

Editorial

The evolution of knowledge systems: ideological confusion or healthy pragmatism?

A *pure* knowledge based system is one in which the semantics of all symbolic elements are explicit and manipulable. The growth of knowledge engineering has produced a trend toward *hybrid* designs, combining explicit representational techniques with traditional algorithmic ones. This trend might be viewed as arbitrary and undesirable (merely reflecting the recent influx of traditionally trained professionals into the AI community) or a healthy development involving pragmatic selection of appropriate, mature techniques uncluttered by ideology.

An example of the trend towards hybrid designs is in the treatment of uncertainty. Techniques of uncertainty management are often viewed as central to the development of expert systems technology, and have attracted great and heated debate over the last few years. The most obvious argument in this debate has been about the choice of quantitative methods for calculating and propagating uncertainty. A second argument has been concerned with whether precise calculation of uncertainty is of importance for practical knowledge systems and whether uncertainty is, in fact, really a quantitative concept or a logical one.

The first argument has tended to eclipse the second. This has been unfortunate because it has distracted attention from important foundational and representational questions in favour of technical issues which may be of little practical importance. The minor debate also seems mistaken in assuming that uncertainty must be viewed as *either* logical *or* quantitative.

In this context Safiotti's balanced review of the field, in which he tries to do justice to the technicalities of many different approaches, while keeping a clear head about the different demands to be satisfied, must mark an important stage in the debate about uncertainty. It is too much to hope for an end to quarrels (there have been quarrels over different intuitions about uncertainty for at least 300 years). However it must be more widely understood that epistemological problems are at least as important as technical problems in AI, and that they must be addressed simultaneously. Safiotti's paper can be regarded as an admirable example of a paper that got the balance about right.

Another way in which knowledge based systems may be evolving in undesirable ways is in the design of user interfaces. In the early, heady days of the expert systems fashion, one of the things that excited newcomers was the idea that software (that autocratic stuff) might actually become intelligible to ordinary folk, and accountable to and controllable by its users. Development of the user interface in general, and the explanations offered by expert systems in particular, looked like liberation.

In the ten years since then progress on the user interface in general, and the "explanations" offered by expert systems in particular, seems conspicuous by its absence. Designers now seem to be returning to their old habit of ignoring (though thankfully no longer denying) the problem of user interface design, or at best assuming that so long as the mouse, menus and graphics are modern and cute there is little more to do. No one can read Stenton's paper and still imagine that the problems of designing flexible, understandable, cooperative systems are so shallow. The paper doesn't attempt to offer final answers, but it does help to understand what the problems are and some of the author's interesting and entertaining views on directions that we might take.