From: Proceedings of the Third International Conference on Multistrategy Learning. Copyright © 1996, AAAI (www.aaai.org). All rights reserved.

Multistrategy Learning to Apply Cases for Case-Based Reasoning

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

Investigations of learning in case-based reasoning (CBR) have traditionally focused on learning two types of knowledge: new cases and new indexing criteria for case retrieval. However, there is increasing recognition that other types of knowledge also play crucial roles in the case-based reasoning process. The effectiveness of a CBR system depends not only on having and retrieving relevant cases, but also on selecting which retrieved cases to apply and determining how to adapt them to fit new situations. Consequently, case-based reasoning can benefit from using multiple learning strategies to acquire, in addition to new cases and indices, new case adaptation strategies and similarity criteria. This paper describes ongoing research that studies how multiple types of learning can improve the case-based reasoning process and examines their interrelationship in contributing to the overall performance of a CBR system.

Introduction

Case-based reasoning (CBR) solves new problems by retrieving records of similar prior problem-solving episodes and adapting their solutions to fit new needs. Learning by acquiring new cases is a fundamental part of case-based reasoning, and the process of learning by case acquisition has been a central focus of CBR research. Recently, however, there has been increasing awareness of the importance of multiple types of knowledge to guide the CBR process. For example, Richter (1995) points out that the knowledge of a case-based reasoner is contained not only in its case base and indexing scheme, but also in the reasoner's similarity metric and in its case adaptation knowledge. Thus an important question is how these types of knowledge may be acquired.

The need to acquire multiple types of knowledge in CBR provides a natural opportunity for multistrategy learning (e.g., Michalski & Tecuci, 1994). Multistrategy learning can enable CBR systems to learn not only new cases, but also how to retrieve cases more reliably, how to judge the relevance of candidate cases more perspicaciously, and how to adapt cases to new situations more effectively. Using different learning strategies for each type of knowledge enables a CBR system to tailor its learning according to the task requirements of different parts of its reasoning process.

Because the component steps of CBR are strongly related, multistrategy learning has an added benefit as well: learning that improves one component may help overcome deficiencies in others. For example, better retrieval can reduce the need for adaptation knowledge, by providing more relevant cases; conversely, better adaptation knowledge can decrease the need for high-quality retrieval, by generating successful solutions even if less than ideal cases are retrieved.

Our research studies how multiple types of learning can improve the case-based reasoning process and examines their interrelationship in contributing to the overall performance of a CBR system. We are investigating how introspective reasoning during CBR can identify needs for information during case adaptation and satisfy them by introspective question transformation, a memory search process involving strategically redescribing needed information to guide memory search. We are also studying how memory search strategies, case adaptation strategies, and similarity criteria can be learned from experience.

This paper first describes our task domain and the basic structure of our testbed system, and summarizes our system's learning strategies and their relationships. It then describes how the system learns to improve its case adaptation process, discusses the effects of case adaptation learning for a small initial set of test examples, and describes how adaptation learning determines new similarity criteria. The paper closes by placing our approach in context of other research on learning from multiple parts of the CBR process.

From: Proceedings of the Third International Conference on Multistrategy Learning. Copyright © 1996, AAAI (www.aaai.org). All rights reserved.

Task and System Overview

Our system's task domain is disaster response planning. Disaster response planning is the initial strategic planning used to determine how to assess damage, evacuate victims, etc., in response to natural and man-made disasters such as earthquakes and chemical spills. There are no hard-and-fast rules for disaster response planning, and human disaster response planners appear to depend heavily on prior experiences when they address new problem situations (Rosenthal, Charles, & Hart 1989).

Our testbed system, DIAL,[1] processes a conceptual representation of a news story describing the initial events in a disaster, and proposes a response plan by retrieving and adapting the response plan for a similar prior disaster. DIAL includes a simple schema-based story understander, a response plan retriever and instantiator, a simple evaluator for candidate response plans, and an adaptation component to adapt plans when problems are found. Its basic processing sequence is as follows:

- A story is input to the system.
- Candidate response plan cases for similar problem situations are retrieved, using coarse-grained static similarity assessment criteria.
- A finer-grained similarity assessment process uses learned information about difficulty of adaptation to select the candidate case whose response plan is expected to be easiest to adapt.
- Problems in the response plan of the selected case are repaired by case adaptation. During adaptation, DIAL learns by storing traces of its case adaptation process and of the memory search process used to find needed information. If its adaptation attempt fails, DIAL can also learn by recording a trace of a user-guided adaptation process.
- The resulting response plan case is stored for future reuse by transformational analogy.

