Gregory Dudek and Michael Jenkin, {Computational Principles of Mobile Robotics}, New York, NY: Cambridge University Press, 2000, xiii + 280 pp., $29.95 (paper), ISBN 0-521-56876-5.

{Computational Principles of Mobile Robotics} comprises 10 chapters, spanning hardware, mathematical formalisms, algorithms, implemented software and systems, and terminology related to autonomous robots that are designed to solve problems in large-scale space, i.e., {mobile robots}. Large-scale space consists of "regions of space substantially larger than those that can be observed from a single vantage point" (p. 1). Robot mobility, large-scale space, and navigational skills go hand in hand. On the one hand, a robot that is mobile needs navigational skills to exploit its movement, and these skills are useful in large-scale space. On the other hand, large-scale space and local sensing imply a need for mobility and navigation to other vantage points. The book provides problems to be worked at the end of each chapter. Six of the chapters also cite further relevant readings. Most of the book is in regards to individual robots, but Chapters 1, 5, 6, and 8 include some consideration of multi-robot issues. I will summarize the contents of the chapters, and then move on to my specific comments about the book.

1. Summary

Chapter 1 introduces the problems facing a robot navigating in large-scale space: dealing with incremental acquisition of knowledge (i.e., because a robot's sensors cannot perceive the entire environment from any specific location), estimation of positional error, recognition of important or familiar objects or places, and real-time response. These problems are framed in the traditional computational trio of input-processing-output: sensing in space, reasoning about space, and movement through space. The chapter also provides some historical context, and goals by contrast—the Stanford cart, for example, "moved roughly 1 m every 10 to 15 minutes," clearly not the real-time response expected of {mobile robots}!

Chapters 2 through 4 take up the sensori-motor aspects of mobile robotics, with {Mobile Robot Hardware} (Ch. 2), {Nonvisual Sensors and Algorithms} (Ch. 3), and {Visual Sensors and Algorithms} (Ch. 4). Chapter 2 is on motor control and covers forward and reverse kinematics for several types of drive systems including differential drive, synchronous drive, steered vehicles, car drive (Ackerman steering), and a constrained type of limbed vehicle. Forward kinematics is the analytic problem of determining the resulting motion produced by a robot given particular control inputs. For example, if I turn on the robot's wheel motors for 5 seconds what will it's resulting (x, y) coordinate be? Inverse kinematics addresses the complementary analytic problem of choosing control inputs given a goal of particular robot motion. For example, if I want my robot to move from an initial (x_1, y_1) coordinate in a room to a new (x_2, y_2) coordinate, how do I use the wheel motors to accomplish the change in position? Chapter 3 considers internal state sensors (analogous to proprioception in biological organisms) such as accelerometers, gyroscopes, compasses, inclinometers, joint angles, and wheel positions, and external state sensors such as contact, infrared, sonar, radar, laser rangefinders, and GPS. Kalman filtering is also considered in the context of combining data from different sources ({data fusion}), and sensing in biological organisms is briefly considered. The introduction to Chapter 4 does a good job in describing the problem of visual sensing, and then launches into a mathematical presentation on modeling pinhole cameras using perspective projection, finding a calibration matrix (which gives a transformation from a 2-D pixel to its 3-D world coordinate as a function of camera and non-camera parameters), reflectance, sampling, detecting the presence of the same feature across images, and computing depth. Section 4.6 addresses active vision, which has a goal of directing a robot's camera(s) to salient events in the world in an effort to simplify the computational vision problem. Section 4.7 considers other light-based sensors (e.g., light striping, single-camera stereo systems).

