

Springer Tracts in Advanced Robotics

Volume 55

Editors: Bruno Siciliano · Oussama Khatib · Frans Groen

Cyrill Stachniss

Robotic Mapping and Exploration

Professor Bruno Siciliano, Dipartimento di Informatica e Sistemistica, Università di Napoli Federico II, Via Claudio 21, 80125 Napoli, Italy, E-mail: siciliano@unina.it

Professor Oussama Khatib, Artificial Intelligence Laboratory, Department of Computer Science, Stanford University, Stanford, CA 94305-9010, USA, E-mail: khatib@cs.stanford.edu

Professor Frans Groen, Department of Computer Science, Universiteit van Amsterdam, Kruislaan 403, 1098 SJ Amsterdam, The Netherlands, E-mail: groen@science.uva.nl

Author

Dr. Cyrill Stachniss
Albert-Ludwigs-University of Freiburg
Institute of Computer Science
Autonomous Intelligent Systems
Georges-Koehler-Allee 079
79110 Freiburg
Germany
E-mail: stachnis@informatik.uni-freiburg.de

ISBN 978-3-642-01096-5

e-ISBN 978-3-642-01097-2

DOI 10.1007/978-3-642-01097-2

Springer Tracts in Advanced Robotics ISSN 1610-7438

Library of Congress Control Number: Applied for

©2009 Springer-Verlag Berlin Heidelberg

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable for prosecution under the German Copyright Law.

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Typeset & Cover Design: Scientific Publishing Services Pvt. Ltd., Chennai, India.

Printed in acid-free paper

5 4 3 2 1 0

springer.com

Editorial Advisory Board

Oliver Brock, Univ. Massachusetts Amherst, USA
Herman Bruyninckx, KU Leuven, Belgium
Raja Chatila, LAAS, France
Henrik Christensen, Georgia Tech, USA
Peter Corke, CSIRO, Australia
Paolo Dario, Scuola S. Anna Pisa, Italy
Rüdiger Dillmann, Univ. Karlsruhe, Germany
Ken Goldberg, UC Berkeley, USA
John Hollerbach, Univ. Utah, USA
Makoto Kaneko, Osaka Univ., Japan
Lydia Kavraki, Rice Univ., USA
Vijay Kumar, Univ. Pennsylvania, USA
Sukhan Lee, Sungkyunkwan Univ., Korea
Frank Park, Seoul National Univ., Korea
Tim Salcudean, Univ. British Columbia, Canada
Roland Siegwart, ETH Zurich, Switzerland
Guarav Sukhatme, Univ. Southern California, USA
Sebastian Thrun, Stanford Univ., USA
Yangsheng Xu, Chinese Univ. Hong Kong, PRC
Shin'ichi Yuta, Tsukuba Univ., Japan

STAR (Springer Tracts in Advanced Robotics) has been promoted under the auspices of EURON (European Robotics Research Network)



For Janna and Maren

Series Editor's Foreword

By the dawn of the new millennium, robotics has undergone a major transformation in scope and dimensions. This expansion has been brought about by the maturity of the field and the advances in its related technologies. From a largely dominant industrial focus, robotics has been rapidly expanding into the challenges of the human world. The new generation of robots is expected to safely and dependably co-habitat with humans in homes, workplaces, and communities, providing support in services, entertainment, education, health-care, manufacturing, and assistance.

Beyond its impact on physical robots, the body of knowledge robotics has produced is revealing a much wider range of applications reaching across diverse research areas and scientific disciplines, such as: biomechanics, haptics, neurosciences, virtual simulation, animation, surgery, and sensor networks among others. In return, the challenges of the new emerging areas are proving an abundant source of stimulation and insights for the field of robotics. It is indeed at the intersection of disciplines that the most striking advances happen.

The goal of the series of Springer Tracts in Advanced Robotics (STAR) is to bring, in a timely fashion, the latest advances and developments in robotics on the basis of their significance and quality. It is our hope that the wider dissemination of research developments will stimulate more exchanges and collaborations among the research community and contribute to further advancement of this rapidly growing field.

The monograph written by Cyril Stachniss is a contribution in the area of self-localization and mapping (SLAM) for autonomous robots, which has been receiving a great deal of attention by the research community in the latest few years. The contents expand the authors doctoral dissertation and are focused on the autonomous mapping learning problem. Solutions include uncertainty-driven exploration, active loop closing, coordination of multiple robots, learning and incorporating background knowledge, and dealing with dynamic environments. Results are accompanied by a rich set of experiments, revealing a promising outlook toward the application to a wide range of

mobile robots and field settings, such as search and rescue, transportation tasks, or automated vacuum cleaning.

