is one paper from Belarus and another with joint authors from Britain, France and Ukraine. In one of the papers with joint American authorship the author having an American affiliation also has a Polish one and is R.S. Michalski, a well established authority on machine learning. Two of the papers having Polish authorship have authors from a Polish-Japanese Institute of Information Technology in Warsaw.

The subject matter obviously ranges widely. In their Preface the editors express their satisfaction that, since the previous events, there is an increase in the diversity of practical applications. The first paper in the book has an important medical application since it refers to use of data mining techniques to help diagnose melanoma. The paper that has R.S. Michalski as joint author, under heading (a), is about the modelling of computer user behaviour. Other papers that mention specific applications refer to e-commerce and deductions from financial data, and another is on the extraction of information from audio signals, where the information includes features relevant to speaker and mood identification.

The seven papers under heading (b) show a fine diversity of approach, since the first refers to a computing technique inspired by an ant colony, and the next two are both about neural nets, but treating them in quite different ways, and the next is about an artificial immune system. These are followed by two papers dealing with genetic algorithms. The final paper under the heading is about an approach using quantitative property structure relations (QSPRs) to infer physical properties of chemical compounds.

The first paper under heading (c) is the one whose authors have British affiliations and is on the intriguing topic of the use of AI techniques for the generation of crowd scenes in TV or film productions. Another paper in this section discusses double clustering, which is essentially the two-stage process in which ordinary clustering is performed first and then the centres of the clusters obtained are subjected to a second stage of clustering. This paper is another having a medical application, since the use of the method is illustrated with reference to a study of cardiac arrhythmia related to meteorological data.

The papers under heading (d) range widely in topic and refer to special algebras and logics and developments in knowledge representation, data mining, reinforcement learning and Bayesian belief networks. Those under heading (e) also range widely and two of them make specific reference to the Internet. One of these discusses an expert system applied to intelligent web search that should enhance the operation of search engines applied to market analysis.

The above attempt to review some contents of the symposium is inevitably patchy and biased towards the contributions mentioning immediate practical applications, though the more abstract treatments may be even more significant. The symposium was undoubtedly productive of many new ideas and methods and these proceedings are a valuable account of them.

Alex M. Andrew
95 Finch Road
Earley
Reading RG6 7JX

DOI: 10.1017/S0263574702244859

SPIKING NEURON MODELS: SINGLE NEURONS, POPULATIONS, PLASTICITY, by Wulfram Gerstner and Werner Kistler, Cambridge University Press, Cambridge, 2002, xiv + 479 pp., ISBN 0-521-89079-9 (Pbk, £24.95, also hardback, ISBN 0-521-81384-0 (£65.00).

Developments in the use of artificial neural nets, in the last few decades, have encouraged the view of the neuron as a continuous device with a sigmoid response function. Such a view is required for both the back propagation method of learning and in Hopfield nets. The transmission of information, at least in peripheral nerves, appears to depend on pulse frequency coding and hence a continuous signalling function. This is in contrast to the earlier theory of McCulloch and Pitts, which has been the basis of much speculation in cybernetics, where the all-or-none character of neural excitation was associated with the true-or-false of formal logic. The McCulloch-Pitts model was put forward as a basis for discussion and not as a detailed representation of real neural activity.

Real neurons (except some in the mammalian retina) do indeed "fire" or "spike" and a variety of neural phenomena can only be explained by taking this into account. It is easy to show that neural channels conveying continuous data by frequency modulation do not account for observed behaviour because the reception of an analogue value requires a certain time of integration, and the responses of people and animals are faster than this would allow. A faster response is possible if the integration is instantaneous over many parallel channels with appropriate randomness, and the analysis of this requires consideration of spiking neurons.

Another demonstration that the spiking behaviour is significant is in connection with stereophony, where location of a source of sound depends on extremely accurate estimation of the difference in times of arrival of a sound at the two ears, or phase differences in a sustained input. For barn owls the time difference can certainly be estimated within 5 microseconds.

