

Book Reviews

ROBOCUP-97: ROBOT SOCCER WORLD CUP 1, edited by Hiroaki Kitano, Lecture Notes in Computer Science Series No. 1395, Springer, Berlin, 1998, xiv + 520 pp., ISBN 3-540-64473-3 (Softcover, £34.00 or \$54.00).

This is the proceedings of the First World Cup Soccer Games and Conference held in Nagoya, Japan, in conjunction with the International Joint Conference on Artificial Intelligence IJCAI-97. It is part of an international initiative termed RoboCup, whose ambitious long-term goal is stated as:

“to build a team of robot soccer players, which can beat a human World Cup champion team.”

The idea of robots playing soccer was first mentioned by Professor Alan Mackworth (University of British Columbia) in a paper in 1992. Independently of this, a group of Japanese scientists organised a Workshop on Grand Challenges in the same year, and decided that robot soccer provided a potentially profitable Grand Challenge. This was followed up a year later by the decision to launch a robotic competition, initially planned to be within Japan. However, the news spread and interest was expressed by workers in other countries, and RoboCup was established as an international project. The first public announcement was made in 1993, and regulations were drafted.

An event called Pre-RoboCup-96 was held during an International Conference on Intelligent Robotics and Systems (IROS-96) in Osaka in November 1996. Eight teams competed in what was termed the ‘simulation league’ and there were demonstrations of robots suitable for the ‘middle size league’ (see below).

In the 1997 event a total of 40 teams participated (real and simulation combined) and over 5000 spectators attended, and the event was judged a great success.

Needless to say, the robotic devices, real or simulated, fall far short of what would be needed to achieve the ambitious long-term goal. The robot players, in the versions where they are physical devices, in fact run on wheels. (There is, in the book, a paper discussing the use of legged players, but suggesting quadrupeds for the time being.) There are three ‘leagues’ in which real robots can compete. One is the small-size league, played on a field the size of a ping-pong table, and another is the middle-size for which the linear dimensions of the table are three times as great. For both of these there are five robots in a team. For the small-size league it is expected that control will be due to global viewing of the field, but for the middle-size entrants the robots should have on-board vision systems and be governed by what is in their field of view. The third league for real robots is the ‘expert robot league’ for which the conditions are left open, allowing the demonstration of special features.

To allow competition in the ‘simulation league’, the competing programs must connect to a ‘soccer server’ which manages the simulation and display. There must be a separate control for each player, though possibly using copies of the same program for each; overall global control is not allowed. Each competing program is able to send messages to its own player, selected from a repertoire of instructions, and ‘sensory data’ from the player is conveyed in the other direction. The soccer server can also let a program play against eleven humans, where the latter control their respective virtual players through appropriate interfaces.

The competitions clearly have entertainment value, but their

significance goes well beyond this. The stimulus provided by the Grand Challenge is real. The Preface and also the back-cover notes emphasis that 1997 was the year in which a computer beat a world master in chess, and also that in which an autonomous vehicle performed successfully on the surface of Mars. These achievements could well have been the aims of earlier Grand Challenges.

It is clear that in the work to date on RoboCup, major problems of balance and dynamics have been bypassed by the use of wheeled robots, but the remaining problems of multi-agent operation and of opponent modelling are enough to be going on with. At the end of one of the papers, in which the outcome of the 1997 event is reviewed, it is argued that a deficiency of the schemes tried was that they did not learn on-line, so that the behaviour did not change radically during a match. Future entries will presumably learn, or at least hone their skills, so as to adapt to their opponents’ special characteristics. Another feature that was said to be missing was communication between members of the same team, like the verbal communication between human players.

The papers in this first publication show that the project has already stimulated valuable effort. They are placed under three main headings of Overview Papers (4 in number), Technical Papers (14), and Team Descriptions. The last group is further subdivided into four papers on small-size teams, five on middle-size ones and no less than eighteen on simulation teams. A Team Description submission was required for each entry in the competitions, but many of these are effectively technical papers emphasising what are felt to be the special features of the respective entries.

Included in the Technical Papers, in addition to that on quadruped robots already mentioned, there is one on a system using the currently-popular *Java* programming language for the controlling software, allowing easy modification. Others review general aspects of multi-agent operation, and of the real-time processing of sensory data, and generation of the required multimedia presentations. Another considers the value of project-based education in robotics, of which RoboCup could be held to be an extreme example. There is also a rather surprising contribution, in the form of a careful analysis of 2402 football articles that have appeared in the *Times* newspaper, all of them from the 1996–1997 football season. The idea seems to be that this will be useful in deciding about the rules and characteristics of future robotic football, so as to preserve the features that make live football interesting.