The system's case-based planning framework is based in a straightforward way on previous case-based planners using transformational analogy, such as CHEF (Hammond 1989). Consequently, we will not discuss DIAL's planning process per se, but instead will focus on how it learns to improve its case adaptation, memory search, and similarity assessment.

Learning Methods and Relationships

Five steps of DIAL's reasoning process involve knowledge transmutations (Michalski 1994) and learning, using a mix of learning strategies:

1. **Response plan learning:** The "baseline" learning method for DIAL is learning by case acquisition, the normal learning of case-based reasoning systems. Response plans are reapplied by **transformational analogy** (Carbonell 1983).

2. **Memory search:** When information is needed from memory (e.g., to adapt a case to a new situation), DIAL identifies the needed information and generates explicit *knowledge goals* (Hunter 1990; Ram 1987) for that information. Its memory search process transforms descriptions of needed information and of known information in memory, using self-knowledge about the system's memory organization, in order to satisfy its knowledge goals. This **introspective question transformation** process is based on the goal-driven learning principle of reasoning strategically about how to satisfy needs for information (desJardins 1992; Ram & Leake 1995).

3. **Memory search strategy learning:** When DIAL generates a memory search plan, it stores a trace of that plan as a *memory search case* for reuse by **derivational analogy** (Carbonell 1986).

4. **Adaptation learning:** DIAL adapts cases by an introspective reasoning process that determines which transformations to apply and which knowledge goals must be satisfied to apply them. If DIAL cannot generate an acceptable adaptation, it asks a user to guide the case adaptation process interactively. Traces of internally-generated adaptations and of user adaptations are stored as *adaptation cases* for reuse by **derivational analogy**.

5. **Similarity learning:** When a case-based reasoning system performs similarity assessment to determine which case is most similar to a new situation, the goal is to select the case that will be easiest to adapt (Birnbaum *et al.* 1991; Leake 1995a; Smyth & Keane 1995). DIAL's similarity assessment process uses prior experiences with case adaptation to estimate case adaptation cost for new problems, by a **transformational analogy** process. Consequently, adaptation learning and similarity learning are coupled: when adaptation cases are learned, that learning provides not only knowledge to use during future case adaptation, but to use during similarity assessment as well.

Table 1 summarizes DIAL's learning processes and the mechanisms used. DIAL's response plan cases, memory search cases, and case adaptation cases can be reused independently, allowing different lessons drawn

[1] For Disaster response with Introspective Adaptation Learning.

From: Proceedings of the Third International Conference on Multistrategy Learning. Copyright © 1996, AAAI (www.aaai.org). All rights reserved.

Process	Learning mechanism
Response plan learning	CBR/transformational analogy
Memory search process	Introspective question transformation
Memory search strategy learning	CBR/derivational analogy
Adaptation learning	CBR/derivational analogy, applied to traces of internal processing or user adaptations
Similarity learning	CBR/transformational analogy

Table 1: DIAL's processes and learning mechanisms.

from a single episode to be reapplied where most appropriate in the future. DIAL's close coupling of adaptation cases and similarity criteria enables the selection of response plan cases based on the state of its case adaptation knowledge. Each aspect of learning complements the others in supporting the CBR process as a whole. The following sections discuss specific aspects of DIAL's learning in more detail.

Learning to Adapt Cases to New Situations

Case adaptation is important to CBR because the ability of CBR systems to solve novel problems depends on adapting prior solutions to fit new circumstances. Unfortunately, hand-coding appropriate case adaptation knowledge has proven to be very difficult, to the point that experts in both CBR research (e.g., Kolodner, 1991) and applications (e.g., Barletta, 1994; Mark et al., 1996) agree that it is not currently practical to deploy CBR applications with automatic adaptation. Thus there is strong practical motivation for developing effective methods for learning to improve case adaptation.

