Towards A Computer Model of Memory Search Strategy Learning

David B. Leake
Computer Science Department
Indiana University
Bloomington, Indiana 47405
leake@cs.indiana.edu

Abstract

Much recent research on modeling memory processes has focused on identifying useful indices and retrieval strategies to support particular memory tasks. Another important question concerning memory processes, however, is how retrieval criteria are learned. This paper examines the issues involved in modeling the learning of memory search strategies. It discusses the general requirements for appropriate strategy learning and presents a model of memory search strategy learning applied to the problem of retrieving relevant information for adapting cases in case-based reasoning. It discusses an implementation of that model, and, based on the lessons learned from that implementation, points towards issues and directions in refining the model.

Introduction

Much recent AI research on memory focuses on analyzing the indices that are relevant to particular classes of retrieval problems (e.g., (Domeshek, 1992; Leake, 1992; Owens, 1991)). The problem of how memory search strategies can be learned and refined has received less attention.[1] Understanding the process of learning memory search strategies is important both for practical reasons, to develop AI systems that can refine memory search criteria as needed, and to extend cognitive models of memory processes. For example, data on childrens' memory strategies show a progression from ineffective to effective memory strategies during early development (Wellman, 1988).

This paper examines the issues involved in modeling the learning of memory search strategies. It focuses on the memory task of retrieving the information that case-based reasoning systems need in order to adapt prior cases to fit new situations. The paper describes a broadly-applicable framework that treats memory search as a planful process and learns the search plans resulting from information search for particular adaptation problems. It discusses the general requirements for appropriate strategy learning within that framework, and presents the lessons learned from ongoing research in applying the model to memory search learning for adapting cases in case-based explanation (Leake, 1993). Based on those lessons, it identifies directions for extending, refining, and validating the performance of the model.

[1] But see Fox & Leake, 1994; Oehlmann, Edwards, & Sleeman, 1994; Redmond, 1992; and Sycara & Navinchandra, 1989, for examples of recent research addressing issues of index learning, refinement and re-indexing.

Memory Search for Case Adaptation

Case-based reasoning (CBR) systems solve problems by *retrieving* similar stored cases and *adapting* their solutions to fit the new situations. The CBR model of reasoning from experience has been successfully applied to many tasks (see (Kolodner, 1993) for a survey), and initial studies are encouraging for its validity as a cognitive model (e.g., (Lancaster & Kolodner, 1988; Ross, 1989; Read & Cesa, 1991)).

A fundamental question in CBR is how to guide the adaptation process that revises a retrieved case to fit a new situation. We view case adaptation knowledge as having two components. The first is knowledge of abstract transformations, such as the transformation *substitute component*. The second is memory search information on how to find the information needed to apply the abstract transformations—for example, in the case of the transformation *substitute component*, information on how to determine which features are relevant to searching for an appropriate substitution and how to conduct the search. Previous research has shown that because of the generality of abstract transformations, a wide range of adaptation problems can be solved using a small set of abstract transformations (e.g., (Carbonell, 1983)). However, a difficult problem is how to determine the features that are relevant in finding the information needed to apply the abstract rules to specific problems. For example, the rule *substitute component* gives no guidance about *what* to substitute.

An example of human learning of relevance criteria for case adaptation is presented in (Gentner, 1988). Gentner's experiments gave children the task of adapting previously-encountered stories to fit new characters. In the experiments, children first acted out stories, using toys to play the roles of the characters. They were then asked to act out the same stories using different toys representing new characters. Although both older children (8-10 years old) and younger children (5–7 years old) were influenced by the transparency of the object mappings between the corresponding characters in the old and new stories, considerations of systematicity—a higher-level feature—aided the older children in making the correct mappings between characters as they adapted the old stories to the new characters. The process we are investigating is (1) how new criteria for determining relevant features for finding adaptation information can be acquired during memory retrieval, and (2) how those criteria can guide the memory search process for the needed information.

