993). From models to cases: where do cases come from and what happens when a case is not available?. In

Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society, pp. 475–480 Boulder, CO. Cognitive Science Society.

Goswami, U. (1991). Analogical reasoning: what develops? a review of research and theory. *Child development*, *62*, 1–22.

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

Hennessey, D. & Hinkle, D. (1991). Initial results from clavier: a case-based autoclave loading assistant. In Bareiss, R. (Ed.), *Proceedings of the Case-Based Reasoning Workshop*, pp. 225–232 San Mateo. DARPA, Morgan Kaufmann, Inc.

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

Hobbs, J., Stickel, M., Appelt, D., & Martin, P. (1993). Interpretation as abduction. *Artificial Intelligence*, *63*(1-2), 69–142.

Hunter, L. (1990). Planning to learn. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, pp. 261–268 Cambridge, MA. Cognitive Science Society.

Kass, A. (1990). *Developing Creative Hypotheses by Adapting Explanations*. Ph.D. thesis, Yale University. Northwestern University Institute for the Learning Sciences, Technical Report 6.

Klein, G. & Calderwood, R. (1988). How do people use analogues to make decisions?. In Kolodner, J. (Ed.), *Proceedings of the Case-Based Reasoning Workshop*, pp. 209–223 Palo Alto. DARPA, Morgan Kaufmann, Inc.

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

Kolodner, J. (1991). Improving human decision making through case-based decision aiding. *The AI Magazine*, *12*(2), 52–68.

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

Koton, P. (1988). Reasoning about evidence in causal explanations. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, pp. 256–261 Minneapolis, MN. AAAI, Morgan Kaufmann Publishers, Inc.

Koton, P. (1989). Evaluating case-based problem solving. In Hammond, K. (Ed.), *Proceedings of the Case-Based Reasoning Workshop*, pp. 173–175 San Mateo. DARPA, Morgan Kaufmann, Inc.

Laird, J. E., Rosenbloom, P. S., & Newell, A. (1990). Chunking in Soar: the anatomy of a general learning mechanism. In Shavlik, J. W. & Dietterich, T. G. (Eds.), *Readings in Machine Learning*, pp. 555–572. Kaufmann, San Mateo, CA.

Lancaster, J. & Kolodner, J. (1987). Problem solving in a natural task as a function of experience. In *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*, pp. 727–736 Seattle, WA. Cognitive Science Society.

Lancaster, J. & Kolodner, J. (1988). Varieties of learning from problem-solving experience. In *Proceedings of the Tenth Annual Conference of the Cognitive Science Society* Montreal. Cognitive Science Society.

Leake, D. (1992a). Constructive similarity assessment: using stored cases to define new situations. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, pp. 313–318 Bloomington, IN. Cognitive Science Society.

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

Leake, D. (1993). Learning adaptation strategies by introspective reasoning about memory search. In *Proceedings of the AAAI-93 Workshop on Case-Based Reasoning*, pp. 57–63 Washington, DC. AAAI. AAAI Press technical report WS-93-01.

Leake, D. (1994a). Towards a computer model of memory search strategy learning. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, pp. 549–554 Atlanta, GA.

Leake, D. (1994b). Workshop report: the AAAI-93 workshop on case-based reasoning. *The AI Magazine*, *15*(1), 63–64.

Leake, D. (1995a). Adaptive similarity assessment for case-based explanation. *International Journal of Expert Systems*, *8*(2), 165–194.

Leake, D. (1995b). Combining rules and cases to learn case adaptation. In *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*, pp. 84–89 Pittsburgh, PA.

Leake, D. (1995c). Representing self-knowledge for introspection about memory search. In *Proceedings of the 1995 AAAI Spring Symposium on Representing Mental States and Mechanisms*, pp. 84–88 Stanford, CA. AAAI. Available from ftp.cs.indiana.edu:/pub/leake/p-95-01.ps.Z.

Leake, D., Kinley, A., & Wilson, D. (1995). Learning to improve case adaptation by introspective reasoning and cbr. In *Proceedings of First International Conference on Case-Based Reasoning* Sesimbra, Portugal. In press.

Leake, D. & Ram, A. (1993). Goal-driven learning: fundamental issues (a symposium report). *The AI Magazine*, *14*(4), 67–72.

LeFevre, J. (1988). Processing instructional texts and examples. *Canadian Journal of Psychology*, *41*, 351–364.

Lesgold, A., Rubinson, H., Feltovitch, P., Glaser, R., Klopfer, D., & Wang, Y. (1988). Expertise in a complex skill: diagnosing X-ray pictures. In Chi, M., Glaser, R., & Farr, M. (Eds.), *The nature of expertise*, chap. 11. Erlbaum, Hillsdale, NJ.

