The Utility Problem in Case-Based Reasoning

Anthony G. Francis, Jr. and Ashwin Ram
College of Computing
Georgia Institute of Technology
Atlanta, Georgia 30332-0280
{centaur,ashwin}@cc.gatech.edu

Case-based reasoning systems may suffer from the utility problem, which occurs when knowledge learned in an attempt to improve a system's performance degrades performance instead. There are two main classes of utility problems: *performance utility problems* and *search-space utility problems*. Search-space utility problems, such as the branching problem in some macro-operator systems [ETZIONI 1992] are a symptom of poorly designed learning algorithms. Performance utility problems, in contrast, are a direct consequence of increased cost of memory accessing and matching as the size of the knowledge base or of learned items increase. Learning algorithms that do not take this overhead into account can cause a system to slow down more than the average speedup provided by individual learned rules. For example, the cost of matching individual items in the Soar system caused the *expensive chunks problem* [TAMBE ET AL. 1990], and the cost of matching a whole rulebase in Prodigy caused the *swamping* problem [MINTON 1988].

Massive parallelism is often offered as a potential solution to the performance utility problem, reducing match cost to nearly constant time. Retrieval of memory items requires both parallel matching of cases and indices *and* selecting the best case. Unfortunately, the cost of matching can grow arbitrarily large as the size of individual cases increases; furthermore, selection algorithms that provide constant-time performance fail to scale up to large case bases. Theoretical considerations of parallel architectures enforce a lower bound of $\Omega(\lg n)$ on the time complexity of ideal selection algorithms [COOK ET AL 1986]; because this lower bound can still cause the utility problem, we must instead turn to various *coping strategies* for solutions. For example, coping strategies for swamping can include deletion policies, restricting search, or restricting learning; coping strategies for expensive chunks can include restricting expressiveness or early termination of long matches.

CBR already incorporates several coping strategies that make it resistant to the utility problem. First, cases have the potential to eliminate vast amounts of problem solving, providing improvements robust enough to survive an architectural slowdown. Second, because the cost of case retrieval is amortized over many adaptation steps, ideal case-based reasoners suffer less severely from the same overhead than conventional problem solvers. Finally, while CBR systems can suffer from the expensive chunks problem, they can easily incorporate a restricted expressiveness policy into the indexing scheme by placing an upper bound on the size of an item that can be matched [DOMESHEK 1992]. However, there have been few attempts to systematically evaluate the cost-utility tradeoffs in CBR systems with very large case libraries.

The authors are currently constructing a memory module called Moore which can be used to represent and access large case libraries without running into the utility problem. Moore uses a combination of restricted expressiveness and asynchronous match policies to limit the expensive chunks problem and a guided search policy called context focusing to limit swamping.

REFERENCES

[COOK ET AL. 1986] Cook, S.; Dwork, C.; Reischuk, R.. "Upper and lower time bounds for parallel random access machines without simultaneous writes." *SIAM Journal on Computing*, 15(1):87-97, 1986.

[DOMESHEK 1992] Domeshek, E. "Do the right thing: A component theory for indexing stories as social advice." *Technical Report #26*, May 1992. The Institute for the Learning Sciences, Northwestern University.

[ETZIONI 1992] Etzioni, O. "An Asymptotic Analysis of Speedup Learning." In *Machine Learning: Proceedings of the Ninth International Workshop*, 1992.

[MINTON 1990] Minton, S. "Quantitative results concerning the utility of explanantion-based learning." *Artificial Intelligence*, 42(2-3), March 1990.

[TAMBE ET.AL. 1990] Tambe, M.; Newell, A.; Rosenbloom, P. S. "The Problem of Expensive Chunks and its Solution by Restricting Expressiveness." *Machine Learning*, v5 p 299-348, 1990.