Invited paper, Proceedings of the Eighth Annual Florida Artificial Intelligence Research Symposium, 1995, pp. 112-116.

BECOMING AN EXPERT CASE-BASED REASONER: LEARNING TO ADAPT PRIOR CASES*

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

Abstract

Experience plays an important role in the development of human expertise. One computational model of how experience affects expertise is provided by research on case-based reasoning, which examines how stored *cases* encapsulating traces of specific prior problem-solving episodes can be retrieved and re-applied to facilitate new problem-solving. Much progress has been made in methods for accessing relevant cases, and case-based reasoning is receiving wide acceptance both as a technology for developing intelligent systems and as a cognitive model of a human reasoning process. However, one important aspect of case-based reasoning remains poorly understood: the process by which retrieved cases are *adapted* to fit new situations. The difficulty of encoding effective adaptation rules by hand is widely recognized as a serious impediment to the development of fully autonomous case-based reasoning systems. Consequently, an important question is how case-based reasoning systems might learn to improve their expertise at case adaptation. We present a framework for acquiring this expertise by using a combination of general adaptation rules, introspective reasoning, and case-based reasoning about the case adaptation task itself.

1 INTRODUCTION

In common parlance, "expertise" may mean a number of different things. It has been noted that sometimes the attribution of expertise is based on social criteria, depending simply on whether society bestows the status of "expert" on a particular person (e.g., Agnew, Ford, & Hayes, 1994; Fuller, 1994; Sternberg & Frensch, 1992). Sometimes the ascription of expertise is based on performance criteria, with the judgment of expertise depending on factors such as problem-solving speed and the quality of the solutions attained. It has also been observed that expertise is sometimes implicitly equated with experience, as in the novice/expert dichotomy often proposed by psychological studies (Campbell, Brown, & DiBello, 1992). This paper focuses on the relationship between performance and experience: how "expert level" performance develops from problem-solving experience.

Case-based reasoning (CBR) (e.g., Kolodner, 1993a; Riesbeck & Schank, 1989) is a paradigm that models the role of experience in refining problem-solving ability. In a case-based reasoner, expert-level performance arises primarily from reuse of cases from a library of previous problem-solving episodes. CBR systems solve new problems by *retrieving* traces of similar previous problem-solving and performing *case adaptation* to fit the previous solutions to the new situation. The adapted cases are then added to the case library for future use, making learning an integral part of the CBR process.

Acquisition of cases is a natural mechanism for knowledge acquisition: rather than needing to extract the principles underlying a domain, a reasoner (or knowledge engineer building a reasoning system) can store a case as a complete unit for future reuse. Many reports suggest that this significantly alleviates the knowledge acquisition problem. Case-based reasoning applications are finding wide acceptance, with over 100 currently fielded applications (Kolodner, 1993b).

However, as Simoudis, Ford, and Canas (1992) point out, CBR does not make all knowledge acquisition problems disappear. Storing and reusing cases obviates having to define domain rules, but new knowledge acquisition problems emerge when developers of CBR systems must define the knowledge needed to retrieve and adapt stored cases.

Although much progress has been made in methods for case retrieval, case adaptation remains poorly understood. The problems with endowing case-based reasoning systems with automatic case adaptation are so acute that leaders in the field have often argued for postponing the inclusion of adaptation components in current CBR applications (e.g., Kolodner, 1991).

We are addressing the problem of acquiring case adap-

*This material is based upon work supported by the National Science Foundation under Grant No. IRI-9409348. The case adaptation research described is being conducted with David Wilson and Andrew Kinley at Indiana University.

tation knowledge by developing methods for learning case adaptation from experience. This is part of a larger effort towards enabling case-based reasoning systems to learn to make better use of their cases—augmenting the domain expertise embodied in their libraries of problem-solving cases with increased expertise at the case-based reasoning process itself. This effort applies introspective reasoning and learning methods to improve memory organization/retrieval criteria (Fox & Leake, 1994, 1995), and case adaptation (Leake, 1993, 1994). After sketching some psychological perspective on case-based reasoning as a model of expertise we summarize ongoing research on how a case-based reasoner can acquire expertise at case adaptation as a natural part of its case adaptation process.