Chapters 5 and 6 describe main aspects of the control software for mobile robots. {Representing and Reasoning About Space} (Ch. 5) suggests that "the most natural representation of a robot's environment is a map" (p. 121). This chapter covers techniques for representing maps using spatial decompositions (including occupancy grids and {quadtrees}), geometric representations (maps in terms of geometric primitives such as points and lines), and topological representations (graphs with nodes and edges). Section 5.2 considers {configuration space}, a method to represent the physical configurations of a robot in its environment, for use in movement planning and also introduces holonomic and nonholonomic constraints (perhaps my favorite terms of the book). Nonholonomic constraints exist when a robot cannot achieve some possible configurations and when these influence the motion the robot can perform (e.g., a car robot with only holonomic constraints would have no difficulty parallel parking). The core of the chapter is Section 5.3, on path planning, which includes visibility graph planning, Voronoi diagrams, potential fields, the problem of detecting a local minimum, vector field histograms, and the bug algorithm. Some of the algorithms in this section can deal with a "major flaw [of] many classical path planners … the assumption that the environment is known in advance" (p. 146). Probabilistic path planning includes a learning phase, and the bug algorithm (which combines circumnavigation of obstacles with movement directly to goal coordinates) can be used without a global map. Chapter 6 ({Operating Environment}) considers different architectural and organizational possibilities for the control software of mobile robots. Approaches include functional decomposition (e.g., perception, cognition, action), reactive control (decomposition of a robot's actions into separate behaviors), high-level languages (e.g., predicate calculus), artificial neural networks, and genetic algorithms.

Chapter 7 describes pose maintenance, the process of maintaining an estimate of the location and heading of a robot with respect to a (given) map of the environment. This chapter presents dead reckoning (navigation without external sensors), landmark positioning, homing ("sensor-based servoing"), model-based localization, use of perceptual structure (computing a sensory signature of a position), and global localization (distinguishing the robots position in the face of sensory ambiguity—when a position could be one of several perceptually similar locations). As the methods of Chapter 7 assume a map of the environment, Chapter 8 considers the problem of having robots generate their own maps. Maps can be of various types including metric maps and topological maps. Metric maps are based on an absolute coordinate frame and contain positional estimates of objects. Techniques covered include image-based mapping (collecting images at specific positions and view angles), spatial occupancy representations (similar to Chapter 5 occupancy grids), use of Markov models (in this context, describing a robot's action as a function of its {current} location and heading in the environment), geometric maps and spiral search, and topological maps and marker-based exploration.

Chapters 9 and 10 close the book with {Practical Mobile Robot Tasks} and {The Future of Mobile Robotics}. Chapter 9 considers applications of delivery (e.g., the Helpmate robot which can move objects from location to location in hospitals), assembly and manufacturing (only non-mobile or industrial robots have thus far made inroads in these areas), intelligent vehicles (e.g., driving assistant systems), survey and inspection, mining, space (e.g., the Mars Sojourner), aircraft, military (including mine disposal), underwater inspection, agriculture and forestry, entertainment, and cleaning. Chapter 10 indicates "that for restrictive environments and for limited tasks, autonomous systems can be readily developed," however, "the inflexibility inherent in … [a restrictive] environment makes the use of robot systems less appealing" (p. 251). These environment restrictions, plus task restrictions, result in mobile robots that are not general purpose. The authors identify open research problems in locomotion (e.g., locomotion over non-smooth terrain), sensing (e.g., sensor fusion and forming a coherent environmental representation), control (e.g., task-level control), and the "formidable challenge" of "integrat[ing] … disparate subsystems for control and sensing at a hardware and software level" (p. 254).

2. Main Comments: Audience & Organization

I have two main comments on the authors' book relating to intended audience and organization and conceptual presentation. I also have two auxiliary comments regarding theoretical perspective, and a lack of treatment of algorithms arising from animal navigation. The authors have certainly provided an extended treatment of navigation in mobile robotics, and I expect to use their book as a reference. However, I think this book does not fully hit its mark for potential audience, and also it could be better conceptually organized.