In the Team Descriptions, quite complex mathematical arguments are used in presenting the playing algorithms used. The currently-fashionable topic of Reinforcement Learning is invoked in several presentations, and is appropriate for this kind of delayed-reward successive-choice activity. Subsumption architecture is also mentioned, again not surprisingly since the requirement is to pursue fixed objectives, with the pursuit interrupted to perform the constituent activities of kicking or dribbling the ball or getting into a favourable position or intercepting an opponent. One of the papers by a Japanese author (“Development of Self-Learning Vision-Based Mobile Robots for Acquiring Soccer Robots Behaviors” by Takayuki Nakamura) combines Reinforcement Learning with an interesting form of self-organisation amounting to automatic partitioning of the state space to suit the control requirement.

Another paper describes the development of playing skills by a Genetic Programming method, where the necessary evolution took place prior to the competition. The evolutionary process is described as being fascinating to watch, since it began with policies such that most of the players never touched the ball during a game, and gradually improved until players could be seen taking on distinct roles as attackers and defenders.

All of the treatments illustrate the shortcomings of what is referred to in the final paper as GOFAIR, or Good Old Fashioned Artificial Intelligence and Robotics. Early work on robotics, still relevant in industrial applications and elsewhere, tended to assume the environment was unchanging or slowly changing, except when operated on by the robot itself. There are many interactive situations, amounting in fact to much of real life, where the GOFAIR approach is inadequate, and the RoboCup environment is a prime example and a useful testbed.

The book, and the whole project, are extremely interesting and potentially valuable. It is hardly necessary to observe that there is a long way to go before we see robot players on a real soccer field, but the rate of technical progress is such that goals should not be set too low.

The plan is announced in the book to hold RoboCup-98 in July 1998 in Paris during the Third International Conference on Multi-Agent Systems (ICMAS-98), and no doubt details of this will appear in a later volume. The third event, RoboCup-99, is announced for Stockholm, as part of the International Joint Conference on Artificial Intelligence (IJCAI-99).

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

STATISTICAL METHODS FOR SPEECH RECOGNITION,

by Frederick Jelinek, Language, Speech and Communication series, MIT Press (Bradford Book), Cambridge, Mass., 1997, xxi+283 pp., ISBN 0-262-10066-5 (Hardback, £29.00).

Speech recognition by machine has become very much a practical proposition in recent decades. This book gives a very thorough account of the methods that have proved successful. The author warns readers that he deals only with what the title indicates, and that successful implementation depends on techniques of signal processing that are not treated here. At a later stage he narrows his field a little further by observing that speech recognition can be divided into small-vocabulary and large-vocabulary versions, and he deals only with the latter. He acknowledges that small-vocabulary schemes can be extremely useful, for example when the only inputs to be recognised are the numeral digits, or even just "yes" and "no". For readers who want to explore related topics, or alternative approaches to those he treats, he gives suggestions for further reading.

The origin of the book is notes for a course given by the author in the Czech Technical University during 1991–92. They were extensively revised for use in teaching at the California Institute of Technology (Spring 1993) and Johns Hopkins University. The methods have mainly been pioneered by members of an IBM Continuous Speech Recognition Group, to which the author belonged, and he pays tribute to the generous administrative attitudes of IBM at that time.

Treatment of the topic is inevitably mathematical, and a flick through the pages shows some of them to be rather dauntingly well-filled with equations. However, the author insists that the treatment makes no great demands on the reader's mathematical knowledge or expertise, and in fact he is right, though perseverance is needed to follow all the arguments through the fifteen chapters.

The author gives a special reason, additional to the obvious practical utility of the result, for his own enthusiasm for the topic. He says he is: "fascinated by the idea that while system structure and parametrisation should come from intuitive understanding of the process, the parameter values are best extracted from actual

data." The method as described starts with a process of 'vector quantisation', by which the incoming speech waveform is converted to digital form, by periodic sampling of the outputs of a set of filters. A form of self-organising clustering is incorporated here. Later stages depend on modelling the source (speaker) as a Hidden Markov Model for which transition probabilities are estimated. By application of Bayes' Theorem these give rise to probabilities of correct identification of, first, phonemes (referred to here as 'phones') and then words.

The book is intended to teach principles that can be applied in a variety of circumstances, and not just to give detailed recipes. In the final chapter some alternatives to the methods are discussed. At an earlier point it is acknowledged that, as in other forms of machine learning, artificial neural nets have found a place in some approaches to the problem.

