Introduction to the *Computational Intelligence* Special Issue to Appear in 2001

Maintaining Case-Based Reasoning Systems

Case-based reasoning (CBR) is the process of reasoning and learning by storing prior *cases*—records of specific prior reasoning episodes—and retrieving and adapting them to aid new problem-solving or interpretation in similar situations (Kolodner 1993, Leake 1996, Watson 1997). Case-based reasoning systems rely on the knowledge contained in multiple "knowledge containers" (Richter 1998), such as the case-base, case adaptation knowledge, and similarity criteria. The contents of each of these knowledge containers may affect system efficiency and the quality of results. Over time, the knowledge containers may need to be updated in order to maintain or improve performance in response to changes in the system's knowledge, task, environment, or user base. This gives rise to the need for strategies to address the problem of maintenance in case-based reasoning systems.

Experience with the growing number of deployed case-based reasoning systems has led to increasing recognition of the value of maintenance to the success of practical CBR systems, as well as the importance of maintenance research. Maintenance issues arise in designing and building CBR systems and support tools that monitor system state and effectiveness in order to determine whether, when, and how to update CBR system knowledge to better serve performance goals. Understanding the issues that underlie the maintenance problem, and using that understanding to develop good practical maintenance strategies, is crucial to sustaining and improving the efficiency and solution quality of CBR systems as their case-bases grow and as their tasks or environments change over long-term use.

Maintaining CBR systems is an active research area that has been well represented at recent conferences. This special issue brings together mature work, focusing on maintaining the essential underlying knowledge of case-based reasoning systems. It provides a snapshot of the state of the art, presenting twelve articles examining core issues, methods, and lessons learned from research and applications. Topics include the foundations of CBR system maintenance—the components of the maintenance process and maintenance goals—as well as proposals for specific maintenance strategies, theoretical and empirical analyses of their performance, and lessons on maintenance arising from practical experience.

In order to understand the issues involved in developing maintenance strategies, as well as to understand maintenance practice and identify opportunities for new research, it is useful to understand the nature of the maintenance process and its relationship to the overall CBR process. The first paper in this issue, Wilson and Leake's "Maintaining Case-Based Reasoners: Dimensions and Directions," provides a characterization of what maintenance is, the components of maintenance policies, and the dimensions along which alternative maintenance policies may differ. It then uses that characterization to examine the state of the art and identify opportunities for future research. Of course, the success of maintenance depends not only on the maintenance policies themselves, but also on how maintenance is integrated with the overall case-based reasoning process. Reinartz, Iglezakis, and Roth-Berghofer's article "On Quality Measures For Case-Base Maintenance" describes an extended, 6-step, CBR cycle that incorporates two explicit maintenance steps into the traditional CBR cycle. The first step, *Review*, makes use of general measures of the quality of a case-base in order to determine whether maintenance is necessary. The second step, *Restore*, makes use of general operators that can be combined to address problems detected in the review phase. The potential of this approach is shown as it is applied to standard and industrial data sets.

A key concern for case-base maintenance is assessing the quality of the case-base, in order to use that assessment to guide refinement of case-base contents. Competence models, which characterize how well a system's case-base covers the set of potential problems, play a vital role in this task. Smyth and McKenna's article, "Competence Models and the Maintenance Problem," provides a survey of their recent research focusing on the role of competence models as a guide to a variety of case-base maintenance strategies. They describe their particular approach to competence modeling and go on to explain and evaluate how this model has driven the development of a range of competence-guided maintenance solutions to problems such as case-base editing, case retrieval, and case-base visualization.

Because the performance of case-based reasoning systems may be adversely affected by the utility problem as the size of the case-base grows, strategies for controlling the size of the case-base are a major focus of current research. Yang and Zhu's paper "A Case-Addition Policy for Case-base Maintenance" presents an algorithm for adding cases incrementally to a case-base. They substantiate their algorithm with a theoretical analysis showing that it provides a well-defined lower-bound on coverage and that it dramatically increases the competence of a case-base when new cases are added, proving properties of case addition that have previously been studied empirically. Portinale and Torasso's "Case-base Maintenance in a Multi-Modal Reasoning System" proposes two policies for maintaining case-base knowledge. One is competence-based, focusing on compacting the case-base when a new case can replace a set of prior cases while retaining roughly comparable coverage. The second is a deletion process guided by tracking the usefulness of stored cases during system problem-solving, which is especially relevant to CBR systems combined with first-principles reasoners that can re-generate deleted cases. The benefits of both approaches are demonstrated empirically.

