Combining Reasoning Modes, Levels, and Styles through Internal CBR*

David B. Leake and Andrew Kinley
Computer Science Department
Indiana University, Bloomington, IN 47405
leake@cs.indiana.edu, akinley@cs.indiana.edu

Abstract

This paper discusses motivations and proposes methods for integrating multiple reasoning modes, styles, and levels within a case-based reasoning system. It describes a CBR system in which rule-based internal processing is augmented with two styles of case-based reasoning, derivational and transformational CBR, and which reasons at both the domain-level and the meta-level, in order to respond to the requirements of different processing tasks. The fundamental principal is for the system to learn by monitoring, capturing, and exploiting multiple types of prior system reasoning. The paper considers the ramifications of this approach and its potential as a strategy for multimodal reasoning in other contexts.

Introduction

The reasoning processes of artificial intelligence systems can be described along multiple dimensions, such as the reasoning mode or paradigm the system uses (e.g., rule-based reasoning or case-based reasoning), the style of reasoning within that paradigm (e.g., transformational or derivational approaches to case-based reasoning), and the level at which that reasoning is applied (e.g., domain-level reasoning or metareasoning). Their combination provides interesting opportunities for multimodal systems.

This paper summarizes a system combining multiple modes, styles, and levels of reasoning, describing its use of multimodal reasoning and considering the potential applicability of similar approaches to other systems. The system described is a case-based planner that uses multiple forms of reasoning to support its domain level case-based reasoning process. The system combines two reasoning paradigms, rule-based and case-based reasoning; two reasoning styles, transformational and derivational CBR; and two levels of reasoning, domain level reasoning (about plans) and metareasoning (about guiding the process for adapting plans to fit new situations).

The system's baseline reasoning process is transformational CBR; it generates new plans by adapting prior plans to fit new circumstances. Initially, this process adapts plans by rule-based reasoning, while internal case-based reasoning components capture the reasoning process used to adapt cases. As case adaptation experience is acquired, internal case-based reasoning supplants the rule-based process for case adaptation and similarity assessment. The case-based case adaptation process uses a different reasoning style from the baseline planner: it uses derivational CBR to compile and replay the reasoning underlying an adaptation.

This paper illustrates the usefulness of this multimodal approach by describing why specific reasoning modes are particularly well-suited to certain system processing tasks, how each approach contributes to the overall function of the system, and how the multiple approaches support each other. Based on experience with this system we make a more general claim: that using CBR components to monitor, capture, and replay a system's reasoning processes is a promising approach to guiding those processes and augmenting their capabilities.

Task and System Architecture

Our testbed system, DIAL (Leake, Kinley, & Wilson 1996), is a case-based planning system. DIAL's domain is disaster response planning, the initial high-level planning involved in deciding, for example, the basic outline for a plan to rescue and relocate the victims of a flood or earthquake. This is a domain for which no hard-and-fast rules exist, and case-based reasoning is often proposed as a reasoning paradigm for such domains. Unfortunately, in such domains it may also be difficult to formulate the knowledge required to guide the application of stored cases to new problems. Our multimodal reasoning approach is aimed at alleviating this problem by adding reasoning components that capture and replay the reasoning done to apply previous plan cases.

DIAL's basic planning process, transformational

*This work was supported in part by the National Science Foundation under Grant No. IRI-9409348. David Wilson has made many contributions to the project described in this paper. We thank the anonymous reviewers for their helpful comments.

case-based planning, uses multimodal reasoning components for both similarity assessment and case adaptation. Each of these uses CBR when possible, falling back on RBR when the CBR component fails. In addition, the rule-based case adaptation component relies on a multimodal memory search process that itself combines CBR and RBR. The results of rule-based case adaptation and memory search are stored in case libraries, to increase the knowledge available to those case-based processes in the future. Likewise, the results of the top-level disaster response planning process are saved for future reuse by case-based reasoning. Both the memory search process and the case-based similarity assessment process reason at the meta-level, the first reasoning about how to search the system's memory, and the second reasoning about the cost of adaptations by examining the reasoning processes involved in prior adaptations.