The treatment is divided into the three parts indicated in the second part of the title. In the first part the structure of a neuron is described, and the nature of its excitability is explained in terms of the well-known Hodgkin-Huxley equations that refer to flows of different ions through the membrane. Since the equations have four variables they are intractable for incorporation in population models, and several more manageable approximations are considered. Synaptic transmission is treated in considerable detail, with accounts of the changes in permeability of postsynaptic membrane. It is acknowledged that neurons may have to be considered as compartmented models, since for example an inhibitory synapse can be at the base of a particular branch of the dendritic tree and it is then more effective in nullifying excitation arising in that branch than in countering excitation generally. Possible sources of noise affecting the output of the neuron are discussed.

In the second part of the book a number of approaches to the modelling and analysis of neuron populations are treated. It is shown that the output of an entirely deterministic net can appear highly irregular, and also that information can be transmitted through an active net much more rapidly than would be expected from the time courses of individual neural responses. This can be attributed to the fact that at any moment there are neurons about to fire and whose firing can be precipitated by the input. Special attention is given to oscillatory behaviour and synchronisation, where one reason for special interest is that the precise synchronisation of impulses has been postulated as part of the means by which sensory stimuli are grouped or "bound" to allow recognition of objects. A theory of complex reverberatory behaviour, first analysed in connection with interactions of fireflies, is shown to be applicable also to neural populations.

The third part of the book, on plasticity and hence possible mechanisms of learning, is specially interesting. The mechanism at the single-cell level is assumed to be essentially as postulated by Hebb but its implications in neural structures and populations are developed in considerable detail. It is shown how "learning to be fast" can occur, so that a response is triggered by the earliest of the events associated with it, an effect that is illustrated by the classical conditioned reflex. A less obvious effect that can also be accounted for is "learning to be precise" where the response comes to be triggered by the event with least time variability. In this part there is discussion of binaural sound localisation and of the localisation abilities of electric fish, as well as other special features of perception.

The book is essentially mathematical and there are few pages that are free of equations. The authors insist, however, that only fairly elementary mathematics is used. Each chapter ends with a summary of its main points and a discussion of the relevant literature. The latter is specially welcome in some of the chapters

where it is easy to get the impression that the emphasis is on the mathematical model, and wider reading is needed to appreciate the biological relevance. It is difficult to see how this could have been avoided without turning this already sizeable work into an encyclopaedia.

The treatment has little immediate relevance to practical robotics, but it undoubtedly holds pointers to future developments that will allow robots to come closer to their biological prototypes and for this it is a very welcome addition to the literature.

Alex M. Andrew
95 Finch Road
Earley
Reading RG6 7JX

DOI: 10.1017/S0263574702254855

2D OBJECT DETECTION AND RECOGNITION: MODELS, ALGORITHMS AND NETWORKS, by Yali Amit, MIT Press, Cambridge, Massachusetts, 2002, xv + 306 pp., ISBN 0-262-01194-8 (Hbk, £26.95).

The treatment here is somewhat more general than the title suggests, since attention is not restricted to images that are intrinsically two-dimensional, but includes those arising as projections of objects or scenes in three dimensions This is a very thorough and practical review of methods for analysis of two-dimensional images, and hence of most work on computer vision, apart from schemes in which the visual input is supplemented by distance estimation by laser or ultrasonic rangefinders. Stereoscopic depth perception is not treated either, though methods are described for location of image features in ways that would seem to be just what is needed as a basis for it.

Images that are intrinsically two-dimensional include printed or hand-formed characters as well as brain scans and other body scans produced by modern techniques. In the book special attention is given to brain scans.

There is an enormous literature on methods for the interpretation of visual images in terms of three-dimensional objects and scenes. Most of it is with the assumption that objects are not deformable, though there are important recognition tasks for which this does not hold, an obvious example being face recognition. The present book is very much concerned with recognition where there may be elastic deformation. Living tissue, as in the brain, is notoriously subject to deformation and the analysis of brain scans has to accept this complication. Similar considerations apply to recognition of hand-formed characters, which are elastically deformed versions of prototypes.

