The Knowledge Engineering Review, Vol. 15:4, 2000, 411–412
Printed in the United Kingdom

Book review

Cambrian Intelligence: The Early History of the New AI by Rodney A. Brooks, MIT Press, ISBN 0-262-52263-2

To paraphrase a colleague, from reading Rodney Brooks's work one gets the strong impression that he is not the kind of person to worry away nights agonising about whether his view of the universe might possibly be flawed. Given the furore that his work first aroused within the AI community, this is probably a good thing—a character of less conviction would probably have abandonned his work on behaviour-based robotics at an early stage, leaving AI as a subject considerably poorer. As it is, some fifteen years have now passed since Brooks first stirred up the storm and, having weathered it, he has collected some of the major papers from the early stages of his work into *Cambrian Intelligence*. Roughly speaking, these run from the first paper on the behaviour-based approach (though it does not use the name) to his Computers and Thought award paper.

One of the interesting things about this collection, for me at least, is the fact that in the year 2000 one has to try very hard to imagine quite how scandalous Brooks's ideas were when they first emerged. Now it is common to find agent architectures which incorporate at least some elements of the behaviour-based approach, and even if Brooks himself might dismiss such architectures (on the grounds that they include elements of more traditional AI), the fact that nobody questions the inclusion of these elements goes to show the extent to which Brooks's ideas have entered the mainstream. However, in 1984 it was a different story. Then the dominant view, if not the only view, in AI was that agents would perceive their environments, translating these perceptions into symbols, reason about these perceptions by manipulating the symbols and then act on their conclusions. This partitioned the world neatly into those who "did" perception, vision in the main, those who did reasoning, in particular planning how to move around the world, and those who built systems which actually did the moving around. Such systems were typified by the SRI robot Shakey (although Shakey was considerably earlier than this the basic picture had not changed), which observed the world through a video camera, analysed the image to determine what it was "looking" at and reasoned by logical theorem proving to determine where it was and where it should go next. Each component was complex, and all (plus the various components which took the plan and converted it into signals to the robot's wheels) needed to be in place before the robot could do anything.

What Brooks suggested, in effect, was that one could proceed in an orthogonal way. First, he suggested, you find the simplest thing that you want your robot to do—wander aimlessly around while not bumping into things, for example. Then you program the robot to do exactly this, to go forward, making changes of direction at random unless its sensors indicate an object ahead, in which case it should veer away from the object. Once this behaviour works, find the next most simple behaviour—exploring by trying to reach distant objects—and do the same. The two behaviours are then run in parallel, though the "higher" level of behaviour is allowed to subsume the "lower" level when it wants to take control. The overall effect of this two-level subsumption architecture is a system which tries to reach distant objects unless it detects something it might crash into, in which case it veers away from it. This, of course, is a combination of the behaviours of the two levels, and the big idea that Brooks was pushing in his early papers was that arbitrarily complex behaviours could be built up from combining simpler behaviours in exactly this way. The attraction of this idea was that it was possible to get robots up and running without having to solve the big problems of vision and reasoning that AI was struggling with. However, because this approach

challenged the conventional position, it seems to have met with a remarkable degree of opposition, documented now with glee by Brooks in his introductions to the various papers in the collection.

When faced with the kind of disbelief and scorn that greeted the early work on the behaviour-based approach, one has the choice of abandonning the ideas which provoked this response, or going away and proving that they work, and the latter was exactly what Brooks did. Across the period documented by *Cambrian Intelligence* his group at MIT built a number of different robots of, broadly speaking, increasing complexity, from a simple machine with the behaviour mentioned above, to one which could roam the lab collecting empty drinks cans (which involves a good number of additional subsumption layers), and this array of robots provided powerful support for the validity of the approach.[1] This weight of evidence, it seems to me, is sufficient to show that the subsumption architecture is one of those ideas which will be found somewhere in any truly intelligent system which has to operate in real-world environments. However, I remain to be convinced that building such systems using only the subsumption architecture is necessarily the easiest way to proceed.

In summary, therefore, I think that the ideas contained in *Cambrian Intelligence* should be part of the armoury of any AI researcher, and it is good to see them collected together in a book. Indeed, my only gripe about the book as a whole is that because the papers it contains were, obviously, meant to be read in isolation, there is a considerable amount of overlap between them and, as a result, the book is not nearly so satisfying as a unified treatment of the material. However, such a treatment would no longer be a "history of the new AI" and so we would lose the rather fascinating overtones of the struggle against the AI establishment which illuminate the book as it stands.

Reviewed by Simon Parsons, Department of Computer Science, University of Liverpool, Liverpool, UK

[1]Though he doesn't suggest this, it seems to me that in this approach Brooks also forshadowed the move towards good empirical science in AI, but that could be because he worked in robotics, and roboticists have always understood the importance of experiments better than the rest of us.