

Linking Adaptation and Similarity Learning

David B. Leake, Andrew Kinley, and David Wilson
Computer Science Department
Lindley Hall 215, Indiana University
Bloomington, IN 47405
{leake, akinley, davwils}@cs.indiana.edu

Abstract

The case-based reasoning (CBR) process solves problems by retrieving prior solutions and adapting them to fit new circumstances. Many studies examine how case-based reasoners learn by storing new cases and refining the indices used to retrieve cases. However, little attention has been given to learning to refine the process for applying retrieved cases. This paper describes research investigating how a case-based reasoner can learn strategies for adapting prior cases to fit new situations, and how its similarity criteria may be refined pragmatically to reflect new capabilities for case adaptation. We begin by highlighting psychological research on the development of similarity criteria and summarizing our model of case adaptation learning. We then discuss initial steps towards pragmatically refining similarity criteria based on experiences with case adaptation.

Introduction

Case-based reasoning (CBR) is a reasoning process that solves new problems by retrieving similar prior problem-solving episodes and adapting their solutions to fit the new situations. Learning by remembering cases is a fundamental part of case-based reasoning: Each problem-solving episode provides a new case for future use. CBR research has also devoted considerable attention to learning by refining the indices used to guide case retrieval. However, little attention has been given to learning how cases should be applied. This paper discusses research modeling how a case-based reasoner can make better use of its prior cases by learning how to adapt them to new circumstances, and by refining the similarity criteria it uses to reflect changes in its adaptation abilities.

Acquisition of case adaptation knowledge is a classic problem for models of case-based reasoning (Kolodner, 1991). CBR systems generally rely on static sets of hand-coded adaptation rules, but developing the needed rules has proven to be a very difficult problem. However, as studies show, human case-based reasoners are adept at applying prior cases (see Kolodner, 1993, for an overview). Consequently, a natural question for CBR as a cognitive model is how the requisite case adaptation knowledge might be acquired.

We are investigating a method for learning specific adaptation knowledge to augment an initial library of

very general case adaptation rules (Leake, 1995b). In our approach, a case-based reasoner begins with a small set of abstract rules for transforming cases and for searching memory to find the information needed to make adaptations. This general knowledge is used to perform adaptations from scratch. The system improves its adaptation capabilities by saving traces of the derivations of new adaptations in a library of *adaptation cases* and reusing them for similar adaptation problems. Thus it makes a transition from rule-based to case-based case adaptation (Leake, 1995b; Leake et al., 1996). By saving memory search cases and adaptation cases, a CBR system can acquire specific adaptation knowledge.

Adaptation learning provides the motivation for another type of learning, learning to refine similarity criteria. A central role of similarity judgments in case-based reasoning is to determine which cases to apply to a new situation and how to adapt them to fit new circumstances. Although CBR systems often base similarity judgments on semantic similarity, the real goal of “similarity assessment” in CBR is to determine *adaptability*: how easily an old case can be adapted to fit the requirements of a new situation (Birnbaum et al., 1991; Smyth and Keane, 1995, 1996). If adaptation knowledge is learned, static similarity criteria may not keep pace with new capabilities for performing adaptations. Thus similarity assessment criteria should change as new adaptation knowledge is acquired.

We first discuss the relationship of our approach to psychological results on the development of similarity criteria and pragmatic influences on case adaptation. We then summarize our model’s approach to improving case adaptation and some preliminary results on the effects of adaptation learning. Finally, we describe the method we are developing for making similarity assessment reflect adaptation experience and relate our approach to other computer models.

Motivations

Our research investigates how case adaptation strategies for case-based reasoning can be learned, and how the similarity assessment process can be refined as adaptation learning makes particular types of differences between old and new situations easier to overcome.

