Integrating CBR components within a Case-Based Planner

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

Abstract

Multimodal reasoning systems can improve the effectiveness of reasoning by integrating multiple reasoning methods, each selectively applied to the tasks for which it is best-suited. One integration approach is to bring CBR into other systems, by developing case-based *intelligent components* (Riesbeck 1996) that collaborate with other reasoning systems, monitoring their successes and failures and suggesting solutions when prior experiences are relevant. Another approach is to bring other reasoning processes into a CBR system's own architecture, to facilitate subprocesses of CBR such as case adaptation and similarity assessment. This paper describes a project combining both approaches: It discusses motivations and methods for a case-based components approach to integrating multiple reasoning modes, styles, and levels within a case-based reasoning system. The fundamental principle is for the system to use case-based components to learn by monitoring, capturing, and exploiting traces of multiple types of prior reasoning within the CBR system. The paper considers the benefits of this approach for improving CBR and its potential applicability to integrations in other contexts.

Introduction

Case-based reasoning provides opportunities for integrations both with other reasoning processes and within the CBR process itself. Case-based reasoning *intelligent components* (Riesbeck 1996), integrated with other reasoning systems, can augment those systems by monitoring their processing and learning from their successes and failures to increase the speed or quality of reasoning. Conversely, other reasoning methods integrated into case-based reasoning systems can help to support the fundamental subprocesses of CBR, such as case adaptation and similarity assessment.

Integrated reasoning systems can be characterized in multiple ways. One characterization, as proposed in the Call for Papers for the AAAI-98 Workshop *Case-Based Reasoning Integrations,* describes the control relationships of the integrated components: master-slave, slave-master, or collaborative. Another characterization focuses on the combinations of reasoning modes or paradigms in the components that are integrated (e.g., rule-based reasoning and case-based reasoning), the style of reasoning within each paradigm (e.g., transformational or derivational approaches to case-based reasoning), the level at which that reasoning is applied (e.g., domain-level reasoning or metareasoning), and the functionality provided by the integration. This paper focuses on this second type of characterization. It describes a project combining multiple modes, styles, and levels of reasoning, illustrating its use of an integrated approach to improve the processing of a CBR system and enable multi-component learning.

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).

This paper illustrates the usefulness of this integration by describing why specific reasoning modes are particularly well-suited to certain system processing tasks, how the processes interact, 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 Methods

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 multi-

modal reasoning approach responds to this problem by augmenting a case-based planner with additional CBR components that capture and reuse the reasoning done to *apply* previous plan cases.

The system's baseline reasoning process is transformational CBR; it generates new plans by adapting prior plans to fit new circumstances. This CBR process serves as a master for component CBR processes that perform case adaptation and similarity assessment. Those two processes in turn each involve two collaborating processes, one case-based and one rule-based. For example, initial case adaptation is done by rule-based reasoning. As case adaptation experience is acquired, internal case-based reasoning supplants the rule-based process for case adaptation and similarity assessment, but the rule-based process can still be used when problems arise that are beyond the scope of the system's stored cases. The following sections summarize DIAL's integrated processes.

The Top-Level CBR Process

Disaster response plans must often be generated without complete information. In practice, human disaster response planners appear 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. This is the process modeled by DIAL's top-level planning process. It is supported by multimodal slave processes for case adaptation and similarity assessment.

Integrations for Case Adaptation

A central problem for case-based planning is adapting prior plans to fit new circumstances. This is illustrated when DIAL processes the story of a 1994 flood in Alakaket, Alaska. It retrieves the disaster response plan for a previous flood in Bainbridge, Georgia. However, the plan cannot be re-applied as-is: the prior plan used volunteers to build levees, while all of the able-bodied inhabitants in the town are away fighting forest fires and therefore unable to perform the required task.

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. In this example, however, no sufficiently 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 a substitution for the previous volunteers.

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). When no applicable memory search case is available, as in this example, DIAL falls back on **rule-based reasoning for memory search**. 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.

DIAL's internal CBR processes supporting case adaptation learn by saving two types of cases: memory search cases and adaptation cases. A memory search case packages a trace of the successful search. This case supplements both a case library of memory search cases and the rules available to rule-based reasoning. If necessary in the future, this case can 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. Additionally, 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 1997b).

Integrations for Similarity Assessment

Learned adaptation cases provide the knowledge needed for another rule-based/case-based internal reasoning process, similarity assessment. As has been pointed out by a number of researchers, useful similarity judgments must reflect "adaptability" (Smyth & Keane 1996). Because DIAL's rule-based case adaptation is augmented with case-based adaptation, the adaptability of cases changes with adaptation learning—so similarity judgments must change as well.

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 adapting prior plans. 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 candidate plan, and 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.

Motivations for DIAL's Integrations

Each of the different reasoning methods applied in DIAL is selected in response to a different set of constraints for a particular system task. Transformational CBR is appropriate for the top-level domain planning task because of the availability of prior examples, the difficulty of developing rules capturing all the interacting factors in the domain, the processing cost of building complicated disaster response plans from scratch, and the lack of planning rationale to enable use of derivational CBR methods.

