The Knowledge Engineering Review, Vol. 9:1, 1994, 61–64

Case-based reasoning

DAVID B. LEAKE
Computer Science Department, Indiana University, Bloomington, IN 47405, USA (email: leake@cs.indiana.edu)

1 Introduction

Case-based reasoning (CBR) systems reason from experience: they solve new problems by retrieving relevant prior cases and adapting them to fit new situations. In 1988 the first case-based reasoning workshop, sponsored by DARPA, identified theoretical foundations and fundamental issues for case-based reasoning research. Since then, much investigation has examined the CBR process itself, the validity of CBR as a cognitive model, and the application of CBR technology. The results of that work include refinements in theories of the case-based reasoning process, psychological evidence for human case-based reasoning, and the fielding of over 100 CBR applications.

At the Eleventh National Conference on Artificial Intelligence, in Washington, DC in July 1993, I chaired a two-day workshop on case-based reasoning. The goal of the workshop was to provide a forum for assessing the state of the field, including both the theory of case-based reasoning and its applications. The workshop enabled researchers addressing theoretical issues and developers building applied case-based reasoning systems to share their findings, to examine current progress, and to define promising areas for future research.

The response to the workshop showed a very active case-based reasoning community: based on paper submissions and statements of interest, over 60 investigators from academia and industry were invited to the workshop, and many other submitters had to be turned away because of workshop size constraints. The workshop included paper sessions discussing both theoretical and applied research, a poster session, invited talks by leaders in case-based reasoning, a panel to discuss the deployment of CBR technology, and much discussion throughout the two days of the workshop.

2 Background

Case-based reasoning systems process new situations by retrieving a relevant prior case from memory, comparing that case to the new situation to determine important similarities and differences, and applying pertinent information from the old case to the new situation. For example, a case-based planning system builds new plans by retrieving prior plans for similar goals and adapting those plans to fit the current goals and circumstances. The results of a CBR system's reasoning are added to the system's memory as a new case that can be retrieved and applied in similar future situations.

Motivations for using case-based reasoning include: increasing efficiency in solving new problems (because relevant reasoning can be re-used rather than having to be re-derived); improving the quality of solutions (because prior cases can guide the reasoner towards successful alternatives and warn of prior problems to avoid); and simplifying knowledge acquisition. CBR facilitates the knowledge acquisition process because cases can be stored and used without having a perfect domain theory; case acquisition can be achieved without analyzing the interactions between individual factors in a case. Also, extensive pre-existing case libraries are available for many tasks, and those libraries can provide a starting point when building CBR systems.

Case-based reasoning has been applied to many tasks and domains, such as message classification, dispute mediation, explanation, language understanding, legal reasoning and planning. For an in-depth discussion of the principles and issues in case-based reasoning, as well as overviews of recent case-based reasoning projects, excellent sources are Riesbeck and Schank (1989) and Kolodner (1993).

The theory and practice of case-based reasoning have been developed through a progression of CBR workshops that started with the 1988 DARPA workshop, which focused on theoretical

foundations and key issues for case-based reasoning (Kolodner, 1988). In 1989, a second DARPA workshop explored those issues in a series of issue-oriented panels (Hammond, 1989), and that workshop was followed by a 1990 AAAI Spring Symposium on case-based reasoning. In 1991, a third DARPA workshop examined progress on central issues in the context of implemented CBR systems (Bareiss, 1991). The aim of the AAAI-93 workshop was to assess theoretical and applied progress to establish the state of the art in case-based reasoning and to identify directions for future progress.

3 Major themes of paper presentations

Paper sessions focused on the central theoretical research areas of indexing and retrieval, representation and learning, and on the issues involved in applying CBR to the task areas of planning and scheduling, teaching, and aiding.

