


INTRODUCTION TO
AI ROBOTICS



ROBIN R. MURPHY

Introduction
to
AI
Robotics

Robin R. Murphy

A Bradford Book
The MIT Press
Cambridge, Massachusetts
London, England

Second printing, 2002

© 2000 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

Typeset in 10/13 Lucida Bright by the author using L^AT_EX 2_ε.
Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Murphy, Robin, 1957–

Introduction to AI robotics / Robin R. Murphy.

p. cm.—(Intelligent robotics and autonomous agents. A Bradford Book.)

Includes bibliographical references and index.

ISBN 0-262-13383-0 (hc. : alk. paper)

1. Robotics. 2. Artificial intelligence. I. Title. II. Series

TJ211.M865 2000

629.8'6263—dc21

00-033251

Introduction
to
AI
Robotics

Intelligent Robots and Autonomous Agents
Ronald C. Arkin, editor

Behavior-Based Robotics, Ronald C. Arkin, 1998

Robot Shaping: An Experiment in Behavior Engineering, Marco Dorigo and Marco Colombetti, 1998

Layered Learning in Multiagent Systems: A Winning Approach to Robotic Soccer, Peter Stone, 2000

Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines, Stefano Nolfi and Dario Floreano, 2000

Reasoning about Rational Agents, Michael Wooldridge, 2000

Introduction to AI Robotics, Robin R. Murphy, 2000

To Kevin

...and Carlyle Ramsey, Monroe Swilley, Chris Trowell

Preface

This book is intended to serve as a textbook for advanced juniors and seniors and first-year graduate students in computer science and engineering. The reader is not expected to have taken a course in artificial intelligence (AI), although the book includes pointers to additional readings and advanced exercises for more advanced students. The reader should have had at least one course in object-oriented programming in order to follow the discussions on how to implement and program robots using the structures described in this book. These programming structures lend themselves well to laboratory exercises on commercially available robots, such as the Khepera, Nomad 200 series, and Pioneers. Lego Mindstorms and Rug Warrior robots can be used for the first six chapters, but their current programming interface and sensor limitations interfere with using those robots for the more advanced material. A background in digital circuitry is not required, although many instructors may want to introduce laboratory exercises for building reactive robots from kits such as the Rug Warrior or the Handy Board.

Introduction to AI Robotics attempts to cover all the topics needed to program an artificially intelligent robot for applications involving sensing, navigation, path planning, and navigating with uncertainty. Although machine perception is a separate field of endeavor, the book covers enough computer vision and sensing to enable students to embark on a serious robot project or competition. The book is divided into two parts. Part I defines what are intelligent robots and introduces why artificial intelligence is needed. It covers the “theory” of AI robotics, taking the reader through a historical journey from the Hierarchical to the Hybrid Deliberative/Reactive Paradigm for organizing intelligence. The bulk of the seven chapters is concerned with the Reactive Paradigm and behaviors. A chapter on sensing and programming techniques for reactive behaviors is included in order to permit a class to get

a head start on a programming project. Also, Part I covers the coordination and control of teams of multi-agents. Since the fundamental mission of a mobile robot involves moving about in the world, Part II devotes three chapters to qualitative and metric navigation and path planning techniques, plus work in uncertainty management. The book concludes with an overview of how advances in computer vision are now being integrated into robots, and how successes in robots are driving the web-bot and know-bot craze.

Since *Introduction to AI Robotics* is an introductory text, it is impossible to cover all the fine work that has been in the field. The guiding principle has been to include only material that clearly illuminates a specific topic. References to other approaches and systems are usually included as an advanced reading question at the end of the chapter or as an end note. *Behavior-based Robotics*¹⁰ provides a thorough survey of the field and should be an instructor's companion.