2 CASES AND EXPERTISE

Reasoning from specific prior experiences has been shown to play an important role both in learning about new domains (e.g., Faries & Schlossberg, 1994; LeFevre, 1988; Pirolli & Anderson, 1985; Ross, 1989), and in performance by expert practitioners. For example, expert decision-makers in complex situations use analogous cases to suggest starting points for problem-solving and to help evaluate candidate solutions (Klein & Calderwood, 1988), and expert mechanics use specific experiences to help generate hypotheses about problems and to help select appropriate tests for diagnosing them (Lancaster & Kolodner, 1987).

However, experience alone does not fully account for expert performance: People with comparable levels of experience may exhibit very different levels of skill (e.g., Bereiter & Scardamalia, 1993; Ericsson & Smith, 1991; Lesgold, Rubinson, Feltovitch, Glaser, Klopfer, & Wang, 1988).

Likewise, in case-based reasoning systems, a sufficient library of cases is only one of the requirements for successful performance; other crucial factors are how those cases are organized, the strategies for retrieving them, and how they are adapted to fit new circumstances.

3 THE KNOWLEDGE ACQUISITION PROBLEM FOR CASE ADAPTATION KNOWLEDGE

Hand-crafting effective case adaptation rules is a difficult and time-consuming task that has proven a serious impediment to developing autonomous case-based reasoning systems. As a result, many CBR systems now function as retrieval-only decision-support systems, relying on their human users to adapt their cases (Blevis, Burke, Glasgow, & Duncan, 1991; Domeshek & Kolodner, 1992; Hennessey & Hinkle, 1991; Kitano, Shibata, Shimazu, Kajihara, & Sato, 1992; Simoudis & Miller, 1991; Slator & Riesbeck, 1991). For example, one such system is an aiding and advisory system for architects (Domeshek & Kolodner, 1992); it presents relevant cases to architects who themselves decide how to apply the cases to new architectural design problems. Such decision-support systems have proven a very useful application of CBR, but developing fully automated CBR systems that can be used by expert and naive users alike depends on providing case-based reasoning systems with appropriate adaptation knowledge.

A few CBR systems have limited capabilities for learning to refine their adaptation capabilities. For example, the case-based planner CHEF (Hammond, 1989) has a static library of domain-independent plan repair strategies, and augments that library with learned *ingredient critics* for suggesting adaptations appropriate to particular ingredients. However, more general methods are needed.

The difficulties faced by developers of case adaptation rules are precisely the classic knowledge acquisition problems of non-learning rule-based systems in general. In other contexts, those problems have been significantly ameliorated by the use of case-based reasoning; this suggests a case-based approach to case adaptation as well (Berger, 1995; Leake, 1994; Sycara, 1988). In order to use case-based case adaptation, however, we must account not only for how adaptation cases are used but how an initial library of cases is acquired. One method is to simply acquire them from human domain experts, either directly as adaptation cases (Berger, 1995) or by an interactive process in which the user is queried for needed information (Bell, Kedar, & Bareiss, 1994). Our method is to build specific adaptation knowledge from introspective reasoning, starting from scratch using general rules and then storing the results for future reuse by case-based reasoning.

4 A FRAMEWORK FOR LEARNING CASE ADAPTATION

Our model of case adaptation learning starts with general rules about information search and about the constraints that must be satisfied by the case adaptation process. As it applies those rules to adapt cases, it stores traces of reasoning done during adaptation as *adaptation cases* to be used in future processing. These adaptation cases are a source of very specific knowledge about how to solve particular adaptation problems. Thus the initial rule-based adaptation process provides (1) flexibility to deal with novel adaptation problems and (2) a way to generate specific adaptation knowledge for future use.

