Robotica (2000) volume 18, pp. 687–689. Printed in the United Kingdom © 2000 Cambridge University Press

Book Reviews

NANOMEDICINE, VOLUME 1: BASIC CAPABILITIES, by Robert A. Freitas Jr., Landes Bioscience, Austin, Texas, 1999, xxi + 509 pp., ISBN 1-57059-645-X Index ((Hardback, $89.000)

The possibilities of using microrobots in medicine have been well explored in the literature, but this book goes further in that it considers manipulation at the molecular level. The prefix nano-, indicating ten to the minus nine, has also been adopted to indicate operation at the molecular level and is appropriate since a nanometre is about the width of six carbon atoms.

Nanomedicine is not with us yet, but the author of the book is confident that it can be expected on the basis of current trends, and in his first chapter, on: 'The Prospect of Nanomedicine' he estimates that molecular technology will be the main paradigm of scientific medicine from the year 2010. The present period, from 1940 to 2010, he terms: "Molecular scientific" in which living processes are understood in terms of molecular structures, but without the means of directly manipulating them by placing individual atoms or atomic groups at will.

The first chapter puts the discussion in context by reviewing the history of medicine from very early times, with some gruesome details of operations before the introduction of anaesthesia, and other intriguing details. Although nanomedicine is still "in the pipeline" it is held to be at a stage where its implications should be seriously considered, including ethical considerations. Some of the ethical aspects are already under discussion in connection with genetically modified crops and cloned animals, and nanotechnology would allow such developments with enormously greater ease and speed. Applied to the nervous system it could allow personality modification, with obvious ethical complications.

The author is in no doubt that these developments are imminent and his affiliation is to an Institute for Molecular Manufacturing in Palo Alto, California. The present volume is the first of three, with the others already in preparation. The aim of this first volume is: "to describe basic capabilities common to all medicinal nanodevices, and the physical, chemical, thermodynamic, mechanical and biological limits of such devices." Its intended audience is a variety of specialists engaged in basic research. The second volume is to deal with: "aspects of device control and configuration, biocompatibility and safety issues, and basic nanomedical components and simple systems." Its primary audience is a variety of specialists engaged in applied research. The third volume is to come still closer to medical practice, with discussion of treatments for specific conditions and injuries and is aimed at clinical specialists and research physicians.

The reasons for believing that molecular manufacturing is not only feasible but just around the corner are given in the second chapter, on: "Pathways to Molecular Manufacturing". Approaches are classed as top-down and bottom-up. The former term is used to refer to successive scaling-down from normal engineering practice, where micromanipulators would be used to build still smaller micromanipulators, and so on. This has resulted in a technique of MEMS (Micro Electro-Mechanical Systems) with some very impressive results but not allowing operation at the truly molecular level.

The bottom-up approaches come under the three headings of biotechnology, supramolecular chemistry and scanning probes. Biotechnology provides almost ready-made solutions. The cell component termed the ribosome can be seen as a programmed protein molecule assembler and it has been found possible to use bacterial ribosomes to produce such molecules to order. Approaches under the heading of supramolecular chemistry are extensions of existing methods for organic synthesis, and a bewildering collection of impressive results is quoted, emphasising the widespread attention being given to these matters. There is overlap between the supramolecular chemistry approach and the biotechnology one, especially since nanodevices for most medical applications will need to be available in large numbers and should be amenable to mass production, or, better still, self replication.

Particularly impressive under this heading are possible structures formed of the "fullerene" or "buckyball" allotrope of carbon, which actually allow meshing gear wheels and rack-and-pinion mechanisms. The pictures shown are from computer simulations of the structures, but presumably using simulation packages that embody appropriate constraints to confirm that the structures are feasible.

The third class of approach is through the use of scanning probes. Scanning probe microscopes, in which the surface to be examined is traversed by a physical probe, are well known and commercially available, and are capable of imaging individual molecules. The technique has been extended to allow the deposition or transport of single atoms or molecules, and nanoassembly robots are visualised.

Later chapters deal with specific aspects, starting with: "Molecular transport and sortation", in which there are proposals for devices to allow the passage of selected molecules, for example by having them adhere to the ends of spokes of molecular wheels that rotate and so convey the molecules through a barrier. The other chapter headings are: "Shapes and Metamorphic Surfaces", "Power", "Communication", "Navigation", "Manipulation and Navigation" and "Other Basic Capabilities".

"Transportation and sortation" are important since important medical applications will come under the heading of scavenging, after the fashion of leucocytes in the blood. Metamorphic surfaces are ones that are able to deform in a useful way by local activation, for example as part of a swimming action. Under "Communication" there is consideration of input and output devices allowing interaction with a human host of nanorobots, or his physician, through display panels made to appear on an appropriate area of skin such as the back of the left hand. In the final chapter possibilities of onboard computation in nanorobots are considered, with a remarkably detailed discussion of theoretical power requirements for computation. Also in this chapter it is acknowledged that nanorobots may have to fight and destroy malignant organisms by mechanical rather than chemical means, using something in the nature of nano-jaws and teeth.

