possibility theory and fuzzy logic), the issue of underlying semantics is ignored. Semantic issues do rarely appear in such papers, because, in my view, these approaches have little or no semantic foundation. I am sure that had these issues been raised in the papers of the non-additive advocates, they would have appeared in this book. Nonetheless, it seems that if the issue of semantics is important in other approaches, the authors might have pointed out the absence of such discussion in the non-additive literature.

Another result of the reportial approach of this book is that the reader might well conclude that probability theory is the most discredited of all the representations of uncertainty. Though they do claim in the beginning of the book that probabilistic techniques represent the "gold standard" of uncertainty, in the remainder of the book probability is simply a target of criticism. Again, this is because of the nature of the uncertainty literature. In this literature, it is fashionable for one to show the value of his or her technique by discussing the aspects in which it overcomes perceived shortcomings of probability. Rarely does one mention the aspects in which probability may be superior, or provide comparisons of one's own theory to other, non-probabilistic, techniques. Likewise in the literature (though hardly in private discussions!), the fashion of probability advocates is to ignore competing theories.

Despite these shortcomings, I heartily recommend this book to those interested in a broad survey of the field. The shortcomings are primarily reflective of the state of the literature in the field. The strength of the book lies in its broad coverage and its understandable presentation, well strewn with good examples.

Received by John F Lemmer, Clinton, NY, USA

From animals to animats 2 edited by Jean-Arcady Meyer, Herbert L Roiblat and Stewart W Wilson, MIT Press, 1993, pp 523, $55, ISBN 0-262-63149-0.

It is time animats lived side by side of animals. They can live, learn, build maps, evolve and be our domesticated pets, too. In the intersection of AI, robotics, philosophy and psychology, traditions of building ever more complex systems that try to exhibit intelligent behaviour or solve puzzles are giving way to systems that live simple lives, without pretending to reason deeply, or to act as if they were real animals. Instead, the system designer builds simple agents and endows them with basic shallow abilities to interact with the world. The point is that the designer doesn't solve all the problems that the system may encounter beforehand. Nor does the designer build a system that acts as if it is trying to talk to itself outloud. Many of the traditional symbolic systems are designed so that they can be made to generate a transcript of human understandable steps they take to decide on performing each and every action. Once deployed in a complex world, the animat (just as a device would) uses its repertoire of interactive behaviours as they become needed, e.g. a microwave oven responds to buttons pushed. The interactions that follow may appear intelligent. Such agents can also be made to learn from these interactions and enrich their range of capabilities.

This volume is an excellent collection of papers to represent this shift in paradigm. Due to space limitations in this review, I will refrain from commenting on each paper. Instead, I will comment section by section. This field hasn't reached maturity and must allow for various approaches. On the whole, the editors have done a good job of setting up sections with some confusion and overlap among action selection and behaviourial sequences, evolution and learning.

The four introductory papers in the animat approach to adaptive behaviour herald the theme of the paradigm shift through the rest of the papers. There is some level of publicising terms and jargon. For example, Maes calls the paradigm shift "behaviour based AI". The paradigm shift is characterized by learning, adaptation and lack of goals directedness.

One of the themes of the paradigm shift is the need to model immediacy of sensing to acting or close knitting of sensations to motor control. Papers in perception and motor control present methods for making these relatively close connections. Presentations here are varied but mostly use neural networks. This seems to be an accurate representation of the concentration of approaches in the research community. However, we believe discrete approaches such as in Horswill's, by the nature of their explanatory nature can help shed light into the circuitry of complex behaviours.

Building internal models of the world and especially acting based on these models is frowned upon by researchers who advocate moment to moment interactions generated through situated activity mechanisms (e.g. Agre & Chapman, 1990; Brooks, 1990). However, we can't deny that humans are capable of building and using internal models. Let's assume that we can set aside the issue of how often humans build internal models and when they use them. Papers in cognitive maps and internal world models describe an approach that allows for adaptive behaviour, develop maps for mobile robots taking into account the robot's experience in navigation, discuss issues related to various ways an agent may represent routes, describe how large scale maps for when places don't share common views can be constructed using a barycentric coordinate system, and describe a neural network approach that simulates how a robot learns to approach desirable places and avoid undesirable places.

The section on learning contains presentations from inductive learning approaches, namely reinforcement based learning (RBL). Several papers discuss various ways of using and scaling RBL to complex domains. The appeal for RBL is that it is an unsupervised learning technique, very little domain knowledge is required, and it is mathematically sound. Nehmzow and colleagues show learning behaviours using associative memories and rules to be maintained by the robot they call *instinct* rules. It is impressive how their Lego robots learn behaviours to avoid obstacles, follow walls, and follow corridors. Having worked with Lego vehicles, one is impressed that such robots with limited structure and computation learn their tasks before self-destructing. Experiments with physical robots are superior to simulations and we are quite impressed with work describing learning using physical robots. Colombetti and Dorigo use classifier systems on a physical robot to learn complex behaviours as composition of simple behaviours.

In the section on evolution, stated by Harvey, Husbands and Cliff, "An animal should not be considered as a solution to a problem posed 4 billion years ago." This position advocates devising evolution mechanisms that use the agent's own physiology to adapt to their world and exhibit intelligent behaviour with respect to their own physiology. Papers in evolution have in common the themes encountered in the artificial life (a-life) community.

Papers in collective behaviour describe experiments with many interacting robots engaged in group activity. These papers explore architecture and/or behaviours of each agent that contribute to collective behaviours. Sometimes, interesting behaviours from local interactions that are not programmed into agents are observed, called emergent behaviours.

In the one page summaries at the end of proceedings we see a diverse array of papers that belong in a-life. Clearly, these papers belong here and it is a nice transition to a-life research. A-life has attracted scientists from varying disciplines. What is in common is the creation of creatures that appear lifelike. These creatures live and may learn, reproduce and die. Many of these creatures are given habitats and studied for their behaviour and evolution. It is yet too early in a-life to identify common techniques for creature building or study of them. However, there is a general emphasis to produce physical systems that interact in realistic situations. There is a trend toward building inexpensive and simple robots with very simple behaviour engines and embed them in a physical environment. In so doing, external observers note interacting behaviours that appear intelligent and unexpected, i.e. emergent behaviours.

References

Agre, P and Chapman, D, 1990. "What are plans for?" In: P Maes (ed.), *Designing Autonomous Agents*, Chapter 2, 17–34. MIT Press.

Brooks, RA, 1990. "Elephants don't play chess". *Robotics and Autonomous Systems* **6** 3–15.

Reviewed by Henry H Hexmoor, Department of Computer Science, SUNY at Buffalo, USA