Robotica (2001) volume 19, pp. 459–462. Printed in the United Kingdom © 2001 Cambridge University Press

Book Reviews

REASONING ABOUT RATIONAL AGENTS, by Michael Wooldridge, MIT Press, Cambridge, Mass., 2000, xi + 227 pp., ISBN 0–262–23213 (Hardback, £23.50).

By an "agent" is meant an entity, possibly a person, capable of independent autonomous action, and the description as "rational" implies that the agent makes good decisions about what to do. Clearly the meaning attached to "good" must be context dependent, but there will usually be a readily-accepted criterion. The book is mainly concerned with principles underlying the design of computer systems that constitute rational agents, but the initial inspiration comes from a treatment that focuses on human performance and consequently comes under psychology or philosophy.

The basis is a Belief–Desire–Intention, or BDI, model. A rational agent has a set of beliefs about the environment, some of them the result of processing sensory input. The formation of beliefs is outside the scope of the present book, but references are made to works dealing with it.

The aim, of course, is to produce artefacts that will operate effectively in unpredictable complex environments which may involve interaction with other agents. There are other well established bodies of theory that go some way towards meeting the requirement, and Games Theory and Decision Theory are mentioned, with the observation that although their results are valuable and may be used they do not in themselves provide the necessary flexibility.

The book introduces and develops a special logic termed LORA, or Logic for Rational Agents. This allows the formal representation of beliefs, desires and intentions, as well as of agents holding or having them. Beliefs give rise to intentions, and a subset of these that are judged to be worth pursuing are termed "desires". (It would seem to me to be more natural to reverse this and to make "intentions" a subset of "desires" but the BDI model has it this way.)

Not everyone would agree that the development of a formal logic is the best way to set about the essentially *AI* problem of constructing rational agents, and the author gives attention to this. He wisely avoids making any assertions about the fundamental nature of biological intelligence, but maintains that a formal logic approach is a powerful means of producing effective artificial systems because of its rigour and transparency.

He also defends the introduction of a new logic, with operators such as *Bel*, *Des* and *Int* (for "believes", "desires" and "intends") that are not found in classical first-order logic, even though some of the requirements could be met using extremely tedious constructions within first-order logic. Purists may object to the fact that the new logic cannot be described as "fully axiomatised". The complexity of LORA, which embodies features from various specialised logics, means that there is little hope of achieving this status, but this is not to deny its usefulness.

The introduction of operators that are extra to classical logic brands this as a "modal" logic. Each of *Bel*, *Des* and *Int* has two arguments, so that for example *Bel*(*x*, *y*) might be used to mean that agent *x* believes *y*, where *y* would be a logical expression. Since the expression for the belief has to be used in ways that do not correspond to an ordinary argument of a function, a notation similar to a LISP statement is preferred, as (*Bel x*, *y*). What is believed by an agent may of course refer to beliefs held by other agents.

As well as being a modal logic, LORA is necessarily a temporal logic, and at the same time it has to be a logic capable of producing actions. For its temporal function, past history is considered to be a fixed sequence of events, but the future branches into a tree of possibilities, with some of the choices influenced by actions of the agent concerned. Variants of the existential quantifier ("there exists, such that . . .") and the universal quantifier ("for all . . .") are introduced to refer to paths in the tree, so that the existential path quantifier is a means of asserting that a statement is true for some path within a given set, while the universal path quantifier allows a similar assertion for all paths.

The author is well aware that he will have a mixed audience, with some readers unfamiliar with formal logic while others may be experts. He handles the situation admirably, with a chapter introducing LORA fairly informally and then another that includes formal proofs, etc. and an even more formal treatment in an Appendix. Then later chapters deal with collective mental states of agents, communication (including the action of initiating or requesting a "speech act"), and co-operation.

All this is presented very clearly and persuasively and the reader is assured right at the beginning that the methods have been used successfully in the control system of an autonomous space probe and in automatic Internet-based systems performing commercial transactions. The basic theory given here clearly needs to be supplemented by a lot of imaginative development to achieve ends such as these, but we are assured it has been done. The final chapter discusses the transition from the logic specification to a working system, with several sections headed "Case Study" but remaining fairly general. A special programming language called *METATEM* is introduced.