Our approach to case adaptation learning models a transition from general (but non-operational) adaptation knowledge to more specific and operational knowledge. The initial knowledge for our approach is a small set of abstract transformation rules and memory search methods. When presented with a new adaptation problem, our system first selects a transformation rule to apply and then performs memory search to find the information needed to operationalize the transformation rule and apply it to the problem at hand (e.g., if a *substitution* transformation is selected, to find what to substitute). The system learns to improve its adaptation capabilities by case-based reasoning applied to the case adaptation process itself: a trace of the steps used in solving an adaptation problem is saved to be reused by derivational analogy when similar adaptation problems arise in the future (Leake 1995b; Leake, Kinley, & Wilson 1995). In this way, a CBR system doing adaptation can acquire specific adaptation procedures starting from domain-independent "weak methods" for adaptation when no specific knowledge is available. At the same time, traces of the memory search process are stored for future reuse, to facilitate building up new adaptations in the future.

DIAL's adaptation component takes two inputs: an instantiated disaster response plan and a description of the problems in the response plan that must be repaired. When presented with an adaptation problem, DIAL's adaptation component performs the following steps:

1. **Case-based adaptation:** DIAL first attempts to retrieve an adaptation case that applied successfully to a similar adaptation problem. If retrieval is successful, the adaptation process traced by that case is re-applied and processing continues with step 3.
2. **Rule-based adaptation:** When no relevant prior case is retrieved, DIAL selects a transformation associated with the type of problem that is being adapted. (E.g., it may decide to *substitute* a new plan step for one that does not apply.) Given the transformation, the program generates a knowledge goal for the information needed to apply the transformation. (E.g., when performing a substitution, the knowledge goal is to find an object that satisfies all the case's constraints on the component being replaced.)

   The knowledge goal is then passed to an introspective planning component that reasons about possible memory search strategies (Leake 1995c) to guide search for the needed information. This search process generates a memory search plan whose operators may include both operators from an initial set of memory search strategies and *memory search cases* stored after solving previous adaptation problems. If the needed information is found, it is used to apply the selected transformation to the retrieved response plan. If it is not found, the process continues with step 4, manual adaptation.
3. **Plan evaluation:** The adapted response plan is evaluated by a simple evaluator that checks the

From: Proceedings of the Third International Conference on Multistrategy Learning. Copyright © 1996, AAAI (www.aaai.org). All rights reserved.
compatibility of the current plan with explicit constraints from the response plan. A human user performs backup evaluation. If the new response plan is not acceptable, other adaptations are tried.

4. **Manual adaptation:** If the previous autonomous case adaptation steps fail to generate an acceptable solution, an interface allows the user to guide the adaptation process, selecting a transformation and suggesting memory search paths 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, so that the system can learn an adaptation case from the interactive episode.
5. **Storage:** When DIAL successfully adapts a response plan, it learns by storing (1) the new *response plan case*, (2) *memory search cases* encapsulating the memory search steps performed during case adaptation, and (3) *adaptation cases,* which encapsulate information about the adaptation problem as a whole—the transformations and memory search cases used when solving the adaptation problem—and its solution.

Thus the system learns not only new response plan cases but also new ways of adapting existing cases to new situations. In addition, new adaptation cases that are learned are used not only to perform new adaptations, but also to estimate adaptation cost during the similarity assessment process for retrieving new cases; learning adaptation cases corresponds to learning new similarity criteria as well.

The Basis of DIAL's Adaptation Learning

DIAL's learning is based on introspective reasoning about requirements for solving adaptation problems. To support reasoning about adaptation problems, a uniform framework is needed for characterizing case adaptation. Following the framework of *adaptation strategies* (Kass 1990), DIAL's rule-based case adaptation treats the case adaptation process as involving two parts: selecting *structural transformations* (e.g., additions, substitutions, and deletions) and performing *memory search* to find the information needed to apply the transformations. Accordingly, two types of case adaptation knowledge are needed: abstract transformations and memory search strategies. It is widely accepted that a small set of transformations is sufficient to characterize a wide range of adaptations (Carbonell 1983; Kolodner 1993), but a large amount of domain-specific reasoning may be required to find the information needed to apply those transformations. The following sections discuss how DIAL's rule-based adaptation process represents the information it needs in the form of knowledge goals, uses the knowledge goals to guide the formation of memory search plans, and packages a trace of its memory search process and other reasoning during adaptation for future use.