The rule-based adaptation process treats the case adaptation problem as the problem of operationalizing a set of very general transformations—substitutions, deletions, and additions to the structure of a retrieved solution—by searching memory for the information needed to apply those transformations. For example, if part of a retrieved case does not apply, substituting an alternative depends on finding that alternative, which depends on identifying the type of information that is needed and then reasoning introspectively about how it might be retrieved. In our framework, the memory search process transforms queries (Kolodner, 1984) using an introspective planning process involving internal "mental" operators (Hunter, 1990) describing how to search for information, such as *to find a cause for an actor's state, search for an action performed by the actor that could cause that state* (Leake, 1993, 1994). The memory search process operationalizes (Mostow, 1983) chains of abstract adaptation rules: it provides the information needed to effectively

apply them to concrete problems.

Because explanation-based generalization is crucially concerned with the problem of operationalization (DeJong & Mooney, 1986; Keller, 1988; Mitchell, Keller, & Kedar-Cabelli, 1986), it might appear that explanation-based generalization is the appropriate method for operationalizing chains of adaptation rules. However, a problem arises: explanation-based generalization assumes a complete and correct domain theory. A domain theory of the memory search task would need to account for each piece of information in memory and how each piece of information is stored. Unfortunately, in general a particular piece of information may be stored in different ways; the problem of memory search is how to find desired information in an idiosyncratic memory whose contents and organization may not be characterized precisely—search rules that work in one instance may not apply to other problems that appear to be within the scope of those rules. Consequently, other methods are needed.

Our solution is to use adaptations done by general rules as cases for later case-based reasoning about the case adaptation process itself. When adaptation using general adaptation rules is successful, a trace of the specific adaptation process used (which reflects details of the current memory organization scheme and the idiosyncratic contents of memory) can be stored in order to provide specific guidance in future adaptations. Our model acquires the knowledge needed for expert adaptation by storing traces of successful adaptations as *adaptation cases* that can be reused in the future by a case-based reasoning process applied to case adaptation. Unlike abstract adaptation rules, adaptation cases encapsulate the system's experience on specific adaptation and memory search problems and reflect the system's specific task, domain, and memory organization.

5 APPLYING THE FRAMEWORK

Our framework is being applied to the domain of *disaster response* for natural and man-made disasters. Examples of natural disasters in the scope of the domain include earthquakes, floods, and typhoons; examples of man-made disasters include events such as toxic chemical spills, the Chernobyl disaster, and the poisonous gas leak at Bhophal.

Reasoning from prior experiences and adaptation of prior solutions play a crucial role in the disaster response domain. In their discussion of human disaster response planning, Rosenthal et al. (1989) observe that because no two disasters are the same, taking into account the differences between old and new cases is essential to applying prior experience successfully. For example, one highly successful episode of disaster response took place in Mississauga, Canada, where 200,000 people had to be evacuated in response to a toxic chemical spill. The success of that episode has been credited to plans developed in a very different crisis situation—a hijacking incident—that occurred in the same community (Scanlon, 1989, pp. 305). Rather than having to reason from scratch, the human disaster response planners reused many aspects of prior reasoning, as well as lessons about prior problems.

We are investigating adaptation learning in the context of the case-based planner DIAL (for **D**isaster response with **I**ntrospective **A**daptation **L**earning). DIAL starts with a library of disaster response cases and learns adaptation strategies as it adapts prior solutions to fit new situations. Its initial adaptation knowledge is limited to a small set of general rules describing possible structural transformations and characterizing how to search memory for needed information. As it successfully adapts its response plans to new situations, it stores not only the problem-solving episode, but also two types of adaptation knowledge for use in similar future adaptation problems: *memory search cases* encapsulating information about the types of operations used in the memory search process and the memory links that were followed, and *adaptation cases* encapsulating information about the adaptation problem as a whole and its solution.

