

A Reply to Reich's Book Review of *Exemplar-Based Knowledge Acquisition*

RAY BAREISS

(BAREISS@ILS.NWU.EDU)

Institute for the Learning Sciences, Northwestern University, Evanston, IL 60201

Yoram Reich has chosen to discuss Protos (Bareiss, 1989) solely as an automated knowledge acquisition system. The Protos project did result in the development and distribution of such a system. However, the motivation for the project was to formulate and test a unified approach to representing, acquiring, and reasoning with polymorphic concepts. I believe this aspect of the research will be of more enduring interest to the Machine Learning community than the Protos knowledge acquisition system itself.

Reich correctly points out the difficulties of evaluating a complex system and understanding its behavior. Our experiments, summarized in Porter et al. (1990), involved five people over the course of two years. Even so, these experiments left many important questions unanswered. I can, however, answer some of the specific questions asked in Section 3.1:

1. **Why was audiological diagnosis chosen as a test domain; is it an instance of a problem class that Protos is explicitly designed to address?** Audiological diagnosis was chosen because it is representative of the general class of tasks for which Protos was designed, an interested expert was available to train Protos, and a group of human subjects was available for a comparative study. As new applications are attempted that do not have the desirable characteristics of audiological diagnosis (Bareiss, 1989; p. 93), it is reasonable to assume that Protos's performance will degrade. The causes of this degradation, and hence, the limitations of a Protos-like approach are important topics for further research.
2. **How does case presentation order affect learning and performance?** Because of the exemplar retention criteria, case presentation order will profoundly affect exemplar retention. Classification efficiency will be affected by an increased number of exemplars, but accuracy should not be.
3. **What is the usefulness of having single-exemplar categories; do these singleton classes affect classification?** Single-exemplar categories were included because they naturally occurred in the training set (as discussed in Question 4). If an exemplar of one of these categories proves to be most similar to a new case during classification, its class is assigned to the case. Note that some categories in Protos's knowledge base are adequately represented by one exemplar because of strong associated domain theory, while others are temporarily singletons because of incomplete learning.
4. **How were the training and test sets selected; why was the largest class not tested?** The training and test sets were selected "pseudo-randomly," i.e., our domain expert obtained cases from a large speech and hearing clinic in the order in which patients

presented themselves for treatment. We expected that this training order would provide a more realistic test of Protos than a handcrafted order. Because test cases were selected in the same way, the most frequently occurring category during training did not occur in the test set. We relied on the expert's judgement to determine the appropriate training and test sets for a convincing experiment in his domain.

The discussion of Protos's explanation capabilities highlights the component of Protos with which we are least satisfied. The intention underlying its design was to permit simple inferences relating features to categories and to other features. Given a large domain theory, computational tractability will be a problem, as will maintaining consistency during learning. The Multifunctional Knowledge Base Project at the University of Texas (Porter et al., 1988), especially Ken Murray's knowledge-integration research (Bareiss et al., 1989), is addressing these issues.

In conclusion, I would like to thank Yoram Reich for his thoughtful review and *Machine Learning* for the opportunity to reply.

References

- Bareiss, R. (1989). *Exemplar-based knowledge acquisition: A unified approach to concept representation, classification, and learning*, Academic Press.
- Bareiss, R., Porter, B., & Murray, K. (1989). Supporting start-to-finish development of knowledge bases. *Machine Learning*, 4, 259-283.
- Porter, B., Bareiss, R., & Holte, R. (1990). Concept learning and heuristic classification in weak-theory domains. *Artificial Intelligence*, 45, 229-263.
- Porter, B., Lester, J., Murray, K., Pittman, K., Souther, A., Acker, L., & Jones, T. (1988). *AI research in the context of a multifunctional knowledge base: The botany knowledge base project* (Technical Report AI88-88). Austin, TX: University of Texas at Austin, Artificial Intelligence Laboratory.