With regard to intended audience, the initial page indicates that the book is suited to a mobile robotics audience, which I assume to mean engineers. As examples of this assumed audience, the treatment of Kalman filters in Chapter 3, was formal and very dense, the definition of convolution in Chapter 4 assumes familiarity with convolution as do the equations for image brightness, optical flow, and transfer functions in the same chapter. The notations appear to assume some specific background as well. For example, in Chapter 8, p. 214 on {Image-based mapping}, I assume that $[I_i (x_i, y_i, \theta_i)]$ means that an image matrix is collected at position x_i, y_i and rotation θ_i. This appears to be the only interpretation that makes sense given the "image" emphasis. While x, y and θ are apparent from the context, it is not clear that I is an image (apparently, a pixel matrix). While the authors state that "the study of mobile robots is an intrinsically interdisciplinary research area involving" (p. 3) at least mechanical and electrical engineering, computer science, cognitive psychology, perceptual psychology, and neuroscience, they have unfortunately made it likely that a good deal of this interdisciplinary audience will be unable to navigate through their book.

In regards to organization and conceptual presentation, I found the sections in the chapters to be missing a level of organization that could assist the reader in making their way from formalism to formalism, comprehension intact. As an example of this, in Chapter 4.1.1, it was not at all clear to me when a {calibration matrix} was needed in a mobile robotics application. In this section, the authors spend four paragraphs getting to a definition of calibration itself, and this and the following presentation lost me in terms of figuring out where calibration could be applied. I found myself getting lost time and time again even when the text wasn't formal. For example, Section 10.1.1 of {Future Research} titled "Smooth Terrain" doesn't state a general research question on this topic. Instead, the authors jump in with the details, and aside from stating in the second paragraph of this section a goal "to overcome the need for a continuous floor surface" (p. 252), they don't follow-up with a bigger picture. Another example of this is in Sections 7.4 and 7.5. Section 7.5 starts with "One of the most difficult aspects of model-based localization" (p. 197) but at this point in reading I had no concept of model-based localization. Presumably this was an anaphoric reference to the immediately previous section (7.4) on Kalman filtering or perhaps it reflects the map assumption of Chapter 7. In general, while focusing on technical content, the book did not paint a broader picture of issues and did not provide a coherent conceptual framework to help the reader organize the topics. Having staked out the territory of navigation as their task domain, the authors could have provided a conceptual view of navigation, and integrated that treatment with the technical content. Instead, the treatment of navigation was piecemeal and difficult to follow. As examples of this, p. 80 uses the terms landmark-based and dead reckoning, but gives no introduction to the terms. Also, the term "metric representation" is used on p. 80, but is not defined. Dead reckoning is not described until p. 175 in Chapter 7. Dead reckoning is navigation without use of external sensors (e.g., finding your way in a room, when you have seen your starting position, but afterwards proceed with your eyes closed), while landmark-based navigation is navigation with external sensors (e.g., finding your way in the room with your eyes open), and these concepts are central to sensori-motor issues in robotic navigation. For example, while dead reckoning can be achieved by keeping track of wheel rotations (odometry) and angle of turns, it is prone to cumulative error, and navigation should also generally be supplemented with landmark-based navigation, which while likely involving more computation (than dead reckoning) can be used to correct for the errors involved in dead reckoning. It is not until over halfway through the book, embedded in Chapter 7, that they finally state a vital conceptual tenet of navigation—that "one approach to dealing with intermittently reliable landmark information is to combine landmark-based position information with position information from dead reckoning" (p. 183). While a knowledgeable reader would have realized earlier that an "assum[ption] of perfect odometry" (p. 214) cannot likely be satisfied, and also that landmark-based information itself is susceptible to sensor error, integrating concepts of navigation earlier and more completely into the book would help.