This approach can be supported directly on functional grounds: The purpose of similarity assessment is to determine the difficulty of adapting cases to new situations, which depends on the reasoner's adaptation knowledge. We are not aware of psychological studies directly examining the connection between adaptability and similarity in case-based reasoning, but psychological studies do provide examples of developmental shift in similarity criteria, of similarity judgments coming to more closely reflect task-relevant features as a task is learned, and of the applicability of an analog to a new problem situation acting as a selection constraint during human analogical reasoning.

Experiments by Gentner & Toupin (1986) demonstrate a developmental shift in the similarity criteria used by children for analogical reasoning, and show that the shift is manifested in how they adapt stories to apply to new characters. Experiments by Suzuki et al. (1992) studying similarity judgments in problem-solving for the Towers of Hanoi problem show that novices' judgments about the similarity of problem states can be characterized by the number of shared surface features, but that experts' judgments are best characterized by the goal-relevant criterion of the number of operators required to transform each problem state to the goal state. Further, Chi et al. (1981) note a dramatic difference between the similarity criteria of novice physics problem-solvers, who rely on surface features, and physics experts, who classify problems according to the underlying methods needed to solve them.

Adaptation factors have also been shown to affect selection of analogues in analogical problem-solving. Experiments by Keane (1994) suggest that when performing analogical problem-solving, subjects favor analogues that are easier to apply to the new problem situation.

Overview of DIAL

We are investigating learning about adaptation and pragmatic similarity in the context of a case-based planner. The planner's task domain is *disaster response planning*, the initial strategic planning used to determine how to assess damage, evacuate victims, etc., in response to natural and man-made disasters (e.g., earthquakes and chemical spills). There are no hard-and-fast rules for disaster response planning; human disaster response planners appear to rely largely on prior cases to guide their decisions (Rosenthal et al., 1989).

Our computer model, DIAL (for **D**isaster response with **I**ntrospective **A**daptation **L**earning), takes as input conceptual representations of news stories describing the initial events in a disaster. It generates candidate response plans by case-based reasoning. The system's case-based planning framework is based in a straightforward way on previous case-based planners such as CHEF (Hammond, 1989).

DIAL's initial knowledge sources are a library of

domain cases—disaster response plans from prior disasters—and general (domain-independent) rules about case adaptation and memory search. When a new story is presented to the system, DIAL uses standard indexing techniques to attempt to retrieve cases representing response plans for similar disasters. This process results in a set of candidate response plan cases. A finer-grained analysis selects the candidate expected to be easiest to adapt, based on similarity criteria learned from experience with prior adaptations. The selected case is provided to the system's case adaptation component, along with a list of differences that must be repaired.

Learning and Reusing Adaptations

The foundation of DIAL's adaptation and similarity learning is case-based reasoning about the adaptation process itself. DIAL's adaptation process begins with the system's adaptation component receiving a description of an adaptation task: a disaster response plan case and a list of the problems that prevent it from applying to a new situation. DIAL first attempts case-based adaptation, searching for an adaptation case that applied successfully to a similar adaptation problem. Adaptation cases are indexed in memory by a vocabulary of categories of problems that may require adaptation (see Leake, 1992). If DIAL succeeds in retrieving a relevant adaptation case, the adaptation process traced by that case is re-applied.

Otherwise, DIAL builds up a new adaptation by a combination of rule-based and case-based reasoning. The system first selects a transformation associated with the type of problem that adaptation must repair. (For example, to *substitute* a new plan step for one that does not apply.) Given the transformation, the program generates a *knowledge goal* (Hunter, 1990; Ram, 1987) for the knowledge needed to apply the transformation. (E.g., when performing a substitution, DIAL needs to identify a good substitute: an object that satisfies the relevant constraints on the object being replaced.) The knowledge goal is used to guide a planning process for how to search memory (Leake, 1994; Leake, 1995c). This process builds a memory search plan, using a small set of built-in memory search strategies (e.g., to perform "local search" for similar objects) and *memory search cases* stored after solving previous adaptation problems. When the needed information is found in memory, it enables DIAL to apply the selected transformation to the retrieved response plan.