Memory search as a planful process

Our approach builds on the model of memory traversal and index elaboration in CYRUS (Kolodner, 1984), and especially on prior proposals for introspective failure-driven learning to repair memory organization problems (Birnbaum, Collins, Brand, Freed, Krulwich, & Pryor, 1991; Ram & Cox, 1994). However, this model differs in treating memory search as a *knowledge planning* process (Hunter, 1990). In the knowledge planning framework, information search is conducted by a planning process based on explicit reasoning about needs for information and how to satisfy them. In that process, a reasoner formulates explicit *knowledge goals* (Leake & Ram, 1993; Ram, 1987) that are pursued by explicit reasoning about the goals it needs to achieve and the available methods for achieving those goals. To allow flexible memory search, we are modeling that process within a similar framework (Leake, 1993).

Our model generates explicit goals to acquire needed information, based on the abstract adaptation rule it is trying to apply and on constraints arising from the case-based reasoner's task and the specific case to which the rule will be applied. The goal reflects a determination of types of features that are relevant for the adaptation problem. Based on the goals that are generated, the model reasons introspectively about its memory organization and uses a planning process to generate a plan for how to search memory for the needed information. This plan reflects a relevance judgement, the judgement of which features of the adaptation problem are relevant to retrieving the needed information. Once a successful memory search plan has been generated, the search information is stored as a new case in memory for future use guiding search in similar adaptation problems. Thus it takes a case-based approach to guiding the memory search process for case-adaptation.

Generating knowledge goals for adaptation

Guiding the memory search process depends on first describing the goal of that search. The model assumes that input to the adaptation component will be expressed in a fixed, structured vocabulary describing the problems that necessitate case adaptation. For the current task, adaptation of implausible explanations, the vocabulary used categorizes problems such as non-normative role fillers in a schema (e.g., in the example we will consider later, the hypothesis that a horse is jogging), ordering problems (e.g., that a hypothesized cause happened before the event it ostensibly caused), or the speedup or delay in an expected sequence of events. For a fuller description of failure categorizations for explanations, see (Leake, 1991, 1992).

Based on the problem description, the adaptation component must select a repair strategy to repair the problem. The choice of the repair strategy for an explanation depends not only on the particular problem—e.g., the implausible event—but also the role it plays in the explanation as a whole, as evidence for other beliefs. Likewise, the intended use of the explanation can place additional constraints on the adaptation. Consequently, the generation of a knowledge goal for adapting an implausible belief depends on three things: the reason the belief is implausible, the repair strategy to be applied, and the context in which it is applied. Each of these properties can constrain the retrieval of stored prior adaptation cases if they are available.

If no stored adaptation cases are retrieved, a new memory search plan must be generated, i.e., an abstract transformation must be selected and needed information for that transformation must be found. What information will be needed depends on the combination of the transformation strategy and constraints arising from relationships between the transformed case and the objects used in the transformation. For example, if the problem to be repaired is that an explanation involves implausible evidence, multiple transformations are possible, such as adding additional support for the implausible evidence, substituting more plausible evidence, or simply deleting the support. If the chosen transformation is to replace the implausible evidence with other evidence, being able to apply that transformation will depend on finding a substitution that both avoids the current problem—i.e., that is itself plausible—and that is evidence for the same conclusion.

Koton (1988) has identified general criteria for deciding which abstract transformations to apply to flawed explanations, and those criteria form a starting point for selection of abstract transformations. However, in general it is difficult to anticipate which types of transformations will apply in a given case, and many conflicting factors may enter into the decision (see (Kolodner, 1993) for an extensive discussion of adaptation issues). The difficulty of balancing conflicting factors provides a strong functional motivation for developing a model in which experience can help to balance those conflicting factors to guide case adaptation.

Introspective reasoning for memory search

Once knowledge goals have been generated, introspective reasoning must be used to generate a plan for finding that knowledge in memory. That reasoning requires self-knowledge of the reasoner's memory organization. In order to perform knowledge planning, initial system knowledge must include information about the relationships between concepts in memory—not just the named links, as in most memory systems, but the *meanings* of those links. The task of providing the system with sufficient information is facilitated by the fact that only local relationships between concepts need to be specified; more distant relationships can be derived on demand using the knowledge planning process. In addition, because this information is task-independent, once it is established it can be applied in multiple contexts and for multiple tasks.