3. Auxiliary Comments: Additional Depth and Breadth

My last two comments reflect on the depth and breadth of the authors' book. First, I would have liked to see a more integrated treatment of issues relating to software and hardware architecture of mobile robots. In their closing chapter (as described above), the authors make a call for additional study of issues related to integrating control and sensing subsystems. The authors seem to have various views on the issues involved. While they figure maps centrally into their consideration of representations suitable for mobile robots (e.g., "the most natural representation of a robot's environment is a map," p. 121), and consider subsumption control (e.g., Brooks, 1999; this method takes a generally non-map view of navigation) to be "difficult" and "problematic" (p. 155), they acknowledge that {a priori} plans also have difficulties in mobile robotics applications: "Because it is often the case that the initial plan will not survive contact with the environment, an alternative to expending considerable effort to generate a complete plan initially is to generate an approximate plan quickly and begin to execute it, refining the plan as the robot moves" (p. 146). There is also an apparent theoretical emphasis in their statement that "discrete and continuous systems are better understood as independent systems rather than in combination" (p. 164-165), where "discrete" refers to higher-level task control, and "continuous" refers to control of low-level systems (e.g., dealing with terrain irregularities).

Second, there was a general lack of treatment related to algorithms arising from (non-human) animal navigation. The authors first sentence in Chapter 1, "The ability to navigate purposefully is fundamental to most animals and to every intelligent organism," while a useful metaphor in this context, is more than a metaphor in computational neuroscience. For example, the hippocampus region of rodent brains has received extensive study and much computational modeling, in large part because of its role in spatial representation and navigation in these mammals (e.g., O'Keefe & Nadel, 1978). While some of these computational models do not emphasize robotic systems (e.g., Reddish, 1999; Schmajuk, 1997; Trullier, Wiener, Berthoz, & Meyer, 1997), the arising spatial navigation and representation algorithms can be applied to robots (e.g., Burgess, Donnett, Jeffery, & O'Keefe, 1997; Franz & Mallot, 2000). On the surface, while it seems possible to relegate treatment of computational and robotic models of animal navigation to a position separate from a book on mobile robotics emphasizing navigation, I think this misses important issues. Perhaps most vital is that, of the biological systems having navigational strategies in their behavioral repertoires, humans are but a single species example. Other species (e.g., rats, bats, bees) have other and varied strategies, and hence may provide us with other and varied computational insights into these problems. And, while topological maps may "have appealing apparent analogies with human spatial perception" (p. 224), this form of map representation may in fact be a relative cognitive and evolutionary aberration—it is far from uncontroversial whether other animals have such strongly abstract spatial skills (e.g., Benhamou, 1996; Macuda & Roberts, 1995; Prince, 1998). In a field which suffers from problems with inflexibility resulting from environmental and task restrictions, why not take insects, rats, birds (to name a few) more seriously? While {Rattus norvegicus} may "suffer" from task restrictions, they are far from inflexible.

4. Additional Notes

(1) The last chapter on {The Future of Mobile Robotics} was too short (5 pages). Does mobile robotics only have a short future? (2) The formatting of tables (e.g., p. 73/74) was difficult to visually separate from the main text. (3) How does this kind of robotics fit in with other aspects of robotics? What I have in mind is again from a biological or comparative psychological perspective. Spatial skills, both small-scale and large-scale serve many animals in good stead. But, these skills integrate with other systems. How can the navigational systems of mobile robots integrate with other types of robotic approaches (e.g., developmental robotics; Prince, 2002; Weng, McClelland, Pentland, Sporns, Stockman, Sur, & Thelen, 2001; Zlatev & Balkenius, 2001)? (4) The authors make only short note of {RoboCup} (Asada, Kitano, Noda, & Veloso, 1999) and entertainment aspects of mobile robotics (e.g., while the Sony AIBO quadruped is on the cover, little mention is made of it). The goals of RoboCup, to have a humanoid-robotic team win against the World Cup human soccer champions by the year 2050, if met, should certainly advance mobile robotics! And, entertainment (e.g., games) is an important market share of computers, so why should mobile robotics be any different? (5) The authors analogize the "software environment within which the control programs of modern mobile robots execute … to mobile operating systems" (p. 119). How do these systems compare to contemporary operating systems? What research and development is needed to further such operating systems? (6) Can the concept of {stability} (p. 127) be used to generate a kind of map representation? I.e., a map should be stable or relatively insensitive to variation in sensor input. (7) What about educational applications of mobile robotics? For example, Lego Mindstorms robotics (http://mindstorms.lego.com) has made strides to providing a low-cost and flexible base of robotic systems to schools.