Minton, S. (1985). Selectively generalizing plans for problem-solving. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pp. 596–599 Los Angeles, CA. IJCAI.

Mitchell, T., Keller, R., & Kedar-Cabelli, S. (1986). Explanation-based generalization: a unifying view. *Machine Learning*, *1*(1), 47–80.

Mostow, J. (1983). Machine transformation of advice into a heuristic search procedure. In Michalski, R., Carbonell, J., & Mitchell, T. (Eds.), *Machine Learning: An Artificial Intelligence Approach*, pp. 367–403. Tioga Publishing Company, Cambridge, MA.

Novick, L. (1988). Analogical transfer, problem similarity, and expertise. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *14*, 510–520.

Piaget, J. (1976). *The grasp of consciousness.* Harvard University Press, Cambridge, MA.

Pirolli, P. & Anderson, J. (1985). The role of learning from examples in the acquisition of recursive programming skills. *Canadian Journal of Psychology*, *39*, 240–272.

Ram, A. & Cox, M. (1994). Introspective reasoning using meta-explanations for multistrategy learning. In Michalski, R. & Tecuci, G. (Eds.), *Machine Learning: A Multistrategy Approach.* Morgan Kaufmann.

Ram, A. (1987). AQUA: asking questions and understanding answers. In *Proceedings of the Sixth Annual National Conference on Artificial Intelligence*, pp. 312–316 Seattle, WA. Morgan Kaufmann Publishers, Inc.

Read, S. & Cesa, I. (1991). This reminds me of the time when ...: expectation failures in reminding and explanation. *Journal of Experimental Social Psychology*, *27*, 1–25.

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.

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

Rissland, E., Skalak, D., & Friedman, M. (1994). Heuristic harvesting of information for case-based argument. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, pp. 36–43 Seattle, WA. AAAI.

Ross, B. (1989). Some psychological results on case-based reasoning. In Hammond, K. (Ed.), *Proceedings of the Case-Based Reasoning Workshop*, pp. 144–147 San Mateo. DARPA, Morgan Kaufmann, Inc.

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

Schank, R. & Leake, D. (1989). Creativity and learning in a case-based explainer. *Artificial Intelligence*, *40*(1-3), 353–385. Also in Carbonell, J., editor, *Machine Learning: Paradigms and Methods*, MIT Press, Cambridge, MA, 1990.

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

Segre, A. (1988). *Machine learning of robot assembly plans.* Kluwer Academic Publishers, Boston.

Seifert, C. (1988). A retrieval model for case-based memory. In Kolodner, J. (Ed.), *Proceedings of the Case-Based Reasoning Workshop*, pp. 120–125. DARPA.

Simoudis, E., Ford, K., & Canas, A. (1992). Knowledge acquisition in case-based reasoning: "...and then a miracle happens". In Dankel, D. (Ed.), *Proceedings of the 1992 Florida AI Research Symposium.* FLAIRS.

Simoudis, E. & Miller, J. (1991). The application of cbr to help desk applications. In Bareiss, R. (Ed.), *Proceedings of the Case-Based Reasoning Workshop*, pp. 25–36 San Mateo. DARPA, Morgan Kaufmann, Inc.

Slator, R. & Riesbeck, C. (1991). Taxops: a case-based advisor. *International Journal of Expert Systems*, *4*(2), 117–140.

Sternberg, R. & Frensch, P. (1992). On being an expert: a cost-benefit analysis. In Hoffman, R. (Ed.), *The psychology of expertise: Cognitive research and empirical AI*, chap. 11, pp. 191–203. Springer-Verlag, New York.

Sycara, K. (1987). *Resolving Adversarial Conflicts: An Approach Integrating Case-based and Analytic Methods.* Ph.D. thesis, School of Information and Computer Science, Georgia Institute of Technology. Georgia Institute of Technology, Technical Report GIT-ICS-87/26.

Sycara, K. (1988). Using case-based reasoning for plan adaptation and repair. In Kolodner, J. (Ed.), *Proceedings of the Case-Based Reasoning Workshop*, pp. 425–434 Palo Alto. DARPA, Morgan Kaufmann, Inc.

VanLehn, K. (1989). Problem solving and skill acquisition. In Posner, M. (Ed.), *Foundations of Cognitive Science*. MIT Press, Cambridge, MA.

Veloso, M. (1994). *Planning and Learning by Analogical Reasoning*. Springer Verlag, Berlin.

Veloso, M. & Carbonell, J. (1993). Derivational analogy in prodigy: automating case acquisition, storage, and utilization. *Machine Learning*, *10*(3), 249–278.