All of the topics are developed in considerable detail, involving a degree of extrapolation that is certainly courageous and perhaps foolhardy. The treatment is supplemented by quite a lot of tabular data, much of it rather unnecessary since it can be found in standard reference works, though including some items that are important in the special context but would be difficult to find elsewhere. References to other published work are indicated by superscripted numbers, most of them four-digit numbers since the total of such citations is **3728**. The actual number of references is substantially greater than this, since many of the numbered items give a primary reference and others in support. Many useful websites are quoted. A curious feature is that the numerical ordering of the references follows no obvious pattern, so that the only systematic way of finding a vaguely-remembered reference is to search for its number in the relevant section of the main text. An index linking authors' names to reference numbers would be a useful addition, which the author may be planning to provide when all three volumes have been completed.

Whether or not the "peek into the future" (as it is described in an Afterword to the book) will turn out to be accurate is impossible to say. The author gives compelling arguments in support of the picture he paints but of course it would be naive to dismiss the possibility of hidden pitfalls. At the very least, though, this is a thorough and fascinating exploration of a field that could bring enormous benefits.

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

UNDERSTANDING INTELLIGENCE, by Rolf Pfeifer and Christian Scheier, MIT Press, Cambridge, Mass., 1999, xx + 697 pp. ISBN 0-262-16181-8 (hardback, £37.50).

Interpreting the title of this book invokes a curious form of recursion, in that the attempt to attach a precise meaning to "understanding" inevitably involves assumptions about the meaning of "intelligence". The treatment in the book starts with what might be termed a common-sense view of intelligence and goes on to assert that the way to gain understanding is to synthesise agents capable of behaviour that would be generally accepted as intelligent, i.e. through *AI*. The emphasis, however, is no longer on the esoteric kind of intelligent agent that can prove mathematical theorems and play board games, but on robots that are "situated agents" and can interact with everyday three-dimensional environments, perhaps operating within societies of similar agents.

The book is, in fact, an enthusiastic introduction and comprehensive review of the new approach to *AI* that has also been treated in a number of other works including *Cambrian Intelligence* by Rodney Brooks, also from MIT Press and recently reviewed in these pages. The new book is based on a course taught by the authors that has been enthusiastically received by the students, though with the curious side effect of encouraging a viewpoint that made it difficult for them to accept the arguments of certain other courses running in parallel, particularly on cognitive psychology, which presumably embodied insupportable principles taken from traditional *AI*.

The book can be thoroughly recommended as a persuasive, comprehensive and lucid treatment of the new approach. It is argued that the way forward is by experimentation with robots capable of interacting with the untidy everyday world, and the authors note with satisfaction that there are now ways of producing these economically. An alternative allowing even more economical experimentation, but with inherent drawbacks that are discussed, is computer simulation, and the associated website: <http://www.ifi.unizh.ch/~pfeifer/mitbook> gives information on simulation packages that can be used.

Whether the new approach to *AI* will fulfil the expectations is another issue. There is little doubt that limitations of traditional *AI* have become apparent, and the emphasis in the last decade or two on expert systems can be seen as a tacit admission of this, since such systems model human performance without deep analysis, essentially as a curve-fitting procedure may model an observed physical phenomenon accurately but as a "black box".

One objection that has been made to the classical *AI* approach is that it ignores continuity (see Mackay,[1] Andrew[2,3] and Churchland.[4] Intelligence has evolved in a continuous 3-D environment in which possibilities of interpolation, extrapolation and smoothing are inherent and are presumably reflected in the behaviours that have emerged. The new approach acknowledges the importance of continuity, as must any approach involving robotics. On the other hand, the full exploitation of continuity has important aspects that have not been touched, particularly the subtle ways in which it enters into the discrete-concept-based reasoning that is the concern of mainstream *AI*. One aspect of this is recognised in Marvin Minsky's principle of "heuristic connection" (see Minsky).[5,6]

Pfeifer and Scheier refer to the success of Big Blue against Garry Kasparov as an example of powerful performance achieved by means that are not readily accepted as "intelligent" since they depend on deep search made possible by enormous computing speed (though since exhaustive search of all game continuations is still not feasible, the performance of Big Blue must also have depended on well-chosen heuristics). On the other hand, Kasparov was able, using his brain (which as far as we know is a vastly slower processor) to play chess of a standard that was only just beaten by the might of Big Blue. It can be argued that this depends on some little-understood combination of derived continuous measures like those of "centre-control" and "mobility" used by Samuel[7] in his famous checker-playing program, with the discrete environment of the chess game.