Figure 1 provides a schematic illustration of how each of the types of reasoning processes fits into our system's overall processing, and each reasoning type is described further in the following section.

Methods and Motivations

Disaster response plans must often be generated without complete information. In practice, human disaster response planners appear often to address this problem by **transformational case-based reasoning for generating disaster response plans**—their planning process is guided by remembering and adapting prior disaster response plans (Rosenthal, Charles, & Hart 1989). This is the process modeled by DIAL's top-level planning process.

An issue that arises in this approach is how to apply prior plans to new situations. Case adaptation is a classic problem in CBR: if a domain is poorly understood, reliable adaptation rules are likely to be unavailable a priori. The potential difficulty of applying prior plans to new situations is illustrated by one of DIAL's examples, the story of a 1994 flood in Allakaket, Alaska. When disaster response personnel selected a plan in which inhabitants of the town would be asked to volunteer to build levees, they discovered that all the able-bodied inhabitants were already away fighting forest fires. To adapt the prior plan, suitable replacements had to be found.

When DIAL processes the Allakaket story it retrieves the disaster response plan for a previous flood in which volunteers built levees, does a coarse-grained evaluation of this plan, and proposes this plan to a human user. The user informs it that the plan cannot be used because no volunteers are available.[1]

This problem triggers DIAL's multimodal case adaptation component. That component first attempts to use **derivational CBR for case adaptation**: to retrieve a trace of reasoning from a similar prior adaptation, in order to replay it. Adaptation cases are more operational, and may be more reliable, than the system's initial set of general rules for guiding the adaptation process. In this example, however, no similar prior adaptation is found, so DIAL falls back on **rule-based reasoning for case adaptation**. It begins by selecting a general adaptation rule for substitutions. Applying this rule depends on searching memory for the specific substitution.

DIAL uses two reasoning modes to search its memory. First, it attempts **derivational CBR for memory search**, attempting to retrieve a stored memory search case describing a similar memory search process (which might have been carried out for a very different adaptation). Derivational CBR can be used for this process because the system has access to all the reasoning underlying its memory search, and has the advantage of increased flexibility compared to transformational CBR: the derivation can be replayed in any situation for which a similar search process is expected to apply.

When DIAL processes this example, no applicable memory search case is available. Consequently, it falls back on **rule-based reasoning for memory search**. This is a metareasoning process that reasons about which strategies the system should use to search its own memory. One memory search rule calls for checking constraints on possible role-fillers for the actors in the prior plan and searching memory for other objects already satisfying those constraints. This check reveals that volunteers were placed under the authority of the police. Searching for others under the authority of the police, it finds prisoners as a possible substitution. Prisoners are judged a reasonable substitution by the system and the human user.

The internal CBR processes learn by saving two types of cases. First, a memory search case packages a trace of the successful search. This case is made accessible both in the case library of memory search cases, and within its set of memory search rules, supplementing the rules available to rule-based reasoning. If necessary in the future, this case may be adapted by extending or revising its search path using rule-based reasoning. Thus just as case acquisition supports the rule-based process by adding operational knowledge for it to apply, the rule-based process supports the application of memory search cases. Second, an adaptation case encapsulating the entire adaptation problem is saved, to supplant rule-based adaptation. When this case is reused, it may also be extended by rule-based reasoning to fit new needs. Our model of how this is carried out emphasizes the use of domain-independent strategies, with the aim of facilitating transfer of the approach to other task domains (Leake, Kinley, & Wilson 1997c).

Learned adaptation cases provide precisely the information needed for another rule-based/case-based inter-

[1] The user inputs this information in a fixed problem-description vocabulary; see (Leake, Kinley, & Wilson 1996).

