The Knowledge Engineering Review, Vol. 8:1, 1993, 1–4

There's no theory better than that of a practical tool

MARK A. MUSEN

Knowledge Systems Laboratory, Departments of Medicine and Computer Science, Stanford University, Stanford, CA 94305-5479, USA

This special issue of *The Knowledge Engineering Review* includes three papers on the subject of automated knowledge-acquisition tools. Gaines and Shaw (1993) review knowledge acquisition tools based on personal-construct psychology; Birmingham and Klinker (1993) review tools based on predefined models of problem-solving methods; and Eriksson and Musen (1993) review metalevel tools that developers may use to build other knowledge acquisition tools. Although this trio of papers by no means summarizes all recent developments in computer-based assistance for knowledge engineering, these contributions provide examples of three important research frontiers that, in many ways, trace their origins to early expert systems research.

The now-legendary difficulties of knowledge acquisition were immediately apparent when people set out to build the first expert system over 25 years ago (Buchanan et al., 1969); research to overcome these difficulties began nearly as soon as the developers realized how labour-intensive is constructing a useful knowledge base. Davis' seminal system, TEIRESIAS (Davis, 1976), was an early proof of concept that demonstrated how an interactive tool might aid system builders in the maintenance of large knowledge bases. In subsequent years, a growing cadré of investigators began to explore new ways for computer systems to aid developers in the creation and support of knowledge-based systems (see Boose & Gaines, 1988; Neale, 1988; Musen, 1993).

Much of the work on automated knowledge acquisition tools has had a strong software-engineering flavour. Knowledge acquisition tools, by definition, achieve success on the basis of whether they can assist their users in practical knowledge-engineering activities. Like tools for computer-assisted software engineering (CASE), knowledge acquisition tools can aid developers with much of the bookkeeping that is required when a large software system is built, and can provide abstractions—often visual—for viewing and manipulating complex data structures. Like CASE tools, knowledge acquisition tools also must maintain an internal model of how the abstractions viewed and manipulated by the user are to be mapped into the data structures that are required by the computer programs that will run within the final system. Knowledge bases are, after all, pieces of software, and principles of software engineering necessarily apply to the construction of knowledge-based systems (Wielinga et al., 1992). Knowledge acquisition tools are more than just specialized CASE tools, however.

Knowledge bases are not just software products; they are models of problem-solving competence (Clancey, 1989). Knowledge bases represent, at the least, models of how problem solving for particular application tasks can take place. Some researchers further suggest that knowledge bases can be models of how *human beings* solve particular tasks.

There is perhaps no greater source of conflict within the knowledge-based systems community than the contention concerning the nature of knowledge-based systems. The debate centers on whether knowledge bases represent merely computational models, or whether they can represent cognitive models as well. Ironically, for the earliest builders of rule-based systems, there apparently was no such controversy; rule bases *were* cognitive models. After all, Newell and Simon (1972) had demonstrated convincingly that production systems could serve as the basis for much of human cognition, and human experts indeed seemed to convey their knowledge in terms of rules. Buchanan and Shortliffe (1984, pp. 55–56), for example, observed during knowledge engineering of the MYCIN system that "by discussing specific difficult patients, and by encouraging our collaborators to justify their questions or decisions, those of us who were not expert in the field began to tease out 'nuggets' of domain knowledge—individual inferential facts that the experts identified as pertinent for problem solving in the domain." System builders consequently identified a correspondence between the data structures in the emerging knowledge base, and the specific heuristics that they presumed experts used during problem-solving.

Eventually, however, both expanding interest in connectionist architectures within the artificial-intelligence community (McClelland et al., 1986) and rumblings of doubt concerning the previously undisputed role of physical symbol systems in human cognition (Clancey, 1991) required developers of traditional knowledge-based systems to step back and to reconsider the nature of their enterprise. Traditional rule bases began to look less and less like 'real' cognitive models. Today, the question of whether knowledge-based systems are models of expertise or merely models of problem-solving competence finds its answer in the eye of the beholder. The implications of the continuing confusion have become particularly significant for those workers building knowledge acquisition tools.

