Machine Learning, 6, 99–103 (1991)


# *Book Review*

## Exemplar-Based Knowledge Acquisition, by Ray Bareiss. Academic Press, 1989.

YORAM REICH (YORAM@CS.CMU.EDU)
*Department of Civil Engineering, Carnegie Mellon University, Pittsburgh, PA 15213*

### 1. Introduction

*Exemplar-Based Knowledge Acquisition* is about the design, construction and testing of the Protos knowledge acquisition system for classification tasks in weak domain theories. It touches issues in knowledge representation, problem-solving, human-computer interaction, machine learning, and psychological validity. This review tries to evaluate the book as well as the Protos system itself from a knowledge acquisition perspective.

### 2. About the book

Bareiss' book starts with a short introduction and a guide to the book. Chapter 2 provides a top-level view of Protos' building blocks: knowledge representation, classification and learning. Chapter 3 provides a detailed description of the Protos system. Each of the three aspects of Protos is described to the lowest level detail, including the variety of algorithms implicit in Protos' behavior. Together with Chapter 2, it is a comprehensive, clear, and well-presented account of Protos. Chapter 4 describes an evaluation of Protos as a knowledge acquisition tool operating in the domain of clinical audiology. Chapter 5 reviews related research with special attention to case-based reasoning and knowledge acquisition; it is more a descriptive account than an analytical comparison. The book concludes by describing the research contribution and outlining a future research plan.

Bareiss' book serves as excellent Protos documentation. The description provided would permit a reconstruction of the original system and continued testing of its behavior (Dvorak, 1988), an important aspect in scientific research. A complementary account of the underlying principles guiding Protos' design and some of its limitations can be found in (Porter et al., 1990). One such limitation is Protos' inability to assess the consequences of its refinement methods. Scaling up to large knowledge bases may require additional tools (i.e., KI, Bareiss et al., 1989) in order to avoid inconsistencies.

For a deeper understanding of the Protos approach I recommend reading Porter et al. (1990), Bareiss (1989), and Bareiss et al. (1989), in this order. A Common-Lisp implementation of Protos, called CL-Protos, with its manual (Dvorak, 1988) rounds out the comprehensive description.

## 3. Methodological issues

The following is this reviewer's interpretation of several general methodological issues for knowledge-acquisition research compiled from the special issue of *Machine Learning* on Knowledge Acquisition (Buchanan, 1989; Boose & Gaines, 1989; Clancey, 1989):

1. Free access to techniques and tools by the research community allows for further testing and dissemination of ideas.
2. Detailed evaluation of applications, techniques and tools are necessary in order to understand their underlying principles, their source of power, and possible weaknesses.
3. Diversity of techniques is important, since premature standardization may impede progress.

The first two issues are related to two styles of research discussed in the next section. The last issue is discussed in Section 3.1, and poses the complex problem of how to integrate this diversity of techniques.

### *3.1. Research styles*

There are two constrasting styles of research in machine learning. The first style favors studying simple problems while trying to understand the principles of learning. Early experimental work as well as current research on formal learning theory fall in this category. This style allows one to build relatively small programs and test them easily on a variety of small-size problems. This style is based on sound scientific methodology; however, it rarely scales-up to dealing with real-world problems.

The second style is characterized by building more complex systems to address real-world problems. Usually, these systems rely on design heuristics imported by their designers. These heuristics are rarely tested experimentally. Demonstrating the ability of these systems is hard and in many cases only a single demonstration is performed. The "scientific nature" of the result is less impressive; indeed, one test cannot substantiate the general utility of the approach or shed light on the principles guiding the system success.

There is a complementary research strategy that alleviates the above difficulty; making the system available to the research community for further experiments. Since the original PROTOS research falls into the second style, and was initially tested only in the clinical audiology domain, the developers of PROTOS have embraced this approach by providing CL-PROTOS to any researcher. Further studies with the PROTOS system (mostly available as technical reports from the AI Lab at the University of Texas at Austin), provide additional information on the PROTOS approach.

Testing a large system on a single domain introduces significant reservations. The first experiment with the PROTOS system left many unanswered questions. For example:

- Why was the clinical audiology task chosen as a test domain? Does it stretch PROTOS' ability to its limits, or is it an instance of a problem class that PROTOS is explicitly designed to address?

- Can the reported learning and performance effects be replicated? How does the case-presentation order affect learning and performance?
- What is the usefulness of having single-case categories if they cannot be tested? Do these singleton classes influence classification?
- Why is the division to training and test sets nonuniform? E.g., why was the largest class not tested?

An interesting analysis of the interaction between the PROTOS system and the expert is reported. Although the conclusions drawn are only preliminary, such analysis is important when dealing with complex large systems, since it can be used to construct a model of the system's knowledge level.

### 3.2. *Integration of tools*

Chandrasekaran (1989) has described the importance of identifying general type-specific architectures, called *generic tasks*, and using them as mechanisms for knowledge acquisition and problem solving. In addition, Buchanan (1989) has critized the traditional syntactic view of machine learning as being "*knowledge-poor*," whereas learning should itself be a problem-solving activity. These two ideas can be merged to serve the task of knowledge acquisition, namely: different learning schemes are good for different tasks and *knowledge about learning* can be used to select between them for the appropriate learning task (Reich & Fenves, 1989).