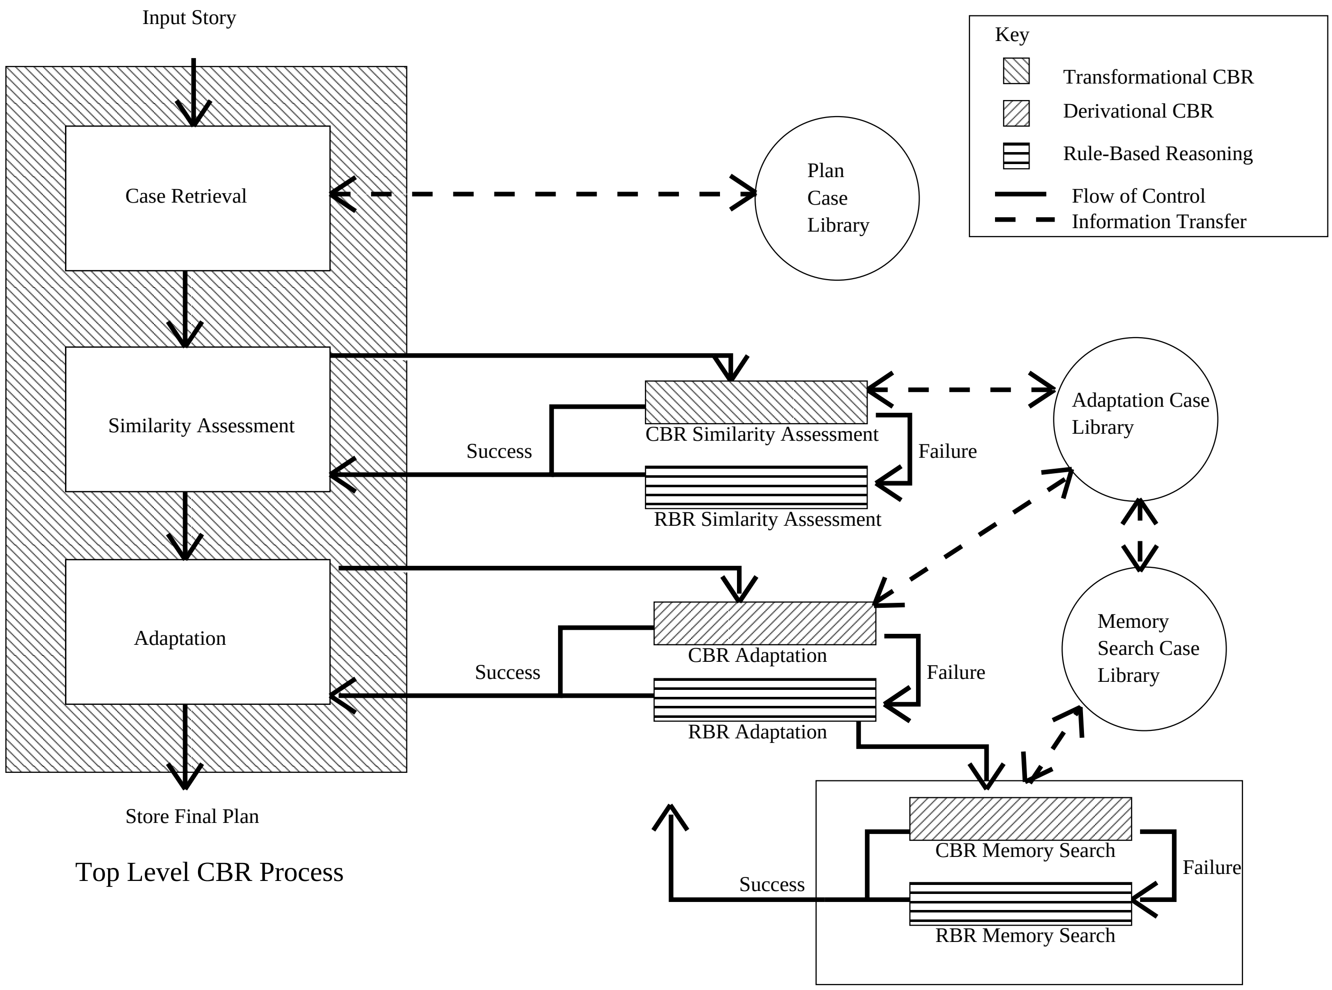

Figure 1: The integration of DIAL's multimodal reasoning processes to support top-level transformational case-based planning.