Knowledge Goals

Knowledge goals provide explicit descriptions of needed information. DIAL's knowledge goals are satisfied by a planning process for how to carry out memory search. Memory search plans are built from simple primitive memory search operators (e.g., to extract slot values or find abstractions) and traces of successful memory searches satisfying similar previous knowledge goals. Knowledge goals represent information about the type of knowledge needed, about the context of the search, about the reasoning giving rise to the knowledge goal, and about what to do with the knowledge, once found. Thus they reflect the basic principle of goal-driven learning, applied to information search in memory: decisions about the knowledge to acquire and how to acquire it should be based on goal-derived criteria and satisfied by a strategic planning process.

In DIAL, initial knowledge goals are generated to obtain information necessary for a specific adaptation, in response to a problem or inconsistency in applying a retrieved response plan to a current disaster situation. For example, a problem applying the response plan for a flood in Bainbridge, Georgia to a flood in Allakaket, Alaska, is that the Salvation Army provided shelter during the Bainbridge flood, but does not exist in Allakaket. A new relief group local to Allakaket must be found. Consequently, a knowledge goal is generated to find a substitute for the Salvation Army that can provide shelter.

Memory search cases

A memory search case consists of a trace of primitive memory search operators (or previously-stored memory search cases) that were used in a successful previous memory search. Initial memory search cases may be built up by applying "weak methods" of memory search, such as "local search," that are built into the system. (Local search is a common strategy for finding substitutions (Kolodner 1993); it attempts to find concepts that are "near-by" in the system's memory, and progressively widens the search until a suitable substitution is found). Memory search cases may also be built up interactively, by recording traces of a user-guided memory search process.

Retrieved memory search cases provide an initial strategy for finding needed information. Memory search cases are indexed both under the adaptation

From: Proceedings of the Third International Conference on Multistrategy Learning. Copyright © 1996, AAAI (www.aaai.org). All rights reserved.
cases that have successfully used them, and under the knowledge goals they satisfy.

Adaptation cases

Adaptation cases package the results of a successful adaptation. An adaptation case consists of three parts: indexing information, adaptation information, and evaluation information. The indexing information includes a representation of the type of problem to adapt and information about the response plan for which the adaptation case was generated. This information guides selection of the adaptation cases to use for new adaptation problems. The adaptation information packages both a transformation type (e.g., substitute, add, delete) and the memory search steps used to find the information needed to apply the transformation. Evaluation information records the cost of applying the adaptation.

Stored adaptation cases are organized in memory by the types of problems they address. The vocabulary of problem types is similar in spirit to the problem vocabularies used to guide adaptation in other CBR systems (e.g., Hammond, 1989; Leake, 1992). For example, a role-filler in a candidate response plan may be inappropriate, and a new role-filler needed, because of problems such as:

FILLER-PROBLEM:UNAVAILABLE-FILLER
The role filler specified in the plan is unavailable. For example, a police commissioner may be out of town and unable to be reached in an emergency situation.

FILLER-PROBLEM:ROLE-MISMATCH
The role filler specified in the plan is incompatible with the given role. For example, a mismatch occurs when a plan for dealing with an industrial disaster is applied to a school disaster (whose victims are children rather than workers), and the industrial response plan calls for notifying the victim's union (instead of parents).

FILLER-PROBLEM:FILLER-UNSPECIFIED
The role specification may be incomplete, simply because of missing information that must be filled in. For example, a plan might call for a rescue without specifying who should carry it out.

The Effects of Learning Plan Cases and Adaptation Cases

An important question is the benefit of augmenting the traditional learning strategy of case-based reasoning—case acquisition—with case adaptation learning. To obtain initial indications of the effects of adaptation learning in DIAL, we performed ablation tests of the system with initial test examples. The system's initial memory included nodes for 800 concepts; the initial case library included 3 disaster response plan cases, and the test examples involved performing a total of 26 adaptations to develop response plans for 6 stories.