It is easy to feel that the new approach to *AI* will meet many of the same difficulties as did the traditional one when the attempt is made to extend its range of application. It will, however, have the advantage of freedom from the bias against continuity that has hampered traditional *AI*, and according to the viewpoint just outlined this could eventually allow its application to top-level chess without the emphasis on computing power. However, the suggestion of a breakthrough in *AI* generally is hardly warranted by developments so far. The book is written with the enthusiasm of the converted, and in the Preface the first-named author recounts the time and place of his conversion, the place being Luc Steels' laboratory in Brussels.

It is encouraging to observe that an approach starting from simple robots in a "real" environment has considerable correspondence to the course of natural evolution and can therefore be expected to find solutions that are similar to those found in living systems. This is not quite the conclusive argument it seems, though, since of course *AI* researchers cannot wait for something like natural evolution to run its course, but must speed things up using their human insights, which amount to "hunches" similar to those underlying traditional *AI*.

Nevertheless, this is an extremely valuable work, being in fact the only comprehensive textbook of the new approach to *AI*, admirably presented. My only reservation is some doubt whether the new approach is quite the breakthrough that is implied.

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

References

1. D.M. MacKay, "On the combination of digital and analogue techniques in the design of analytical engines" **In**: *Mechanisation of Thought Processes* (HMSO, London, 1959) pp. 55–65.
2. A.M. Andrew, *Continuous Heuristics: the prelinguistic basis of intelligence* (Ellis Horwood, Chichester, 1990).
3. A.M. Andrew, "Logic and continuity – a systems dichotomy", **In**: *Cybernetics and Systems Research* (ed. R. Trappl) (North-Holland, Amsterdam, 1982) pp. 19–22.
4. P.S. Churchland, *Neurophilosophy – Toward a Unified Science of the Mind/Brain* (MIT Press, Cambridge, Mass., 1986).
5. M.L. Minsky, "Contribution to discussion" **In**: *Mechanisation of Thought Processes* (HMSO, London, 1959) p. 71.
6. M.L. Minsky, "Steps towards artificial intelligence" **In**: *Computers and Thought* (ed. E.A. Feigenbaum and J. Feldman) (McGraw-Hill, N.Y., 1963) pp. 406–450.

7. A.L. Samuel, "Some studies in machine learning using the game of checkers", **In**: *Computers and Thought* (ed. E.A. Feigenbaum and J. Feldman) (McGraw-Hill, N.Y., 1963) pp. 71–105.

AN INTRODUCTION TO SUPPORT VECTOR MACHINES AND OTHER KERNEL-BASED LEARNING METHODS by Nello Christianini and John Shawe-Taylor, Cambridge University Press, Cambridge, 2000, xiii + 189 pp., ISBN 0–521–78019–5 (Hbk, £27.50).

This is the first textbook to appear on a new and enormously important development in learned pattern classification. The basic ideas were put forward in the early nineteen-nineties by the Russian workers Vapnik and Chervonenkis, though some aspects can be traced further back. The application is to the classification of input patterns represented as high-dimensional vectors, amounting essentially to the supervised-learning task associated with devices of the perceptron types.

The new techniques are therefore distinct from the other recent groundbreaking developing in machine learning, described under the heading of "reinforcement learning" by Sutton and Barto.[1] The latter deals with dynamic situations, but in practical applications such as robotics must depend also on static classification, so future systems will certainly combine these two powerful innovations.

The book begins by discussing the learning of linear discriminations such as are possible with a "simple perceptron" or by statistical techniques devised, like so much of statistical mathematics, by R.A. Fisher. A simple perceptron operating directly on its unmodified sensory inputs is limited to linear discriminations, and usually such devices operate on a set of functions of their inputs that can be termed features. The success of many practical devices depends critically on the initial manual choice of features. The authors of the present book take the view that the choice of features should be a part of the automated learning process.

Another point they make about the perceptron is that it does not converge on a unique solution. No further adjustments are made once a set of weight values has been found that performs the required discrimination within the training set of inputs. This set of weights can be different depending on the order of presentation of the patterns used for training. Variants of the method are however possible that do converge to a unique solution. They can operate to maximise the margin of the discrimination, i.e. the minimum distance from the separating hyperplane to any point on either side representing a vector in the training set. It is also possible to incorporate criteria favouring minimisation of the norm of the weight vector, which amounts to a kind of Occam's razor principle that should improve the chance of the solution being generally applicable outside the training set.

