Available online at www.sciencedirect.com

SciVerse ScienceDirect

Int. J. Human-Computer Studies 71 (2013) 195–199

International Journal of Human-Computer Studies

www.elsevier.com/locate/ijhcs

The knowledge acquisition workshops: A remarkable convergence of ideas

Mark A. Musen

Stanford Center for Biomedical Informatics Research, Stanford University, Stanford, CA 94305-5479, USA

Accepted 25 September 2012
Available online 16 October 2012

Abstract

Intense interest in knowledge-acquisition research began 25 years ago, stimulated by the excitement about knowledge-based systems that emerged in the 1970s followed by the realities of the "AI Winter" that arrived in the 1980s. The knowledge-acquisition workshops that responded to this interest led to the formation of a vibrant research community that has achieved remarkable consensus on a number of issues. These viewpoints include (1) the rejection of the notion of knowledge as a commodity to be transferred from one locus to another, (2) an acceptance of the situated nature of human expertise, (3) emphasis on knowledge acquisition as the modeling of problem solving, and (4) the pursuit of reusable patterns in problem solving and in domain descriptions that can facilitate both modeling and system implementation. The Semantic Web community will benefit greatly by incorporating these perspectives in its work.

Keywords: Knowledge acquisition; Knowledge-based systems; Semantic web; Workshops and conferences

1. An academic agenda responding to practical problems

Brian Gaines and John Boose convened the first Knowledge Acquisition Workshop in Banff in 1986 at the height of the era of expert-systems hype. When the call for participation was released, the International Joint Conference on Artificial Intelligence (IJCAI '85) had recently concluded in Los Angeles—the largest single event that the city had seen since the 1984 Olympic games. At IJCAI '85, companies with names such as *Intellicorp* and *Teknowledge* were outdoing one another to attract the attention of eager customers, in one case bussing attendees in droves from the conference venue on the UCLA campus to an exclusive ice-cream store in Beverley Hills so that they could enjoy lavish sundaes that somehow might entice them to purchase expensive "shells" for building knowledge-based systems.

The expert-systems companies gave the illusion that engineers, with the right set of skills, could take the knowledge of professionals and bottle it up for reuse. The operative question was only one of *which* expert-system shell offered the greatest number of capabilities. Knowledge acquisition was known to be a hard problem, but it was seen as difficult merely because application specialists and system builders did not speak the same language at first.

Gaines and Boose had the radical idea that knowledge acquisition was something that actually could be studied. Conferences such as IJCAI that summer and books that were popular at the time (Hayes-Roth et al., 1983) largely took it for granted that, metaphorically, when enough pressure was applied, knowledge inevitably would flow from the heads of experts into the electronic knowledge bases of intelligent computer systems. Suddenly, here was a call for participation for a workshop to be held in the Canadian Rocky Mountains that would put that assumption under a microscope.

The first Knowledge Acquisition Workshop was held in November of 1986. Investigators from computer science, psychology, anthropology, and linguistics—most of whom had never met one another before—convened for an entire week. This widely interdisciplinary group of participants found it remarkable that they had so much in common. They also found it significant that they had never gotten together before. The Knowledge Acquisition Workshop suddenly made the study of knowledge engineering a legitimate scholarly discipline. It initiated the growth of a coherent scientific literature. The workshop

E-mail address: musen@Stanford.EDU

1071-5819/$ - see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.ijhcs.2012.10.011

also catalyzed the creation of an international community that remains cohesive to this day.