This important notion of integration is addressed by the PROTOS system in its start-to-finish support of knowledge-base construction. The main method is the assimilation of cases into its knowledge base by creating appropriate indices. The second method is a simple form of learning-by-completing-explanations. In this case, PROTOS elicits additional domain knowledge from the expert in order to complete a chain of reasoning to establish correspondences between cases. Another learning process refines inconsistencies by adding various types of indices (e.g., difference and censors indices). Finally, a simple form of generalization of cases occurs when the relative importance of properties of a case are modified.

The PROTOS approach is to integrate knowledge acquisition with performance. The main trigger for learning is failure in classification problem-solving. This failure driven approach to eliciting new knowledge in the appropriate place side-steps the credit-and-blame assignment problem.

## 4. Flexibility: The source of power and the root of possible problems

One of PROTOS' most appealing characteristics is its flexible interaction with a domain expert. This, in turn, increases the complexity of the PROTOS system's operation in three ways:

1. The language used by PROTOS admits only nominal properties.[1] This may constrain experts (or the system) if the domain cannot be fully captured by such a set of nominal terms; indeed the domain of clinical audiology seems to belong to this class. Terms

such as *s-neural* (*mild*, *greater-than-2k*) and *s-neural* (*mild*, *greater-than-3k*) are treated as two different properties. This results in the loss of implicit ordering information and complicates the description of cases and concepts.
2. The interactive nature of PROTOS' operation and the desire to avoid restricting the expert allows the expert to alias existing terms (e.g., *notch-at-4k* equals *notch-4k*).
3. A case may be described by properties having multiple values. For example *history* can take the values *vomiting* and *dizziness* for the same case.

The first two complications are handled by the explanation facility which relates terms using a pre-defined syntax and semantics. As the language grows, these explanations become harder to maintain. The third complication will cause difficulties in matching terms when creating and using explanations.

There is a similarity between PROTOS' explanations (which are products of both problem-solving and external expert guidance) and SOAR's chunking mechanism (Laird & Newell, 1986). In SOAR, the tradeoff between expressive language and computational efficiency (known as the problem of *expensive chunks* (Tambe et al., 1990)) can be avoided by reducing the expressiveness of the language. Since PROTOS favors a flexible language, it uses a second approach which performs an heuristic, possibly incomplete, match operation when indexing and explaining the correspondence between exemplars.

The elaborate expressive language used by the PROTOS system might prove to be computationally expensive. It is not clear what the limitations of this heuristic approach are, since there is no complexity analysis presented. The effectiveness of the heuristics should be explored in larger domains and their semantics should be better understood. This reviewer suspects that the same tradeoff of expressiveness vs. efficiency as occurs in SOAR will arise in PROTOS when used in more complex domains, especially due to the elaborate language introduced in CL-PROTOS.

The PROTOS system appears to achieve an excellent performance despite (or maybe because of) the flexible language used. This suggests that PROTOS may have captured certain important intrinsic properties of the domain. It would be interesting to solicit a domain expert's assessment of PROTOS' internal knowledge representation in the clinical audiology domain. This might lead to some deeper insights in this domain.

## 5. Summary

*Exemplar-Based Knowledge Acquisition* is an excellent reference document describing a promising knowledge acquisition tool, the PROTOS system. Since it is publicly available, PROTOS provides a testbed for a variety of techniques and issues in machine learning: (1) the integration of similarity-based and explanation-based approaches, (2) the transition from user guidance to autonomy, and (3) the relation between knowledge representation and efficiency. Much work remains to be done to assess PROTOS' scaleability and to uncover and repair any possible existing complexity problems.

## Acknowledgments

I would like to thank Bruce Porter for providing additional material on PROTOS and to Alberto Segre for his comments on a draft of this review.

## Notes

1. CL-Protos allows a much more elaborate language, including recursive definitions of terms with both nominal and numeric components, quantifiers and relations.

## References

Bareiss, R. (1989). *Exemplar-based knowledge acquisition*. Boston, MA: Academic Press.

Bareiss, R., Porter, B.W., & Murray, K.S. (1989). Supporting start-to-finish development of knowledge bases. *Machine Learning, 4*, 259–283.

Boose, J.H., & Gaines, B.R. (1989). Knowledge acquisition for knowledge-based systems: notes on the state-of-the-art. *Machine Learning, 4*, 377–394.

Buchanan, B.G. (1989). Can machine learning offer anything to expert systems? *Machine Learning, 4*, 251–254.

Chandrasekaran, B. (1989). Task-structures, knowledge acquisition and learning. *Machine Learning, 4*, 339–345.

Clancey, W.J. (1989). The knowledge level reinterpreted: modeling how systems interact. *Machine Learning, 4*, 285–291.

Dvorak, D.L. (1988). *Guide to CL-Protos: An exemplar-based learning apprentice* (Technical Report AI88-87). Austin, Texas: University of Texas at Austin, Artificial Intelligence Laboratory.

Laird, J.E., Rosenbloom, P.S., & Newell, A. (1986). Chunking in Soar: the anatomy of a general learning mechanism. *Machine Learning, 1*, 11–44.

Porter, B.W., Bareiss, R., & Holte, R.C. (1990). Concept learning and heuristic classification in weak-theory domains. *Artificial Intelligence*, (to appear).

Reich, Y., & Fenves, S.J. (1989). *Integration of generic learning tasks* (Technical Report EDRC 12-28-89). Pittsburgh, PA: Carnegie Mellon University, Engineering Design Research Center.

Tambe, M., Newell, A., & Rosenbloom, P.S. (1990). The problem of expensive chunks and its solution by restricting expressiveness. *Machine Learning, 5*, 299–348.