

# Review of *Artificial Intelligence and Mobile Robotics: Case Studies of Successful Robot Systems*

Chris Brown

In the 1970s, a schism developed between (symbolic) AI and robotics (including computer vision). Today, mobile robotics is an increasingly important bridge between the two areas. It is advancing the theory and practice of cooperative cognition, perception, and action and serving to reunite planning techniques with sensing and real-world performance. Further, developments in mobile robotics can have important practical economic and military consequences.

For some time now, amateurs, hobbyists, students, and researchers have had access to how-to books on the low-level mechanical and electronic aspects of mobile-robot construction (Everett 1995; McComb 1987). The famous Massachusetts Institute of Technology (MIT) 6.270 robot-building course has contributed course notes and hardware kits that are now available commercially and in the form of an influential book (Jones 1998; Jones and Flynn 1993). The MIT-inspired materials concentrate mainly on Rodney Brooks's subsumption architecture. This approach to robot control relies on simple reactive behaviors, notionally running in parallel, with only simple inhibitory connections between them. Actions are determined by the sensed state of the external world, but internal representation of the world is held to be unnecessary. With these tenets, it is possible to produce robots with reliable, robust, compelling, but arguably limited, repertoires of behavior.

Indeed, pure subsumption robots

have not yet demonstrated the ability to perform complex real tasks. Practical autonomous systems tend to use more complex, multilevel organizations of robot capabilities. Common components are a deliberative, strategic, or cognitive level (in the domain of symbolic AI); a tactical action-sequencing level; and a real-time control level (which often can be implemented as reactive behaviors).

Papers in the robotics literature often concern specific technical aspects of robot research and develop-

---

**Artificial Intelligence  
and Mobile Robotics,**  
*eds. David Kortenkamp,*  
*R. Peter Bonasso, and*  
*Robin Murphy, AAI*  
*Press, Menlo Park, Cali-*  
*fornia, 1998, 390 pp.,*  
*ISBN 0-262-61137-6.*

---

ment. At the same time, several robot competitions (run by the American Association for Artificial Intelligence [AAAI], the International Joint Conferences on Artificial Intelligence [IJCAI], and the Department of Transportation [DOT], for example) have emphasized integrated, practical approaches to increasingly general versions of real-world problems such as driving, inspection, and environ-

mental cleanup. However, with a few exceptions (for example, Dorigo and Colombetti 1998), it is rare to find thorough, holistic descriptions of the techniques and software organization used in these practical robots.

Enter *Artificial Intelligence and Mobile Robots: Case Studies of Successful Robot Systems*, which, as the title says, is about AI algorithms, representations, and software for perception, planning, and robot control. As the subtitle says, the chapters often describe practical, successful, and competition-winning robots. The book concentrates on the organization of high-level functions in mobile robots, not mechanical or electronic design: Most of the robots mentioned are commercial products, and robot construction is not usually addressed. The book is likewise not concerned with traditional control theory for low-level systems such as direction or speed control. Its strength is its description of software, which is at a level of detail and breadth that one does not find in conference papers or usually in the archival literature. Coupled with an existing text on mobile robot construction and programming, this book would be an excellent primary or supplementary text for the next classroom stage beyond a successful subsumption robot. It is also a source book for professionals who must build software architectures for robot perception and control because it contains detailed, current best-practice solutions from proven authorities.

The book begins with an editorial introduction and intellectual history of robotics, a brief description of robot competitions, and an overview of the volume. Three main sections follow: (1) four chapters on navigation and mapping, (2) three on vision, and (3) six on architecture. Throughout, the main theme is the organization of robot intelligence: what abstractions, connections, representations, and techniques have proven successful in robustly carrying out practical tasks under a wide variety of conditions. Because the chapters are often descriptions of entire working systems, these topical divisions are suggestive but not mutually exclusive. There is no real intellectual dependence between

## ***KDD-98 Now Available!***

*KDD-98—New York, New York*

### **Proceedings of the Fourth International Conference on Knowledge Discovery and Data Mining**

*Edited by Rakesh Agrawal and Paul Stolorz*

As an interdisciplinary field, KDD has made important strides by generating cross-fertilization of ideas between individual disciplines. In keeping with this theme, KDD-98 was collocated with the Conference on Very Large Databases (VLDB-98) with the goal of catalyzing discussions and interactions between researchers in the two fields. Paper topics in this proceedings include theory and foundational issues in KDD, data-mining methods and algorithms, database-centric data mining, and KDD process and human-interaction applications..

ISBN 1-57735-070-7      392 pp., index, \$60.00 softcover

*KDD-97—Newport Beach, California*

### **Proceedings of the Third International Conference on Knowledge Discovery and Data Mining**

*Edited by David Heckerman, Heikki Mannila, Daryl Pregibon, and Ramasamy Uthurusamy*

ISBN 1-57735-027-8      320 pp., index, \$60.00 softcover

**To order, call 650-328-3123 or visit the  
AAAI Press web site at [www.aaai.org/Press/](http://www.aaai.org/Press/)**

Illah Nourbakhsh formalizes and recognizes progress in a plan as the evolution of a set of state variables. The planner interleaves maximum-likelihood assumptions with execution and state monitoring to achieve high-level control. This robot won a difficult office-delivery AAI robot competition.