The knowledge of the meanings of links can be augmented with domain knowledge, as well as the knowledge of relevance built through experience with adaptation problems, in order to determine the search strategies to apply to identify particular types of information. It should be noted, however, that although that domain knowledge is useful, it is not essential. In the absence of domain knowledge, the initial search for information becomes an unguided *local search* process of the sort already commonly used to support case adaptation (Kolodner, 1993, pp. 407-410), but the process can still learn to improve its search process by favoring search plans that were successful in similar prior situations.

Treating the memory search process this way increases flexibility in selection, application and refinement of memory search strategies. For example, components of effective strategies can be combined as needed, making it possible to build complex memory search strategies given only local information about memory organization. In addition, because standard plan learning methods can be applied to the memory search plans that are generated, retrieval strategies can be refined with experience.

Credit assignment and adaptation cost issues

Two credit assignment problems arise from this model. The first is the problem of knowing whether, if no previous adaptation strategies are sufficient to adapt a prior case, the problem should be addressed by attempting to find a new and closer case (e.g., by querying a human user), or instead by learning new adaptation procedures. Our work has not attempted to address this question: The model simply attempts to generate an adaptation strategy for each case for which no stored strategy exists or for which strategies that are indexed as relevant fail to apply.

The second credit assignment problem concerns memory search failures. When memory search fails, there is no way (short of exhaustive search, which the model does not use) to determine whether the search plan itself was flawed or the needed information was simply missing from memory. A limit on the amount of search effort (measured in memory links followed) provides a criterion for when a particular search plan should be terminated, and can also provide a cut-off point for the set of search plans to be attempted to find a given piece of information. Reaching that limit is not in itself enough to determine whether the wrong search strategy was chosen, but if one memory search plan fails and a later plan succeeds in finding the needed information, the failure of previously-attempted search plans *can* be ascribed to inapplicability of the search plans that were tried initially. In that case, if the search plan was retrieved from memory it can be re-indexed to avoid being retrieved for similar knowledge goals in the future.

Note that even if no search plan succeeds in finding the needed information within the search limit, when solving a future adaptation problem it may be possible to find that information using new memory search plans that were discovered and stored in the interim.

The Computer Model

The theory is being tested in the computer system AL (Adaptation Learner). That system applies the planning framework described in the previous sections to the task of learning the search strategies hand coded as part of the *adaptation strategies* (Kass, 1990) in the case-based explanation systems SWALE (Schank, Riesbeck, & Kass, 1994) and ABE (Kass, 1990). To provide context and illustrate the role of adaptation strategies in those systems, we sketch SWALE's processing of its namesake example, the story of the racehorse Swale.[2]

Swale was a 3-year-old superstar racehorse who died unexpectedly at the peak of his career. When SWALE detects the anomaly of Swale's death, one of the explanatory cases SWALE retrieves is the episode of the death of the runner Jim Fixx, who, like Swale, died when in peak physical condition. The retrieved explanation for Fixx's death is that Fixx was doing recreational jogging, leading to a high exertion level that overtaxed a hidden heart defect, leading to a fatal heart attack. That explanation does not apply directly to Swale: Swale was not a recreational jogger.

In a system using abstract adaptation rules, the adaptation rule to apply would be *substitute evidence*, which can be used in any domain but gives no guidance as to how to find the evidence to substitute. In a system using very specific rules tailored to the domain of horse-racing, an applicable rule for finding the evidence might be *when a racehorse's exertion must be substantiated, horse racing is a likely cause of the exertion*—a rule that is easy to apply but that has no applicability to other types of actions and actors. Instead of either rule, SWALE and ABE use an adaptation strategy called *Replace-action: Use agent theme links,* which suggests trying to find substitute actions by examining actions habitually associated with the actor. This search strategy is implemented as a procedure that retrieves the filler of the role-theme role in the schema describing the actor.[3] One of the role themes stored in the system's memory is that racehorses run in races. Consequently, when the strategy is applied, it finds running in races as a candidate for the substitution.[4] Replacing jogging with horse racing as the cause of exertion that overtaxed a heart defect leads to a plausible explanation for why Swale died despite appearing to be in outstanding physical condition.