Acknowledgments

It would be impossible to thank all of the people involved in making this book possible, but I would like to try to list the ones who made the most obvious contributions. I'd like to thank my parents (I think this is the equivalent of scoring a goal and saying "Hi Mom!" on national TV) and my family (Kevin, Kate, and Allan). I had the honor of being in the first AI robotics course taught by my PhD advisor Ron Arkin at Georgia Tech (where I was also his first PhD student), and much of the material and organization of this book can be traced back to his course. I have tried to maintain the intellectual rigor of his course and excellent book while trying to distill the material for a novice audience. Any errors in this book are strictly mine. David Kortenkamp suggested that I write this book after using my course notes for a class he taught, which served as a very real catalyst. Certainly the students at both the Colorado School of Mines (CSM), where I first developed my robotics courses, and at the University of South Florida (USF) merit special thanks for being guinea pigs. I would like to specifically thank Leslie Baski, John Blitch, Glenn Blauvelt, Ann Brigante, Greg Chavez, Aaron Gage, Dale Hawkins, Floyd Henning, Jim Hoffman, Dave Hersherberger, Kevin Gifford, Matt Long, Charlie Ozinga, Tonya Reed Frazier, Michael Rosenblatt, Jake Sprouse, Brent Taylor, and Paul Wiebe from my CSM days and Jenn Casper, Aaron Gage, Jeff Hyams, Liam Irish, Mark Micire, Brian Minten, and Mark Powell from USF.

Special thanks go to the numerous reviewers, especially Karen Sutherland and Ken Hughes. Karen Sutherland and her robotics class at the University of Wisconsin-LaCrosse (Kristoff Hans Ausderau, Teddy Bauer, Scott David Becker, Corrie L. Brague, Shane Brownell, Edwin J. Colby III, Mark Erickson, Chris Falch, Jim Fick, Jennifer Fleischman, Scott Galbari, Mike Halda, Brian Kehoe, Jay D. Paska, Stephen Pauls, Scott Sandau, Amy Stanislawski, Jaromy Ward, Steve Westcott, Peter White, Louis Woyak, and Julie A. Zander) painstakingly reviewed an early draft of the book and made extensive suggestions and added review questions. Ken Hughes also deserves special thanks; he also provided a chapter by chapter critique as well as witty emails. Ken always comes to my rescue.

Likewise, the book would not be possible without my ongoing involvement in robotics research; my efforts have been supported by NSF, DARPA, and ONR. Most of the case studies came from work or through equipment sponsored by NSF. Howard Moraff, Rita Rodriguez, and Harry Hedges were always very encouraging, beyond the call of duty of even the most dedicated NSF program director. Michael Mason also provided encouragement, in many forms, to hang in there and focus on education.

My editor, Bob Prior, and the others at the MIT Press (Katherine Innis, Judy Feldmann, Margie Hardwick, and Maureen Kuper) also have my deepest appreciation for providing unfailingly good-humored guidance, technical assistance, and general savvy. Katherine and especially Judy were very patient and nice—despite knowing that I was calling with Yet Another Crisis. Mike Hamilton at AAAI was very helpful in making available the various “action shots” used throughout the book. Chris Manning provided the L^AT_EX 2_ε style files, with adaptations by Paul Anagnostopoulos. Liam Irish and Ken Hughes contributed helpful scripts.

Besides the usual suspects, there are some very special people who indirectly helped me. Without the encouragement of three liberal arts professors, Carlyle Ramsey, Monroe Swilley, and Chris Trowell, at South Georgia College in my small hometown of Douglas, Georgia, I probably wouldn’t have seriously considered graduate school in engineering and computer science. They taught me that learning isn’t a place like a big university but rather a personal discipline. The efforts of my husband, Kevin Murphy, were, as always, essential. He worked hard to make sure I could spend the time on this book without missing time with the kids or going crazy. He also did a serious amount of editing, typing, scanning, and proofreading. I dedicate the book to these four men who have influenced my professional career as much as any academic mentor.

Brief Contents

I Robotic Paradigms 1

- 1 *From Teleoperation To Autonomy* 13
- 2 *The Hierarchical Paradigm* 41
- 3 *Biological Foundations of the Reactive Paradigm* 67
- 4 *The Reactive Paradigm* 105
- 5 *Designing a Reactive Implementation* 155
- 6 *Common Sensing Techniques for Reactive Robots* 195
- 7 *The Hybrid Deliberative/Reactive Paradigm* 257
- 8 *Multi-agents* 293

II Navigation 315

- 9 *Topological Path Planning* 325
- 10 *Metric Path Planning* 351
- 11 *Localization and Map Making* 375
- 12 *On the Horizon* 435

Contents

Preface xvii

I Robotic Paradigms 1

1 *From Teleoperation To Autonomy* 13

- 1.1 Overview 13
- 1.2 How Can a Machine Be Intelligent? 15
- 1.3 What Can Robots Be Used For? 16
 - 1.3.1 Social implications of robotics 18
- 1.4 A Brief History of Robotics 19
 - 1.4.1 Industrial manipulators 21
 - 1.4.2 Space robotics and the AI approach 26
- 1.5 Teleoperation 28
 - 1.5.1 Telepresence 32
 - 1.5.2 Semi-autonomous control 33
- 1.6 The Seven Areas of AI 34
- 1.7 Summary 37
- 1.8 Exercises 37
- 1.9 End Notes 39