Stored cases and new stories were based on the 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, and a flood in Bainbridge, Georgia. The tests considered only the efficiency of performing an adaptation, but as described later in this paper, efficiency is only one possible dimension for measuring the value of case adaptation learning.

In the baseline condition, the system performed no learning of either cases or adaptations. In addition, it performed all memory search during case adaptation by "local search." The second condition added learning of response plan cases, but no learning of adaptations; this reflected the "standard" configuration of most CBR systems. The third condition included learning of adaptation cases, but not of response plan cases. The fourth condition included learning of both response plan cases and adaptation cases. The fifth and sixth conditions replaced "local search" with other memory search strategies, such as attempting to extract constraints on acceptable role-fillers of a schema and using them to define knowledge goals to be satisfied by a knowledge planning process. The fifth condition involved response plan learning only, and the sixth involved learning of both response plans and adaptation cases. We expected that either learning of response plan cases alone (conditions 2 and 5) or adaptation cases alone (condition 3) would improve performance over no learning (condition 1), and that the performance would be better with both types of learning (conditions 4 and 6) than with either individually.

Because most of the system's adaptation cost comes from memory search, efficiency was estimated by two different criteria reflecting memory search cost: the number of primitive memory operations performed and the number of memory nodes visited. In each case, lower values suggest less effort expended, but the two numbers can vary significantly as multiple operations can be applied to a single memory node, and, conversely, many nodes may be examined but never have operations applied directly to them. In general, variations in the order of presenting the problems may also have an effect on overall performance, but for the sample set changes in problem order did not appear to

From: Proceedings of the Third International Conference on Multistrategy Learning. Copyright © 1996, AAAI (www.aaai.org). All rights reserved.

	Nodes Visited					Memory Ops				
	Max	Min	Avg	Med	Dev	Max	Min	Avg	Med	Dev
Using "local search" to find needed information										
1. *No learning*	252	8	77	41	41	164	6	45	28	40
2. *Plan learning only*	252	8	55	36	64	164	6	32	22	38
3. *Adaptation learning only*	159	5	45	14	50	86	2	25	10	27
4. *Plan + Adaptation learning*	159	5	38	11	49	86	2	22	6	27
Using multiple strategies to find needed information										
5. *Plan learning*	1270	111	369	217	360	284	72	119	90	65
6. *Plan + Adaptation Learning*	1268	5	147	42	281	238	1	36	4	59

Table 2: Effort expended adapting the five sample cases.

have a significant effect. Table 2 shows the results for a single problem order. The table shows the maximum, minimum, average, median and standard deviations of the number of operations applied when processing the test stories in each condition.

In trials using local search, results were as predicted, with a combination of response plan learning and adaptation learning performing best overall. When multiple search strategies were used as opposed to local search, similar results were achieved, however, both the number of memory nodes visited and the number of operations performed were significantly higher than in tests using only local search. While initially surprising, this result can be explained by the fact that initial selection of the multiple search strategies is largely arbitrary. The dramatic decrease of the median in condition 6 suggests that adaptation learning is resulting in much more effective strategy selection.

The benefit of using multiple search strategies is evident in problems which are very difficult using local search methods alone. For example, adapting a Los Angeles earthquake response plan to generate a response plan for an earthquake in Liwa, Indonesia, requires adapting the means of transportation for relief supplies: The Los Angeles response plan involves the Red Cross sending supplies in by truck, but the roads to Liwa are impassable. In the real episode, the solution was a military airlift. The Red Cross and the military are distant in the system's memory and are each characterized by different constraints, making local search extremely costly, but performing a more strategic search process, characterizing the problem as "lack of access" and searching for actions to overcome that impediment and actors who could carry out those actions, suggests a more direct solution. In the tests, such problems did not arise often, however. When adaptation cases based on both local search and other strategies are saved and reused, average performance is better than for either method individually.

We note that, consistent with predictions, response plan learning did improve performance, as did adaptation learning. Interestingly, adaptation learning alone was somewhat more effective than case learning alone. 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. This is the foundation for the benefits of learning found in most CBR systems. Adding adaptation learning to response plan learning produced a moderate drop in cost 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.