The adapted response plan is evaluated by a simple evaluator that checks the compatibility of the current plan with explicit constraints from the response plan. A human user performs backup evaluation, detecting more subtle problems. If problems are found, DIAL attempts other adaptations. If the autonomous case adaptation process fails to generate an acceptable solution, an inter-

face allows the user to guide the adaptation process, selecting a transformation and suggesting features to consider. During the adaptation, the system records a trace of the adaptation process. The trace is represented in the same form as the traces of system-generated adaptations and is added to the adaptation case library for future use.

When adaptation is successful, the resulting response plan, adaptation case, and memory search plan are stored for future use.

The Effects of Adaptation Learning

Although our computer model is still under development, we have conducted initial ablation tests studying the benefit of adaptation learning and its relationship to case learning in the initial model. In these tests, starting from an initial memory of 870 concepts and case library of 6 initial cases, DIAL performs a total of 30 adaptations to develop response plans for 5 stories. Stored cases and new stories were based on Clarinet News Service newswire and the *INvironment* newsletter for air quality consultants; stored cases involved an earthquake in Los Angeles, an air quality disaster at a manufacturing plant, a flood in Bainbridge, Georgia, a chemical disaster at a factory, a flood in Izmir, Turkey, and an air quality disaster in a rural elementary school.

In these tests, response plan learning did improve performance, as did adaptation learning. Interestingly, by the measure of memory operations performed, adaptation learning alone was more effective than case learning alone, although both required comparable numbers of memory nodes to be visited. As was also expected, when no adaptation cases are learned, learning additional response plan cases enables the system to solve new problems with less adaptation effort—more similar cases are available. Adding adaptation learning to response plan learning produced insignificant benefits when memory search during adaptations was based on local search. There was much greater benefit when response plan learning was combined with adaptation learning using other memory search strategies. Details on the adaptation learning process and this preliminary test can be found in Leake, Kinley, and Wilson (1996).

We are now “scaling up” the system for additional tests. One particular concern is the potential for a “utility problem” (Francis and Ram, 1995; Minton, 1988) as large numbers of adaptation cases are learned.

Learning Similarity from Adaptability

To realize the full benefits of adaptation learning, similarity learning is needed as well. Learning about how difficult it is to repair particular differences helps to decide which cases are most usefully similar—which will be most easily adapted. For example, initially it might be assumed that the locale of a disaster is comparatively unimportant when deciding similarity. However, adapt-

ing the response for a small town disaster into the response plan for a big city disaster may be quite difficult, because of the added need to work out arrangements for a large-scale evacuation. From experience with this adaptation, a disaster response planner can learn to consider the locale of the disaster when assessing similarity between a new situation and prior disasters.

DIAL improves its similarity assessment process by using learned adaptation cases to provide estimates of the cost of adapting particular types of problems. In order to facilitate later processing, response plans that require less adaptation effort to apply are considered more similar than those requiring expensive adaptation. In DIAL’s pragmatic similarity judgment, the “most similar” case is the one expected to be easiest to adapt. This approach to similarity follows the same principle as Leake’s (1992) *constructive similarity assessment* and Smyth and Keane’s (1995, 1996) *adaptation-guided retrieval*.

If similarity judgments are to be based on adaptability, two questions are how to estimate the cost of adaptation and how to make a reasonable tradeoff between accuracy of adaptation cost estimates and the cost of the estimation process itself. DIAL’s retrieval uses two types of similarity assessment in a two-step process. The first step retrieves a rough “first-pass” set of cases based on static semantic similarity criteria applied to the type of disaster (e.g., flood, earthquake, etc.) and its attributes. The second step prioritizes these candidate cases for adaptation according to estimates of their adaptability.

