

speed of a steam engine: "[It] is preferable to the Turing machine as an archetype for models of cognition." I would like to believe that this turns out to be a barren program.

To summarize: Aaron Sloman and John McCarthy, two major proponents of fundamental research in AI, demonstrated that realistic AI requires arming a program with a philosophy. Their IJCAI '95 papers ("A philosophical encounter," available at http://www.cs.bham.ac.uk/~axs/cog_affect/ijcai95.text, and "What has AI in common with philosophy?" available at <http://www-formal.stanford.edu/jmc/aiphillaiphil.html>, respectively) consolidated the need for deeper inquiries into foundations by suggesting numerous pathways for growth. In the words of McCarthy, "AI needs many ideas that have hitherto been studied only by philosophers." But McCarthy also notes that "many philosophical problems take new forms when thought about in terms of how to design a robot." He then issues a warning: "Some approaches to philosophy are helpful and others are not." As I have already stated, *Mind Design II* places heavy — some would think unusual, even biased — emphasis on connectionism, neurocomputing, situated AI, and yes, dynamic systems theory. In the words of Margaret A. Boden (from the back cover of *Mind Design II*), "which of these is the most stable, and whether the foundations need to be re-worked, are questions readers will be eager to explore." Haugeland uses the umbrella term NFAI (New-Fangled AI) to mean scientific mind design that is not GOFAI (Good Old-Fashioned AI, that is, thinking as internal symbol manipulation). Connectionism directly falls under NFAI, and several other possibilities such as the embodied mind design present themselves as worthy candidates.

A few observations on the structure of the book: There are 22 pages of references in *Mind Design II* compared to 14 pages in *Mind Design*. I found them

to be quite complete. I also liked the fact that each reference has pointers indicating in which chapters it was cited. All the figures in *Mind Design II* have been redrawn, rendering a more harmonious look. There are quite a few typos, but they are not of the confusing kind. A valid complaint could be made that there are no name or subject indices; these would indeed be very useful in a book of this size. Finally, while the cover illustration is fine, both the color and the quality of the cover could be improved.

I would like to end this review by explaining the relevance of the opening stanza of Perloff as follows. *Mind Design*, contained the seeds of a second edition, and *Mind Design II* shows that at least in this particular case the destruction didn't take place; what we have is as accomplished as *Mind Design*. And clearly in *Mind Design II* there are the seeds of a possible "Mind Design III." I, for one, will look forward to reading that third revision.


References

- Churchland, Paul M. 1989. *A Neurocomputational Perspective: The Nature of Mind and the Structure of Science*. Cambridge, MA: MIT Press.
- Dennett, Daniel C. 1987. *The Intentional Stance*. Cambridge, MA: MIT Press.
- Dreyfus, Hubert L. 1979. *What Computers Can't Do: A Critique of Artificial Intelligence*. New York: Harper and Row. Second edition.
- Guttenplan, Samuel, ed. 1994. *A Companion to the Philosophy of Mind*. Oxford: Blackwell.
- Haugeland, John, ed. 1981. *Mind Design*. Cambridge, MA: MIT Press.
- Haugeland, John. 1985. *Artificial Intelligence: The Very Idea*. Cambridge, MA: MIT Press.
- Haugeland, John. 1997. *Having Thought*. Cambridge, MA: Harvard University Press.
- Nadel, Lynn, Cooper, Lynn A., Culicover, Peter, and Harnish, R. Michael, eds. 1989. *Neural Connections, Mental Computation*. Cambridge, MA: MIT Press.
- Posner, Michael I., ed. 1989. *Foundations of Cognitive Science*. Cambridge, MA: MIT Press.
- Ramsey, William, Stich, Stephen P., and

Rumelhart, David E., eds. 1991. *Philosophy and Connectionist Theory*. Hillsdale, NJ:

Lawrence Erlbaum.

Winston, Patrick H., ed. 1975. *The Psychology of Computer Vision*. New York: McGraw Hill.



Artificial Intelligence and Mobile Robots: Case Studies of Successful Robot Systems

D. Kortenamp, R.P. Bonasso, and R. Murphy

The MIT Press
Cambridge, MA, 1998

390 pp.