This approach using formal logic, with many references back to earlier work in *AI*, is in sharp contrast to other developments in *AI* and Robotics which favour relatively simple autonomous devices. The arguments of Brooks[1] and those advanced by Warwick[2] in connection with simple robots at Reading University are examples of this alternative trend. A common feature of the approaches is that both emphasise the need to consider "situated" agents rather than such lone performers as mathematical theorem-provers or chess machines. Work on situated agents is clearly directly relevant to advanced robotics.

Another common feature of the two approaches is that both have been applied to space exploration, the "simple robot" one to the design of planetary rovers and the formal logic one to control of an autonomous space probe. It is clear that both approaches have their own fields of application, and it would for example be likely to prove expensive to put a simple robot in control of automatic financial transactions.

The present book is a lucid and authoritative treatment of the logic approach, which could be worth dipping into purely for its relatively painless introduction to formal logic as such. Its topics are linked to a bibliography of no less than 251 items.

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

References

1. Rodney A. Brooks, *Cambrian Intelligence: The Early History of the New AI* (MIT Press, Cambridge, Mass., 1999).

https://doi.org/10.1017/S0263574701253491 Published online by Cambridge University Press

2. Kevin Warwick, *In the Mind of the Machine: The Breakthrough in Artificial Intelligence* (Arrow Books, London, 1998).

HETEROGENEOUS AGENT SYSTEMS, byV.S. Subrahmanian, Piero Bonatti, Jürgen Dix, Thomas Eiter, Sarit Kraus, Fatma Ozcan and Robert Ross, MIT Press, Cambridge, Mass., 2000, xiv + 580 pp., ISBN 0-262-19436-8 (Hardback, £39.95).

This very substantial work introduces and develops a new approach to software and computer system development, where an "agent" is a module of code, capable of performing a specific task or set of tasks, and interacting with other agents to achieve some overall result. The first-named author, as well as the two last-named, belong to the University of Maryland and the treatment is associated with a project denoted by the acronym IMPACT, or Interactive Maryland Platform for Agents Collaborating Together. (I find the redundancy of the last two words slightly irritating, but have not been able to think of non-redundant alternatives that would suit the acronym.) The project is clearly international since the four other authors are respectively from Italy, Germany, Austria and Israel.

The agents are visualised as providing some service that may be called upon by other agents. It is necessary that an agent has an associated interface giving information on the services offered and how they can be invoked. The principle is illustrated in the book with reference to three application areas. The first is to the operation of a department store offering web-based services and dealing with customer enquiries as well as trying to promote business by informing customers about merchandise likely to interest them. The operation is described as depending on a total of five agents, assumed to be implemented in software although any or all of them could be human.

One agent in this example is the Interface Agent, which receives inputs from customers and generates the multimedia output to them. This has two-way communication with a Content Determination Agent, which interacts with a Product Database Agent to produce answers to enquiries and also promotional material. The Interface Agent also informs a Profiling Agent about the customer input, and this agent prompts the Content Determination Agent about targeted promotional activity. The Profiling Agent takes account of customer credit information which it requests and receives from the Credit Agent.

Of the other two application areas, one is to navigation of an aircraft, with the acronym CFIT, for Controlled Flight Into Terrain. The objective is to integrate and use position data from different sources, each operating through an agent, with the possibility that some sources may malfunction, perhaps due to interference by an enemy or saboteur. The possibility of unreliable agents is a new feature. The third application area is to management of the supply chains for a manufacturing process, where the reliability of suppliers has to be judged and possibly new suppliers requested from a Yellow Pages Agent.

The introduction of a Yellow Pages Agent, and also of a Thesaurus Agent to be called in when agents try to communicate using words not mutually understood, give the scheme an intriguing anthropomorphic character and it is possible to imagine a host of agent types offering a variety of brokerage, consultancy and possibly insurance and even espionage services. These wild speculations (not from the book) should not obscure the fact that this is a serious work in which conventions for agent communication are worked out in detail. Belief logics are invoked, and temporal aspects have to be taken into account, so that, for example, agents involved in the Controlled Flight Into Terrain application do not continue to sound warnings that the flight is heading for a mountain when in fact the planned path has been modified and an appropriate change of direction will occur later but in good time.