Yet another STAR volume on SLAM, a very fine addition to the series!

Naples, Italy
February 2009

Bruno Siciliano
STAR Editor

Foreword

Simultaneous localization and mapping is a highly important and active area in mobile robotics. The ability to autonomously build maps is widely regarded as one of the fundamental preconditions for truly autonomous mobile robots. In the past, the SLAM has mostly been addressed as a state estimation problem and the incorporation of control into the map learning and localization process is a highly interesting research question. In this book by Cyrill Stachniss, the reader will find interesting and innovative solutions to the problem of incorporating control into the SLAM problem. I know Cyrill since over eight years and I still appreciate his enthusiasm in developing new ideas and getting things done. He has been working with a large number of different robots, participating in several public demonstrations, and has gained a lot of experience which can also be seen from his large number of papers presented at all major robotic conferences and in journals. His work covers a variety of different topics. He has acquired several project grants and received several awards. He furthermore is an associate editor of the IEEE Transactions on Robotics. It's safe to say that he is an expert in his field.

This book is a comprehensive introduction to state-of-the-art technology in robotic exploration and map building. The reader will find a series of solutions to challenging problems robots are faced with in the real world when they need to acquire a model of their surroundings. The book focuses on autonomy and thus the robot is not supposed to be joysticked though the world but should be able to decide about his actions on its own. I regard the ability to learn maps by making own decisions as a key competence for autonomous robots. Cyrill rigorously applies probabilistic and decision-theoretic concepts to systematically reducing the uncertainty in the belief of a robot about its environment and its pose in the environment.

The book contains impressively demonstrates the capabilities of the described solutions by showing results obtained from real robotic datasets. A further strength lies in the sound and thorough evaluation of all presented techniques going beyond the world of simulation. At this point, I would like to encourage the reader to follow Cyrill's example to take real

robots and data obtained with real robots to demonstrate that novel approaches work in reality. For readers not in possession of particular sensors or for comparison purposes, Cyrill and colleagues have created a Web site (<http://www.openslam.org/>) in which the community can share implementations of SLAM approaches and where the reader will find links to datasets to support future research.

Freiburg, Germany
February 2009

Wolfram Burgard

Preface

Models of the environment are needed for a wide range of robotic applications including search and rescue, transportation tasks, or automated vacuum cleaning. Learning maps has therefore been a major research topic in the robotics community over the last decades. Robots that are able to reliably acquire an accurate model of their environment on their own are regarded as fulfilling a major precondition of truly autonomous agents. To autonomously solve the map learning problem, a robot has to address mapping, localization, and path planning at the same time. In general, these three tasks cannot be decoupled and solved independently. Map learning is thus referred to as the simultaneous planning, localization, and mapping problem. Because of the coupling between these tasks, this is a complex problem. It can become even more complex when there are dynamic changes in the environment or several robots are being used together to solve the problem.

This book presents solutions to various aspects of the autonomous map learning problem. The book is separated into two parts. In the first part, we assume the position of the robot to be known. This assumption does not hold in the real world, however, it makes life easier and allows us to better concentrate on certain aspects of the exploration problem such as coordinating a team of robots. We describe how to achieve appropriate collaboration among exploring robots so that they efficiently solve their joint task. We furthermore provide a technique to learn and make use of background knowledge about typical spatial structures when exploring an environment as a team.

In the second part, we relax the assumption that the pose of the robot is known. To deal with the uncertainty in the pose of a robot, we present an efficient solution to the simultaneous localization and mapping problem. The difficulty in this context is to build a map while at the same time localizing the robot in this map. The presented approach maintains a joint posterior about the trajectory of the robot and the model of the environment. It produces accurate maps in an efficient and robust way. After addressing step-by-step the different problems in the context of active map learning, we integrate the main techniques into a single system. We present an integrated approach

that simultaneously deals with mapping, localization, and path planning. It seeks to minimize the uncertainty in the map and in the trajectory estimate based on the expected information gain of future actions. It takes into account potential observation sequences to estimate the uncertainty reduction in the world model when carrying out a specific action. Additionally, we focus on mapping and localization in non-static environments. The approach allows a robot to consider different spatial configurations of the environment and in this way makes the pose estimate more robust and accurate in non-static worlds.

