"engineering" view, which tends to concentrate on what is technically possible. Kolodner attempts to satisfy both communities by providing a core of practical detail, interspersed with references to cognitive modelling. This works well—making it possible to navigate through the chapters according to taste. Given this subtheme, the main structure of the book is based around the architectural components of case-based systems. An early chapter takes us through some of the representational decisions which must be made in structuring a case library. This is followed by a discussion of the choice of vocabulary for indexing cases. The narrative then moves to considering how features may be organized to make retrieval of cases easier, then introduces heuristic algorithms for assigning "best" matches on the basis of strength of similarity. Having chosen a best matching case, the next step is usually to adapt it to match the current problem, so we find chapters on common adaptation strategies and on using explanations of failed cases to guide the repair process. The sequence of chapters on representation, indexing, matching and adaptation are bolstered by introductory chapters which set the general context for case-based reasoning, and a penultimate chapter on the practicalities of making the various components work effectively when combined.

Looking at the list of computational mechanisms used in the book (frames, discrimination networks, causal models, semantic networks, model-based reasoning, and so on), one is struck by the diversity of implementational techniques, and also that most of them are familiar from other KBS areas. Thus, if we simply look at the implementational components of many case-based systems we will find little of novelty. However, this is to miss the point. The contribution of case-based reasoning is not in the implementation methods themselves, but in the way in which these are employed to tackle problems in difficult domains. The reuse of previous examples for problem solving can often be achieved, in appropriate domains, using comparatively standard techniques. The trick, which Kolodner attempts to teach us, is to approach domain problems in a way which makes this form of attack successful. In her concluding chapter, Kolodner cites "scaleup" (by which she means increasing the size of case libraries) as the major technological issue for future research. However, of at least equal importance is the need to understand more fully the formal basis on which case matching is undertaken, and to develop a body of experience which permits us to judge reliably when a given problem warrants a case-based approach. These problems are not purely technical but, as Kolodner points out, the production of improved tools and design methodologies can help. What also helps is a thorough understanding of the extent to which case-based reasoning actually works in realistic domains; the book provides a detailed summary of over 40 case studies but (since each is tightly limited in application), it would be instructive to know what the limiting factors were in each domain.

Kolodner's book is invaluable as a guide to the variety of case-based reasoners which have been built. It contains a wealth of practical advice on how to approach the design of such systems, and tempers the obvious enthusiasm of the author for her subject with warnings of potential pitfalls. As the foundation for a teaching course, its only failing may be that it contains no demonstration programs—but perhaps these will appear in a second edition. As it stands, this is the most comprehensive and best written survey of case-based reasoning which I have seen. I recommend it as an addition to any KBS library.

Reviewed by Dave Robertson, University of Edinburgh, UK

**Decentralized A.I. 3** edited by E. Werner and Y. Demazeau, North-Holland, Amsterdam, 1993, pp 356, hardback, ISBN 0 444 89661 9, Dfl 250.00.

There has been an increasing amount of research in multi-agent systems (MAS), including human, automated and hybrid systems. These systems may cooperate or compete with each other, or they may simply be indifferent about each other's activities. However, if agents have inter-related individual goals, which may contribute to or contradict each other, or if they have joint goals, issues

such as communication, cooperation and coordination are extremely important for coherent social behaviour. To investigate these issues, the research community in MAS has devoted considerable effort to the modelling of agent internal states (beliefs, goals, intentions and commitments, desires, and so on) as well as organizational control structures.

This book includes a collection of papers presented at the *Third European Workshop on Modelling Autonomous Agents in a Multi-Agent World.* Papers are grouped under various headings largely based on an agent-oriented viewpoint: Communication, Psychology of Agent Plans and Intentions, Agent Psychology of Evaluations, Multi-Agent Organization, Reactive and Deliberative Agents, Emergent and Real-Time Systems. Based on this view, DAI issues such as cooperation, coordination and coherence are explored throughout the book.

The editor first gives an introduction to MAS issues and overview of the authors' contributions. He presents a multi-layer view of agent design: an object level of agent system and meta-observer levels of user, application programmer, system programmer and designer. Corresponding modes of interaction are identified and each mode has its own language.

Chang and Woo's paper describes a communication protocol for negotiation, based on speech act theory. Rao, Georgeff and Sonenberg then propose a notion of social plan and a formalism for representing social plans, assuming that agents have a joint goal and a joint plan before forming a joint intention. Burmeister and Sundermeyer present a "generic" agent model, a taxonomy of behaviour and argue that control and cooperation should be based on intentions, resources and behaviour. Jennings then proposes a notion of joint commitment as the basis for joint responsibility, which is an extension of joint intentions. A joint responsibility defines preconditions for joint actions and how agents should behave once such an action is started (committed). The issue of belief revision in MAS is investigated by Dragoni. Information is associated with a source and each source has a credibility factor.

