of some value to researchers, especially as background, but it lacks the polish or single-mindedness to make it a seminal work. Overall, it probably meets the worthwhile objectives of the Council for Science and Society.

Reference

International Electrotechnical Commission, 1989. "Draft—Software for computers in the application of industrial safety-related systems", *Geneva: IEC Report 65A(Secretariat)94-1.*

Reviewed by Richard Thomas, AI Group, School of Computer Studies, University of Leeds, UK.

Benefits and risks of knowledge-based systems Report by a Working Party, Council for Science and Society, OUP Oxford 1989, pp 76, £6.95.

The Council for Science and Society (CSS) is a charitable institution dedicated to the study of the social impact of science and technology. In particular, the CSS promotes the early discussion of topics of future significance not yet the subject of serious public debate. The stated aim of this report is to increase general awareness of knowledge-based systems (KBSs) and to alert the public at large to the potential benefits and possible dangers arising from the widespread introduction of this technology.

The report opens with a brief history of artificial intelligence and the development of knowledge-based system technology. Unfortunately, considering the lay audience at which this report is aimed, the explanations are couched in rather technical terms which may hinder understanding. A wide range of applications is discussed with illustrations drawn from the fields of finance, law, education, the military and health.

The authors don different "hats" to present the benefits and risks of KBSs. In the guise of "the believer", the benefits of the knowledge-based filtration and delivery of information are expounded, while the "social critic" warns of the hazards of uncritical acceptance of a KBSs advice, misplaced belief in a system's wisdom, the risks to civil liberties, the ethics of cultural imperialism and, most ominously, the terrifying dangers threatened by the military use of autonomous decision systems. Even the "disbeliever" has his say, but gets short shrift, as the authors clearly believe in the technical capability of KBSs.

The report observes that many of the risks arising from the use of KBSs are shared with computer systems in general: the threat to privacy and civil liberties, for example. The dangers of autonomous decision making systems (systems where an action of some consequence is taken by the system directly, without a human authority in the decision loop), particularly those in the military field, are heavily underlined, and rightly so, for not only are the consequences of error most serious, but military secrecy shrouds their development from public scrutiny. But here again, the risks are shared with "traditional" computing. As long ago as 1980, Professor C. A. R. Hoare implored American military programmers not to use Ada for safety-critical software as he believed it to be too complex a language ever to be properly verified. This report levels the same criticism at military KBSs.

Finally, the report makes thirteen recommendations of measures to enhance the benefits and minimise the risks of such systems, ranging from education and the initiation of socially desirable applications to the adoption of industry standards and codes of practice.

Does this report achieve its objectives? The answer must be a qualified yes. Yes, because the benefits to society of appropriate systems, together with the dangers inherent in the inappropriate application of the technology, are clearly spelled out. Qualified, because discussion is largely limited to what could be termed "public" systems: systems in domains such as health, education and law with which all members of society will come into contact. There is no discussion of the myriad "private" systems, each affecting only a small group of people, which can be expected to outnumber the "public" systems by orders of magnitude. Qualified, too, because there is no discussion of the risks of failure by business to adopt the technology in the face of its adoption by our competitors.

The social impact of loss of industrial competitiveness has been huge, as continuing unemployment testifies, and failure to capitalise on the commercial advantages offered by KBSs will only compound Britain's industrial problems. It is important that Britain's industrial managers are made aware of the significance to their businesses of KBSs. The report does not address this issue.

There may be more subtle social consequences of KBS implementation in industry. Will KBSs bring about a change in the power hierarchy of the workplace, as workers are provided with decision support systems to enable them to make decisions once referred to a manager? Will management find many of their responsibilities for planning and budgeting taking over by planning systems? Will the corporate mainframe pass its processed sales figures to an expert system for analysis, rather than the present legions of accountants? How will industry cope with these changes?

Excepting the lack of material of direct industrial relevance, this is a valuable report which will interest anyone who wishes to be acquainted with the forces which will shape their future. It should be compulsory reading for all those who have a responsibility for controlling those forces, especially those in politics and the military, whose ignorance of the consequences of their actions could be catastrophic.

Reviewed by J. A. G. Hawkins, Johnson Matthew Technology Centre, Blount's Court, Sonning Common, Reading RG4 9NH

Benefits and risks of knowledge-based systems Report by a Working Party, UK Council for Science and Society, OUP, Oxford 1989, pp 76, £6.95.

The aim of this report, sponsored by the Council for Science and Society, is to alert the public to the consequences of working with knowledge-based systems. The council was right to select this as a topic for study, since much of the work in this field is likely to affect the way in which society will react to advances in science and technology.

Because of the high incidence of errors in computer systems, we now all know that both hardware and software can give incorrect or misleading answers. It is obvious that people will be concerned when they realise they have been under the influence of mechanisms they do not understand and can not control.

When a *knowledge-based system* goes wrong, the effects are likely to be much more serious than for conventional applications. One reason is that knowledge-based systems are deliberately designed to appear natural and human—this makes their output more plausible than conventional programs, and hence more disturbing when they are in error. Another reason has to do with our still imperfect ideas of what is correct (and incorrect) about such programs.

Anyone who has written a simple program in Basic or Pascal knows that computer arithmetic can result in errors, that mistakes in addressing can result in the wrong variable of an array being accessed, and so on. It is worth considering what kinds of error can appear in a knowledge-based system. They include the same kinds of error that appear in conventional programs, but in addition there are the more subtle and pervasive errors that come from the contributions of an "expert". These may be errors of judgement, and perhaps encoding, and are much more subtle than the errors we have become used to in computer systems. Some systems, when they give results, allow the user to see the chain of reasoning used in the system. Few users will be able to evaluate this kind of output, and it will serve merely to re-enforce the apparent correctness of the system, rather than its fallibility.

The report starts with a short and necessarily sketchy account of the working of knowledge-based systems. This is followed by a survey of five application areas: finance, law, education, military and health. This is a potentially valuable part of the book, since it covers the selection of applications, one of the most critical parts of building a system. The accounts of what is being attempted, and in particular how these applications can fail, will be alarming to those unaware of directions in this field.

AI people have always seemed to try and separate themselves from the rest of the computing community. This can be seen from the very terms they use: Artificial Intelligence, Expert Systems, Knowledge-based, all seem to be making a point. Who could be *against* intelligence, experts,