

Introduction: A Sampler in Knowledge Acquisition for the Machine Learning Community

This special issue is devoted to invited editorials and technical papers on knowledge acquisition. In the past, special issues have been devoted to recognized subfields of machine learning, where a subfield might be characterized by a particular method of machine learning, such as genetic algorithms. The relationship between machine learning and knowledge acquisition is not so clearcut as the field-subfield one. Neither are the methods of knowledge acquisition so homogeneous and easily characterized as for genetic algorithms.

Just as with machine learning as a whole, people who consider themselves to be working in the field of knowledge acquisition are identified more by the goal of their work than by any particular methodology they apply. I think most of us in knowledge acquisition would agree that our goal is to make expert systems easier to build and maintain and, along the way, to make the expert systems built more explainable, more robust, and so on. This aim has given the knowledge acquisition field an engineering flavor. Our tools and methodologies are judged by their suitability for their user community and the environment in which they are used, and by the performance of the knowledge bases they produce and the expert systems they support. We tend to place a heavy emphasis on fielding systems and evaluating the practicality of our approaches.

In the quest for this goal, knowledge acquisition workers have employed a diversity of methods. Some work in knowledge acquisition, including some of the earliest work, was done using traditional machine learning techniques (e.g. [Michalski and Chilausky, 1980; Quinlan, 1986]), but other knowledge acquisition approaches fall outside of what is generally viewed as machine learning. In general, machine learning techniques are appreciated in the knowledge acquisition field only according to their ability to achieve practical results. In turn, the machine learning community seems to show little interest in tools or methodologies that have great practical benefits to today's expert systems but don't exhibit "real" learning.

In part, a goal of this issue is to extend a hand from the knowledge acquisition community to the machine learning community. To a large degree, knowledge acquisition and machine learning share a common goal: We are each trying to improve the performance of some automated system. Knowledge acquisition is very eager to use what it can from machine learning. It is both rewarding and challenging to put one's work into practice—even more challenging as we raise our ambitions for the performance of our application systems. We hope you'll accept the challenge, or at least keep it in mind. Furthermore, because of our shared goal, we have encountered and will continue to encounter some of the same problems. Issues such as appropriateness of domain models, design of lexicons for exchange between teacher and pupil, credit assignment in debugging, and sources and use of bias cut across both fields. Work in either community can reveal issues and approaches that should be of interest to the other community, even when the results are not directly importable.

What follows is a sampler of work in knowledge acquisition. It comprises three technical papers and six guest editorials. The technical papers give an in-depth look at some of the important issues and current approaches in knowledge acquisition. The editorials were produced by authors who were basically invited to sound off. I've tried to group and order the contributions somewhat coherently. The following annotations emphasize the connections among the separate pieces.

Buchanan's editorial starts on the theme of "Can machine learning offer anything to expert systems?" He emphasizes the practical goals of knowledge acquisition and the challenge of aiming for them.

Lenat's editorial briefly describes experience in the development of CYC that straddles both fields. He outlines a two-phase development that relies on an engineering approach early on and aims for a crossover to more automated techniques as the size of the knowledge base increases.

Bareiss, Porter, and Murray give the first technical paper. It comes from a laboratory of machine learning researchers who have taken an interest in supporting the development of knowledge bases, with an emphasis on how development changes with the growth of the knowledge base. The paper describes two systems. The first, Protos, adjusts the training it expects and the assistance it provides as its knowledge grows. The second, KI, is a system that helps integrate knowledge into an already very large knowledge base. It is designed to help acquire knowledge at the fringes of a knowledge base built using the CYC knowledge-base editing tools.

Clancey's editorial reflects on the enterprise of building expert systems. He debunks a popular myth about the nature of knowledge and provides a thoughtful perspective on the goals and challenges of the knowledge acquisition process.

Gruber's is the second technical paper of the collection. He analyzes the problems of knowledge acquisition in terms of representation mismatch, taking up some of the modeling themes in Clancey's editorial. These ideas are applied in the development of ASK, an interactive knowledge acquisition tool that elicits strategic knowledge from people. He defines strategic knowledge to be that knowledge used by an agent to decide what action to perform next, where actions have consequences external to the agent.

McDermott's editorial focuses on the aim of one particular line of work in knowledge acquisition—one that makes use of task-specific architectures. The practical aim is conveyed in the title: "The World Would Be a Better Place if Non-Programmers Could Program."

Chandra's editorial also explores the issue of how task-specific architectures support knowledge acquisition and emphasizes the connection to machine learning approaches as well. He supports his ideas with a description of some recent work in the OSU Generic Tasks group.

Musen's is the final technical paper. His work embraces the aim that McDermott describes—that is, to provide a tool that allows non-programmers to program. He places his work in the framework of constructing and extending domain models. This paper describes PROTÉGÉ, and interactive knowledge-acquisition system that addresses the two activities individually.

Finally, Boose and Gaines' editorial comments on the state-of-the-art in knowledge acquisition, based on their experience as the organizers of a series of workshops on the subject. It gives a little flavor of the variety of work being done in knowledge acquisition, but

emphasizes the development of interactive knowledge acquisition tools. For those interested in more information on knowledge acquisition, their editorial concludes by listing sources of information on past workshops and by announcing future events.

References

- Michalski, R.S., and Chilausky, R.L. 1980. Learning by being told and learning from examples: An experimental comparison of the two methods of knowledge acquisition in the context of developing an expert system for soybean disease diagnosis. *Policy Analysis and Information Systems*, 4. 125–160.
- Quinlan, J.R. 1986. Induction of decision trees. *Machine Learning*, 1. 81–106.

—Sandra Marcus
Advanced Technology Center
Boeing Computer Services