- *Indexing and retrieval:* The effectiveness of CBR systems depends upon their ability to retrieve the right case at the right time. Presentations on indexing included a case study of the iterative process involved in indexing cases for a system to teach architectural design (Zacherl & Domeshek); a method for facilitating generation of indices for multipurpose case libraries (Goldstein, Kedar & Bareiss); and a study of how factors such as imagery contribute to human case retrieval when solving creative design problems (Kolodner & Wills). Presentations on the retrieval process discussed methods for making the process more flexible by combining traditional indexing with best-first search (Rissland, Skalak & Friedman), and for guiding retrieval with goal-based feedback (Rissland, Daniels, Rubinstein & Skalak).
- *Representation:* Depending upon the task being performed, a case-based reasoning system may be able to function successfully using cases that have little internal structure. Other tasks may require cases that include more complete representations, or may require cases that can be manipulated at varying levels of granularity according to current needs. Papers in the representation session discussed the tradeoffs involved in using minimalist representations and their relationship to the issues of accuracy, efficiency and difficulty of indexing (Hinrichs, Bareiss & Slator); discussed the use of multi-layer representations in legal reasoning (Cuthill & McCartney); and presented a vocabulary for describing multi-agent interactions (Goldweic & Hammond).
- *Learning:* Learning by storing new cases is fundamental to the case-based reasoning process: CBR systems learn by adding the results of current processing to their case libraries as cases for future use. The learning session of the workshop examined how additional types of learning could be productively applied to CBR systems. For example, the rules that CBR systems use to adapt cases to fit new situations are pre-defined in most current CBR systems, and defining appropriate adaptation rules has proven to be an important practical problem. A method was proposed for acquiring adaptation rules by reasoning about the information needed for different types of adaptations and for learning how to seek that information (Leake). Likewise, it was shown that prototypical cases from a case library may be learned by genetic algorithms (Skalak), and preference predicates for guiding search through alternative cases may be acquired by compositional instance-based learning (Broos & Branting).

In addition to the paper sessions centring on fundamental theoretical issues, the workshop included paper sessions focused on issues in applying CBR to particular task areas:

- *Planning and scheduling:* Results in this session included methods for integrating search with case-based reasoning for path planning (Haigh & Veloso), methods for applying case-based reasoning to continuous control tasks, such as control of a robot's sensorimotor interactions with its environment (Ram & Santamaria), techniques for planning from partial and multiple plans (McCartney), and methods for applying CBR to scheduling (Miyashita & Sycara). Experimental results were presented to substantiate the benefits of CBR for the path planning and scheduling tasks.
- *Teaching and aiding:* A trend revealed by the workshop was a considerable emphasis in current research on case-based systems for teaching and aiding. Applications discussed at the workshop included teaching of design in beginning architecture classes, by allowing a student to browse

through a library of prior cases with the teaching system explaining, justifying and critiquing design decisions based on those cases (Goel, Malkawi, Pearce & Liu); teaching social skills, by presenting educational stories that can make appropriate points about social interactions (Burke); and teaching business decision-making (Slade). Systems for aiding in the development of X-ray treatment plans (Berger) and for aiding the development of computerized decision aids (Simoudis & Miller) were also described.

A poster session included 16 additional papers on theoretical issues and applications. The abstracts of those posters, as well as all presented papers, have been published in the workshop proceedings (Leake, 1993) which may be ordered from AAAI Press, 445 Burgess Drive, Menlo Park, CA, 94025, USA.

4 Invited addresses

The invited addresses provided four individual perspectives from leaders in case-based reasoning. The first address was by Kris Hammond, of the University of Chicago, who discussed the basic framework for case-based reasoning and examined the types of issues that must be addressed in CBR research. He proposed that case-based reasoning efforts be divided into three categories according to their fundamental goals: "true faith" CBR, "hard core" CBR and "CBR lite". In "true faith" CBR, the primary goal is addressing fundamental theoretical issues (both as they raise questions from a purely AI perspective and as cognitive models). In "hard core" CBR, the goal is testing and refining "true faith" ideas by applying them to build practical systems for performing challenging tasks. In "CBR lite", the goal is simply achieving a high level of performance, using selected ideas from CBR to further that end. These categories were widely adopted by workshop participants to describe the goals of CBR projects and for evaluating their contributions.