In a group of papers on Agent Psychology of Evaluations, Kiss and Reichgelt first present a model theory for agent hedonic states, preferences, goals and values based on a physical world view, instead of a possible world view. Levy and Rosenschein then propose a method for incorporating global goals into the local interests of agents, based on game theory techniques. The method is applied to solve the pursuit problem. Kraus and Rosenschein suggest that agents discover focal points to coordinate their activities. They argue that different ways of representing knowledge affect agent ability to find focal points. They then propose a focal-point algorithm based on step logic and analyse its advantages over other conventional approaches such as game theory, standard logic and decision theory.

Under the heading of Multi-Agent Organisation, Osawa and Tokoro extend the contract net protocol to allow contracting agents to construct collaborative plans from their individual plans based on their own, possibly incorrect beliefs. Areas of collaboration such as the removal of obstacles are identified by the manager and communicated to appropriate contractors. Papers by Numaoka and Gasser then discuss issues in dynamically forming organizations, and propose ways of doing so. Castelfranchi, Miceli and Cesta argue that dependency relations are the basis for agent influence and interaction.

In the Reactive and Deliberative Agents section, the relationships and tradeoffs between low-level reactive agents and high-level deliberative agents are investigated by Kiss, Ferguson and Latombe in separate papers. Kiss and Ferguson propose a multi-layered agent architecture to combine closely coupled, situation-specific agents with loosely coupled generic agents. Latombe finds a middle ground between fully reactive planning and centralized planning through the exchange of individual plans and coordination prior to plan execution.

In the last section, Emergent Systems are explored by Drogoul and Dubreuil, Wavish, Boissier and Demazeau, Gambardella and Haex in three separate papers. Colliot and Hayes-Roth then present an algorithm for achieving real-time performance.

On the whole, the book provides a good collection of research papers in various areas of MAS. It is not a textbook for beginners in the field, but a highly technical and informative book for researchers who have made a reasonable start in MAS.

Reviewed by Jun Huang, Advanced Computation Laboratory, Imperial Cancer Research Fund, London, UK

**Probability and the art of judgement** by Richard Jeffrey, Cambridge University Press, 1992, pp 244, £12.95 (p/back) or £37.50 (h/back), ISBN 0-521-39770-7.

For anyone who is interested in decision theory, and like me finds the name of Richard Jeffrey far more familiar than his work, this volume of essays is essential reading. Despite being a collection of pieces, and therefore lacking a common narrative thread to tie the various ideas tightly together, it nevertheless manages to give a comprehensive treatment of the topics with which he is concerned. Broadly speaking, these are all aspects of the use of probability and utility theory to make decisions in the face of uncertain information, but Jeffrey's treatment differs from that of other Bayesian probabilists in two important ways. First, Jeffrey approaches the problem very much from the philosophical rather than technical point of view; he is, after all, a philosopher. Second, he gives considerable space to discussing whether the theories provide a reasonable model of human, rather than purely rational, decision making. As a result, the book is always very readable, and the theories he advances more convincing than those built from purely technical considerations.

In total, the book contains 16 essays, the earliest of which was written in 1956, and the most recent of which dates from 1991. Thus, what the book provides is a historical record of Jeffrey's ideas gathered together in one handy volume and ordered so that similar topics are covered in adjacent essays. In addition, through the new introductory essay, and the postscripts to a couple of the others, recent developments are also acknowledged. As a result, this is more than just a collection of papers, and is instead a coherent whole, albeit one that is more of a palimpsest than a conventional monograph.

In brief, the topics of the individual essays are as follows. The introduction argues that the position established by the whole collection is but a development of Carnap's views on probability, and represents a generalization of logic deductive methods in which the concepts of true and false are replaced by a continuum of states of belief. The second essay then addresses the question of scientific reasoning in the light of this position, and reaches the conclusion that the best decisions will be reached only if hypotheses are not merely accepted or rejected, but are assessed probabilistically—exactly the kind of conclusion one would expect from an advocate of Bayesian techniques. However, to be fair, it should be noted that Jeffrey also admits that there are considerable problems with this view, not least among which is the fact that this is clearly not what happens in practice.

The third essay introduces the important idea of probability kinematics. This is a means of updating probabilistic information in the light of uncertain information, and is a generalization of conditioning, which only permits updating in the light of certain information. Clearly the use of probability kinematics makes if possible to carry out probabilistic updating in a much broader set of situations than conditioning, and Jeffrey uses this fact to argue that probability kinematics provide the "probabilism of the commonplace". Then, in the fourth essay, he uses it to form the basis of an account of scientific method. This account is extended in the fifth essay with a description of how scientific theories can become more probable by explaining previously known facts (the tricky problem of old evidence), and how incomplete and indefinite probabilistic information can be usefully employed. Further work on probability kinematics is contained in the sixth and seventh essays. The first of these discusses the order of revisions, and proposes a solution to the problem of revision in the face of evidence in which, unlike other proposals, the order of two successive revisions does not effect the final probability. The second looks at the connections between probability kinematics and subsequent developments in Bayesian statistics, showing various equivalences between the approaches, and further developing the theme of successive updating.

The eighth essay introduces Jeffrey's third major theme after scientific reasoning and probability kinematics—the study of preferences and their representation. Indeed, with the exception of the eleventh essay, which is a convincing demolition of the frequency view of probability, and the