Sven Koenig and Reid Simmons provide a comprehensive application of partially observable Markov decision-process (POMDP) models to navigation, including the learning of the models and the compiling of them from prior but uncertain maps and sensor models, the application of the models to state and state-sequence estimation, sensor interpretation, obstacle avoidance, and environment modeling. A POMDP is a probabilistic mechanism that estimates and uses probability distributions instead of point probabilities. This chapter, as well as several others, illustrates the power of simulation in developing and quantifying control algorithms.

### **Vision**

Ian Horswill describes how a low-cost vision system and an organization of parallel strategies work together to take advantage of prior constraints of a particular environment (office building) to produce complex and robust behavior (giving a laboratory tour). The vision system's reliable production of percepts (such as *person-ahead?* and *blocked?*), as well as sequencing logic and navigational algorithms, is of particular interest.

Robin Murphy compares two generations of entries to a DOT obstacle-course competition and shows the evolution of their high-level behavioral control scheme. This chapter gives a clear description of reactive, hierarchical, and hybrid control styles and relates them to current theories from cognitive psychology. The performance of the basically reactive robots is examined in terms of these theories, and recommendations for successful sensing strategies are given.

Bill Schiller and several colleagues participate in the SAFETRUCK program, which is to use vision to help vehicles stay in their lane. This chapter gives technical detail about multitasking

chapters, so they can be read in any order. Each of the chapters has its own charm: Some explore general questions, some present enough detail to let the reader understand a system in depth. The pages are rich with block diagrams, code fragments, graphics, photographs, graphs, and maps showing experimental results.

The following brief paragraphs do not do full justice to the chapters, which often tell complete integrated stories of complete integrated systems.

### **Navigation and Mapping**

Sebastian Thrun and several colleagues produce geometric and topological (adjacency) maps using several sensors and a well-founded probabilistic framework. The beneficial effects of position estimation on map making

are illustrated. Maps are used for jobs such as route planning, exploration, and obstacle avoidance. Details of the computer vision and obstacle avoidance during fast driving are provided.

David Kortenkamp and several colleagues describe, in useful technical detail, their entrant to a three-stage AAI robot competition (static and dynamic obstacle avoidance, object location, and a race to visit previously located objects). A vector field histogram representation of space (a discrete two-dimensional [2-D] grid with probabilistic contents) stores robot and goal locations and the locations of obstacles sensed by ultrasound. This 2-D world representation is used for global path planning and obstacle avoidance. Good descriptions of obstacle avoidance, vision, and self-localization algorithms are given.

implementations of robot control, pipelined real-time vision hardware and algorithms, and a comparison of edge and region vision techniques. There is a detailed description of lane keeping and obstacle-avoidance control using range sensors, local maps, computed potentials, and the transformation of potentials into control algorithms to steer around obstacles.

## Architecture

Erann Gat gives a readable history, philosophy, and anatomy of the canonical three-layer system—(1) low-level controller, (2) intermediate-level sequencer, (3) high-level deliberator—with a case study. This chapter could well be read first of the volume as well as first in this section.

Kurt Konolige and Karen Myers describe SAPHIRA, a flexible architecture capable of controlling any robot server by means of a standard interface. SAPHIRA is designed to coordinate activity, maintain a coherent model of local perceptual space and prior information, and communicate with other agents. Low-level control is reactive, initiated by fuzzy-control ideas applied to the local perceptual space. Computed potentials are used in an interesting way. A high-level procedural reasoning system includes goal-oriented analysis of activities and goal sequencing. An overview of communication modalities for use with mobile robots is given. This detailed chapter is motivated by a challenging real-life scenario specification, but (rather frustratingly) we never find out about performance.

James Firby, Peter Prokopowicz, and Michael Swain use a two-layer architecture to control their entrant in the IJCAI-95 competition (locate, classify, and put away trash). Low-level skills are controlled by a special robot language, and task sequencing of high-level skills is controlled by the reactive action plan (RAP) system. The details of skills, plans, and their connection are described along with several good examples of RAP programming and details of world map representations and the vision system.

Ron Arkin and Tucker Balch's work is in controlling cooperating robots in

Insert Mitre Ad

two domains: (1) the Defense Advanced Research Project Agency's unmanned ground vehicle (UGV) DEMO II and (2) the AAAI-94 office-cleaning competition (which their three-robot team won). These examples illustrate schema-based systems, which control behavior through various biologically motivated methods such as vector summation for action fusion, context-dependent behavior instantiation and deletion, and event-triggered sequencing (Arkin 1998).