Roger Schank, of the Institute for the Learning Sciences at Northwestern University, presented a case-based model of human learning and examined the ramifications of that model for education. In light of the model, he argued that the aim of education should be to facilitate appropriate case acquisition. He argued that to encourage and facilitate acquisition of needed cases, instructional computer systems should teach by using "goal-based scenarios", rich learning environments designed so that students learn skills and conceptual knowledge as part of the pursuit of compelling goals.

Janet Kolodner, of the Georgia Institute of Technology, took stock of the field in an address entitled "A Case-Based View of Case-Based Reasoning—What Have We Wrought?" She discussed the state of the art in CBR applications, highlighting some of the over 100 CBR applications now fielded. Reports from the field show considerable acceptance of CBR from the user community. One reason is development time: reports suggest that CBR systems can be developed more rapidly than conventional systems. Another reason is that users seem to interact naturally with case-based advisory systems, which support the user's own decision-making process by suggesting relevant cases to consider. Her address also identified misconceptions about CBR and CBR issues (for example, the widespread belief that CBR systems must always have large case bases). In her view, overcoming these misconceptions will require a concerted effort to clarify the CBR paradigm and to better disseminate the nature of case-based reasoning.

Chris Riesbeck, of the Institute for the Learning Sciences, discussed possible next steps beyond the current CBR paradigm. In his view, performance of current CBR systems is sometimes impaired by their "fussiness": they insist on generating a perfect solution, which may result in their failing to generate any solution at all. He proposed an alternative framework, nicknamed the "shoot first and ask questions later" model, in which the CBR system presents tentative results and continues to search for a better solution. If it finds a better solution, it revises its response. As the search continues, the user may offer guidance and additional clarification based on the results that are currently available, facilitating further processing.

5 Final panel: Deploying CBR

The workshop closed with a panel to examine issues in the deployment of CBR systems. The panel was chaired by Ray Bareiss of The Institute for the Learning Sciences, with participants Kris

Hammond, Alan Meyrowitz, of the Navy Artificial Intelligence Center, Hiroaki Kitano, of NEC and Carnegie Mellon University, and Evangelos Simoudis, of Lockheed AI Center. Panelists pointed to many CBR systems that are succesfully deployed, and especially to the usefulness of CBR-based advisory systems and decision support systems. For example, Kitano described SQUAD, a large-scale case-based advisory system for software quality control that has been successfully deployed as a corporation-wide information source at NEC.

6 Conclusions

The large response to the workshop reflected a high level of activity in case-based reasoning, and the workshop itself showed a dynamic CBR community. Active questioning and followup discussion raised and clarified many important issues, sharpening the ideas presented in the workshop. The workshop demonstrated that there is a significant research focus on the problems of indexing and retrieval—on indexing vocabularies, on retrieval strategies, and on the development of tools to facilitate the indexing process—and showed a trend towards applying CBR for design and design support tasks, and towards case-based learning environments.

The workshop also revealed that case adaptation is the least understood aspect of case-based reasoning. Because of the difficulty of developing appropriate adaptation rules, many current CBR systems are designed to function primarily as memories, retrieving relevant cases that the user must apply. Such systems have proven very useful, but fully autonomous CBR systems must perform their own case adaptation. Therefore, a key future focus for CBR research is to deepen understanding of the case adaptation process.

Acknowledgements

I would like to thank the organizing and program committees of the workshop for making the workshop possible through their efforts, and to thank them and the workshop participants for the workshop's success. I would also like to thank AAAI for its sponsorship of the workshop.

References

Bareiss, R (ed.), 1991. *Proceedings of the DARPA Case-Based Reasoning Workshop*. San Mateo, CA: Morgan Kaufmann.

Hammond, K (ed), 1989. *Proceedings of the DARPA Case-Based Reasoning Workshop*. San Mateo, CA: Morgan Kaufmann.

Kolodner, J (ed.), 1988. *Proceedings of the DARPA Case-Based Reasoning Workshop*. San Mateo, CA: Morgan Kaufmann.

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

Leake, D (ed.), 1993. *Proceedings of the AAAI-93 Workshop on Case-Based Reasoning*. Menlo Park, CA: AAAI Press.

Riesbeck, C and Schank, RC, 1989. *Inside Case Based Reasoning*. Hillsdale, NJ: Lawrence Erlbaum.