money had been well spent, how could Congress fund post-war research into nuclear energy?

As we know, the products of Oak Ridge, Hanford and Los Alamos were dropped on Japan in August 1945. Rhodes not only tells of the missions, but he presents about 18 numbing pages, mostly of eye-witness accounts, of the destruction of Hiroshima and Nagasaki. Except for an epilogue that sketches the development of the hydrogen bomb, the story ends here. The book is less scholarly than the magisterial *The New World, 1939/1946*, by Richard Hewlett and Oscar Anderson, yet more accurate than Robert Jungk's well-written but flawed *Brighter Than a Thousand Suns*. Altogether Rhodes has

produced the most readable, exciting and just book to date that covers both the bomb and the preceding four decades.

The book is not polemical, but Rhodes does gently advance a message. The nation-state of modern times, he says, has become a "death organization". Science, the other major force to arise in the past few centuries, is the only structure strong enough to challenge the nation-state, for the bomb has made it vulnerable. Our visionaries, Bohr and Szilard, appear to have been more practical than the so-called hard-headed politicians. □

Lawrence Badash is Professor of History of Science in the Department of History, University of California, Santa Barbara, California 93106, USA.

Clever computers

Bernard Meltzer

Artificial Intelligence and Natural Man, 2nd Edn. By Margaret A. Boden. *Basic Books, New York/MIT Press, London: 1987. Pp. 576. Hbk £25.95; pbk \$14.95, £10.25.*

THIS is the second edition of a book which, first published ten years ago, was one of the best non-technical introductions to artificial intelligence. It deals also with psychological, philosophical and social issues related to the subject, and differs from the original only in having an added postscript chapter dealing with some of the developments of the intervening period such as logic programming, basic robotics, naive physics, non-monotonic logic, creativity and parallel distributed processing (PDP).

A putative third edition — in ten years' time, say — would probably have to be much more radically modified because, in Kuhn's terminology, a paradigm shift in the subject seems to have started already, with the promise of fast-gathering momentum. It is too early to give a precise definition of this shift, but a good starting point for indicating its character is to consider the title of a section of the author's first chapter, namely, "Computers Manipulate Symbols". Here Professor Boden's intention was to demolish a common illusion of the time, that computers can process only numbers, but the assertion in the title as such is possibly misleading and bears some examination.

A Martian, examining — let us say — the behaviour and inner workings of a robot actuated and controlled by an internal "symbol-manipulating" computer, might very well fail to discover any symbols whatsoever in the system, but only certain patterns of movement, sounds, connections, electric currents and so on, just as we would do in the case of a dog or a

lion. In other words, the symbols are not inherent in the computer, they arise only from our interpretation of certain patterns of physical processes. But even to say this is rather misleading. For it is not the case simply that we observe such patterns and then interpret them as symbols, but that the particular form the physical processes in the computer take is determined by design (of the hardware and software of the system), and such design might be done (and in the past has mainly been so done) using the high-level notion of symbols and symbol structures.

Today, however, impressive computer models of mental processes have been designed which make much less use of the concept of symbols. For instance, at a workshop in Geneva last year Sejnowski described a model implemented on a PDP computer, which from the reading of children's stories learned rules of English pronunciation (as well as their exceptions!). Models of this kind work mainly by interactive activations and inhibitions of a very large number of connected, simple units; so far from symbols being manipulated in the system, it appears to have been a positively mind-cracking job for Sejnowski, from observation of the activity of the model, to come up with any sort of detailed account of the processes going on in high-level terms.

The term "subsymbolic computation" has been coined for this kind of processing, though it is clear from the discussion above that in a sense all digital electronic computation is "subsymbolic". The distinction has to do only with the character of the design of the system: for instance, a learning program which resulted in the addition of a new item to the knowledge base in a serial computer would be predominantly of type "symbolic", while one which varied a set of connection strengths in a PDP computer would be predominantly of type "subsymbolic". It is probable that the functional design of our own brain and nervous systems is subsymbolic in this sense.

It might appear that the new paradigm is entirely the result of the advent of PDP computers. Certainly, the effort to exploit these massively parallel systems with their promise of very much greater speed, robustness and flexibility for cognitive modelling, compared with the von Neumann (serial) computer, has been determinative in this development.

But there have been other, less palpable, factors involved. With its emphasis on symbolic representation, the von Neumann computer encouraged certain restrictive attitudes in the development of cognitive modelling. For example, from its beginnings at the end of the Second World War, artificial intelligence was dominated by the issue of the representation of knowledge, which — when one comes to think of it — is rather peculiar: an encyclopaedia is a good (symbolic) representation of knowledge, but hardly an intelligent system. The central issue surely is the process of thinking. Knowledge on the one hand is the result of such processes, and on the other hand is able to supplement sensory and motor experience in providing the raw material for them.

As for the symbolic representations themselves, they were in the main those suggested by natural, mathematical and logical languages: these are dominantly what Sloman termed "Fregean", in that predicates and functions are represented by the application of function symbols to argument symbols. However a map is also a symbolic representation, but one in which relations are represented by other relations, and it is only comparatively recently that the power of such non-linguistic "analogical" representations has been demonstrated on serial computers by cognitive models of naive mechanics and physics. PDP-type systems would implement such models much more effectively.

Finally, the "actor" model of computation proposed many years ago, in which computation is viewed not as the action of a central processor, but as the exchange of messages between data structures, was an anticipation of the new paradigm.

We see then that artificial intelligence, the computational modelling of cognitive processes, is still very much an exploratory and experimental discipline. This is the reason why it remains both exciting and, on occasion, frustrating — one can seldom can fall back on well-established principles and laws to guide research, as in some of the older sciences. However, Boden's book is, for the layman, still one of the most readable and certainly the most scholarly initiations to the subject. □

Bernard Meltzer, Via G. Galilei 5, Ispra (Va), Italy, was formerly Head of the Department of Artificial Intelligence at the University of Edinburgh, UK, and is now a consultant to the European Commission on basic research in artificial intelligence and signal processing.