DIAL's adaptation component takes as input an instantiated response plan and a description of an applicability problem to repair. The basic sequence of processing steps is as follows:

1. **Case-based adaptation:** DIAL attempts to retrieve an adaptation case for a similar prior problem. If it succeeds, the previous adaptation case is re-applied and processing continues with step 3.
2. **Rule-based adaptation:** Based on the problem type, the system selects a transformation (e.g., substitution), and generates a *knowledge goal* (Leake & Ram, 1993; Ram, 1987) for the information needed for the transformation. Introspective planning (Hunter, 1990) is done to search for the needed information, re-using any applicable memory search cases generated for previous adaptation problems.
3. **Plan evaluation:** The disaster response plan that would result from the candidate adaptation is evaluated by a human user who inputs to the system whether the adaptation is acceptable. If not, other adaptations are tried.
4. **Storage:** When adaptation is successful, the resulting response plan, the *adaptation case*, and the *memory search cases* reflecting current processing are stored to be accessible for future use. This information is then available to apply to future case adaptation problems, supporting the transition from adaptation rules that are general but had to apply, to more specific adaptation knowledge.

The rules provided for use by DIAL's rule-based adaptation process are relatively unguided, depending on "weak methods" of memory search (e.g., ascending and descending abstraction hierarchies to find related nodes) with the results filtered by constraints from the particular adaptation problem. However, traces of this relatively unguided process become cases that can provide more precise guidance in similar future situations.

6 CONCLUSION

The case-based reasoning paradigm is receiving attention both as a technology for developing intelligent systems and

as a cognitive model of expert human reasoning. However, developing case-based reasoning systems that can generate complete solutions on their own—which is crucial for making the systems usable by naive users as well as experts—depends on endowing them with the capability for automatic case adaptation, which in turn depends on solving the problem of acquiring case adaptation knowledge.

We are developing a model that combines rules and case-based reasoning to build up expertise at case adaptation. This approach provides a way for case-based reasoning systems to make the transition from adaptation by general rules (which have wide applicability, but may be unreliable and expensive to apply) to an adaptation process that reflects specific lessons acquired from experience doing case adaptation, without sacrificing the flexibility to perform novel adaptations when needed. Thus our approach is a flexible model of how a case-based reasoner can increase expertise at the case-based reasoning process itself.

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.

Bell, B., Kedar, S., & Bareiss, R. (1994). Interactive model-driven case adaptation for instructional software design. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, pp. 33–38 Atlanta, GA.

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

Berger, J. (1995). Using past repair episodes. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence* Montreal, Canada. IJCAI. Submitted.

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.

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.

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. (1995). Planning for repairing reasoning failures. In *Proceedings of the 1995 AAAI Spring Symposium on Representing Mental States and Mechanisms* Stanford, CA. AAAI. In press.

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

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.

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.

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

Kitano, H., Shibata, A., Shimazu, H., Kajihara, J., & Sato, A. (1992). Building large-scale and corporate-wide case-based systems: integration of organizational and machine executable algorithms. In *Proceedings of the Tenth National Conference on Artificial Intelligence*, pp. 843–849 San Jose, CA. AAAI.

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 making. *The AI Magazine*, *12*(2), 52–68.

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

Kolodner, J. (1993b). A case-based view of case-based reasoning—what have we wrought?. Invited talk at the AAAI-93 Workshop on Case-Based Reasoning.

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.

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. (1994). 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. & 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.

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.

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. (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.

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

Rosenthal, U., Charles, M., & Hart, P. (Eds.). (1989). *Coping with crises: The management of disasters, riots, and terrorism*. C.C. Thomas, Springfield, IL.

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.

Scanlon, T. (1989). Toxic chemicals and emergency management: the evacuation of mississauga, ontario, canada. In Rosenthal, U., Charles, M., & Hart, P. (Eds.), *Coping with crises: The management of disasters, riots, and terrorism*, chap. 11, pp. 303–322. C.C. Thomas, Springfield, IL.

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. (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.