

Technology transfer and the evolution of knowledge engineering

Professor F of Pittsburgh asserts that today's sophisticated environments which represent the current state of the art are of limited value to those developing serious AI systems. These tools, he says plausibly, suffer the inevitable fate of the generalist competing with the specialist. For any single application only ten per cent of what they provide is needed—the rest is unacceptable overhead. Consequently serious designers will simply write for themselves the bits they need for specific applications.

Professor F of Palo Alto responds that his colleague should examine what is happening in the real, commercial world; to his certain knowledge, company after company has successfully used the present generation of tools to develop advanced and successful applications. In many cases these are already making a serious contribution to the competitive position of their sponsors. As yet we hear little of these merely because of the companies' secrecy policies, but many of these successes will be described in detail in his forthcoming book.

The fact is that no one has any clear idea of the overall balance of success and failure in transferring expert systems, knowledge engineering tools, applied AI etc., from the academic laboratory to the commercial marketplace. Some people have tales of triumphs, while others have their war stories. Many (whether because their companies want their successes to be kept secret, or their failures to be kept secret) don't have any stories at all.

The reason for planning this special issue was that it was obvious that whatever is going on it is complicated, untidy, and most of it goes unreported. Different organizations have tried to transfer string-and-sealing-wax prototypes from academia to unforgiving commerce in different ways; commercial products have evolved into many different forms to exploit distinct potential markets, and the pattern of these developments has been different in different fields and countries.

Striking examples of rapid evolution of applied AI and knowledge engineering can be found in tools and applications. Among the most obvious, and the most predictable, of these trends has been the willingness to move away from a purist AI position about the design of practical systems to a pragmatic compromise with established computing technology. This can be seen in the way that the shell suppliers have moved away from LISP and Prolog in favour of C for product implementation. Shell suppliers and language suppliers have wisely embarked on a continuing programme of development to provide hooks and facilities of wholly conventional kinds; recognizing that knowledge based systems have to be able to talk to a wide range of devices, conventional applications, databases and so on. (In passing, it might be remarked however, that wise managers are continuing to use and develop these languages in their product research and prototyping, otherwise their research costs would rise and they would risk losing touch with the developments in AI which continue at a rapid rate.)

Perhaps a little less predictable has been an apparent divergence of developments in Europe and North America, in the scale of applications which have been attempted. This may have something to do with the higher European commitment to Prolog rather than LISP. No doubt this was partly due to Europeans being loyal to homegrown technology, but it was also the fact that useful Prologs appeared on PCs before useful LISPs did. As I observed in a state-of-the-art report written in 1983 "It is interesting to note that the leading American company, Teknowledge, has built its large expert system package S.1. in LISP on a LISP machine, its smaller system for the IBM personal computer is implemented in Prolog. The typical price of an expert system package in the US is in the thousands, even tens of thousands, of dollars; although they are in many ways less sophisticated the typical price of commercial packages in Europe is a few hundred dollars!" Since then small systems have proliferated in North America, and the sophistication (and price) of the leading European packages

has risen substantially, but the pattern has been established. Ed Feigenbaum has caricatured the difference as the European (and particularly the UK) landscape dominated by large numbers of “rabbits” (small expert systems) in contrast to the buffalo which roam the North American plains (and apparently a growing number of hunting parties in pursuit of elephants). Which of these forms of exploitation will prove to offer the sounder business remains to be seen.

Two aspects of developments to date must be regarded with some disappointment. One was predicted by many, the other wasn't (and still isn't) even widely understood.

The first issue concerns the degree to which academia, let alone industry, is engaged in developing some form of artificial intelligence. It is now generally apparent, and was apparent to those experienced in AI from the beginning, that none of the products on offer in the near future will have the intelligence of rabbits, let alone that of buffalos or elephants. Indeed most of the traditional subfields of AI have yet to offer a commercializable product, which is a rather more modest goal than the creation of intelligence. However the prospects continue to be bright in some areas.

Many of those arguing for expert and knowledge based systems a few years ago tried to avoid being identified with AI, at least as a short term commercial venture. They had more modest though still, as it turned out, extremely difficult goals. As it was put at the time “. . . the technology of expert systems is as much about knowledge as it is about computing”. This was thought to mean that—among the great potential benefits of knowledge based systems—they cause us to reflect upon, articulate and write down, what we know but have traditionally left tacit, unanalysed—and unquestioned. The great topics like medicine or the law have seen little progress (in the sense of revolutionary epistemologies of medical theory or jurisprudence; see the comments of Susskind and of Leith in this issue) but the proliferation of the rabbit population may be an optimistic sign for the future.

The collection of articles in this special issue on technology transfer and the evolution of knowledge engineering reflects a heterogeneity of development, and progress at different rates in different fields. At a time when many people in many countries are taking stock of progress since the excitement of the fifth generation announcements, and prospects for the future, this collection of observations, reflections, and assessments may make interesting reading. However, we see this issue as only an initial foray into the general topics of technology transfer and the evolution of knowledge engineering. The *Review* will continue to monitor the development of the field and we invite further surveys, analyses and even polemics for discussion in its pages.

References

Expert systems. State of the Art Report 12:7, Maidenhead, Berkshire: Pergamon Infotech, 1984.