The essential components of a large-vocabulary recogniser are covered in the first five chapters, and the remaining ten are concerned with refinements. Some of the refinements depend on considerations of entropy, and hence on information theory, and chapter seven is devoted to an account of the latter, with some special twists adapting it to this requirement.

The book is admirably comprehensive (though no claims are made for exhaustive coverage of the literature) and well presented and will certainly be a standard text and reference work.

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

PRACTICAL MOTION PLANNING IN ROBOTS: CURRENT APPROACHES AND FUTURE DIRECTIONS,

edited by Kamal Gupta and Angel P. del Pobil, Wiley, Chichester, 1998, xi+356 pp., ISBN 0-471-98163-X (Hardback, £65.00).

The papers in this volume come from a one-day workshop having the same title, organised by the editors within the 1996 IEEE International Conference on Robotics and Automation, held in Minneapolis. According to their Preface, the status of motion-planning theory is controversial, with some people seeing it as a mature field while others feel that the theoretical studies have made little inroad into real-world applications. The aim of the workshop was to focus on approaches that were particularly amenable to such applications.

A distinction is made between gross and fine motion control. The former is associated with non-contact tasks with relatively high tolerances, while the latter requires high precision, and usually some involvement of sensory input. The papers in this book are exclusively concerned with gross motion control.

A further distinction is between motion planning for articulated arms and that for mobile robots. Attention is given mainly to the former, except in one chapter.

Yet another distinction is between planning in known and unknown environments. In the former, the environment is assumed to be known in advance, often in the form of polyhedral models (though the alternative of models formed of multiple spheres is considered in one of the papers). In an unknown environment, operation must rely on sensors. In the book, operation in a known environment is referred to as classical motion planning, while the other is termed sensor-based motion planning.

The book is divided into five Parts. The first, and largest with seven papers, is on Practical Model-Based Approaches to Motion Planning. It begins with an overview of the state of the art by one of the editors. The second Part, with two papers, is on Applications in Industry. Interestingly, one of the papers is not about robotics as such, but about the application of the same principles in the design of complex mechanisms where some components have to be replaced regularly (like the sparking plugs and filters of a car engine). In modern jet engines numerous components, some of them subject to periodic replacement, have to be housed in a limited space and there is a real problem in ensuring that there are

feasible pathways for the removal and insertion of the renewable items.

The third Part, with four papers, is on Collision Detection and Geometric Complexity, and the fourth, with three papers, has some overlap of subject matter as it is on Sensor-Based Approaches. One paper is on Harmonic Functions for Path Planning and Control, and is a development of the use of potential fields for path planning with obstacle avoidance, but using a field function of a particularly useful kind. Another paper describes an articulated arm controlled by an operator as far as its end-point movement is concerned, but nevertheless covered in its entirety with touch sensors. The touch sensors are used to ensure obstacle avoidance for the arm as a whole, allowing the operator to concentrate on the end-point. It has been found that human operators have great difficulty in planning obstacle-avoiding paths for a complete tele-operated arm (even though we do it unconsciously for our own limbs).

The final Part is useful for a reader trying to get a quick impression of the main points and of their wider implications. It consists of a single paper by the editors in which they try to summarise the contents of round-table discussions during the event. There is general agreement that the reason some of the techniques are not more widely used is that they are tedious to set up. This is likely to be true where they depend on digital representation of the work environment. The development of convenient user interfaces would be helpful. The final remark in the paper and hence in the book is to the effect that motion planning cannot be treated in isolation and an important question for future work is how to integrate the motion planning approaches with related problems such as grasp planning, manipulation, and fine motion planning. There is clearly much food for thought here.

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

ARTIFICIAL NEURAL NETWORKS – ICANN'97, edited by Wulfram Gerstner, Alain Germond, Martin Hasler and Jean-Daniel Nicoud, Lecture Notes in Computer Science Series No. 1327, Springer, Berlin, 1997, xix + 1274 pp., ISBN 3-540-63631-5 (Softcover, £81.50 or \$129.00).

This is the proceedings of the Seventh International Conference on Artificial Neural Networks, held in Lausanne, Switzerland, during 8–10, October 1997. The number of pages must be close to the limit that is practicable for a single volume in this format. In the Preface it is stated that, of 365 submitted papers, 190 were accepted for publication, and in addition there were several invited papers. The three days of the meeting were obviously fully occupied. These conferences have been held annually since 1991, and have become the major European meetings dealing with the topic. A flick through the pages of these proceedings shows that interest is not confined to Europe, with at least one paper from New Zealand and one with joint authorship from South Africa, and substantial participation of Israel, Japan and U.S.A.

