Mental simulation

Geoffrey Hinton

An Introduction to Neural Computing. By Igor Aleksander and Helen Morton. *Chapman and Hall: 1990. Pp. 240. Pbk £15.95*.*

RESEARCHERS in the fields of physics, computer science, electrical engineering, statistics, neuroscience, psychology and others are currently investigating the computational abilities of networks of nonlinear neurons. Their scientific aim is to understand how the brain computes, and their technological goal, which is the main concern of this book, is to produce useful adaptive devices.

The central problem in the field of 'neural networks' is to find a learning algorithm that will adjust the connection strengths in the network so that the neural activities represent important underlying properties of the network's input. By associating its responses with these underlying properties, the network can behave in much more complex ways than a system that simply associates responses with the raw input. Much of the recent interest has been stimulated by the discovery of new learning procedures such as Kohonen's method of forming topographic maps or the back-propagation procedure described by Rumelhart *et al.* (*Nature* **323**, 533; 1986). These procedures are relatively painless to program on scientific workstations and work well for training small networks to solve small tasks, so curious researchers can get hooked in a few days.

For those entering the field of neural networks, there is an acute need for an authoritative textbook that explains the current models, points out their weaknesses, and relates them to their roots in established disciplines. Aleksander and Morton fail in this task because they assume a complete ignorance of calculus, linear algebra and statistics. They do a good job of explaining the basic models by using detailed examples, so their book will be valuable to those who want to get the flavour of the ideas without understanding the relevant mathematics, but they fail to provide a solid foundation that will transcend the specific models they describe. This is unfortunate, as any introductory book in this fast-moving field is inevitably out of date by the time it appears.

Aleksander and Morton cover most of the important models up until 1987 and their descriptions are fairly accurate. They do, however, make a few careless blunders, such as confusing the number of times a neuron fires with the number of times it is given an opportunity to decide whether or not to fire (page 95). They also choose to describe the unsupervised version of the Boltzmann machine learning procedure instead of the more easily understood supervised version, and they seem unaware that the 'mean field' approximation used by physicists leads to an equivalent but much faster learning procedure that can be applied to analogue deterministic networks of the kind that Hopfield and Tank introduced for optimization tasks. The authors devote only 15 pages to applications of neural networks in speech, language and vision, so they inevitably omit some of the most significant practical successes such as the use of multilayer time-delay networks for phoneme recognition or the use of constrained multilayer networks for hand-written character recognition. Both these successes stem from ideas about how neural networks can achieve perceptual invariance so that their recognition is not affected by changes in the onset time of a phoneme or the position of a character.

Aleksander and Morton do not attempt to relate neural-network models to important ideas in other fields. For example, competitive learning can be viewed as a crude and inefficient approximation to a maximum likelihood fit of a mixture of gaussians, and back propagation can be viewed as a multilayer extension of logistic regression, or as a nonlinear generalization of methods used in optimal control. Hopfield networks, particularly the analogue version, are closely related to iterative techniques used within computer vision for processes such as stereo-fusion and surface interpolation.

A significant limitation of the book is that the authors do not seriously discuss the issue of generalization. There is no mention of theoretical work by Baum and Haussler (published in *Neural Computation* in 1989) which establishes that if a small network can be trained to give the correct response for a large number of examples, it will almost certainly generalize correctly to further examples. Nor do the authors mention that generalization can be improved by shrinking the weights, eliminating unnecessary connections or using a validation set to decide when to stop learning on the training set.

An even more significant omission is a discussion of the speed of learning. This is now the crucial issue in the practical application of neural networks because current algorithms such as back propagation do not scale up well to large networks. Greedy algorithms that add hidden units one at a time can greatly improve learning speed, but for really large networks it will probably be necessary to develop more modular systems in which individual modules can learn independently because they have local, internally generated goals that allow them to build appropriate representations of the underlying causes of the sensory input without requiring an external teacher. ▶

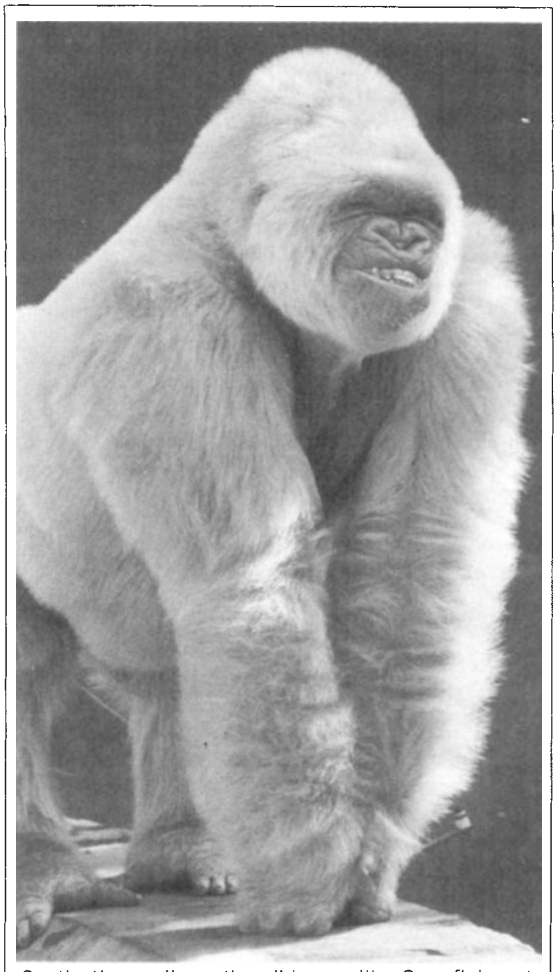

Captivating smile — the albino gorilla, Snowflake, at the Barcelona Zoo attracts many fans. Traditional menageries were once the domain of the rich, where people went to be amused or amazed; modern zoos have been derided as souless prisons. For many species zoos are now the last refuge and their only hope for survival. *Zoo. The Modern Ark* by Jake Page with photos by Franz Maier looks at zoos throughout history and examines their role in protecting endangered wildlife. Published by Facts On File, price is £18.95. □

* A software package, *Cortex*, is also available from Unistat Ltd, PO Box 383, London N6 5UP, UK. This package allows users to construct their own neural networks and is for use in conjunction with the book reviewed here. Price is £100.