Developers of knowledge acquisition tools based on personal-construct psychology (Gaines & Shaw, 1993), for example, attempt to relate the knowledge bases that their tools construct to a well-described theory of how people make distinctions about the world, of how people anticipate events, and of how people select appropriate actions (Kelly, 1955). Such tools have as their basis Kelly's theory of human behavior, and generate knowledge-based systems that necessarily are consonant with that theory. Tools such as ETS (Boose, 1985) and KSSn (Gaines & Shaw, 1993) embody Kelly's psychological theory, and validate that theory each time they create a new knowledge base that has expert-level performance.

Other developers of knowledge acquisition tools, on the other hand, are at best agnostic regarding any relationship between the knowledge-based systems that their tools can generate and human problem-solving activities. The models of problem-solving methods that underlie the tools discussed by Birmingham and Klinker (1993) are useful because they clarify the roles in which an inference engine uses knowledge during problem solving (McDermott, 1988), and because they can facilitate the generation of more specialized knowledge acquisition tools (Musen, 1989a). The motivations for using method-oriented models—although substantial—are primarily related to the desire to improve the engineering of large knowledge bases. Knowledge acquisition tools based on problem-solving methods thus embody a theory of knowledge-base construction, rather than a theory of cognition. The practical success of these model-based tools in some sense validates their underlying software engineering theory.

Interactive knowledge acquisition tools thus automate specific theories of knowledge engineering. The papers in this special issue by Gaines and Shaw and by Birmingham and Klinker illustrate two extremely different theories, although these theories are not necessarily mutually exclusive or competitive. Personal-construct psychology itself does not address issues of software engineering—although knowledge acquisition tools based on personal-construct psychology can be viewed from a software engineering perspective (Musen, 1989b). Similarly, models of problem-solving methods are artifacts that reflect few, if any, suppositions about the cognitive processes of the human beings who may use knowledge acquisition tools based on those methods—although the cognitive processes themselves can be effectively studied (Allemang & Rothenfluh, 1992). The many dimensions of knowledge engineering make it possible for knowledge acquisition tools to address multiple, discrete aspects of the life cycles of knowledge-based systems (Motta et al., 1990). The theory incorporated within any particular tool should not generally be construed to be making a claim about the comprehensive process of knowledge acquisition and system maintenance.

The first two papers in this special issue thus present cognitive and engineering theories of knowledge acquisition, and show how these theories can be translated into practical knowledge acquisition tools. The third paper discusses how metalevel tools may generate custom-tailored knowledge acquisition tools—each with its own theory—from abstract models of the knowledge to be acquired and of the means to acquire it (Eriksson & Musen, 1993). The argument for interactive systems that developers can use to design better-adapted knowledge acquisition tools makes sense from a software-engineering perspective: Computer programs are most usable when they permit their users to think about the tasks at hand in natural ways. There also is an important scientific argument in favour of such metalevel knowledge acquisition environments: When the theories that underlie computer-based knowledge acquisition tools can be represented as explicit models, with terms and relations that are editable, not only can these theories be examined more directly, but also they can be readily tested and revised.

From the rule-based theory of knowledge associated with TEIRESIAS, to the theory of personal constructs incorporated in tools such as KSSn, to the theories of software engineering advanced by contemporary task-specific architectures, interactive knowledge acquisition tools embody theories of cognition, of knowledge representation, of problem solving, and of human–computer interaction. The validity of these theories must be ascertained from the pragmatic effectiveness of the various tools. Whereas there may be nothing more practical than a good theory, it is only through the practice of building real knowledge bases that we can acquire the empirical data that can test our theories in the first place. The challenge, of course, is to build our knowledge acquisition systems so that we may critically evaluate our work. When such experiments are possible, the validation of our research hypotheses—both through controlled experiments and through anecdotal experience—often has the gratifying side effect of producing useful knowledge bases that can aid other people in their professional activities.