These data are only suggestive; they involve a very small set of examples and the typicality of the examples is unclear. We plan to follow up on these initial data by performing a more controlled analysis of the effects of learning for a larger set of problem examples, and also to examine the potential utility problem (Francis & Ram 1993; Minton 1988) as the number of adaptation cases grows.

Learning similarity from adaptability

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. As Smyth & Keane (1995) observe, CBR systems often base similarity judgments on semantic similarity, but the real goal of their "similarity assessment" is to determine *adaptability:* how easily an old case can be adapted to fit the requirements of a new situation. If new adaptation strategies are learned, static similarity criteria do not keep pace with new capabilities for performing adaptations. When adaptation learning makes it easier to apply particular cases, those cases should be judged more relevant to a new situation. Thus similarity assessment criteria should change as new adaptation knowledge is acquired.

We have begun to study methods aimed at enabling

From: Proceedings of the Third International Conference on Multistrategy Learning. Copyright © 1996, AAAI (www.aaai.org). All rights reserved.

DIAL to improve its similarity assessment process by using learned adaptation cases to provide estimates of the cost of adapting particular types of problems. After retrieving an initial set of candidate response plan cases using traditional indexing techniques, DIAL compares their adaptability. For each problem to be repaired by adaptation, DIAL retrieves the adaptation case for the most similar prior problem (using the description of the adaptation problem). If the adaptation case was generated to solve an identical adaptation problem, the solution to that previous adaptation can be reapplied directly, resulting in very low adaptation cost: the cost is simply the cost to perform the needed transformation. If the adaptation case dealt with an adaptation problem that was similar but not identical, the cost of applying it to the new problem is estimated as the cost of the transformation used, plus the cost of the primitive memory search operations used in the prior adaptation. The rationale is based on the principle of derivational analogy: If a previous adaptation for a similar problem had to extract certain features and constraints, and transform them in certain ways to generate an appropriate adaptation, the process for the current situation should follow the same basic steps, even if the specifics of the situation are different.

If no similar adaptation case is found, DIAL uses a crude estimate of the cost: it maintains a record of the average cost (measured in primitive memory search operations) of adapting problems in each problem category, starting from scratch, and estimates the cost of the current problem using that average. By basing similarity assessment directly on the current state of its changing adaptation knowledge, DIAL's similarity assessment process reflects its adaptation knowledge. The aim is to improve the overall CBR process by favoring cases that are likely to be easier to adapt.

Perspective

Motivations for multistrategy learning during case-based reasoning

Our application of multistrategy learning to CBR is motivated by a number of potential benefits. One of these is that learning new similarity criteria and storing and "replaying" adaptations, by derivational analogy, will make it possible both to select better (more easily adaptable) prior cases and to expedite the adaptations that are performed, providing speedup learning.

In domains such as disaster response planning, for which no hard-and-fast rules are available to characterize what constitutes a good plan, an equally important potential benefit is increasing the quality of the solutions generated. Part of the appeal of reasoning from prior cases is to reflect regularities of a situation that may not be explicitly represented in a reasoner's domain theory. Thus storing and replaying successful adaptations may help to generate better adaptations than would be generated by reasoning from scratch. In the examples we have considered so far, the results generated by reusing adaptation cases are reasonable, and much more reliable than the results of, for example, simply selecting candidate substitutions by "local search." However, although we see this as an important potential benefit, the comparative effects on quality remain to be tested.

A final motivation for studying this multistrategy learning during CBR comes from cognitive modeling. Although the previous discussion provides functional arguments for learning to improve case adaptation skills and for adjusting similarity criteria as adaptation knowledge is learned, some psychological studies point to related aspects of human reasoning. Gentner & Toupin (1986), for example, 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 adults' similarity judgments 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. 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. Finally, Keane (1994) has shown that when selecting analogues for an analogical problem-solving task, subjects favor analogues that are more readily adaptable to the new problem situation.