To estimate adaptability, DIAL first identifies inapplicable aspects of the retrieved response plans, using procedures for pattern-based anomaly detection, and describes the problems according to a vocabulary of problem types based on Leake (1992). This vocabulary includes, for example, categories to describe the problem when objects specified by the plan are unavailable, or when role-fillers of a schema have been left unspecified and need to be selected. Associated with each category is a frame structure to be filled in with the specifics of the current problem. That frame structure, instantiated with the particulars of the current situation, describes the problem to be repaired by adaptation.

For each problem to be repaired by adaptation, DIAL searches its memory of prior adaptations, and retrieves the adaptation case addressing the most similar prior problem. That previous adaptation case provides information about how to repair the problem. The information is used to estimate the cost of its repair, as follows:

- If the retrieved adaptation case was generated to solve an identical adaptation problem, the solution to that previous adaptation can be reapplied directly, so adaptation cost will be minimal. For example, if a previous adaptation involved adapting the response plan for an American flood to a flood in Turkey, and the Red Cross was involved in the original plan, a new relief agency

would have to be found to apply the plan in Turkey, where the Red Cross does not exist. However, replacing the Red Cross by the Red Crescent is a reasonable adaptation. Once that adaptation is learned, the problem of adapting the Red Cross to a relief organization to apply to Turkey is trivial.

- If the retrieved adaptation case dealt with an adaptation problem that was similar but not identical, the cost of adapting the new problem is estimated from the cost of the prior adaptation. How best to perform this estimation is still an open issue, but DIAL's current method is to focus on the cost of performing the memory search needed to find the information to allow the previous adaptation to be performed. To illustrate, if finding a substitution—for example, an appropriate evacuation method—required considerable effort for a previous adaptation, it is assumed that finding a new evacuation route will take considerable effort in the current situation. Consequently, if one of the retrieved plans requires finding an evacuation method, while another avoids evacuation by containing the disaster, it may be reasonable to favor the plan for containing the disaster (if containing the disaster is practical).¹ The rationale for this cost estimation criterion is based on the idea of derivational analogy (Carbonell, 1986; Veloso, 1994): If a previous adaptation for a similar problem had to infer certain features and constraints from the plan, and transform them in certain ways to generate an appropriate adaptation, the process for the current situation is expected to follow analogous steps, even if the specifics of the situation are different.
- If no similar adaptation case is found, DIAL uses an estimate based on the average cost (measured in primitive memory search operations) of adapting problems in each problem category. Kass (1990) proposes a similar method for coarse-grained estimates of adaptation cost.

By basing similarity assessment directly on the current state of its changing adaptation knowledge, DIAL's similarity assessment process reflects information about the actual difficulty of adapting to repair certain types of problems. We are now designing experiments to examine the performance of these simple strategies, to test how they affect the needed adaptation effort and to guide their refinement.

Relationship to Other Computer Models

Some early case-based reasoning systems included components for learning limited forms of adaptation knowl-

¹It should be noted that in general, the different plans may each contain useful parts of the solution to a problem (e.g., Ram and Francis, 1996, Redmond, 1992). Extending our model to consider relevant pieces of candidate plans is a topic for future research.

edge. For example, CHEF (Hammond, 1989) bases its adaptations on both a static library of domain-independent plan repair strategies and a library of special-purpose *ingredient critics* that are learned; PERSUADER (Sycara, 1988) uses previously-stored adaptation episodes to suggest adaptations. In both examples, the learned adaptations can only be reused in highly similar situations. However, the adaptation cases learned by DIAL can be reused more flexibly. Rather than learning only by storing specific adaptation episodes, DIAL stores both the specific adaptation and its derivational trace. In very similar situations, the adaptation can be reapplied directly; in less similar situations, the steps used to determine the previous adaptation can be replayed, taking into account differing circumstances.