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.

More generally, the combination of these processes provides important benefits to the overall system. In CBR, 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 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.

Collaborative interaction of DIAL's case-based and rule-based methods 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 the system's 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. In this case, the CBR component calls upon the RBR component to complete the solution.

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 1997a). 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. However, further work is needed to examine the potential utility problem as the case library grows and to develop methods for controlling retrieval costs of internal cases (e.g., by selective forgetting).

Issues for Case-Based Components

We believe that the integration of case-based components, to capture and reuse experiences of a system's own reasoning process, has broad applicability for learning useful reasoning paths and operationalizing general knowledge. General issues in integrations of 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 an internal planning process to guide memory search, 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. Likewise, 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.

Relationship to Prior Research

A number of CBR systems integrate CBR and other reasoning methods, both as a means of increasing efficiency and as a way of alleviating problems of imperfect or incomplete domain theories (e.g., (Branting & Porter 1991; Goel 1989; Portinale & Torasso 1995;

Rissland & Skalak 1991; Surma & Vanhoof 1995)). Rather than focusing on multimodal reasoning at the domain level, the focus of our research is on multimodal *internal* reasoning. Early CBR research proposed internal CBR to guide adaptation (Sycara 1988); recent projects have examined the use of learning processes to improve adaptation knowledge (Hanney & Keane 1996) and similarity criteria (Ricci & Avesani 1995; Veloso 1994). DIAL's use of internal derivations CBR is most in the spirit of Oehlmann's metacognitive adaptation (Oehlmann 1995). DIAL's use of CBR for memory search to build up memory search information usable by case-based adaptation can also be seen as a type of multilayer CBR (Grimnes & Aamodt 1996).

Our approach also differs from other projects in emphasizing 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 can provide each other with information as needed during processing. The rule-based reasoner can use derivational cases as composite 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 help 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 its integration 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.

Acknowledgments

This paper is an adapted and revised version of "Combining Reasoning Modes, Levels, and Styles through Internal CBR," published in *Proceedings of the 1998 AAAI Spring Symposium on Multimodal Reasoning.* This work was supported in part by the National Science Foundation under Grant IRI-9409348. We thank David Wilson for his contributions to this project.

References

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. 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.

Grimnes, M., and Aamodt, A. 1996. A two layercase-based reasoning architecture for medical image understanding. In Smith, I., and Faltings, B., eds., *Advances in case-based reasoning*, 164–178. Berlin: Springer Verlag.

Hanney, K., and Keane, M. 1996. Learning adaptation rules from a case-base. In *Proceedings of the Third European Workshop on Case-Based Reasoning.* Berlin: Springer Verlag.

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. 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. 1997b. A case study of case-based CBR. In *Proceedings of the Second International Conference on Case-Based Reasoning*, 371–382. Berlin: Springer Verlag.

Oehlmann, R. 1995. Metacognitive adaptation: Regulating the plan transformation process. In *Proceedings of the Fall Symposium on Adaptation of Knowledge for Reuse.* AAAI.

Portinale, L., and Torasso, P. 1995. ADAPtER: An integrated diagnostic system combining case-based and abductive reasoning. In *Proceedings of First International Conference on Case-Based Reasoning*, 277–288.

Ricci, F., and Avesani, P. 1995. Learning a local similarity metric for case-based reasoning. In *Proceedings of the First International Conference on Case-Based Reasoning*, 301–312. Berlin: Springer Verlag.

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.

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.

Sycara, K. 1988. Using case-based reasoning for plan adaptation and repair. In Kolodner, J., ed., *Proceed-*

ings of the DARPA Case-Based Reasoning Workshop, 425–434. San Mateo, CA: Morgan Kaufmann.

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

Appendix

1. **Integration name**: DIAL
2. **Performance Task**: Top-level: Case-based planning. Internal: Case adaptation and similarity assessment
3. **Integration Objective**: Increase processing speed and solution quality compared to initial rule-based processing; simplify knowledge acquisition
4. **Reasoning Components**: CBR for top-level planning; RBR for adaptation, memory search, and similarity assessment; Derivational analogy for adaptation and memory search; CBR for similarity assessment
5. **Control Architecture**: Case-based planning as master for CBR/RBR adaptation and similarity assessment. Initial control in CBR/RBR processes is sequential (CBR first; RBR as fallback), with RBR supporting CBR
6. **CBR Cycle Step(s) Supported**: Retrieval and revision
7. **Representations**: Cases for disaster response plans; semantic net (to find substitutions and estimate similarity); rules for similarity assessment, memory search, and adaptation; cases for memory search paths and adaptations
8. **Additional Components**: Interactive adaptation
9. **Integration Status**: Initial evaluation of efficiency and coverage effects by lesion studies
10. **Priority future work**: Scaleup and further evaluation.