Relationship to other computer models

A number of previous CBR systems learn by both case acquisition and refining their indexing criteria (e.g., Hammond, 1989; Veloso & Carbonell, 1994). A few include restricted mechanisms for learning limited forms of adaptation knowledge. 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*, which suggest steps that must be added to any recipe using particular ingredients (e.g., that shrimp should be shelled before being added to a recipe). CHEF uses special-purpose procedures to learn new ingredient

From: Proceedings of the Third International Conference on Multistrategy Learning. Copyright © 1996, AAAI (www.aaai.org). All rights reserved.

critics when omitted preparation steps cause recipes to fail. However, the learned adaptations can only be reused in very similar situations, while the adaptation cases learned by DIAL can be reused more flexibly because they are derivational traces of the results of a general introspective reasoning mechanism.

Learning to refine similarity criteria has been investigated in Prodigy/Analogy (Veloso & Carbonell 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. Our emphasis on adaptation cost is shared by Smyth and Keane (1995) , who have developed a CBR system that ties similarity judgments directly to adaptability, using heuristics coded to recognize the difficulty of performing particular types of adaptations. In their system, adaptation-guided retrieval results in significant improvements in overall problem-solving cost. In their work, however, similarity and adaptation knowledge are static. As our method learns new adaptations, it derives similarity criteria directly from its own experience with adaptation problems, changing both as it acquires adaptation experience.

DIAL's approach to memory search by question transformation is inspired by the memory search process of CYRUS (Kolodner 1984), and is similar to recent research on applying heuristic search to gathering information for argumentation (Rissland, Skalak, & Friedman 1994) and on strategic methods for information retrieval (Baudin, Pell, & Kedar 1994). Neither of these methods, however, learns from the search process. Our use of analogical techniques for internal reasoning is related to Ram and Cox's (1994) theory of *meta explanation patterns*, Kennedy's (1995) *internal analogy,* and Oehlmann's (1995) *metacognitive adaptation.*

Our use of transformational analogy for case-based planning, and derivational analogy for case-based reasoning applied to case adaptation, combines benefits of both learning strategies. Transformational CBR approaches store and adapt a *solution* to a problem, while derivational approaches store and replay a *derivational trace* of the problem-solving steps used to generate a previous solution. For CBR tasks such as disaster response planning, derivations of solutions are not generally available, and planning from scratch is not satisfactory because domain theories are inaccurate and intractable. However, examples of prior solutions are readily available in news stories and casebooks used to train disaster response planners (e.g., Rosenthal et al., 1989). This favors a transformational approach to reusing disaster response plans. On the other hand, derivational approaches can simplify the reapplication of a case to a new situation, and the rationale for the system's choice of particular steps during adaptation of prior cases *is* available. This makes it possible to use derivational analogy for learning about case adaptation and memory search.

Conclusion

We have described ongoing research on multistrategy learning within a case-based reasoning context, focusing 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 learning which types of adaptations are difficult to perform. Our approach to learning response plan cases uses transformational analogy; our approach to learning about case adaptation uses derivational analogy, which it applies both to memory search during adaptation and to the adaptation process as a whole. Our approach to learning 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 effects of adaptation learning on our system are encouraging for decreasing memory search cost, but more thorough tests are needed, both to study how the process "scales up" when large numbers of adaptations are learned and to determine effects on the quality of the response plans generated. Tests are also needed to examine how well current estimates of adaptation cost predict the difficulty of future adaptations. Our model is now being refined in preparation for more extensive tests of the system as a whole and the effects of its multiple forms of learning.

Acknowledgments

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

References

Barletta, R. 1994. A hybrid indexing and retrieval strategy for advisory CBR systems built with ReMind. In *Proceedings of the Second European Workshop on Case-Based Reasoning*, 49–58.

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*, 553–559.

Birnbaum, L.; Collins, G.; Brand, M.; Freed, M.; Krulwich, B.; and Pryor, L. 1991. A model-based

From: Proceedings of the Third International Conference on Multistrategy Learning. Copyright © 1996, AAAI (www.aaai.org). All rights reserved.
approach to the construction of adaptive case-based planning systems. In Bareiss, R., ed., *Proceedings of the DARPA Case-Based Reasoning Workshop*, 215-224. San Mateo: Morgan Kaufmann.