In sum, the contributions of this book are solutions to various problems of the autonomous map learning problem including uncertainty-driven exploration, SLAM, active loop closing, coordination of multiple robots, learning and incorporating background knowledge, and dealing with dynamic environments.

A lot of the work presented in this book has been done in collaboration with other researchers. It was a pleasure for me to work with all the wonderful people in the AIS lab in Freiburg. First of all, I thank Wolfram Burgard for his tremendous support, his inspiration, and for providing a creative atmosphere. My thanks to my friends and colleagues for the great time in the lab, especially to Maren Bennewitz, Giorgio Grisetti, Dirk Hähnel, Óscar Martínez Mozos, Patrick Pfaff, Christian Plagemann, and Axel Rottmann for the great collaboration on the topics addressed in this book. It was a pleasure to work with all these people and to benefit from their knowledge. My thanks also to Mark Moors and Frank Schneider for the collaboration on multi-robot exploration. Special thanks to Nick Roy and Mike Montemerlo who did a great job in developing and maintaining the Carnegie Mellon Robot Navigation Toolkit. It was a pleasure for me to work together with all of them.

Additionally, I thank several people, who published robot datasets and in this way helped to make mapping approaches more robust and more easily comparable. In this context, I would like to thank Patrick Beeson, Mike Bosse, Udo Frese, Steffen Gutmann, Dirk Hähnel, Andrew Howard, and Nick Roy.

Freiburg, Germany
December 2008

Cyrill Stachniss

Contents

1	Introduction	3
2	Basic Techniques	7
2.1	Introduction to Particle Filters	7
2.1.1	Mobile Robot Localization Using Particle Filters	11
2.2	Grid Maps	13
2.2.1	Occupancy Probability Mapping	14
2.2.2	Sensor Model for a Laser Range Finder	16
2.2.3	Sensor Model for a Sonar Sensor	16
2.2.4	Reflection Probability Mapping	18

Part I: Exploration with Known Poses

3	Decision-Theoretic Exploration Using Coverage Maps	23
3.1	Introduction	23
3.2	Definition of Coverage Maps	24
3.3	Updating Coverage Maps Upon Sensory Input	25
3.4	Decision-Theoretic Exploration with Coverage Maps	29
3.4.1	Choosing the Closest Target Location	29
3.4.2	Exploration Using the Information Gain	30
3.4.3	Using IG in a Local Window	32
3.4.4	Combination of IG and CL	32
3.5	Exploration Using Occupancy Grid Maps	32
3.6	Experimental Results	33
3.6.1	Mapping with Noisy Sensors	33
3.6.2	Comparing the Viewpoint Selection Strategies	34
3.6.3	Advantage over Scan Counting	37
3.7	Related Work	38
3.8	Conclusion	40

4 Coordinated Multi-Robot Exploration	43
4.1 Introduction	43
4.2 Coordinating a Team of Robots during Exploration	45
4.2.1 Cost of Reaching a Target Location	46
4.2.2 Computing Utilities of Frontier Cells	47
4.2.3 Target Point Selection	48
4.2.4 Coordination with Limited Communication Range	49
4.3 Collaborative Mapping with Teams of Mobile Robots	50
4.4 Experimental Results	51
4.4.1 Exploration with a Team of Mobile Robots	51
4.4.2 Comparison between Uncoordinated and Coordinated Exploration	52
4.4.3 Simulation Experiments	54
4.4.4 Exploration with Limited Communication	58
4.5 Comparisons to Other Coordination Techniques	59
4.5.1 Target Assignment Using the Hungarian Method	59
4.5.2 Using a Priorization Scheme to Coordinate a Team of Robots	64
4.5.3 Coordination of a Team of Robots by Solving a TSP	66
4.6 Related Work	66
4.7 Conclusion	71
5 Multi-Robot Exploration Using Semantic Place Labels	73
5.1 Introduction	73
5.2 Semantic Place Labeling	74
5.3 Estimating the Label of a Goal Location	77
5.4 Using Semantic Place Information for Efficient Multi-Robot Exploration	80
5.5 Experimental Results	81
5.5.1 Performance Improvement Using Semantic Place Information	81
5.5.2 Influence of Noise in the Semantic Place Information	84
5.5.3 Applying a Trained Classifier in New Environments	84
5.5.4 Improvements of the HMM Filtering and Error Analysis of the Classifier	85
5.6 Related Work	88
5.7 Conclusion	89