2 *The Hierarchical Paradigm* 41

- 2.1 Overview 41
- 2.2 Attributes of the Hierarchical Paradigm 42
 - 2.2.1 Strips 44
 - 2.2.2 More realistic Strips example 46
 - 2.2.3 Strips summary 52

2.3	Closed World Assumption and the Frame Problem	53
2.4	Representative Architectures	54
2.4.1	Nested Hierarchical Controller	54
2.4.2	NIST RCS	57
2.4.3	Evaluation of hierarchical architectures	59
2.5	Advantages and Disadvantages	61
2.6	Programming Considerations	62
2.7	Summary	62
2.8	Exercises	63
2.9	End Notes	64
3	<i>Biological Foundations of the Reactive Paradigm</i>	67
3.1	Overview	67
3.1.1	Why explore the biological sciences?	69
3.1.2	Agency and computational theory	70
3.2	What Are Animal Behaviors?	73
3.2.1	Reflexive behaviors	74
3.3	Coordination and Control of Behaviors	75
3.3.1	Innate releasing mechanisms	77
3.3.2	Concurrent behaviors	82
3.4	Perception in Behaviors	83
3.4.1	Action-perception cycle	83
3.4.2	Two functions of perception	85
3.4.3	Gibson: Ecological approach	85
3.4.4	Neisser: Two perceptual systems	90
3.5	Schema Theory	91
3.5.1	Behaviors and schema theory	92
3.6	Principles and Issues in Transferring Insights to Robots	97
3.7	Summary	99
3.8	Exercises	100
3.9	End Notes	102
4	<i>The Reactive Paradigm</i>	105
4.1	Overview	105
4.2	Attributes of Reactive Paradigm	108
4.2.1	Characteristics and connotations of reactive behaviors	110
4.2.2	Advantages of programming by behavior	112
4.2.3	Representative architectures	113

4.3	Subsumption Architecture	113
4.3.1	Example	115
4.3.2	Subsumption summary	121
4.4	Potential Fields Methodologies	122
4.4.1	Visualizing potential fields	123
4.4.2	Magnitude profiles	126
4.4.3	Potential fields and perception	128
4.4.4	Programming a single potential field	129
4.4.5	Combination of fields and behaviors	130
4.4.6	Example using one behavior per sensor	134
4.4.7	Pfields compared with subsumption	136
4.4.8	Advantages and disadvantages	145
4.5	Evaluation of Reactive Architectures	147
4.6	Summary	148
4.7	Exercises	149
4.8	End Notes	152
5	<i>Designing a Reactive Implementation</i>	155
5.1	Overview	155
5.2	Behaviors as Objects in OOP	157
5.2.1	Example: A primitive move-to-goal behavior	158
5.2.2	Example: An abstract follow-corridor behavior	160
5.2.3	Where do releasers go in OOP?	162
5.3	Steps in Designing a Reactive Behavioral System	163
5.4	Case Study: Unmanned Ground Robotics Competition	165
5.5	Assemblages of Behaviors	173
5.5.1	Finite state automata	174
5.5.2	A Pick Up the Trash FSA	178
5.5.3	Implementation examples	182
5.5.4	Abstract behaviors	184
5.5.5	Scripts	184
5.6	Summary	187
5.7	Exercises	188
5.8	End Notes	191
6	<i>Common Sensing Techniques for Reactive Robots</i>	195
6.1	Overview	196
6.1.1	Logical sensors	197
6.2	Behavioral Sensor Fusion	198