The UGV demonstration illustrated multirobot formation keeping along with communication of a commander's intentions. The three-robot cleaning team is described in considerable technical detail, including several pseudocode illustrations of controlling schemas.

Huosheng Hu and Michael Brady use the domain of manufacturing to motivate a description of the planning, sensing, and representational components of a typical robot. Their

NEW!

ISMB-98—Montréal, Quebec, Canada

## Proceedings of the Sixth International Conference on Intelligent Systems for Molecular Biology

Edited by Janice Glasgow, Tim Littlejohn, François Major, Richard Lathrop, David Sankoff, and Christoph Sensen

ISBN 1-57735-053-7 234 pp., index, \$45.00 softcover

This volume comprises the formal proceedings for the Sixth International Conference on Intelligent Systems for Molecular Biology (ISMB-98) held in Montréal, Canada from June 28 – July 1, 1998. As with the previous ISMB conferences, this meeting provides a general forum for disseminating the latest developments in bioinformatics. ISMB is a multidisciplinary conference that brings together scientists from computer science, molecular biology, mathematics and statistics. Its scope includes the development and application of advanced computational methods for biological problems. Relevant computational techniques include, but are not limited to: machine learning, pattern recognition, knowledge representation, databases, combinatorics, stochastic modeling, string and graph algorithms, linguistic methods, robotics, constraint satisfaction, and parallel computation. Biological areas of interest include molecular structure, genomics, molecular sequence analysis, evolution and phylogenetics, metabolic pathways, regulatory networks, developmental control, and molecular biology generally. Emphasis is placed on the validation of methods using real data sets, on practical applications in the biological sciences, and on development of novel computational techniques.

To order, call AAAI Press at 650-328-3123  
or visit the AAAI Press web site at [www.aaai.org/Press/](http://www.aaai.org/Press/)

Prices subject to change without notice

locally intelligent control agent (LICA) architecture, like SAPHIRA, transcends any particular implementation, but it is particularly well suited to embedded microprocessor control. Its modularity and scalability are demonstrated, and details of a vision system built of LICAs are given. A probabilistic cost-function approach to navigation and a Kalman filter formalization of dynamic localization are thoroughly described. The multilayered perception and multilayered action components are implemented as a multiprocessor, distributed real-time system.

Don Brutzman and colleagues highlight the difficulties of underwater robotics and the potentially crucial contributions of sophisticated sensor and control simulation. Their underwater robot uses the three-level rational behavior model (RBM) for control: The lowest level is commanded by a mission-script language that is readable by robots and humans and is similar to orders given down the command chain by officers. The middle tactical layer features multiprocessing and interprocess communication to implement cooperative navigation, sonar sensing, and replanning. The RBM for the highest strategic level can

take the form of a deterministic finite automaton or can be as general as a Prolog program; it can also be generated by nonprogrammers using a graphic interface combined with an expert system.

The fairly detailed index and the unified bibliography add to the book's usefulness. It is priced to be accessible: Although the graphics and photographs are clear and the book is quite usable, I feel it deserved a better-looking typeface, more attractive (perhaps hard) cover, and a more modern layout. Also, a chapter on some of the commercially practical (though possibly less AI-intensive) robots emerging from places such as Carnegie Mellon University would have been useful. Still, this very welcome collection is well placed in a large, empty niche. It is a fine resource for anyone teaching about, thinking about, designing, or producing sophisticated mobile robot perception, reasoning, and control software.

### References

- Arkin, R. C. 1998. *Behavior-Based Robotics*. Cambridge, Mass.: MIT Press.
- Dorigo, M., and Colombetti, M. 1998. *Robot Shaping: An Experiment in Behavior Engineering*. Cambridge, Mass.: MIT Press.

Everett, H. R. 1995. *Sensors for Mobile Robots: Theory and Application*. Wellesley, Mass.: A. K. Peters.

Jones, J. L. 1998. *RUG WARRIOR PRO: Hardware and Assembly Guide*. Wellesley, Mass.: A. K. Peters.

Jones, J. L., and Flynn, A. M. 1993. *Mobile Robots: Inspiration to Implementation*. Wellesley, Mass.: A. K. Peters.

McComb, G. 1987. *The Robot Builder's Bonanza: 99 Inexpensive Robotics Projects*. New York: McGraw-Hill.

**Chris Brown** (B.A., Oberlin College, 1967; Ph.D., University of Chicago, 1972) is professor of computer science at the University of Rochester, where he has been since finishing a postdoctoral fellowship at the School of Artificial Intelligence at the University of Edinburgh in 1974. He is coauthor of the influential text *Computer Vision* with his Rochester colleague Dana Ballard and is coeditor of *Real-Time Computer Vision* (Cambridge University Press, 1994), among other volumes. His current research interests are computer vision and robotics; integrated parallel systems performing animate vision (the interaction of visual capabilities and motor behavior); and the integration of planning, learning, sensing, and control. His e-mail address is [brown@cs.rochester.edu](mailto:brown@cs.rochester.edu).