5. In Closing

This book can certainly serve as a technical reference to those who are already familiar with the area of mobile robotics. Engineers should certainly benefit from this treatment as well. I have strong reservations, however, about the use of this book for people not familiar with engineering and an interdisciplinary audience. It should be noted that at least part of this audience consideration may reflect "a completely new research environment developing" where "mobile robotics audience" now means an interdisciplinary audience (U. Nehmzow, Pers. Comm.). For example, it may become feasible in the near future for neuroscientists and psychologists to run studies with mobile robots prior to work with live animal subjects (e.g., Sporns & Alexander, 2002).

References

Asada, M., Kitano, H., Noda, I., and Veloso, M. (1999), 'RoboCup: Today and Tomorrow—What We Have Learned', {Artificial Intelligence} 110, pp. 193-214.

Benhamou, S. (1996), 'No Evidence for Cognitive Mapping in Rats', {Animal Behaviour} 52, pp. 201-212.

Brooks, R. A. (1999), {Cambrian Intelligence: The Early History of the New AI}, Cambridge, MA: MIT Press.
Burgess, N., Donnett, J.G., Jeffery, K.J., and O'Keefe, J. (1997), 'Robotic and Neuronal Simulation of the Hippocampus and Rat Navigation', {Phil. Trans. Roy. Soc. B}, 352, pp.1535-1543.
Franz, M. O. and Mallot, H. A. (2000), 'Biomimetic Robot Navigation', {Robotics and Autonomous Systems} 30, pp. 133-153.
Macuda, T. and Roberts, W. A. (1995), 'Further Evidence for Hierarchical Chunking in Rat Spatial Memory', {Journal of Experimental Psychology: Animal Behavior Processes} 21, pp. 20-32.
O'Keefe, J., and Nadel, L. (1978), {The Hippocampus as a Cognitive Map}, London: Oxford University Press.
Prince, C. G. (1998), {Rats' Capabilities With Relations Between Locations in Space}, Unpublished Doctoral Dissertation, University of Louisiana at Lafayette.
Prince, C. G. (2002), 'Introduction: The Second International Workshop on Epigenetic Robotics', {Proceedings of the Second International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems}, pp. XX-YY. Lund University Cognitive Studies, Volume 94.
Reddish, A. D. (1999), {Beyond the Cognitive Map: From Place Cells to Episodic Memory}, Cambridge, MA: MIT Press.
Schmajuk, N. A. (1997), {Animal Learning and Cognition: A Neural Network Approach}, New York, NY: Cambridge University Press.
Sporns, O. and Alexander, W. H. (2002). 'Dopamine, Reward Conditioning, and Robot Behavior', In {Proceedings of the 2nd International Conference on Development and Learning} (pp. 265-270). IEEE Computer Society Press. Los Alamitos, CA.
Trullier, O., Wiener, S. I., Berthoz, A., and Meyer, J.-A. (1997), 'Biologically Based Artificial Navigation Systems: Review and Prospects', {Progress in Neurobiology} 51, pp. 483-544.
Weng, J., McClelland, J., Pentland, A., Sporns, O., Stockman, I., Sur, M., and Thelen, E. (2001), 'Autonomous Mental Development by Robots and Animals', {Science} 291, pp. 599-600.
Zlatev, J. and Balkenius, C. (2001), 'Introduction: Why "Epigenetic Robotics"?',{Proceedings of the First International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems}, pp. 1-4. Lund University Cognitive Studies, Volume 85.

CHRISTOPHER G. PRINCE
{Department of Computer Science}
{University of Minnesota Duluth}
{Duluth, MN 55812 U.S.A.}
{E-mail: chris@cprince.com}