One chapter is devoted to Meta Agents, by which is meant agents that are part of the model of the environment on which decision making depends. The example used is a military one, where a variation of the Controlled Flight Into Terrain application is applied to a group of helicopters attempting to destroy a group of tanks on the ground. The tanks can be assumed to be using a CFIT system essentially similar to that of the helicopters, and deductions can be made by sirnulating this. The authors of the book seem to consider that such modelling is appropriate only in warlike adversarial situations, but in a milder form it could presumably be part of the operation of the Profiling Agent in the department store application. The idea is particularly interesting since it can be argued that a survival advantage of the evolution of consciousness in people and animals is that it allows the use of one's own reactions to an imagined situation to be used as a model of the responses of others with whom there is interaction.

Two potentially important features seem to be missing from the treatment, though I may have overlooked references to them in the very detailed development that occupies the bulk of the book. One is the assignment of priority of access to agents. It is possible to imagine, for example in the department store example, that a forward-planning agent might decide to run a long series of simulations of situations that might arise, in order to compile statistical data. It would be counterproductive if this were allowed to monopolise the Content Determination Agent and other essential agents to the extent that current customer enquiries were delayed. There also appears to be no provision for automatic self-improvement by learning, either of the internal workings of agents or of their interactions.

Both of these features indicate the need for a feedback of "worth" of the kind visualised by Selfridge[2] in his "Pandemonium" scheme. A generalisation has been introduced by this reviewer as "significance feedback" and will be reviewed in a forthcoming paper.[1]

A large number of AI projects in the past have considered interaction between agents, but the present book introduces and develops this general approach to system design in essentially practical contexts, and is an important contribution. It is surprising to find everything specified in such detail in an introductory text and the book's long-term value must depend on whether the hunches of these authors have been right, but in the meantime it will be accepted as essential reading. A minor criticism of the presentation is that rather a lot of acronyms are sprinkled around without interpretation. A glossary would have been useful.

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

References

1. A.M. Andrew, "Backpropagation", *Kybernetes 30* (2001) (in preparation).
2. G.O. Selfridge, "Pandemonium: a paradigm of learning", **In**: *Mechanisation of Thought Processes* (HMSO, London, 1959) pp. 511–531.

INDEPENDENT COMPONENT ANALYSIS: PRINCIPLES AND PRACTICE, by Stephen Roberts and Richard Everson, Cambridge University Press, 2001, xii + 338 pp., ISBN 0-521-79298-3 (Hardback, £45.00).

This deals with an important new development in statistical mathematics, related to the much older topic of Principal Component Analysis. Each of them operates on a set of data inputs, and PCA provides a means of attributing these to components that are uncorrelated, where correlation is defined as usual in terms of second-order statistical characteristics. As the

https://doi.org/10.1017/S0263574701253491 Published online by Cambridge University Press

editors of this work put it in their introductory chapter: "ICA aims at a loftier goal: it seeks a transformation to coordinates in which the data are maximally statistically independent, not merely decorrelated".

The theory and applications of ICA are developed in twelve chapters, from different authors or groups of authors. The contributions have been specially commissioned with an attempt at unitorm nomenclature throughout, and with a common set of references that is surprisingly lengthy for a relatively new topic. The longest of the chapters, by a substantial margin, is the first in which the two editors give a valuable Introduction, and there is a useful Preface giving an overview of the coverage of each chapter.

In the first chapter the general idea is introduced by reference to the "cocktail party problem" which is that of understanding how a person can follow a conversation under cocktail-party conditions. Where inputs from a number of sources are mixed, to give several mixture signals containing the original components in different proportions, the method is remarkably effective in recovering the original signals from the mixtures. Results of a test are shown, where there were three original components, two of them music of very different kinds and the third white noise.

The introductory example is rather simpler than the full cocktail-party problem since the components are mixed without the time differences that would be introduced by the spatial separation of ears or microphones. The full problem with time separation also receives attention. The methods utilise the fact that the probability density functions of the components may be other than Gaussian. In this they contrast with standard statistical methods in which Gaussian distributions are assumed (and methods are then rated for their robustness against violation of assumptions, particularly of Gaussianity).