nal reasoning process, for similarity assessment. As has been pointed out by a number of researchers, the goal of "similarity assessment" in CBR is to select the prior cases that can be most easily applied to the new situation. As a result, useful similarity judgments must reflect "adaptability" (Birnbaum *et al.* 1991; Leake 1992; Smyth & Keane 1996). Because DIAL's rule-based case adaptation is augmented with case-based adaptation, which enables learning how to perform adaptations, adaptability of cases changes with adaptation learning—so similarity judgments must change as well. This led us to investigate case-based methods for augmenting rule-based similarity assessment. The key principle is that when similar prior adaptations have been done in the past, the difficulty of performing those adaptations should be predicted from the difficulty of the prior adaptations: by CBR.

DIAL initially selects a disaster response case to apply using **rule-based similarity assessment**, according to pre-defined domain-specific criteria. As adaptation cases are learned, it replaces this process with a **case-based similarity assessment** process that estimates the cost of new adaptations based on examining the reasoning process involved in similar prior adaptations. After an initial processing phase that retrieves a small set of candidate disaster response plans, DIAL's case-based similarity assessment component retrieves the adaptation cases DIAL would apply to adapt each problem in each plan, and—by examining its prior reasoning process and estimating the amount of effort the retrieved adaptations required when last applied, measured by the lengths of their derivational traces—estimates the total cost of adapting each candidate response plan to the new situation. This process not only judges similarity, but also provides the information needed for future adaptation. Retrieved adaptation cases for the best plan are passed on to the adaptation component, in the spirit of Smyth and Keane's (1996) adaptation-guided retrieval. Details are available in (Leake, Kinley, & Wilson 1997a).

Why Multimodal Reasoning

Each of the different reasoning methods applied in DIAL is selected in response to a different set of constraints. Case-based reasoning is appropriate for the domain planning task because of the availability of prior examples, the difficulty of developing rules capturing all the interacting factors in the domain, and the processing cost of building complicated disaster response plans from scratch. The transformational style of CBR is appropriate because the rationale for prior plans is seldom available.

Rule-based reasoning, using very general rules, enables initial adaptation or memory search with minimal knowledge acquisition effort. However, general adaptation rules are neither operational nor reliable. This supports using case-based reasoning when possible. Derivational CBR is practical for this task because the system can store the rationale for its successful adaptation decisions. This trace also provides an object for introspective reasoning, for example when predicting the cost of adaptation during similarity assessment.

The use of adaptation cases for similarity assessment as well as adaptation shows another interesting benefit of the approach in this system: the ability of case-based components to share knowledge, using cases from a common case base in different ways.

Contributions of the Multimodal Process to the System as a Whole

In case-based reasoning, the basic knowledge sources—cases and adaptation knowledge—are overlapping in the sense that each can compensate for weaknesses in the other. For example, a large case library can compensate for limited adaptation knowledge, by providing cases that require less effort to adapt. Conversely, good adaptation knowledge enables successful reasoning with a smaller case library, by facilitating the reuse of existing cases. The internal CBR components make it possible for the system to learn either domain cases or adaptation knowledge (or both), learning multiple lessons from its experiences.

The interaction of DIAL's methods also helps each part to perform its processing. When the rule-based memory search process uses a case to suggest a search path, the case focuses its processing on a sequence of steps that—because it was useful in the past—might be expected to be useful again. In turn, DIAL's rule-based reasoning can be called upon by its internal CBR components. The case-based components of DIAL are intentionally limited to using very simple CBR processes, to simplify knowledge acquisition for these components. Consequently, reapplication of a single case may result in only a partial solution, which is then augmented by RBR.

We are now gathering quantitative data on the benefits of this multimodal processing. An initial set of ablation studies of the contributions of different combinations of reasoning methods—no learning, domain-level CBR and RBR alone, adaptation CBR alone, and the combination—is described in (Leake, Kinley, & Wilson 1997b). In these tests, the overall processing speed of the combined system is superior to that of the "standard" CBR system on which it is based, as is the range of problems the system can solve.