The learning process

The method for adaptation strategy learning starts with a small library of abstract adaptation rules giving generalized coverage of possible adaptations, and operationalizes them by combining them with memory search plans and storing them for future use. The basic process of generating a new adaptation strategy involves four steps:

1. Input a case and a description of a problem to be solved by adaptation.
2. Attempt to retrieve relevant existing adaptation strategies. If success, done—no new strategy is needed. If failure, retrieve abstract adaptation rule.
3. Use an introspective knowledge planning process to generate memory search plans to operationalize the abstract rule.
4. Generalize and package the search plan with the rule as a new adaptation strategy.

This process is summarized in figure 1.

A program example

AL is implemented as a new component of the program microSWALE (Schank et al., 1994), a distillation of the SWALE system. The following section illustrates the algorithm with a sketch of how it applies to the problem of adapting the Jim

[2] A more detailed discussion of this example is contained in (Schank et al., 1994), which addresses additional issues of case retrieval and evaluation that are irrelevant to the current discussion.

[3] *Role themes* (Schank & Abelson, 1977) represent stereotyped knowledge about the plans and goals associated with actors in certain societal roles, such as the knowledge that a racehorse runs in races or that a policeman performs actions such as directing traffic and investigating crimes.

[4] It also finds the irrelevant theme action of eating oats, which leads to the explanation that Swale's death was caused by the exertion of eating oats. That explanation is rejected as implausible.

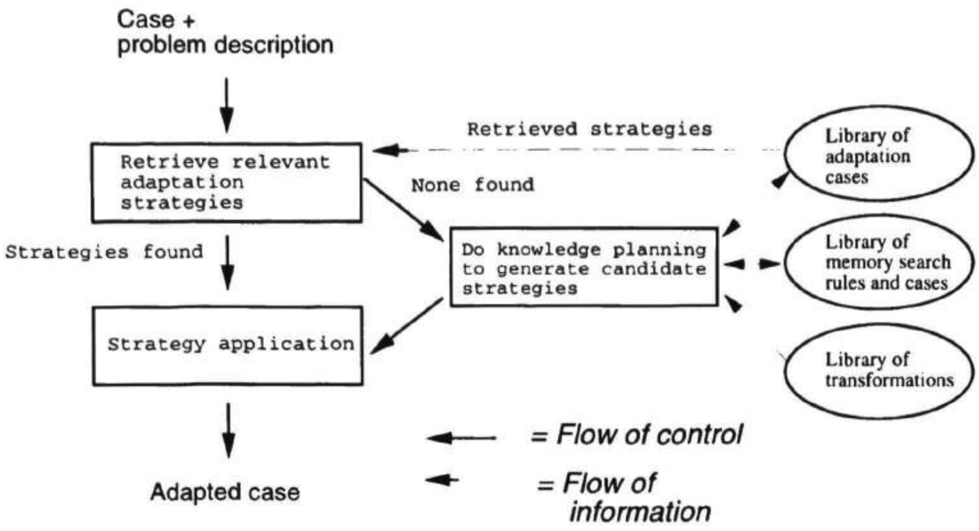

Figure 1: The adaptation strategy learning process.

Fixx explanation to Swale. In the process, the system learns an adaptation strategy corresponding to *Replace action: Use agent-theme links,* starting with only abstract adaptation rules and knowledge of the system's memory organization.

In attempting to apply the Jim Fixx explanation to Swale's death, microSWALE detects a problem and provides the adaptation component with a description of the problem: Swale is a *non-normative-role-filler* for the role of *actor* in the program's schema for *jogging*. It attempts to retrieve existing adaptation strategies indexed as relevant to the problem description. The adaptation system starts only with abstract adaptation rules without memory search information. In the current implementation, adaptation rules are stored as in a hierarchy of MOPs (Schank, 1982) containing (a) a description of the classes of problems that they address and (b) procedures that take a problem characterization and explanation as input, generate knowledge goal descriptions based on those inputs, and call a memory search module with the knowledge goal to attempt to retrieve or generate an appropriate memory search plan to find the needed information for the transformation to be applied.