\$37.50

ISBN 0-262-61137-6

Reviewed by: John Clarke

Instituto de Sistemas e Robotica
Universidade de Coimbra
Polo II
Coimbra 3030 Portugal

Artificial Intelligence and Mobile Robots is a compilation of 13 chapters by well respected robotics researchers describing some of their most recent work. The book aims to provide a "how-to" guide to building the control system of a mobile robot. These case studies are definitely worth studying: seven of the 13 have placed highly in various robotics competitions. In the introduction, the editors give a good, yet brief, history of

intelligent robots. The book is then divided into three sections.

The first section, "Mapping and Navigation," could easily be titled "Sonar." It is packed with ideas and pseudo-code and is a superb guide to adding this type of sensing to a robot platform. Map making and navigation are classical problems in mobile robotics and a number of theoretical solutions exist (e.g. Lumelsky's "Cocktail party" model). This section presents methods which are both well founded and practical. For example, both occupancy grids and topological maps are discussed in some detail.

The first chapter, "Map Learning and High-Speed Navigation in RHINO" (Thrun et al.), gives a comprehensive account of an effective map learner and path planner. Koenig and Simmons in their description of the "Xavier" robot take a similar approach, but they use a partially observable Markov model to solve the localization ("Where am I?") problem. The "Dervish" robot (Nourbakhsh) with its novel sonar layout and purely topological mapping is also described, as is the "Carmel" robot (Kortenkamp et al.), which employs a reactive path planner. In short, this section is absolutely indispensable reading for anyone new to mobile robots.

The second section, "Vision for Mobile Robots," is a useful introduction to the types of vision algorithms that have proved themselves useful in practice. The methods described are not state-of-the-art computer vision algorithms. They are, however, simple and effective. The primary goal of the vision systems described here is to (help) define a driveable region. This is principally achieved by color segmentation, but also by blob and edge detection. These algorithms lack the sophistication of the advanced structure from motion algorithms. Indeed, they lack the sophistication of Moravec's Stanford cart. However, they make up for this by achieving high frame rates.

The two questions of how much of the world to represent, and in what form to represent that world, have long challenged robotics researchers. Early attempts to faithfully model the global environment floundered and were soon challenged by the fast but low-level "reactive planners." The current consensus is that both types of control are necessary and must be somehow integrated into a coherent whole. The third section, "Mobile Robot Architectures," is devoted to this topic. The first chapter, "Three-Layer Architectures" (Gat), sets the scene masterfully with a full discussion of why three kinds of control (fast obstacle avoidance, task sequencing, and overall planning) are necessary as well as the role of the robot's internal state or memory. While the remaining chapters in this section are interesting as descriptions of working fielded systems, Gat's chapter also provides a clear and succinct theoretical analysis.

The final chapter, "The Phoenix Autonomous Underwater Vehicle" (Brutzman et al.), is also noteworthy. Not only is the underwater domain an especially demanding one, but even the physics of simple motions are non-intuitive. There is a great opportunity to learn about the fundamental aspects of robot control architectures by comparing systems operating in very different environments. It would have been nice if the authors had included such an explicit comparison. For example, the Phoenix controller is a particularly clear example of the three layer architecture, and while the authors state that the architecture allows responses on a range of time scales, it is the analogy with the organization of the crew on naval submarines that is pursued at length. Your reviewer was left wondering whether the architecture of this robot, and by implication other three-layer robots, has been unduly influenced by the sociology of human hierarchies. To summarize this section: it succeeds as a how-to manual, but one is left wondering "Why do these robots

book reviews

work?" The reader will have no trouble duplicating the machines described, but there are few clues as to how one might improve on them.

In summary, *Artificial Intelligence and Mobile Robots* is a thought provoking, if somewhat eclectic, snapshot of current robotics, concentrating on the development of competition-winning robots. (Will robots ever do anything other than clean up trash?) It does not try to be, and is not, a history of intelligent robots, although at times I wondered if the successes and lessons learned might not be more clearly presented set against past failures and dead ends. I found the book a pleasure to read. Judging by the way the review copy mysteriously found its way to desks far across the lab, this is a book to own.

Mathematical Methods for Neural Network Analysis and Design

Richard M. Golden

A Bradford Book

The MIT Press

Cambridge, MA, 1996

432 pp.

\$67.50

ISBN 0-262-07174-6

Reviewed by: Toshiro Kubota

Department of Computer Science

University of South Carolina

Columbia, SC 29208

kubota@cs.sc.edu

This book describes various neural networks and their behavior based on David Marr's theory of vision and information processing systems [Marr82]. The author, Richard Golden, is an Associate Professor in the Program of Cognition and Neuro-science at the University of Texas at Dallas. He has published a number of