With such a wealth of material to draw on, an impression can be gained of the current state of the art, or at least the conference organisers' view of it, by examining the list of contents and the headings used to classify it. The papers are grouped in eight Parts, with the respective headings:

- (1) Coding and Learning in Biology (26 papers)
- (2) Cortical Maps and Receptive Fields (15 papers)
- (3) Learning: Theory and Algorithms (41 papers)
- (4) Signal Processing: Blind Source Separation, Vector Quantization, and Self-Organization (25 papers)
- (5) Robotics, Adaptive Autonomous Agents, and Control (32 papers)
- (6) Speech, Vision, and Pattern Recognition (16 papers)
- (7) Prediction, Forecasting, and Monitoring (25 papers)
- (8) Implementations (21 papers)

The headings are obviously not mutually exclusive, and for example several of the papers in the final Part, on Implementations, describe schemes for pattern recognition that could have been allocated to Part 6. An interesting and welcome aspect is that the first two Parts are concerned primarily with real nervous systems, and their papers qualify for inclusion here either because artificial nets have been constructed as models of real ones or because methods of analysis are applied that stem from neuromodelling rather than directly from physiological experiments. The attention to biology is genuine and constructive, and in the references the biological journals (*J. Neurophysiol.*, *Exp. Brain Res.* etc.) are well represented.

Many of the papers the first two Parts refer to quite specific neural structures and phenomena. Papers with biological relevance are not restricted to these two Parts, and in Part 6, on Pattern Recognition, there is a paper by Stephen Grossberg, whose pioneering work in neurobiology is enshrined in a collected works.¹ His paper in the present volume gives a theory accounting for some puzzling aspects of processing in the visual cortex. The paper following it has as joint author the well-known Christian von der Malsburg, and refers to the part played in pattern recognition by the 'complex cells' of the visual cortex.

It is interesting to see that, in Part 3, on Learning, there is a paper by Richard Sutton, one of the main instigators of the current developments in Reinforcement Learning. In his contribution here he discusses and defends the representation of the task environment in terms of Markov decision processes, a fundamental aspect of the recent important developments.

The phrase 'Blind Source Separation' in the heading of Part 4 is not self-explanatory, though it is sufficiently well-known in some contexts to be replaced by the initials BSS. In one of the papers it is defined: "In blind source separation the problem is to recover independent sources given sensor outputs in which the sources have been mixed by an unknown channel. The problem has become increasingly important in the signal and speech processing area due to its prospective application in speech recognition, telecommunications and medical signal processing." In most work it is assumed that the unknown mixing process is linear, but the extension to the nonlinear case is considered in one of the papers here.

One of the papers in Part 5 is on 'Learning Soccer Strategies' and so has a connection with another Springer conference proceedings to be reviewed in these pages. In other papers in this Part, neural nets are applied to robotics in a wide variety of different ways, including applications to path planning, force feedback control, visual tracking and various kinds of modelling. One paper is on 'Robot-Animal Interaction' and describes a robot that models the reaction of a cricket to sound and can be observed interacting with real crickets as confirmation of the accuracy of the model.

The subject-matter of the contributions ranges widely, and includes some down-to-earth practical applications of artificial neural nets, such as forecasting the loading of power generating stations, and fault detection in power systems. Speech recognition and the recognition of cursive handwriting continue to receive attention, as does the recognition of facial features. A paper in Part 7 is on 'Neural Combustion Control' and discusses the use of a neural net to control an internal combustion engine for efficiency and low emissions.

The topics of the 201 papers range so widely that these comments inevitably have a 'lucky dip' character. All of the papers appear to be of high quality and well presented and there is a mine of valuable material here.

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

Reference

1. S. Grossberg, *Studies of Mind and Brain*, (Reidel, Dordrecht, 1982).

ACTIVE CONTOURS, by Andrew Blake and Michael Isard, Springer-Verlag London, 1998, 352 pages including index and bibliography., (\$74.95).

"The aim of this book is to show how prior knowledge can be applied in machine vision at the lower level of shapes and outlines" (pp. 1–2). The prior knowledge in question is that which people acquire of the world through interacting with it. The book is, therefore, about contour-based object tracking algorithms.

The book is divided into two parts. Part One deals with the representation of curves using splines. It then moves to a design methodology for linear image-based parametric models of shape. Methods for fitting splines to image features and image processing are considered. Part Two deals with motion and deformation. Gaussian and non-Gaussian models are used as the basis for the treatment.

This is highly technical material, aimed at the specialist researcher (which I am not, unfortunately). It is also highly mathematical in content. This should serve as a warning to those with an interest in new developments in vision but who do not particularly care for differential equations. The book is also somewhat expensive. It contains a large number of illustrations, some in colour.