For the example of Swale's jogging, possible candidate abstract adaptations to repair the problem include *add support,* to add additional supports for the hypothesized event of the racehorse jogging, *remove evidence,* to simply delete the offending belief, and *substitute evidence,* to find a replacement support for the act of jogging. For this example, *add support* does not apply—no supports exist for a racehorse jogging—and *remove evidence* immediately results in an explanation with insufficient support. This leaves *substitute evidence* as the only remaining candidate. The role of jogging in the explanation imposes an additional constraint: That the substituted component must support the hypothesized exertion.

The system generates a *knowledge goal* to find evidence to replace "Swale jogging," the problem in the Fixx explanation, and support Swale's exertion. After generating that goal, it attempts to generate a memory search plan for satisfying it. The following output shows highlights from this process. (Minor editing has been done for readability.)

```
Generating knowledge goal to guide
search for substitute evidence:
"Find substitute support for
M-EXERT-EVENT-FOR-M-SWALE-22".

Generating search plans for satisfying
"Find substitute support for
M-EXERT-EVENT-FOR-M-SWALE-22".
```

The knowledge goal is passed to the memory search planner, which generates a chain of steps for finding the needed information. The current process is hierarchical: The plan is first characterized in terms of abstract operators which are specified until the plan is described in terms of directly-executable steps. It generates a plan, applies it successfully, and packages it as a new memory search rule:

```
Packaging search plan
NEW-M-MEMORY-SEARCH-RULE-27:

"To find a cause for an actor's state,
  search for an action performed
  by the actor that could cause that
  state"

"To find actions performed by an actor,
  check the actor's theme actions"

"To find themes of an actor, retrieve
  the value of the 'theme' slot for the
  MOP for that actor"

"To find the value of a slot of a MOP,
  apply the function 'get-slot-value'
  to the MOP and slot name".
```

That plan, combined with the abstract rule *substitute evidence,* provides a result equivalent to SWALE's adaptation strategy *Replace-action: Use agent theme links.* Thus the results of the process are both a solution to a particular adaptation problem and a new adaptation strategy that can be applied to a wide range of future situations. The adaptation strategy is stored in memory, indexed by microSWALE's initial description of the adaptation problem for which it was generated.

Lessons Learned

Although the described system is in preliminary form, it has brought to light some notable points. The first concerns the types of learning mechanisms that apply to the task. The system was initially envisioned as performing explanation-based generalization (EBG) (Mitchell, Keller, & Kedar-Cabelli, 1986) on the search plans that it generated. That approach seemed reasonable because the model can be viewed as operationalizing a set of abstract transformation rules with search procedures, and EBG is a standard operationalization method (Keller, 1988). However, EBG is not appropriate to the memory search task because a deductive explanation of why a memory search plan succeeds would be required to guide generalization. The memory search process reasons using heuristics; whether they succeed in a given instance depends not only on the rules themselves but on the idiosyncratic contents of memory. Thus unlike the explanation-based generalization process, the learning problem for memory search strategies involves specifying unreliable general rules, in light of experience, to learn specific information that is more reliable.

This concern suggests applying case-based reasoning to the memory search task itself: Case-based reasoning is a method for learning in imperfectly-understood domains. Consequently, developing a case-based model of the learning process, to automatically build up a library of useful strategies from specific experiences applying memory search heuristics, is an important current direction of this research. Like the work on derivational analogy described in (Veloso & Carbonell, 1993), the process will store and re-play successful memory search plans. Because strategies will be applied to similar future situations, the case-based memory search approach can build up knowledge reflecting the idiosyncratic contents of a particular memory.

A related issue is how to organize memory search plans in memory. In the initial system design, memory search information was packed with the transformations it operationalized to form adaptation strategies along the lines of those developed in (Kass, 1990). It has become clear that fully exploiting prior learning depends on being able to apply memory search knowledge in new contexts, necessitating storing and indexing memory search plans independently from the adaptations for which they are used, even though the combinations are useful for memory search problems similar to previous tasks. Flexible application of memory search plans in novel situations requires developing a general vocabulary to characterize the information needed for different types of case adaptation problems, in order to use that characterization to select memory search plans indexed by the information that they provide. Additional questions raised by the current model and being investigated include the level at which to represent the information in memory search rules initially provided to the system, and how to detect, learn from, and recover from particular classes of memory search failures.

