

Knowledge versus Knowledge Representation

Daniel Roy

1 Introduction

The conception of knowledge by the engineers charged with building expert systems contrasts sharply with that of social scientists. However, the narrow focus of Diana Forsythe on knowledge engineers missed a respect of the effects of representing knowledge in the larger AI community [FOR2]. In “What Is a Knowledge Representation” [DAVIS], Davis et al. describe the five roles of a knowledge representation, observing that a knowledge representation requires that a set of ontological commitments as well as theories of reasoning be embodied, necessarily biasing the knowledge and resulting inferences. I argue that the failure of knowledge engineers to express their notion of knowledge is largely due to the restrictive apparatus of the rule-based expert systems in which they operate. Instead, I argue that a truly “expert” system will require numerous knowledge representations, just as a human uses, and that a true, automated expert will most likely need to be trained in much the same way that a human expert is: through life experiences as well as expert training. Expertise is not solely the formal knowledge in the expert domain, but the entire collection of social and cultural experiences.

2 Background

Knowledge-based, or expert, systems were built throughout the last forty years, growing in popularity in the 1980’s as AI scientists explored practical applications of their techniques and algorithms. Systems such as DENDRAL, MYCIN, RA-1, and CYC were built upon “inference engines” that operated on databases of formal rules that embodied a set of knowledge. Knowledge was encoded by “knowledge engineers,” usually graduate students who sat, face-to-face, with experts, slowly encoding the complicated inferences of the expert into simple, mechan-

ical rules. These systems were met with resistance by end-users, despite, in many cases, achieving very good technical performance.

3 Survey

Diana Forsythe observed and interviewed knowledge engineers at several academic labs in order to understand the knowledge elicitation process and to determine how assumptions made in this process affected the resulting expert systems [FOR2]. Forsythe clearly lays out the distinctions between how social scientists and knowledge engineers view the nature of knowledge [FOR2, 463].

First, while knowledge engineers consider the definition of knowledge to be straightforward (almost obvious), social scientists consider the nature of knowledge to be incredibly hard to define. Immediately, knowledge engineers underestimate their difficult task.

Second, knowledge engineers treat knowledge in very binary terms; An expert has knowledge while a novice lacks knowledge. On the other hand, social scientists tend to believe that knowledge is highly situational and cultural. The result is that engineers misunderstand that their systems encode a certain knowledge bias that is social and cultural.

Third, knowledge engineers consider that reasoning occurs in much the same way that their expert systems are built—via formal rules. In contrast, social scientist believe that, instead of a universal formal logic, reasoning is highly dependent on social and cultural factors. This restrictive setting results in an overly narrow definition of knowledge, and diminishes the power of the resulting expert systems.

Fourth, knowledge engineers build their systems with the assumption that all useful knowledge is stored

in the mind of the expert, while social scientists attribute a large amount of knowledge to “cultural, social and organizational order” [FOR2, 464].

Fifth, while knowledge engineers tend to ignore tacit knowledge both because their techniques only illicit “conscious” knowledge and because they tend to ignore social and common-sense knowledge, social scientists have developed techniques to acquire tacit knowledge because they believe that such knowledge is crucial to the overall understanding of knowledge.

Sixth, knowledge engineers ignore the problems of relying solely on interviews to extract knowledge. In contrast, social engineers use observation in the actual work place to observe action and compare this to the informants’ perceptions’ of their actions. The result is that the encoded knowledge is fundamentally flawed and does not correspond to actual practice. This problem is further exacerbated by the fact that most systems employed only a single expert and considered the resulting knowledge to be representative of the entire field.

Finally, knowledge engineers consider knowledge to be universally applicable, across domains, and devoid of social or cultural bias. Social scientists, instead, stress the social and cultural aspects of knowledge, observing that knowledge is highly local and situational.

4 Analysis

Forsythe’s observations are very important in that they point out systematic omissions in the work of knowledge engineers. By ignoring tacit knowledge, social and cultural bias, and by using single experts to represent entire fields, the knowledge at the foundation of these expert systems is fundamentally flawed. Such problems, no doubt, have had an effect on the usefulness of these systems for end users [FOR1].

Forsythe objects to the incredibly narrow view of knowledge that the engineers hold. However, her informants seemed to have been entirely ignorant of conceptions of knowledge held by the rest of their field outside expert systems. In fact, there is some discussion of these same issues Forsythe brings up in [DAVIS] (however, there is little overlap in nomencla-

ture). That these specific knowledge engineers considered knowledge to be straightforward is most likely a product of their dealing with a system in which the definition of knowledge is a formal rule: it is no wonder that, when forced to work within such restrictions, knowledge engineers mold their definition of knowledge into equally restrictive domains.

For many scientists in the field of AI, the notion of knowledge is best understood by observing how to represent knowledge. According to Davis et al., knowledge representation is defined by the five roles it plays [DAVIS]:

1. It is a replacement for the real thing: a surrogate.
2. It is a set of ontological commitments.
3. It is an embodiment of a conception of reasoning, a set of inferences that it supports and a set of inferences it recommends.
4. It is a medium for efficient computation.
5. It is a language that allows humans to exchange knowledge.

Unlike the notion of knowledge presented by the knowledge engineers, this notion of a knowledge representation does observe that the act of representing knowledge results in difficult and unavoidable omissions. For example, the representation may encode certain social and cultural biases as it is used as a surrogate. In the above roles, there is a danger that the choice of the embodied “conception of reasoning” overrules other types of reasoning.

5 Conclusion

My conception of knowledge follows that presented in [DAVIS] and outlined above. However, while Forsythe expresses little hope for the ability of AI scientists to build knowledge-based systems, I am less pessimistic. I do not believe that the methods employed by expert systems in any way support a realistic model of knowledge or reasoning. One of the largest problems is that expert systems attempt to compress the innumerable representations that ex-

perts use to reason about their domain into a single knowledge representation: formal rules. I believe that a truly “expert” system will require an entire model of conception, certainly requiring natural language and common-sense knowledge. I agree with Forsythe that the act of extracting knowledge from an expert results in omission and erroneous translation of that knowledge, and I believe that any system based on the extraction and encoding of knowledge is doomed to failure. For the time being, expert systems will be relegated to a advisory role.

References

[DAVIS] Davis, R., et al. “What Is a Knowledge Representation,” *AAAI Journal*. Spring 1993.

[DUDA] Duda, Richard O., and Edward H. Shortliffe. “Expert Systems Research.” *Science* 220, no. 4594 (April 15, 1983): 261-268.

[FOR1] Forsythe, Diana E. “Engineering Knowledge: The Construction of Knowledge in Artificial Intelligence.” *Social Studies of Science* 23 (1993): 445-477.

[FOR2] Forsythe, Diana E. “Blaming the User in Medical Informatics: The Cultural Nature of Scientific Practice.” In *Studying Those Who Study Us: An Anthropologist in the World of Artificial Intelligence* (Stanford, CA: Stanford University Press, 2001), pp. 1-15.