DIAL also differs from previous approaches to adaptation learning in its emphasis on learning about the required memory search. Its characterization of adaptations is inspired by the *adaptation strategies* in SWALE (Kass et al., 1986) and ABE (Kass, 1990), which combine transformations with domain-independent memory search information, and its approach to memory search is inspired by the memory search process of CYRUS (Kolodner, 1984). However, those systems did not learn to improve their search processes. In reasoning about the information needed to carry out the adaptation task, our model also relates closely to Oehlmann's (1995) metacognitive adaptation. It is in a similar spirit both to recent research on applying heuristic search to gathering information for argumentation (Rissland et al., 1994) and to work in information retrieval on strategic reasoning about where to search for needed information (Baudin et al., 1994).

Smyth & Keane (1995, 1996) have developed a CBR system that ties similarity judgments directly to adaptability, using heuristics coded to reflect the difficulty of performing particular types of adaptations. They also demonstrate that their adaptation-guided method produces significant improvements in the cost of performing adaptations. Our approach to similarity judgments is strongly in the spirit of their approach, and their results are encouraging for the potential benefit of adaptation-based criteria compared to traditional semantic similarity criteria. However, in their work, similarity and adaptation knowledge are static.

Learning to refine similarity criteria has been investigated in Prodigy/Analogy (Veloso, 1994). That system's "foot-print" similarity metric focuses consideration on goal-relevant portions of the initial state, in order to retrieve cases that refer to the prior problem situations with the most relevant similarities. Our adaptability-based similarity method focuses on a different issue, estimating the costs of repairing relevant differences that have been found. Finally, two-stage retrieval processes, such as that used in DIAL's initial filtering of retrieval

candidates followed by a deeper but more computationally expensive analysis, have been advocated by many previous models (e.g., Bareiss & King, 1989) not only on functional grounds (to restrict processing effort) but on cognitive grounds as well (Gentner and Forbus, 1991). The new contribution of our approach is to tie similarity criteria directly to learning about the relative importance of different types of differences when adapting cases to new situations.

Conclusion

We have described ongoing research on how case-based reasoners can learn to apply cases more effectively, both by learning how to adapt prior cases to new situations and by refining similarity criteria according to experience concerning which types of adaptations are difficult to perform. Our approach to learning case adaptation models the acquisition of specific adaptation knowledge starting from “weak methods” for case adaptation; our approach to learning useful similarity criteria builds on the adaptation learning process, to consider cases “usefully similar” if they are expected to be easy to adapt, given experience with prior adaptations. Preliminary trials of the adaptation learning system are encouraging, but further tests are needed, especially to study how the process “scales up” when large numbers of adaptations are learned. Tests are also needed to examine how well current similarity estimates predict the difficulty of future adaptations. The model is now being refined and extended in preparation for tests of the effects of similarity learning and more extensive tests of the system as a whole.

