Chapter.6, on reasoning about actions, is not much better. Thanassas tries to distinguish between facts and events, and exemplifies the former by *loud*(*noise*, t_1) (which is presumably meant to say that there is a loud noise at time t_1, though the only way I can interpret it is as saying that the entity "noise" was loud at t_1, which is rather different), and the latter by, *inter alia*, *love* (x, y, t_1, t_2). Of the latter he says that "since if x loves y in [t_1 . . . t_2] then x loves y at any time point t between t_1 and t_2, this event conforms to the definition of a fact but does not conform to the definition of an event", which makes me wonder why he wants to classify it as an event in the first place.

The confusions pile up thick and fast. Thanassas speaks of "the fact *loud* (*noise*, t_1)" but he also refers to a fact's "changing its truth value". How can the fact that there was a loud noise at t_1 change its truth value? This is what Quine (1960) calls an *eternal sentence*: its truth value does not depend upon the time at which it is uttered. Some of the problems arise, perhaps, from thinking that there is a meaningful dichotomy between facts and events. To my mind, the significant distinction is between *states* and events; and *both* of these can participate equally in facts, for example the fact that it is now dark (a fact about a state) and the fact that the sun set an hour ago (a fact about an event).

At this point I did not see much point in reading further. I hesitate to condemn the *substance* of the book as worthless, although as I have indicated it contains many gross conceptual confusions. The problem is that the material is so badly presented that it is sometimes impossible to see what the author is trying to say. There are faults both in the way the material is organized and at the level of detail; there are distracting misprints and spelling errors on every page. This reflects badly on author and publisher alike.

Indeed, I believe that the publisher should accept the major part of the blame on this score: the author is evidently not a native English speaker, and is inexperienced in publishing. Yet there is no sign that the book has been read through by a copy editor: the impression one has is that the author has submitted camera-ready copy and the publisher has printed it without further ado. Whether I am right or wrong in this surmise makes little difference: a book with this many errors should never have been published, and it is surely in the publisher's own interests to impose some more rigorous quality control than appears to have been exercised in this case.

Speaking of errors, and returning now to the material itself rather than the presentation: who originated this nonsense about Australian penguins being able to fly (p.42)? How can a scientific researcher seriously work with a sentence such as "red birds are usually Australian penguins"? Of course, these are only light-hearted illustrative examples and not really about penguins. But one of these days, someone is going to read this sort of stuff and come away believing that Australian penguins really do fly, or that they are red. There is an intellectual irresponsibility here that should not be allowed to pass without comment. It betrays a cavalier approach to reality which indicates that its perpetrator is working in a fantasy world where real facts do not matter: it is all just an intellectual game.

References

Galton, A, 1988. "Review of Y. Shoham *Reasoning about Change*". *Artificial Intelligence Review* **2** 285–288.
Galton, A, 1991. "A critique of Yoav Shoham's theory of causal reasoning". In: *Proceedings of the Ninth National Conference on Artificial Intelligence* AAAI Press/MIT Press, pp. 355–359.
Moore, RC, 1985. "Semantical considerations on nonmonotonic logic". *Artificial Intelligence* **25** 75–94.
Quine, WVO, 1960. *Word and Object*. MIT Press.

Reviewed by Antony Galton, Department of Computer Science, University of Exeter, UK

Do the right thing—studies in limited rationality by Stuart Russell and Eric Wefald, MIT Press, Cambridge, MA, £24.75, ISBN 0-262-18144-4.

As the cover blurb of this book points out, most intelligent systems face the same dilemma as the hero of Spike Lee's film—knowing just what is the right thing to do. Previous theories of rationality

are fine as far as they go, but they clearly do not go far enough. They assume that the decision maker has perfect information, and, even more unrealistically, has unlimited resources that may be applied to the problem of determining that elusive "right thing". In contrast, Russell and Wefald address the far more challenging problem of establishing the best course of action under severe resource constraints, one of their applications being the construction of systems capable of reasoning about the best way to proceed in a game of table-tennis, in real time.

Such an aim is clearly very ambitious, and it would be unrealistic to claim that the book does more than lay some of the groundwork for such an endeavour, but even the groundwork makes fruitful reading. What Russell and Wefald do is to take the domain of two player games with perfect information, that is games such as chess, go and othello, which have been widely studied already in artificial intelligence, and look to see what computational advantage can be gained by replacing blind heuristic search with a search based upon what they call "rational metareasoning". What this means is that before spending any time evaluating the usefulness of a particular node in the search tree, the system deliberates a little to ensure that it is worth doing. Now, of course, in any normal search there is some such deliberation when algorithms such as A* are used to try to ensure that only nodes that are liable to lead most cheaply to a goal are expanded. However, simple algorithms such as A* are easily confused, and the main idea behind "Do the right thing" is that more deliberation, and thus more expenditure of resources upfront, leads to much better selection of nodes to expand, and a much better payoff in terms of the decisions that it is possible to reach in a limited time.

This point is borne out in practice. The algorithms developed in the book have been tested out playing othello, on classic search problems such as the 15-puzzle, and (by Dickson, 1991, at Oxford) on a problem from computer vision. In the first case, the method outperformed existing search algorithms, and produced similarly encouraging results in the other domains.

Before I get too dewy-eyed about this book, or at least give the impression that I am, I should point out that I have at least one reservation about the approach. That is, that the kind of deliberation that the method applies in order to choose the best alternatives is based on classical decision theory. To me this seems rather counter-intuitive. Granted, classical decision theory, in which every alternative is assigned a probability of happening and has its utility of occurrence assessed, is a well established model for making "rational" decisions under conditions of uncertainty, but it is also well established that it has a number of problems which largely relate to the establishment of these numerical probabilities and utilities. Indeed, establishing the numbers is in many ways very similar to the problem of deciding what the best node to expand is—it is very easy to do if you assume an omniscient agent with unlimited time (or, alternatively, compile the answer into your intelligent system), but is very difficult to achieve on the fly with limited resources.

Having said this, it should be noted that Russell and Wefald do acknowledge that there is a problem, considering ways in which their method can be augmented to learn the best parameters when given some initial distribution that enables it to muddle its way, albeit less than perfectly, to some kind of initial solution. This is a thoroughly sensible approach and cannot really be faulted. I would rather see an approach that admitted the flaws in classical decision theory and tried to overcome them, but then perhaps I should go out and pursue that line of work myself.

References

Dickson, W, 1991. "Image structure and model-based vision". DPhil thesis, Department of Engineering Science, University of Oxford.

Reviewed by Simon Parsons, Advanced Computation Laboratory, Imperial Cancer Research Fund, London, UK.