Carbonell, J. 1983. Learning by analogy: Formulating and generalizing plans from past experience. In Michalski, R.; Carbonell, J.; and Mitchell, T., eds., *Machine Learning: An Artificial Intelligence Approach*. Cambridge, MA: Tioga. 137-162.

Carbonell, J. 1986. Derivational analogy: A theory of reconstructive problem solving and expertise acquisition. In Michalski, R.; Carbonell, J.; and Mitchell, T., eds., *Machine Learning: An Artificial Intelligence Approach*, volume 2. Los Altos, CA: Morgan Kaufmann. 371-392.

Chi, M.; Feltovich, P.; and Glaser, R. 1981. Categorization and representation of physics problems by experts and novices. *Cognitive Science* 5(2):121-153.

desJardins, M. 1992. Goal-directed learning: A decision-theoretic model for deciding what to learn next. In *Proceedings of the Machine Discovery Workshop, Ninth International Machine Learning Conference*. San Mateo: Morgan Kaufmann.

Francis, A., and Ram, A. 1993. Computational models of the utility problem and their application to a utility analysis of case-based reasoning. In *In Proceedings of the Workshop on Knowledge Compilation and Speed-Up Learning.*

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*. San Diego: Academic Press.

Hunter, L. 1990. Planning to learn. In *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, 261-268. Cambridge, MA: Cognitive Science Society.

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

Keane, M. 1994. Adaptation as a selection constraint on analogical mapping. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, 490-495.

Kennedy, A. 1995. Using a domain-independent introspection mechanism to improve memory search. In *Proceedings of the 1995 AAAI Spring Symposium on Representing Mental States and Mechanisms*, 72-78. Stanford, CA: AAAI Press. Technical Report WS-95-05.

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

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*. San Mateo, CA: Morgan Kaufmann.

Leake, D.; Kinley, A.; and Wilson, D. 1995. Learning to improve case adaptation by introspective reasoning and CBR. In *Proceedings of First International Conference on Case-Based Reasoning*, 229-240. Sesimbra, Portugal: Springer Verlag.

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

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*, 84-89.

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*, 84-88. Stanford, CA: AAAI Press. Technical Report WS-95-05.

Mark, W.; Simoudis, E.; and Hinkle, D. 1996. Case-based reasoning: Expectations and results. In Leake, D., ed., *Case-Based Reasoning: Experiences, Lessons, and Future Directions*. Menlo Park, CA: AAAI Press. In press.

Michalski, R., and Tecuci, G., eds. 1994. *Machine Learning: A Multistrategy Approach*. San Mateo, CA: Morgan Kaufmann.

Michalski, R. 1994. Inferential theory of learning: Developing foundations for multistrategy learning. In Michalski, R., and Tecuci, G., eds., *Machine Learning: A Multistrategy Approach*. Morgan Kaufmann. chapter 1, 3-61.

Minton, S. 1988. *Learning Search Control Knowledge: An Explanation-Based Approach*. Boston: Kluwer Academic Publishers.

Oehlmann, R. 1995. Metacognitive adaptation: Regulating the plan transformation process. In *Proceed-*

From: Proceedings of the Third International Conference on Multistrategy Learning. Copyright © 1996, AAAI (www.aaai.org). All rights reserved.

ings of the Fall Symposium on Adaptation of Knowledge for Reuse. AAAI.

Ram, A., and Cox, M. 1994. Introspective reasoning using meta-explanations for multistrategy learning. In Michalski, R., and Tecuci, G., eds., *Machine Learning: A Multistrategy Approach.* Morgan Kaufmann. 349–377.

Ram, A., and Leake, D. 1995. Learning, goals, and learning goals. In Ram, A., and Leake, D., eds., *Goal-Driven Learning.* MIT Press.

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

Richter, M. 1995. The knowledge contained in similarity measures. Invited talk, the First International Conference on Case-Based Reasoning, Sesimbra, Portugal.

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*, 36–43. Seattle, WA: AAAI.

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

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.*

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*, 343–348. Bloomington, IN: Lawrence Erlbaum.

Veloso, M., and Carbonell, J. 1994. Case-based reasoning in PRODIGY. In Michalski, R., and Tecuci, G., eds., *Machine Learning: A Multistrategy Approach.* Morgan Kaufmann. chapter 20, 523–548.