Acknowledgments

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

References

- Bareiss, R. and King, J. (1989). Similarity assessment in case-based reasoning. In Hammond, K., editor, *Proceedings of the DARPA Case-Based Reasoning Workshop*, pages 67–71, San Mateo. Morgan Kaufmann.
- Baudin, C., Pell, B., and Kedar, S. (1994). Using induction to refine information retrieval strategies. In *Proceedings of the twelfth national conference on artificial intelligence*, pages 553–559, Seattle, WA.
- Birnbaum, L., Collins, G., Brand, M., Freed, M., Krulwich, B., and Pryor, L. (1991). A model-based approach to the construction of adaptive case-based planning systems. In Bareiss, R., editor, *Proceedings of the DARPA Case-Based Reasoning Workshop*, pages 215–224, San Mateo. Morgan Kaufmann.
- Carbonell, J. (1986). Derivational analogy: A theory of reconstructive problem solving and expertise acquisition. In Michalski, R., Carbonell, J., and Mitchell, T., editors, *Machine Learning: An Artificial Intelligence Approach*, volume 2, pages 371–392. Morgan Kaufmann, Los Altos, CA.
- Chi, M., Feltovich, P., and Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5(2):121–153.
- Francis, A. and Ram, A. (1995). A comparative utility analysis of case-based reasoning and control-rule learning systems. In *Eighth European Conference on Machine Learning*, Crete, Greece.
- Gentner, D. and Forbus, K. (1991). MAC/FAC: A model of similarity-based retrieval. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*, pages 504–509, Chicago, IL. Cognitive Science Society.
- Gentner, D. and Toupin, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, 10(3):277–300.
- Hammond, K. (1989). *Case-Based Planning: Viewing Planning as a Memory Task*. Academic Press, San Diego.
- Hunter, L. (1990). Planning to learn. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, pages 261–268, Cambridge, MA. Cognitive Science Society.
- Kass, A. (1990). *Developing Creative Hypotheses by Adapting Explanations*. PhD thesis, Yale University. Northwestern University Institute for the Learning Sciences, Technical Report 6.
- Kass, A., Leake, D., and Owens, C. (1986). SWALE: A program that explains. In *Explanation Patterns: Understanding Mechanically and Creatively*, pages 232–254. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Keane, M. (1994). Adaptation as a selection constraint on analogical mapping. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, pages 490–495, Atlanta, GA.
- 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.
- 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*, pages 313–318, Bloomington, IN. Lawrence Erlbaum.
- Leake, D. (1992b). *Evaluating Explanations: A Content Theory*. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Leake, D. (1994). Towards a computer model of memory search strategy learning. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, pages 549–554, Atlanta, GA. Lawrence Erlbaum.
- 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*, pages 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*, pages 84–88, Stanford, CA. AAAI Press. Technical Report WS-95-05.
- Leake, D., Kinley, A., and Wilson, D. (1996). Acquiring case adaptation knowledge: A hybrid approach. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, Menlo Park, CA. AAAI Press. In press.
- Minton, S. (1988). *Learning Search Control Knowledge: An Explanation-Based Approach*. Kluwer Academic Publishers, Boston.
- Oehlmann, R. (1995). Metacognitive adaptation: Regulating the plan transformation process. In *Proceedings of the Fall Symposium on Adaptation of Knowledge for Reuse*. AAAI.
- Ram, A. (1987). AQUA: Asking questions and understanding answers. In *Proceedings of the Sixth Annual National Conference on Artificial Intelligence*, pages 312–316, Seattle, WA. Morgan Kaufmann.
- Ram, A. and Francis, A. (1996). Multi-plan retrieval and adaptation in an experience-based agent. In Leake, D., editor, *Case-Based Reasoning: Experiences, Lessons, and Future Directions*. AAAI Press, Menlo Park, CA. In press.
- Rissland, E., Skalak, D., and Friedman, M. (1994). Heuristic harvesting of information for case-based argument. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, pages 36–43, Seattle, WA. AAAI.
- Rosenthal, U., Charles, M., and Hart, P., editors (1989). *Coping with crises: The management of disasters, riots, and terrorism*. C.C. Thomas, Springfield, IL.
- Smyth, B. and Keane, M. (1995). Experiments on adaptation-guided retrieval in case-based design. In *Proceedings of First International Conference on Case-Based Reasoning*, Sesimbra, Portugal.
- Smyth, B. and Keane, M. (1996). Design à la déjà vu: Reducing the adaptation overhead. In Leake, D., editor, *Case-Based Reasoning: Experiences, Lessons, and Future Directions*. AAAI Press, Menlo Park, CA. In press.
- Suzuki, H., Ohnishi, H., and Shigermasu, K. (1992). Goal-directed processes in similarity judgment. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, pages 343–348, Bloomington, IN. Lawrence Erlbaum.
- Sycara, K. (1988). Using case-based reasoning for plan adaptation and repair. In Kolodner, J., editor, *Proceedings of the DARPA Case-Based Reasoning Workshop*, pages 425–434, San Mateo, CA. Morgan Kaufmann.
- Veloso, M. (1994). *Planning and Learning by Analogical Reasoning*. Springer Verlag, Berlin.