The table below shows the average number of CPU seconds required to process each adaptation in a set of 118 adaptations required to generate response plans for 18 disasters, starting from a case library of 5 response plans. To test the extent to which internal CBR can compensate for deficiencies in the top-level CBR process, we tested the system both with regular plan learning (PL) of response plan cases by the top-level CBR system, and with that learning disabled. Interestingly, the processing speed with both adaptation learning and case learning was actually slower than with either alone. We hypothesized that the problem might be caused by a mismatch between the system's fixed rule-based similarity assessment criteria and the system's changing case adaptation abilities. In fact, when case-based similarity assessment was added in the complete system, the best average was achieved.

Test condition	**Time**
No learning	4.5
PL; internal RBR only	4.3
No PL; RBR/CBR for adaptation	3.0
PL + RBR/CBR for adaptation	3.2
All processes	2.2

Internal CBR as a General Strategy

We consider the use of case-based components that capture and reuse derivations of a system's own reasoning to be a promising approach that could be applied within a broad range of systems to learn useful reasoning paths and operationalize general knowledge. The basic strategy is to augment AI systems by embedding within them case-based "intelligent components" (Riesbeck 1996) that learn during normal processing. When similar problems arise in the future, the intelligent component furnishes a solution based on the stored case, to replace the initial reasoning process with CBR. Thus the intelligent component is seamlessly integrated with the initial system to improve its performance by building a case library covering actual problems the system encounters. DIAL's contributions to this area are both to investigate this general approach and to apply case-based components to capturing and reusing the *rationale* underlying the system's own reasoning processes.

Issues in applying case-based components include how the components' knowledge must be represented and organized, how much specialized knowledge must be provided to support component CBR processes (and how this effort compares to hand coding rules for these processes, given that the motivation for the component is to increase system performance while alleviating the knowledge acquisition burden), and the effects on overall performance. Experiences with our system provide only one data point about these issues, but the results are encouraging. Improvements were achieved despite the fact that very little effort was expended to tune the internal CBR systems.

It should be noted, however, that some of our specific methods depend on particular properties of the underlying system. For example, DIAL uses derivational analogy for its case adaptation process. In order to use this type of CBR, it is necessary to have access to a derivational trace of the underlying process that can be captured and reused. Because DIAL uses a "planful" memory search process (Leake 1995), it is practical to capture a trace of that process for reuse. This would not be possible in systems with a more opaque reasoning process.

The internal CBR process for case adaptation also benefits from knowledge already used for the top-level CBR process as the basis of its indexing. Standard CBR systems, such as DIAL's baseline CBR system with rule-based adaptation, must include some sort of indexing scheme to associate problems requiring adaptation to adaptation rules. The indexing scheme used for this purpose in DIAL's top-level CBR process is also used by its internal case-based adaptation component to index stored adaptation cases, decreasing the knowledge acquisition burden for this process. If case-based adaptation were added to a retrieval-only CBR system, for example, this information would not be available. Thus this strategy is appropriate for improving the performance of an existing adaptation component, but would be more difficult to use to provide an adaptation component starting from scratch. This is consistent with Riesbeck's general intelligent component strategy: the aim of case-based intelligent components is to improve the performance of existing processes, rather than to provide entirely new capabilities.

Relationship to Prior Research

Case-based reasoning systems commonly use rule-based systems for internal tasks such as guiding the application of cases (Kolodner 1993). Less common are CBR systems that integrate CBR and rule-based reasoning for a single task, though this has been investigated by a number of researchers, for example, for design (Goel 1989) and for guiding the application of ill-defined terms in the legal domain (Branting & Porter 1991; Rissland & Skalak 1991). CBR has also been used in multimodal systems for tasks such as making decisions when rules conflict (An, Cercone, & Chan 1997), as a strategically selected method within multimodal reasoning systems (Goel *et al.* 1994), and to address problems that cannot be solved by rule-based techniques (Surma & Vanhoof 1995).

