Book Review

Josh Bongard*
University of Vermont

Probabilistic Robotics. Sebastian Thrun, Wolfram Burgard, and Dieter Fox. (2005, MIT Press.) 647 pages.

It's a wild world out there. The most striking pattern one can observe in the history of robotics (since its beginnings in the 1950s) is its staggering successes in completely revolutionizing heavy industry through automation, and its equally spectacular failure to produce robots that work alongside us in the home, or out of doors. Like the artificial life community, roboticists have struggled to develop ways to enable their creations to deal with the constantly changing demands of the real world. In the first attempts, robots were provided with internal models crafted for them by their creators, but this limited their usefulness: They slowly evaluated their options against these models, and became useless (or dangerous) if their models became inaccurate through environmental change. One of the very first autonomous robots used this approach: Shakey the Robot [1], developed in the late 1960s, could reason using internal models, but it shook and hesitated as it used them to plan actions. In the 1980s Rod Brooks of MIT fomented a rebellion in the field by stating that robots may not require models at all in order to exhibit useful behavior [2], and (along with others that followed) loosed upon the research community a series of scrambling, bounding, and otherwise fast-moving robot critters. The debate between classical robotics and behavior-based robotics continues today.

With the release of *Probabilistic Robotics*, Sebastian Thrun and his coauthors have laid down another, equally large gauntlet. If the only constancy in life is change, then robots should be able to deal with the uncertainties around them by taking them into account when operating. Thrun et al. replace the complete (and therefore computationally intensive) internal models from classical robotics with statistical ones that honestly reflect the uncertainty out there in the world, from the point of view of a robot. With precision, elegance, and depth, the authors indicate the deep philosophical and methodological differences that distinguish deterministic and model-free robotics from probabilistic robotics. (It is interesting to note that Thrun is currently director of the artificial intelligence laboratory at Stanford University, which, coincidentally, is also the birthplace of Shakey and where Brooks received his Ph.D.)

As the authors state at the outset of the book, uncertainty does not just surround the robot in the form of environmental noise, but exists within the robot as well. Any real-world robot can safely assume that there are gaps in its knowledge about the environment (*is this a door I'm seeing?*), but also in regard to its sensors (*how well did that last measurement actually indicate that the door is open?*), the effect of an action (*did I actually close the door?*), and its local position within a larger environment (*have I seen this door before?*). Chapter by chapter, the authors systematically reveal the challenges related to equipping a robot with the wherewithal to deal with increasing uncertainty, and along the way introduce algorithms designed to handle them.

This book, however, is not for the faint of heart, as the mathematics required to model and contain this uncertainty may be daunting for some. This book serves well as a graduate textbook, and as a mandatory reference and handbook for those working in the field. That being said, the authors provide a very thorough treatment of the mathematics required in the first part of the book. For those who are not working directly in robotics, such as artificial life researchers, machine learning researchers, and biologists, these initial chapters alone are invaluable as an accessible introduction to probability theory.

The structure of the book is excellent, as the mathematics is interspersed with examples provided at varying levels of detail, from simple "imagine a robot attempting to ..." scenarios, to visual

* Department of Computer Science, University of Vermont, Burlington, VT 05405. E-mail: josh.bongard@uvm.edu

depictions of the algorithms, to descriptions of actual robot implementations. For instance, on p. 150 the authors provide an overhead view of a robot attempting to locate itself within an office corridor using an ultrasound scan. The image indicates that, depending on the geometry and texture of the corridor's walls and doors, many of the outgoing sonar signals will reflect in such a way that the robot does not receive them, and therefore erroneous, infinite readings occur. This example provides a simple illustration of why a robot must approach sensor readings statistically, rather than deterministically. Many of the chapters and sections also start off with motivational examples that indicate why a particular approach is needed, and serve to highlight the advantages and disadvantages of each approach.

After the mathematical development, a series of increasingly challenging tasks (for a robot) are introduced, along with the probabilistic algorithms developed to solve them. A thorough treatment of localization is provided first, followed by mapping. The former challenge involves finding oneself within a known environment; the latter involves building a description of the environment based on knowledge of one's own whereabouts. The authors then guide the reader through the various approaches for simultaneously locating oneself while mapping an unknown environment, one of the most thoroughly investigated but still challenging problems in mobile robotics. The book concludes with an investigation of robot planning and exploration.

Not only is the treatment of this large and important body of work in robotics (much of it carried out by the authors themselves) deep, systematic, accessible, and broad, it is also honest. For most of the algorithms presented, the authors describe their limitations, thereby guiding practitioners in selecting the appropriate approach for their robot and problem domain. In addition, each chapter ends with a number of methodological and material-expanding exercises, as well as a summary of the relevant literature.

The preface might have benefited from indicating what level of mathematical background a reader should have. Granted, the authors do indicate in the preface that each algorithm is accompanied by an example implementation in pseudocode, results generated either by hypothetical, simulated, or real robots, and a discussion of the algorithm's pros and cons, as well as a detailed mathematical derivation from first principles, which the reader may choose to skip over during a first reading. The book is accompanied by a Web site (www.probabilistic-robotics.org) that provides all of the figures in the book, as well as ready-made Power Point slides that can be used to present material from the book. The Web site might have benefited from providing some (or all) solutions to the book's exercises, as it is sometimes difficult, even after going back through the chapter, to self-verify whether one's solution is correct.

As already mentioned, this book successfully clarifies a "third way" in robotics, in addition to model-based and model-free robotics. In the authors' own words: "In contrast with traditional programming techniques in robotics—such as model-based motion planning techniques or reactive behavior-based approaches—probabilistic approaches tend to be more robust in the face of sensor limitations and model limitations. This enables them to scale much better to complex real-world problems. . . ."

This issue of scalability is of paramount importance in robotics, and many promising approaches and impressive robot demonstrations have failed to scale into practical, tough, and sophisticated machines capable of really helping or replacing humans in dangerous, remote, or boring task environments. And although not mentioned explicitly in the text, the main author led a team of students at Stanford to create a robot capable of doing just this. Stanley [3], a driverless Volkswagen Touareg, won the 2005 Defense Advanced Research Projects Agency (DARPA) Grand Challenge: to drive, as fast as possible and without any human intervention, across 142 miles of the Mojave desert. Using some of the methods outlined in the book, Stanley took home the gold medal—along with the $2 million in prize money—after finishing nearly seven hours ahead of the competition. With more than 100 fatalities from auto accidents a day, worldwide, and the promise that driving automation could alleviate this, one is hard pressed to imagine a more resounding, dramatic, and needed demonstration that the theoretical content of *Probabilistic Robotics* can be translated into real-world solutions to real-world problems.

References

1. Nilsson, N. J. (1984). *Shakey the robot* (Technical Report 323). Menlo Park, CA: Artificial Intelligence Center, SRI International.

2. Brooks, R. A. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, *2*(1), 14–23.

3. Thrun, S., Montemerlo, M., Dahlkamp, H., Stavens, D., Aron, A., Diebel, J., Fong, P., Gale, J., Halpenny, M., Hoffmann, G., Lau, K., Oakley, C., Palatucci, M., Pratt, V., Stang, P., Strohband, S., Dupont, C., Jendrossek, L.-E., Koelen, C., Markey, C., Rummel, C., van Niekerk, J., Jensen, E., Alessandrini, P., Bradski, G., Davies, B., Ettinger, S., Kaehler, A., Nefian, A., & Mahoney, P. (2006). Stanley, the robot that won the DARPA Grand Challenge. *Journal of Field Robotics*, *23*(9), 661–692.