Detection and recognition when there may be elastic deformation can be treated as applications of Bayesian statistical methods. Given an input or "instantiation", it is usually possible to estimate the probability with which it could have come from a number of prototypes, and Bayes' rule then allows posterior probabilities to be assigned to these. This becomes impracticable when the number of prototypes, is large, and more economical alternatives are described, some employing Dynamic Programming. The treatment is entirely practical and computational complexity is kept within bounds.

In reviews quoted on the back of the dust cover, the book is welcomed by leading workers in computer vision as presenting a novel synthesis of earlier strands of research, and as presenting efficient and well-motivated algorithms that have fundamental as well as practical implications for the study of the topic. There seems to be no doubt that this is a major contribution and will be a standard work of reference. Like other books from MIT Press it is well produced with helpful illustrations and an extensive bibliography.

Computer programs implementing the methods can be downloaded from an associated website. They are in $C++$ and will compile under Linux. It is mentioned that familiarity with $C++$ is needed to understand the programs, and some familiarity with Unix will be needed to make them run.

This is an area of investigation that has significance for biological studies as well as for artificial systems, and in one of the chapters a possible neural network implementation is discussed in considerable detail It is linked convincingly to the known structure of the mammalian visual cortex and will certainly be of interest to neurophysiologists.

Alex M. Andrew
95 Finch Road
Earley
Reading RG6 7JX
(UK)

DOI: 10.1017/S0263574702264851

FAULT DIAGNOSIS AND FAULT TOLERANCE FOR MECHATRONIC SYSTEMS: RECENT ADVANCES, edited by Fabrizio Caccavale and Luigi Villani, Springer, Berlin, 2003, pp. xii, 191. ISBN 3-540-44159-X. Springer Tracts in Advanced Robotics (STAR), Vol. 1 (Hardback, £49.00).

In a Foreword to this first volume in a new series, the series editor indicates the purpose as being to disseminate information on new developments rapidly and informally but with high quality, as soon as a topic has reached the stage of maturity where such presentation is appropriate. Since this series is concerned with robotics he takes satisfaction in noting that this first volume has wide-ranging applicability, including underwater vehicles and aircraft.

The term "mechatronics" denotes the integration of mechanical, electronic and information technologies. In many applications, the consequences of faults are likely to be serious and automatic fault diagnosis is valuable. It is particularly valuable if the location and nature of the fault are indicated, and even more so if operation can be maintained despite it.

Application areas in which automatic fault diagnosis is specially valuable include unmanned vehicles for underwater exploration and for operation in conditions that are dangerous for humans because of radioactivity or other conditions, as well as in space craft and planetary rovers. The breakdown of a robot in an area inaccessible to humans can obviously be costly, and similar considerations apply to military aircraft, ships and vehicles. Prompt fault diagnosis and remediation is also important in industrial plants and it is mentioned that there is a plan to have motor vehicles in North America fitted with a form of automatic fault diagnosis in order to protect the environment. The authors of one of the papers visualise this being extended to boats and lawn mowers.

The six chapters of the book originated as presentations at a workshop on *Fault Diagnosis and Fault Tolerance for Dynamic Systems*, held in Vancouver in October 2002 in conjunction with an IEEE Symposium, and have been revised and expanded by the authors. All of them are fairly mathematical and acquaintance with modern control theory is assumed.

The first chapter is on "Sliding Mode Observers and Their Application in Fault Diagnosis", where the mention of Sliding Mode Observers (SMOs) indicates a previously-existing body of theory that is introduced here extremely sketchily. A useful reference seems to be a book by Utkin, about which notes can be found at: http://www.eleceng.ohio-state.edu/~utkin/ee894v.html

The second chapter is also strongly theoretical and refers to "Fault Diagnosis and Fault Tolerant Control for Non-Gaussian Stochastic Systems with Random Parameters".

The four remaining chapters refer to specific areas of application. One is on "Fault Diagnosis for Industrial Robots", and another is a "Survey of Fault Detection/Tolerance for AUVs and ROVs" where the acronyms stand for Autonomous Underwater Vehicle and Remotely Operated Vehicle. In the chapter the reference to ROVs is also to underwater operation. The discussion is highly practical with a listing of specific failures that might