Rather than focusing on multimodal reasoning at the domain level, the focus of our research is on multimodal *internal* reasoning; we primarily examine how internal derivational CBR guides the system's reasoning processes and augments their capabilities. The resulting cases increase reasoning efficiency, in the spirit of domain-level CBR systems that use cases to capture the domain-level results of a generative problem-solver (Koton 1988) or its reasoning process (Veloso 1994). In addition to focusing on internal reasoning issues, our approach emphasizes the ongoing integration of case-based and rule-base processing. Rather than simply falling back on rule-based reasoning when no case is available, its case-based and rule-based reasoning provide each other with information during ongoing processing. The rule-based reasoner can use derivational cases as collections of rules; the case-based reasoner can call on the rule-based reasoner as it refines a retrieved case.

Conclusion

Multimodal reasoning approaches have the potential to develop robust AI systems combining the strengths of multiple reasoning paradigms, but also raise questions of how to integrate the reasoning approaches. This paper illustrates an integration into a CBR system of internal CBR processes supporting—and supported by—rule-based reasoning, to guide the case application process of the main CBR system. The paper justifies this multimodal approach by discussing why each reasoning method used is particularly well-suited to its task and how different reasoning methods support each other and contribute to the overall function of the system. From this example it makes a more general claim: that building multimodal systems that use CBR components to monitor, capture, and replay the reasoning of other reasoning processes is a promising approach to supporting and augmenting their reasoning.

References

An, A.; Cercone, N.; and Chan, C. 1997. Integrating rule induction and case-based reasoning to enhance problem solving. In *Proceedings of the Second International Conference on Case-Based Reasoning*, 499–508. Berlin: Springer Verlag.

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., ed., *Proceedings of the DARPA Case-Based Reasoning Workshop*, 215–224. San Mateo: Morgan Kaufmann.

Branting, K., and Porter, B. 1991. Rules and precedents as complementary warrants. In *Proceedings of the Ninth National Conference on Artificial Intelligence*, 3–9. Menlo Park, CA: AAAI Press.

Goel, A.; Ali, K.; Donnellan, M.; de Silva Garza, A.; and Callantine, T. 1994. Multistrategy adaptive path planning. *IEEE Expert* 9(6):57–65.

Goel, A. 1989. *Integration of case-based reasoning and model-based reasoning for adaptive design problem solving*. Ph.D. Dissertation, The Ohio State University. Computer Science Department.

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

Koton, P. 1988. Reasoning about evidence in causal explanations. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, 256–261. San Mateo, CA: AAAI.

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*, 684–689. Menlo Park, CA: AAAI Press.

Leake, D.; Kinley, A.; and Wilson, D. 1997a. Case-based similarity assessment: Estimating adaptability from experience. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*. AAAI Press.

Leake, D.; Kinley, A.; and Wilson, D. 1997b. Learning to integrate multiple knowledge sources for case-based reasoning. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, 246–251. Morgan Kaufmann.

Leake, D.; Kinley, A.; and Wilson, D. 1997c. A case study of case-based CBR. In *Proceedings of the Second International Conference on Case-Based Reasoning*, 371–382. Berlin: Springer Verlag.

Leake, D. 1992. Constructive similarity assessment: Using stored cases to define new situations. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, 313–318. Hillsdale, NJ: Lawrence Erlbaum.

Leake, D. 1995. 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.

Riesbeck, C. 1996. What next? The future of CBR in postmodern AI. In Leake, D., ed., *Case-Based Reasoning: Experiences, Lessons, and Future Directions*. Menlo Park, CA: AAAI Press.

Rissland, E., and Skalak, D. 1991. CABARET: Rule interpretation in a hybrid architecture. *International Journal of Man-Machine Studies* 34(6):839–887.

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. 1996. Design à la Déjà Vu: Reducing the adaptation overhead. In Leake, D., ed., *Case-Based Reasoning: Experiences, Lessons, and Future Directions*. Menlo Park, CA: AAAI Press.

Surma, J., and Vanhoof, K. 1995. Integrating rules and cases for the classification task. In *Proceedings of the First International Conference on Case-Based Reasoning*, 325–334. Berlin: Springer Verlag.

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