Part II: Mapping and Exploration under Pose Uncertainty

6 Efficient Techniques for Rao-Blackwellized Mapping	93
6.1 Introduction	93
6.2 The Concept of Rao-Blackwellized Mapping	94
6.3 Improved Proposals and Selective Resampling	96
6.3.1 Using Laser Range Data to Compute an Improved Proposal Distribution	97
6.3.2 Selective Resampling	100
6.4 Complexity	101
6.5 Experimental Results	102
6.5.1 Mapping Results	103
6.5.2 Quantitative Results	105
6.5.3 Effects of Improved Proposals and Adaptive Resampling	107
6.5.4 Situations in Which the Scan-Matcher Fails	109
6.5.5 Computational Cost	112
6.6 Related Work	112
6.7 Conclusion	115
7 Actively Closing Loops During Exploration	117
7.1 Introduction	117
7.2 Active Loop-Closing	118
7.2.1 Detecting Opportunities to Close Loops	119
7.2.2 Representing Actions under Pose Uncertainty	121
7.2.3 Stopping the Loop-Closing Process	122
7.2.4 Reducing the Exploration Time	124
7.2.5 Handling Multiple Nested Loops	126
7.3 Experimental Results	126
7.3.1 Real World Exploration	127
7.3.2 Active Loop-Closing vs. Frontier-Based Exploration	127
7.3.3 A Quantitative Analysis	128
7.3.4 Importance of the Termination Criterion	129
7.3.5 Evolution of N_{eff}	130
7.3.6 Multiple Nested Loops	132
7.3.7 Computational Resources	132
7.4 Related Work	133
7.5 Conclusion	134
8 Recovering Particle Diversity	135
8.1 Introduction	135
8.2 Recovering Particle Diversity after Loop Closure	137
8.3 Experimental Results	138

8.4 Related Work	141
8.5 Conclusion	142
9 Information Gain-based Exploration	143
9.1 Introduction	143
9.2 The Uncertainty of a Rao-Blackwellized Mapper	145
9.3 The Expected Information Gain	147
9.4 Computing the Set of Actions	151
9.5 Experimental Results	152
9.5.1 Real World Application	152
9.5.2 Decision Process	154
9.5.3 Comparison to Previous Approaches	154
9.5.4 Corridor Exploration	157
9.6 Related Work	157
9.7 Conclusion	160
10 Mapping and Localization in Non-Static Environments	161
10.1 Introduction	161
10.2 Learning Maps of Low-Dynamic Environments	163
10.2.1 Map Segmentation	163
10.2.2 Learning Configurations of the Environment	163
10.2.3 Map Clustering	164
10.3 Monte-Carlo Localization Using Patch-Maps	166
10.4 Experimental Results	168
10.4.1 Application in an Office Environment	168
10.4.2 Localizing the Robot and Estimating the State of the Environment	171
10.4.3 Global Localization	173
10.5 Related Work	173
10.6 Conclusion	175
11 Conclusion	177
A Appendix	181
A.1 Probability Theory	181
A.1.1 Product Rule	181
A.1.2 Independence	181
A.1.3 Bayes' Rule	181
A.1.4 Marginalization	182
A.1.5 Law of Total Probability	182
A.1.6 Markov Assumption	182
References	183
Index	195

Notation

Throughout this book, we make use of the following notation:

variable	description
x_t	pose of the robot at time step t . This pose is a three dimensional vector containing the x, y -position and the orientation θ of the vehicle
$x_{1:t}$	sequence of poses of the robot from time step 1 to time step t
z_t	sensor observation obtained at time step t
u_t	odometry information describing the movement from x_t to x_{t+1}
a	action or motion command
w	importance weight
$w_t^{[i]}$	importance weight of the i -th particle at time step t
m	grid map
c	grid cell
r	resolution of a grid map. Each cell covers an area of r by r .
\mathcal{G}	topological map
$E[\cdot]$	expectation
$\mathcal{N}(\mu, \Sigma)$	Gaussian with mean μ and covariance Σ
H	entropy
I	information gain
U	utility function
V	cost function
η	normalizer, typically resulting from Bayes' rule
N_{eff}	effective number of particles