Independence of components is defined in terms of mutual information between them, but with a specially rigorous definition of this taking account of probability density functions. Remarkably, efficient algorithms have been developed that operate with modest sample sizes. In the later chapters variations of the method are described that operate in non-stationary situations.

Apart from its use in isolating sources in cocktail parties and other acoustically noisy environments, the technique has various important areas of application. A fairly obvious extension of the application to cocktail parties is to the elimination of crosstalk in communication networks. Another area of application is in biomedicine, with electroencephalography particularly mentioned. In electroencephalography, signals are picked up from a number of electrodes placed over the head and the technique allows isolation of distinct sources of electrical variations contributing to these. Applications have also been made to financial analysis, to identify sources of variation in investment returns.

An application with a less parametric character is to the classification of text documents into groups with independent themes. There are also various applications to scene analysis, including a controversial attempt to analyse human faces as aggregates of a number of basic types.

The method also allows the removal of noise from signals and images in ways that seem to contravene basic ideas of information theory since they operate without a correction channel. The auxiliary information allowing this refers only to an assumed nature of the original signal or image. An example is shown where an image containing printed text and a well-defined decorative pattern is first corrupted by random substitution of pixels and then substantially restored by an application of ICA.

The book is intended to be a self-contained introduction and overview of this important development and it appears to meet the requirement admirably.

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

GEOMETRIC PARTIAL DIFFERENTIAL EQUATIONS AND IMAGE ANALYSIS, by Guillermo Sapiro, Cambridge University Press, 2001, xxv + 385 pp., ISBN 0-521-79075-l (Hardback, £40.00).

As its title suggests, this introduces a rather formidable body of mathematical theory having valuable applications in image processing and computer vision. The first chapter of 70 pages is devoted to providing basic mathematical background, starting with nine recommendations for books that should be on the shelves of anyone involved in this topic. Everything is presented in a lucid style, but it is clear that the topic has many ramifications, indicated by the numerous references to the bibliography for further development of particular aspects. The total number of references listed is 430.

Despite this, in the first part of the chapter the treatment is self-sufficient and accessible without great demands on mathematical background, but when it comes to a section on "Differential Invariants and Lie Group Theory", the author acknowledges that the reader needs to be familiar with concepts of "manifold" and "smooth function", and that in this section results are stated without proofs (but with copious references to where proofs can be found). The chapter ends with an overview of numerical computational techniques, and the orientation is firmly towards producing practical algorithms.

The methods described depend on the "evolution" of geometric curves (in two dimensions) and surfaces (in three dimensions). The curves or surfaces may be "level sets" of higher-dimensional surfaces, and results obtained by J.A. Sethian are included. (Sethian's book, *Level Set Methods and Fast Marching Methods*, was reviewed in *Robotica*, **volume 18**, p. 89.) An impressive example of the use of such a method is illustrated on page 93, where the starting data is the coordinates, in three dimensions, of unconnected points on the surfaces of two interlinking tori, and the method allows a surface to collapse from being an enveloping sphere to one conforming to the surfaces of the tori. The shapes of 3-D objects may be smoothed according to either volume-preserving or surface-preserving flows.

Methods are described for segmentation of images, with an example of a picture of two spanners that is the sort of thing often shown in connection with visually guided industrial assembly robots. An algorithm is described that will maintain a segmentation, originally introduced manually, while the segregated parts move and change shape. An example is the following of two hands moving around, but a more serious application is to the following of structures through successive microscope images of slices of biological tissue.

There are many references to biomedical image processing, and one application is mentioned where a particular object (the image of a neuron) is separated out from a complicated messy image. Other methods are known for this, but the one described allows the user to guide the process by indicating some points on the outline of the required object. This is a useful feature when the image contains uninteresting blobs that could otherwise be reported.

Various ways of treating photographic images are described, some of them describable as "denoising". A way of converting coloured pictures to monochrome versions is discussed, with setting of brightness levels to indicate the 3-D structure much better than could be done by straight colour conversion. In the final chapter on "Additional Theories and Applications", methods are described for removing superimposed text, or unwanted objects, from photographs, with the previously obscured parts of the image appropriately filled. An example is shown of someone diving from a high building, where the original picture shows that his legs are securely bound and fastened to a bungee, but in the "doctored" version no trace of the binding or bungee can be seen.

