BOOK REVIEWS

and distribute many mutant mice; but without such a facility, the mutants run the risk of being irremediably lost. The recent decision to create the European Mutant Mouse Archive is an example of the application of the rule proposed above.

Existing cooperative programmes could be improved by keeping past experience in mind. The administration in Brussels must not be left alone with such a task, not least because the language used in its documents is vastly different from that of contemporary science. (I am of the opinion that the responsibility for the administration's isolation lies mostly on the shoulders of European scientists.)

The recently instituted European Sciences and Technologies Assembly (ESTA) could become the tool for improving the European Union's research and technological development policies while strengthening coordination of the research policies and activities of European countries. The role played by ESTA, *mutatis mutandis*, should be analogous to that of the US National Research Council. ESTA brings together scientists who are supposed to guarantee contacts with the frontiers of research and development, the directors of European research organizations such as EMBL, ESO and so on, the presidents of national research organizations and the heads of the research departments of large industrial concerns. It is not at all rash to conclude that the possibility of realizing a Unitary European Research and Technology Policy depends entirely on how ESTA handles the tasks assigned to it. As ESTA was created during Ruberti's term as commissioner, how it actually fares in practice will determine the outcome of Ruberti's and André's reflections. □

Glauco P. Tocchini-Valentini is in the CNR Institute of Cell Biology, Viale Marx 43, 00137 Rome, Italy.

Fleshing out intelligent machines

Igor Aleksander

The Handbook of Brain Theory and Neural Networks. Edited by Michael A. Arbib. *MIT Press: 1995. Pp. 1,118. $150, £95.*

IT must say something about the science of neural networks that it has gone from being a fringe activity for the misguided few to meriting Michael Arbib's encyclopaedic treatment in the space of 12 or so years. At more than three kilograms, this handbook is twice the weight of an adult brain and far beyond the size of the average book one might read on the train. It claims to tackle the questions "How does the brain work?" and "How can we build intelligent machines?" through its 266 articles written by individuals and a few guiding introductions written by the editor and his advisers. It should immediately be said that the articles on brain modelling are of a 'computational' kind: that is, they concentrate on theoretical models rather than reporting on neurobiological findings.

Can one really learn about brains by doing mathematics and simulating artificial neural networks? In a very broad sense the answer must be yes. A neural network is a system containing elements that adjust their function according to some external requirement. The brain too is a network of variable-function cells and in this sense has the artificial neural network as its distant cousin. Studying the artificial network is more tractable than studying the brain, and the former leads to articles in the handbook on epileptic nets, nets that sort visual background from foreground and nets that act like the hippocampus. The credibility gap comes from the fact that many artefacts such as error back-propagation are unlikely to exist in the brain. Also, many theoretical models are used because of their mathematical elegance rather than their biological relevance. Personally, I would be happier if the neural network community were more realistic about the true biological relevance of their work and stopped believing that if an artificial network behaves in a brain-like way, the exact underlying mechanism must somewhere to be found in the flesh. The true biological significance of neural modelling lies in the discovery of broad principles and not detailed function. My main criticism of the handbook is that it makes little mention of what is known of such principles — the role of network topology, the structure-function problem, the richness and plasticity of mental states, the constraints of evolutionary development, the difference between learning, adaptation and instruction, the role of modularization, and so on.

So what of intelligent machinery? I happen to believe that this is really where neural networks are now making an impact. Sadly, it too is not a strong point of this handbook. Although applications are mentioned in a sensible article by Françoise Fogelman-Soulie, the handbook gives the impression that the biological effort is paramount and that the engineering of neural systems is a lesser pursuit. Many important engineering contributions such as the work on sparse memory by the Finnish engineer Pentti Kanerva and the UK-based work on logical nets receive no mention at all.

The overall coverage can be best understood by looking at the main subheadings: connectionism and artificial intelligence, dynamics, the mathematics of learning, applications, models of biological neurons, sensory systems, biological learning and motor control. As with many handbooks of this kind, the contributions are useful but should be treated with care. The detailed articles, while providing a reasonably extensive coverage, are written by experts describing their own research. In many, an attempt has been made to provide a good perspective and there are some good, informative reviews such as Paul Werbos's treatment of back-propagation (training of multi-layer networks) and Jeffrey Elman's article on language processing. On the other hand, Eduardo Sontag's paragraphs on automata and neural networks hardly touch on the considerable literature now available on the topic. There are sins on the neurobiological side too. Material on models of the hippocampus makes no reference to important work at the University of Oxford in the United Kingdom and at Rutgers University in the United States. Driven by my own interest I looked up the article on "Theories of Consciousness" by Max Velmans — interesting but hardly comprehensive. There is virtually no reference to the intricacies of the current debate on this topic or to the vast philosophical history of the subject. This is merely an example of the real dangers of trying to gain entry into a complex element of the widespread field of neural networks by reading an article of three or so pages.

How permanent is the content of the handbook? Regrettably, the field is not yet sufficiently mature to provide stability, for reasons already mentioned. But, despite this negative impression, for a year or two this book will be a handy addition to the library shelves of any institution in which research is being done on neural networks. The number of contributions and the expertise of the contributors are reasonable. It is the patchiness, the lack of treatment of general principles and the lack of reference to many applied machine-building techniques that make the project less effective than it might have been. Many publishers have embarked on similar projects and other tomes in the pipeline may plug some of the holes. So, librarians, strengthen your shelves in the vicinity of the "neu-" areas. □

Igor Aleksander is in the Department of Electrical Engineering, Imperial College of Science, Technology and Medicine, Exhibition Road, London SW7 2BT, UK.