As the first Knowledge Acquisition Workshop (KAW) was taking place in Banff, the early frosts of the "AI Winter" were already becoming apparent. Attendance at meetings such as ECAI '86 and AAAI '86 were down from that of IJCAI '85, and the trade shows associated with those conferences were noticeably smaller. (In future meetings, the extravagant trade shows would disappear entirely.) It was becoming clear that the excitement that the commercial sector had displayed for expert systems in the early part of the 1980s was beginning to wane. Simply put, building and—more important—*maintaining* electronic knowledge bases was not a simple task, and the corporate world was starting to question the return on investment for its knowledge-engineering activities. Although the emergent disenchantment with AI technology in commercial settings would later cause many academic researchers to move away from symbolic AI, the real-world problems associated with expert systems would only redouble the determination of the KAW community to understand the nature of elicited knowledge and of methods for imbuing computer systems with the capacity for intelligent behavior.[1]

Although the KAW community was quite heterogeneous, its scholarly diversity was actually an asset that helped the workshop participants to synthesize a very broad set of research contributions and to develop important unifying ideas that facilitated both the theoretical study of knowledge elicitation and the pragmatic development of knowledge-engineering methods and tools. Although it is perhaps dangerous for one person to assume that there was consensus within a broad community, it did appear that the researchers who attended KAW converged on a common mode of thought and a set of shared assumptions. These positions included (1) the rejection of the prevailing notion of knowledge as a commodity to be transferred from one locus to another, (2) an acceptance of the situated nature of human expertise, (3) a renewed emphasis on knowledge acquisition as the modeling of problem solving, and (4) the pursuit of reusable patterns in problem solving and in domain descriptions that could facilitate both modeling and system implementation. In retrospect, it is remarkable that the KAW community converged around these ideas as quickly and as consistently as it did. These core positions represent enduring products of the interactions that took place at the knowledge-acquisition workshops.

1.1. The end of the "transfer" metaphor

When the expert-system paradigm emerged in the 1970s, knowledge engineering was viewed as a problem in "mining." In the parlance of the day, application experts were viewed to have "nuggets" of knowledge in their heads and the goal of knowledge engineering was to mine those nuggets and to cast them in the form of a computer-interpretable knowledge base. Rule-based architectures caught on because developers could imagine nearly a one-to-one correspondence between the nuggets of knowledge in the heads of the experts and the rules in the rule base. Just as professional knowledge was construed as nearly self-contained heuristics waiting to be applied at the right time, production rules were viewed as modular chunks of knowledge that could be applied opportunistically and independently at run time (Buchanan and Shortliffe, 1984). Knowledge engineers were the skilled professionals who knew how to "transfer expertise" bit by bit from the experts to the emerging rule bases.

While knowledge engineers were busy trying to mine their nuggets of knowledge, psychologists were learning that there really are no such nuggets (Nisbett and Wilson, 1977). Humans are very good at offering post-hoc rationalizations of their behavior, but their ability to introspect and to provide truthful descriptions of their problem-solving knowledge is extremely limited. Thus, it is impossible simply to transfer knowledge from informants to a computer, as the informants can never really know what they know (Johnson, 1983). Bill Clancey (1997) said it best in a keynote talk at the second KAW workshop: Knowledge acquisition is not hard because it is difficult to get domain experts and knowledge engineers to speak the same language; knowledge acquisition is hard because the domain experts and the knowledge engineers together must create a language to define a model of professional expertise—a model that never existed previously in any formal sense.

With the AI Winter setting in, and much of the corporate world showing increasing intolerance for the complexities of building useful intelligent systems, this shift in perspective was extraordinarily useful. The KAW community recognized that enhancing the communication channel between subject-matter experts and system builders by itself would not ease the development of intelligent systems. What was needed were frameworks to assist everyone concerned with the construction of electronic knowledge bases with the specification of new models of expertise and the means to transform those models into useful implementations.

1.2. The rise of situated cognition

While the KAW community was wrestling with its basic premise concerning the "transfer of expertise," the cognitive-science community was dealing with its own internal battles. There was growing insistence that much (perhaps all) of human decision making was critically dependent on the context in which it took place, suggesting that human problem solving could never be divorced from myriad environmental variables at every turn. Cognition, it was argued, could not be understood in terms of disembodied processes; cognition and action are necessarily situated (Clancey, 1997).

Although the KAW community generally watched this debate from the sidelines, the implication for the construction

[1]The KAW community includes not only the regular attendees of the Banff workshops, but also the participants of recurring meetings in Europe (EKAW) and the Pacific Rim (PKAW).