Computer-based knowledge acquisition tools are important both from theoretical and from practical perspectives. Perhaps the most exciting element of the work presented in this issue is that these two dimensions are so thoroughly intertwined.

References

Allemang, D and Rothenfluh, TE, 1992. "Acquiring knowledge of knowledge acquisition: A self-study of generic tasks" In: T Wetter, K-D Althoff, J Boose, BR Gaines, M Linster and F. Schmalhofer (eds.), *Current Developments in Knowledge Acquisition—EKAW '92.* Springer-Verlag, pp 353–372.

Birmingham, W and Klinker, G, 1993. "Knowledge-acquisition tools with explicit problem-solving models" *Knowledge Engineering Review* **8** pp 5–25.

Boose, JH, 1985. "A knowledge acquisition program for expert systems based on personal construct psychology" *International Journal of Man-Machine Studies* **23** 495–525.

Boose, JH and Gaines, BR, 1988. *Knowledge Acquisition Tools for Expert Systems,* Academic Press.

Buchanan, BG and Shortliffe, EH (eds.), 1984. *Rule-based Expert Systems: The MYCIN experiments of the Stanford Heuristic Programming Project,* Addison-Wesley.

Buchanan, BG, Sutherland, G and Feigenbaum, EA, 1969. "Heuristic DENDRAL: A program for generating explanatory hypotheses in organic chemistry" In: B Meltzer and D Michie (eds.), *Machine Intelligence vol 4*, Edinburgh University Press, pp 209–254.

Clancey, WJ, 1989. "Viewing knowledge bases as qualitative models" *IEEE Expert* **4** 9–23.

Clancey, WJ, 1991. "Why today's computers don't learn the way people do". In: PA Flach and RA Meersman (eds.), *Future Directions in Artificial Intelligence,* North Holland, pp 53–62.

Davis, R, 1976. Applications of meta level knowledge to the construction, maintenance, and use of large knowledge bases. PhD dissertation, Department of Computer Science, Stanford University, California, Report STAN-CS-76-564.

Eriksson, H and Musen, MA, 1993. "Conceptual models for automatic generation of knowledge-acquisition tools" *Knowledge Engineering Review* **8** pp 27–47.

Gaines, BR and Shaw MLG, 1993. "Knowledge acquisition tools based on personal construct psychology" *Knowledge Engineering Review* **8** pp 49–85.

Kelly, G, 1955. *The Psychology of Personal Constructs,* Norton.

McClelland, JL, Rumelhart, DE and the PDP Research Group (eds.), 1986. *Parallel Distributed Processing: Explorations in the microstructure of cognition,* MIT Press.

McDermott, J, 1988. "Preliminary steps toward a taxonomy of problem-solving methods" In: S Marcus (ed.), *Automating Knowledge Acquisition for Expert Systems,* Kluwer, pp 225–256.

Motta, E, Rajan, T and Eisenstadt, M, 1990. "Knowledge acquisition as a process of model refinement" *Knowledge Acquisition* **2** 21–49.

Musen, MA, 1989a. *Automated Generation of Model-based Knowledge Acquisition Tools,* Pitman.

Musen, MA, 1989b. "Conceptual models of interactive knowledge-acquisition tools" *Knowledge Acquisition* **1** 73–88.

Musen, MA, 1993. "An overview of knowledge acquisition" In: J-M David, J-P Krivine and R Simmons (eds.), *Second Generation Expert Systems*, Springer-Verlag.

Neale, IM, 1988. "First generation expert systems: A review of knowledge acquisition methodologies" *Knowledge Engineering Review* **3** 105–145.

Newell, A and Simon, H, 1972. *Human Problem Solving*, Prentice-Hall.

Wielinga, B, Schreiber, AT and Breuker, J, 1992. "KADS: A modeling approach to knowledge engineering" *Knowledge Acquisition* **4** 5–53.