The utility of adaptation learning

The previous issues concern simply building a model of case adaptation learning, but another important concern is the effectiveness of such a model. In principle an enormous number of alternative memory searches could be tried for a particular case adaptation, making the process potentially very costly, but the limit on adaptation effort imposes some control on the cost of adaptation search. Nevertheless, the model may still require considerable effort to learn new adaptation strategies, and even after they have been learned, it is not guaranteed that they will improve overall performance, because of the *utility problem:* It has been shown that the learning of control knowledge may actually degrade the performance of the system using that knowledge (Minton, 1990), due to increased overhead costs overwhelming the savings from individual rules. The intention of this model is to address that problem by indexing learned memory search cases for efficient access, but the effectiveness of that approach must be validated.

Another aspect of the utility of adaptation learning, however, is the quality of the result of adaptation. In many of the domains used by case-based reasoning, system knowledge is incomplete, making adaptation rules unreliable. Applying similar adaptations to similar problems may help to improve the quality of the solutions generated by the adaptation system, just as case-based reasoning can improve the quality of solutions in imperfectly-understood domains (Kolodner, 1993). An analysis of the effectiveness of this model must also examine its effect on the quality of adaptations.

Conclusions

This paper proposes treating memory search as a planful process guided by explicit reasoning about needs for information and the organization of memory. The memory search plans resulting from that process are then learned for future use. Current results suggest using case-based reasoning as the basic for this memory search process, and point to key issues to be addressed in future research on treating memory search as a planful process.

Although our memory search framework has been illustrated in the context of case adaptation for case-based explanation construction, its approach to planful memory search has wider applicability. For example, abstract adaptation rules of the type that this framework requires as a starting point also apply to other tasks such as case adaptation for case-based planning. Much more generally, the planful memory search process has wide applicability to memory search problems outside of the context of case adaptation. The development of an introspective model of planful memory search is important for enabling memory systems to refine their performance by generating, re-using and refining memory search strategies in response to their needs. It is also a first step towards a cognitive model of the developmental process for relevance criteria and memory search strategies.

Acknowledgements

I would like to thank the anonymous reviewers for their helpful comments on a draft of this paper.

References

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.

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.

Domeshek, E. (1992). *Do the Right Things: A Component Theory for Indexing Stories as Social Advice*. Ph.D. thesis, The Institute for the Learning Sciences, Northwestern University.

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* Atlanta, GA.

Gentner, D. (1988). Metaphor as structure mapping: the relational shift. *Child Development*, *59*, 47–59.

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.

Keller, R. (1988). Defining operationality for explanation-based learning. *Artificial Intelligence*, *35*(2), 227–241.

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

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.

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. (1991). An indexing vocabulary for case-based explanation. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, pp. 10–15 Anaheim, CA. AAAI.

Leake, D. (1992). *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. & Ram, A. (1993). Goal-driven learning: fundamental issues (a symposium report). *The AI Magazine*, *14*(4), 67–72.

Minton, S. (1990). Quantitative results concerning the utility of explanation-based learning. *Artificial Intelligence*, *42*, 363–391.

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

Oehlmann, R., Edwards, P., & Sleeman, D. (1994). Changing the viewpoint: re-indexing by introspective questioning. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society* Atlanta, GA.

Owens, C. (1991). *Indexing and retrieving abstract planning knowledge*. Ph.D. thesis, Yale University.

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. In Press.

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. American Association for Artificial Intelligence, 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.

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. (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 New Jersey.

Sycara, K. & Navinchandra, D. (1989). Index transformation and generation for case retrieval. In Hammond, K. (Ed.), *Proceedings of the Case-Based Reasoning Workshop*, pp. 324–328 San Mateo. DARPA, Morgan Kaufmann, Inc.

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

Wellman, H. (1988). The early development of memory strategies. In Weinert, F. & Perlmutter, M. (Eds.), *Memory Development: Universal Changes and Individual Differences*, chap. 1. Lawrence Erlbaum, Hillsdale, NJ.