Iain Craig
Department of Computer Science,
Warwick University,
Coventry CV4 7AL
(UK)

GENETIC PROGRAMMING AND DATA STRUCTURES, by William B. Langdon, Kluwer Academic Publishers, 1998, 278 pages including bibliography and index.

Genetic programming is an outgrowth of genetic algorithms. The idea behind genetic algorithms is that computational entities compete in a simulated evolutionary setting, and the fittest (according to some *a priori* criterion supplied by the experimenter) survive to be used in either another round of evolution or, if the evolutionary process is judged to have completed, in the computation of some function. Since their first introduction in the 1960s by Holland, genetic algorithms have been employed in a number of practical as well as purely theoretical settings. It appeared natural to apply the concepts of genetic algorithms to the development of program code.

Koza produced some of the first genuine examples of genetically produced LISP programs a few years ago and this book develops the ground which he prepared. The book shows how advanced data structures such as lists, stacks and queues can be evolved. As part of this, the author first introduces the area and relates it to the more familiar field of genetic algorithms. He explains how genetic programs are constructed from components and how they rely upon such external entities as fitness functions, control parameters (parameters that affect the control of growth and fitness ratings of populations of program candidates) and criteria for determining when the evolutionary process should terminate.

This is interesting stuff, but I am not persuaded to engage in it, nor do I think that it will provide us with the answer to the software problem (compare with genetic algorithms, neural nets, etc.). Nevertheless, it is still worth keeping an eye on.

The book is nicely produced (I suspect from camera-ready copy) and well illustrated. The text is well written and is easy to read and understand. The concepts in genetic programming are well presented, as are those in genetic algorithms. The references are selected to be useful. People of a non-mathematical disposition will find few problems in reading this book.

Iain Craig

Department of Computer Science,
Warwick University,
Coventry CV4 7AL
(UK)

UNITED NATIONS, WORLD ROBOTICS 1998 (Statistics, Market analysis, forecasts, case studies and profitability of robot investment), co-authored by The International Federation of Robotics, UN Economic Commission for Europe, Palais des Nations, CH-1211 Geneva 10, Switzerland, 1998, 299 pp. (\$120). This compilation, the only of its kind, presents comprehensive global statistics on industrial and service robots for 20 countries. The data cover application areas, industrial branches, types of robots, worldwide diffusion of industrial robots, their applications in various countries (ranging from Australia to the USA), forecasts of worldwide investment in industrial robots in the period 1998–2001, robotics in the food and agricultural industries, case studies of the profitability of industrial robots (Germany, Sweden and the USA), and other relevant information. There is also a valuable summary of 12 pages at the beginning of the publication. In all, this collection of worldwide data is an indispensable aid to anyone seeking cogent and reliable information on robotics in various parts of the world, to researchers and development engineers, industrialists and economics experts.

J. Rose

OPTICS OF LIGHT SCATTERING MEDIA, Problems and Solutions by Alex A. Kokhanovsky, J. Wiley and Sons, Chichester, 1999, ISBN 0-471-97260-6, x + 217 pp., index (Hbk, £75.00).

This book is a part of the Wiley-Praxis series in atmospheric physics and climatology which reflects the major developments of the physical and dynamic processes occurring within the Earth's atmosphere. There are five chapters (Introduction; Single Light Scattering; Multiple Light Scattering; Light Scattering and Radiative Transfer in Densely Packed Dispersed Media; Applications) and five appendices dealing with refractive indices, solutions of light-scattering problems, special functions, light scattering on the Internet, and phase functions. This is followed by an extensive list of references (19 pages) and an index.

The work summarises current knowledge of the optical properties of single small particles and natural light scattering media, viz. aerosols, clouds, foam, snow, etc. The introduction is a valuable description of the light field, particles, radiative transfer and scattering theory. Chapter 2 describes various particles, small and large compared to the wavelength, particles with a refractive index close to that of the host medium, and layered particles. The next chapters 3 and 4 consider multiple light scattering and radiative transfer. Chapter 5 focuses on applications in geophysical optics, image transfer, remote sensing and inverse problems, bio-optics and planetary optics. The emphasis of the text is to deal with modern analytical methods used in light scattering optics. Simple equations are derived relating the size, shape and concentration of particles, and their optical constants and intensity of the reflected, internal or transmitted light.

Because of its sound treatment of a difficult topic, the book will be of interest to undergraduate and postgraduate students of atmospheric physics, optical engineering, biological optics and ocean optics. The work, however, is of marginal value to students of robotics and artificial intelligence, though those concerned with the application of vision in robotics may find it of use.

T.E. Sutton
(c/o Editor)