Other papers in this issue examine the benefits of applying maintenance to additional types of knowledge. Muñoz-Avila's "Case-Base Maintenance by Integrating Case Index Revision and Case Retention Policies in a Derivational Replay Framework" examines two maintenance policies for derivational replay planners. One guides case retention based on the efficiency benefits of using the retrieved case as a start for new reasoning; the other guides index revisions, based on whether the results of retrievals can be extended to new problems without revising the planning decisions suggested by the retrieved case. Shiu, Sun, Wang and Yeung's article "Transferring Case Knowledge to Adaptation Knowledge: An Approach for Case-base Maintenance" presents a methodology for extracting adaptation knowledge from a case-base automatically. Their system uses a fuzzy decision-tree learning approach to transform a given case-base into a representation that is more conductive to case adaptation. This approach provides for the transfer of knowledge between two knowledge containers, the case-base and adaptation knowledge.

The problem of supporting case authoring is receiving increasing attention from the CBR community as a whole. Ferrario and Smyth's "Distributing Case-Base Maintenance, the Collaborative Maintenance Approach" describes and evaluates an approach which allows the case authoring task to be distributed across a variety of authors, and which provides support mechanisms to manage and review author submissions. It provides a method for managing the ongoing maintenance of case-bases in dynamic domains where traditional human-based or automatic maintenance strategies prove too costly or ineffective. McSherry's "Intelligent Case-Authoring Support in CaseMaker-2", also focuses on the case acquisition task, and presents a system that performs background reasoning on behalf of the case author while new cases are being added, in order to help the user determine the best cases to add in light of their competence contributions. The system uses its evaluations of the contributions of potential cases to suggest cases to add to the case library.

Three articles illustrate the practical benefits of applying general maintenance strategies in the context of ongoing applications. Craw, Jarmulak, and Rowe's "Refining Retrieval Knowledge After Case-Base Maintenance" applies a genetic algorithm approach to maintaining retrieval knowledge, refining the parameters and feature weights used for indexing and similarity matching. The approach is demonstrated on a pharmaceutical tablet formulation application, with benefits shown as the case-base expands and in response to changes in the manufacturer's tablet formulation policy. Nick, Althoff and Tautz's article "Systematic Maintenance of Corporate Experience Repositories" investigates organizational learning issues, and in particular, use of the Experience Factory approach from software engineering along with case-based reasoning, as a framework to support repository-based organizational learning. This article explores the issue of maintenance in an experience-base and proposes a flexible and practical maintenance and evaluation framework that is well suited to organizational learning and corporate knowledge scenarios. Watson's "A Case Study of Maintenance of a Commercially Fielded Case-Based Reasoning System" examines maintenance practice in the context of Cool Air, a system to support engineers in creating HVAC installation specifications. Maintenance support in Cool Air addresses issues of redundancy and obsolescence in the case-base, based on the practical needs of engineers using the system, as well as on the organizational requirements for managing their knowledge. Its approach is then placed in the larger context of maintenance work.

These papers bring together a range of issues, methods, and lessons learned from research and applications in building systems and support tools that explicitly address the maintenance aspects of CBR systems. We hope that this collection will provide a valuable resource for researchers, implementers, maintainers, and users of CBR systems, as well as a useful starting point for others interested in this area.

David B. Leake, Indiana University

Barry Smyth, University College Dublin

David C. Wilson, University College Dublin

Qiang Yang, Simon Fraser University

References

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

Leake, D., ed. 1996. *Case-Based Reasoning: Experiences, Lessons, and Future Directions.* Menlo Park, CA: AAAI Press/MIT Press.

Richter, M. 1998. Introduction. In Lenz, M.; Bartsch-Spörl, B.; Burkhard, H.-D.; and Wess, S., eds., *CBR Technology: From Foundations to Applications.* Berlin: Springer. chapter 1, 1–15.

Watson, I. 1997. *Applying Case-Based Reasoning: Techniques for Enterprise Systems.* San Mateo, CA: Morgan Kaufmann.