Once the learning process has been expressed in terms of extremum seeking, a dual representation can be obtained by equating to zero the derivatives of the relevant function with respect to the weights. The resulting expression, normally in matrix form, is termed a kernel. A kernel corresponds uniquely to a set of features, but for many purposes the kernel representation can be manipulated without explicit reference to the features. Mathematical criteria for a function to be acceptable as a kernel are derived, as well as means of combining existing kernels to produce new ones.

Support vectors are introduced fairly well on in the book, and refer to the fact that, where a separating hyperplane is chosen for maximal margin, it is effectively a function of only those points that are distant from it by exactly the width of the margin. The vectors defining these points are termed support vectors.

Operating the learning process in its dual, kernel-based version brings a remarkable package of benefits. The effects of the "curse of dimensionality" in high-dimensional features spaces lose their force, and computations depend on the computationally-inexpensive evaluation of inner products of vectors. The definition of a kernel includes a convexity condition that ensures that the optimisation process cannot stick on local optima. It becomes feasible to accept training sets with as many as 100,000 elements. The method performs well when the elements of the training set to be classified in one way are extremely sparse. Alternative methods tend to settle for the trivial solution in which all inputs are attributed to the non-sparse category, a solution which gives a good percentage score if the other category is sufficiently sparse, though obviously no useful discrimination is achieved.

Another nice feature is that the method is applicable to non-Euclidean pattern spaces, and in particular kernels have been devised that apply to the recognition of substrings in sequences of symbols, and hence in DNA sequences. This area of application is of particular current interest.

In the final chapter, a number of specific applications are described, and in all of them the new approach performs significantly better than alternatives. The applications include text categorisation, with applications in e-mail filtering, web searching, office automation, sorting of documents by topic, and classification of news agency stories. There are also applications to straightforward image recognition, including some in medicine as well as the familiar one of handwritten character recognition. Applications in bioinformatics include the examination of gene expression in DNA sequences, and the detection of protein homologies, or similarities in properties, from examination of the amino acid sequences of the molecules.

The book is of course essentially mathematical, but with clear and readable explanatory text and notes to help with difficult points in the mathematics. Since Support Vector Machines are topics of ongoing research, the book cannot possibly be comprehensive in its coverage, but an associated website is quoted as: <http://www.support-vector.net> and the reference list as it appears there will be duplicated, as will also a list of relevant items of software available for downloading.

Ideas about feature selection and refinement were visualised by, notably, Oliver Selfridge in an early paper,[2] and it is difficult to believe that all the problems associated with his "mutated fission" are solved just by shifting to the dual representation in terms of kernels. On the other hand, quite enough problems are solved or mitigated by the shift to make this an extremely important development. The book is an admirable presentation of this powerful new approach to pattern classification.

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

References

1. R.S.Sutton and A.G. Barto, *Reinforcement Learning: An Introduction*, MIT Press, Cambridge, Mass., 1998.
2. O.G. Selfridge, "Pandemonium: A paradigm for learning" **In:** *Mechanisation of Thought Processes* (National Physical Laboratory Conference, HMSO, London, 1959), pp. 511–531.

https://doi.org/10.1017/S0263574700212824 Published online by Cambridge University Press

KYBERNETES

International Journal of Cybernetics and Systems
(a quarterly)

OFFICIAL JOURNAL OF THE WORLD ORGANISATION OF SYSTEMS AND CYBERNETICS

This journal is concerned with the interdisciplinary study of cybernetics and systems in the widest sense. Among the extensive range of topics covered are the following: *Artificial Intelligence, Automation, Cybernetic Modelling, Computer Simulation, Biocybernetics, Economic and Social Systems, Nature and Validation of General Systems, Ecosystems, Adaptive Systems, Philosophy of Cybernetics, the Interrelation between Cybernetics and other Sciences.*

Kybernetes, now in its twenty seventh year of publication, aims to endow cybernetics and general systems with an authoritative voice of its own and to establish a competent forum for the exchange of knowledge and information in relevant fields. One of the primary objectives of this publication is to weed out pseudo-cybernetics claims and to base cybernetics on a sound foundation; hence, highly speculative papers that lie on the fringe of cybernetics and systems sciences will not be published. The journal contains only *original papers*; it is the official publication of the World Organisation of General Systems and Cybernetics.

Enquiries and orders to:

MCB UNIVERSITY PRESS
62 Toller Lane
Bradford
BD8 9BY
England
Tel. 01274 499821
Telex. 51317 MCBUNI G
Fax. 547143

For Subscription Rates, Advertising Rates and Back Issues enquire at MCB University Press, as above.

A PRESTIGE PUBLICATION, INDISPENSABLE TO YOUR ORGANISATION AND TO YOU NOW IN ITS TWENTY NINTH YEAR

https://doi.org/10.1017/S0263574700212824 Published online by Cambridge University Press

https://doi.org/10.1017/S0263574700212824 Published online by Cambridge University Press