6.3	Designing a Sensor Suite	202
6.3.1	Attributes of a sensor	203
6.3.2	Attributes of a sensor suite	206
6.4	Proprioceptive Sensors	207
6.4.1	Inertial navigation system (INS)	208
6.4.2	GPS	208
6.5	Proximity Sensors	210
6.5.1	Sonar or ultrasonics	210
6.5.2	Infrared (IR)	216
6.5.3	Bump and feeler sensors	217
6.6	Computer Vision	218
6.6.1	CCD cameras	219
6.6.2	Grayscale and color representation	220
6.6.3	Region segmentation	226
6.6.4	Color histogramming	228
6.7	Range from Vision	231
6.7.1	Stereo camera pairs	232
6.7.2	Light stripers	235
6.7.3	Laser ranging	239
6.7.4	Texture	241
6.8	Case Study: Hors d'Oeuvres, Anyone?	242
6.9	Summary	250
6.10	Exercises	251
6.11	End Notes	254
7	<i>The Hybrid Deliberative/Reactive Paradigm</i>	257
7.1	Overview	257
7.2	Attributes of the Hybrid Paradigm	259
7.2.1	Characteristics and connotation of reactive behaviors in hybrids	261
7.2.2	Connotations of "global"	262
7.3	Architectural Aspects	262
7.3.1	Common components of hybrid architectures	263
7.3.2	Styles of hybrid architectures	264
7.4	Managerial Architectures	265
7.4.1	Autonomous Robot Architecture (AuRA)	265
7.4.2	Sensor Fusion Effects (SFX)	268
7.5	State-Hierarchy Architectures	274
7.5.1	3-Tiered (3T)	274

7.6	Model-Oriented Architectures	277
7.6.1	Saphira	278
7.6.2	Task Control Architecture (TCA)	280
7.7	Other Robots in the Hybrid Paradigm	283
7.8	Evaluation of Hybrid Architectures	284
7.9	Interleaving Deliberation and Reactive Control	285
7.10	Summary	288
7.11	Exercises	289
7.12	End Notes	291
8	Multi-agents	293
8.1	Overview	293
8.2	Heterogeneity	296
8.2.1	Homogeneous teams and swarms	296
8.2.2	Heterogeneous teams	297
8.2.3	Social entropy	300
8.3	Control	301
8.4	Cooperation	303
8.5	Goals	304
8.6	Emergent Social Behavior	305
8.6.1	Societal rules	305
8.6.2	Motivation	307
8.7	Summary	309
8.8	Exercises	310
8.9	End Notes	312
II	Navigation	315
9	Topological Path Planning	325
9.1	Overview	325
9.2	Landmarks and Gateways	326
9.3	Relational Methods	328
9.3.1	Distinctive places	329
9.3.2	Advantages and disadvantages	331
9.4	Associative Methods	333
9.4.1	Visual homing	334
9.4.2	QualNav	335

9.5	Case Study of Topological Navigation with a Hybrid Architecture	338
9.5.1	Path planning	339
9.5.2	Navigation scripts	343
9.5.3	Lessons learned	346
9.6	Summary	348
9.7	Exercises	348
9.8	End notes	350
10	<i>Metric Path Planning</i>	351
10.1	Objectives and Overview	351
10.2	Configuration Space	353
10.3	Cspace Representations	354
10.3.1	Meadow maps	354
10.3.2	Generalized Voronoi graphs	357
10.3.3	Regular grids	358
10.3.4	Quadtrees	359
10.4	Graph Based Planners	359
10.5	Wavefront Based Planners	365
10.6	Interleaving Path Planning and Reactive Execution	367
10.7	Summary	371
10.8	Exercises	372
10.9	End Notes	374
11	<i>Localization and Map Making</i>	375
11.1	Overview	375
11.2	Sonar Sensor Model	378
11.3	Bayesian	380
11.3.1	Conditional probabilities	381
11.3.2	Conditional probabilities for $P(H s)$	384
11.3.3	Updating with Bayes' rule	385
11.4	Dempster-Shafer Theory	386
11.4.1	Shafer belief functions	387
11.4.2	Belief function for sonar	389
11.4.3	Dempster's rule of combination	390
11.4.4	Weight of conflict metric	394
11.5	HIMM	395
11.5.1	HIMM sonar model and updating rule	395
11.5.2	Growth rate operator	398

11.6	Comparison of Methods	403
11.6.1	Example computations	403
11.6.2	Performance	411
11.6.3	Errors due to observations from stationary robot	412
11.6.4	Tuning	413
11.7	Localization	415
11.7.1	Continuous localization and mapping	416
11.7.2	Feature-based localization	421
11.8	Exploration	424
11.8.1	Frontier-based exploration	425
11.8.2	Generalized Voronoi graph methods	427
11.9	Summary	428
11.10	Exercises	431
11.11	End Notes	434
12	<i>On the Horizon</i>	435
12.1	Overview	435
12.2	Shape-Shifting and Legged Platforms	438
12.3	Applications and Expectations	442
12.4	Summary	445
12.5	Exercises	445
12.6	End Notes	447
	<i>Bibliography</i>	449
	<i>Index</i>	459