All of these intriguing and valuable possibilities, and more, stem from the mathematical theory mentioned earlier and this book will certainly be accepted as a standard text. It is written with course teaching in mind, with exercises for the student following each chapter. A student who had worked through everything here

https://doi.org/10.1017/S0263574701253491 Published online by Cambridge University Press

would be very well equipped to develop powerful image processing software.

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

STATISTICAL OPTICS, by Joseph W. Goodman, Wiley, New York, 2000, xvii + 550 pp., ISBN (Classics Library Edition, 2000) 0-471-39916-7, (Paperback, £38.95) (Earlier edition, 1985, ISBN 0-471-01502-4).

SCATTERING OF WAVES FROM LARGE SPHERES, by Walter T. Grandy, Jr., Cambridge University Press, 2000, ix + 370 pp., 0-521-66126-9 (Hardback, $45.00).

Each of these is a comprehensive and authoritative treatment of an area of optical theory that tends to receive little attention in other texts. Neither makes specific reference to application areas having obvious immediate relevance to robotics, but both present fundamental theory that must become relevant as robotic devices depending on visual perception or signalling are pushed towards theoretical limits of performance.

As pointed out by Goodman in the first of the two books, optical theory is often presented as though all relevant phenomena were deterministic, with statistical treatment introduced only reluctantly to allow for such things as atmospheric inhomogeneity or manufacturing defects in lenses. Such an approach ignores important aspects of optical phenomena, including the basic nature of light from a thermal source, which is a superposition of unrelated photons. Light from a laser has a different character, but from real lasers it also has variations of phase and intensity that demand statistical treatment.

The author argues that the statistical approach is the more fundamental and useful, with deterministic theory derivable as a special case. The book is written in textbook style, and originated as notes for a class taught at Stanford University. Each chapter is followed by suggestions for further reading and exercises for the student. Following an introductory chapter, two chapters are devoted to foundations of the necessary statistical theory.

In the following chapters the theory is applied to the analysis of image formation and interference with light of varying degrees of coherence, taking account also of polarisation. Numerous optical experiments and effects are explained in terms of second-order coherence functions, and the difficult concept of high-order coherence is introduced, with applications to interferometric imaging as widely practised in radio astronomy.

The last but one of the nine chapters is on imaging in the presence of randomly inhomogeneous media, with particular reference to the effect of the atmosphere on images formed by optical instruments. The term “speckle” is introduced to refer to a particular kind of granularity that appears in images formed by coherent light when reflecting surfaces are rough on the scale of the optical wavelength, or when transmission media have corresponding defects. A technique termed Stellar Speckle Spectrometry allows partial escape from the effects of atmospheric turbulence on imaging and is described in detail.

The final chapter is on fundamental limits in photoelectric detection of light, and different kinds of interferometry, including speckle interferometry, are discussed in this context.

The book by Grandy is also essentially about optics, and in the reference to “large spheres” the description as “large” is relative to the wavelength of the scattered waves. The book is almost entirely about light waves and water droplets in the atmosphere, and the formation of all the features of the rainbow and of the less common phenomenon referred to as a “glory”, a name given to a colourful halo effect around the shadow of a person or object in mist. The front cover of the book has an attractive detailed picture of a rainbow and on the back is one of the shadow of an airliner, taken from the plane, and showing the “glory” effect.

The glory is a feature of the phenomenon known to climbers as the Brocken spectre, and it is speculated that it could be the origin of the use of halos in religious art. An amusing point is that, when several people stand side by side, each can see all the shadows, but the “glory” effect is so precisely directed that only the observer’s own shadow is seen with a halo.

A possible application of the theory in the robotics field (not from the book) could be to the design of reflective devices having the “corner reflector” property of directing incident light back in the direction of its source. Reflective adhesive tapes are available having this property, achieved by incorporation of transparent spheres with appropriate refractive properties. However, the detailed mathematical treatment in the book certainly goes well beyond what would be needed to analyse this effect.

Both books are authoritative and thorough and well presented.

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

https://doi.org/10.1017/S0263574701253491 Published online by Cambridge University Press