

Springer Tracts in Advanced Robotics

Volume 52

Editors: Bruno Siciliano · Oussama Khatib · Frans Groen

Andreas Nüchter

3D Robotic Mapping

The Simultaneous Localization and Mapping
Problem with Six Degrees of Freedom



Springer

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, Robotics 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. Andreas Nüchter
University of Osnabrück
Institute of Computer Science
Albrechtstraße 28
Room: 31/506
49069 Osnabrück
Germany
E-Mail: Andreas@nuechti.de

ISBN 978-3-540-89883-2

e-ISBN 978-3-540-89884-9

DOI 10.1007/978-3-540-89884-9

Springer Tracts in Advanced Robotics ISSN 1610-7438

Library of Congress Control Number: 2008941001

©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

Herman Bruyninckx, KU Leuven, Belgium
Raja Chatila, LAAS, France
Henrik Christensen, Georgia Institute of Technology, USA
Peter Corke, CSIRO, Australia
Paolo Dario, Scuola Superiore Sant'Anna Pisa, Italy
Rüdiger Dillmann, Universität Karlsruhe, Germany
Ken Goldberg, UC Berkeley, USA
John Hollerbach, University of Utah, USA
Makoto Kaneko, Osaka University, Japan
Lydia Kavraki, Rice University, USA
Sukhan Lee, Sungkyunkwan University, Korea
Tim Salcudean, University of British Columbia, Canada
Sebastian Thrun, Stanford University, USA
Yangsheng Xu, Chinese University of Hong Kong, PRC
Shin'ichi Yuta, Tsukuba University, Japan

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



To Jutta

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 Andreas Nüchter is yet another title in the series belonging to the fertile research area of self-localization and mapping (SLAM) for mobile robots. The main contribution of the work is certainly in the derivation of SLAM algorithms for building 3D maps, as opposed to the majority of existing algorithms for 2D maps. When tackling the problem for a 3D geometry, all six degrees of freedom are considered to reconstruct the robot pose. Results span from theory through a rich set of experiments, revealing a promising outlook toward benchmarking SLAM robot maps in a wide range of applications.

Expanded from the author's doctoral dissertation (finalist for the Seventh Edition of the EURON Georges Giralt PhD Award) and overall eight years of research and teaching on SLAM, this volume constitutes a very fine addition to STAR!

Naples, Italy,
August 2008

Bruno Siciliano
STAR Editor

Foreword

Have you ever jumped to bang on the emergency stop button of your mobile robot, to prevent it from crashing into an open drawer above its collision avoidance laser scan plane? That was the moment when you learned that the world is a 3D place. Well, maybe you guessed it before, which is the reason why you would painstakingly avoid to let the robot get close to staircases: You do know, but your robot may not know, that this wide open area so tempting in the 2D scan on which its heading direction control is based, is right above the flight of stairs down.

While it has always been possible to calculate 3D geometry from stereo images, recent years have witnessed creative usage of 2D laser scanners to get it – be it by having scanners rotate or pitch, or by mounting them at angles between the scan plane and the robot motion direction. Given the rocketing in performance of SLAM for 2D maps during the same period, it is only natural to think about building 3D geometry maps using these sensor data, so: How turn into a consistent 3D geometry map the point clouds generated by a 3D laser scanner on-board a mobile robot?

The obvious approach to do this is: Take your favorite SLAM algorithm for building 2D maps, say, a Particle Filter algorithm, and apply it to 3D rather than 2D points. That could work, but there is a problem. When thinking about 3D geometry, it makes sense to consider for the robot rigid body motion in 3D, i.e., in six degrees of freedom (6 DoF: x, y, z , yaw, pitch, roll). A typical SLAM algorithm for building 2D maps from 2D laser scans would consider only robot motion on a plane, i.e., 3 DoF poses. If it uses for building the map an approximation of the full posterior distribution of robot poses in pose space, then changing from 3 to 6 DoF poses lets explode the distribution. Put briefly, not all that works well in 3 DoF pose space is practical in 6 DoF.

Andreas Nüchter's book is a comprehensive introduction into approaches that do work well on 3D scan data in 6 DoF – in fact, it is the most comprehensive that I know of. It is based on Andreas's own research and teaching on the topic since about the year 2000. This has involved a number of different robots, and a healthy number of applications; it has earned him a diploma degree and a doctorate; it has yielded for his respective working groups grants and awards, such as a vice-championship in the RoboCup real rescue league (2004 in Lisbon), not to mention papers in all major robotics conferences and

journals. So one can safely say, this author knows his topic well. Moreover, Andreas could exploit various teaching events for annealing the presentation.

The basic technique that underlies 6DoF SLAM here is the good old ICP (iterative closest/corresponding points) algorithm for registering 3D point clouds. If you don't know ICP yet, you will learn it from Chapter 4; read Chapter 4, too, if you ever wanted to know how to implement and use ICP really efficient, or how to choose between different versions of the math behind the required 6 DoF transformations or of the algorithms behind finding corresponding points. Then, scan registration is the basis of mapping here, but left alone would normally produce maps with unbounded errors. To build good maps, SLAM needs information about closed loops in the robot trajectory – this is no different for 6 DoF than it is for 3 DoF. Chapter 5 describes how this works in 6 DoF. The noblest technique here is a version of the Graph-SLAM algorithm, maximizing the joint probability of 6 DoF scan poses. So we are back at a probabilistic approach, yet a uni-modal one, for the good reason stated above.

The book also contains a healthy amount of results on data from real applications, where many of the data sets can be downloaded from an online 3D scan repository, to reproduce the results or do own experiments. Here, I think the topic of Benchmarking is particularly interesting in terms of methodology. You find many papers with SLAMmed robot maps that “look good”; you find few papers, at least if it comes to 3D maps, that question explicitly whether their maps are also truthful. In fact, that question is often hard to answer if the mapped area is thousands of square meters in open terrain with significant elevation differences. Section 6.4 tells you how it could be answered anyway. Finally, there is a section about Semantic Mapping in 3D: If benchmarking is a hot topic for applications, bringing semantics into robot-generated maps and into the mapping process is most probably the emerging hot topic for AI related robotics research.

So if you are interested in getting serious with 3D geometry in robotics, I recommend strongly that you read and understand this book. If you aren't interested in robotics, but only in 3D maps, there may be good reasons to read it anyway: The map building algorithms presented here don't really care whether their data is taken on-board a mobile robot or on a surveyor's tripod or on some hand-held scanning device. Give them 3D point clouds with rough 6 DoF pose estimates, and they allow consistent maps to be built without further manual interference. I am looking forward to seeing a lot of applications of that – in robotics and beyond.

Osnabrück, Germany
Juli 2008

Joachim Hertzberg

Preface

Intelligent autonomous acting in unstructured environments requires 3D maps with labelled 3D objects. 3D maps are necessary to avoid collisions with complex obstacles and to self-localize in six degrees of freedom (x -, y -, z -position, roll, yaw and pitch angle). Meaning becomes inevitable, if the robot has to interact with its environment. The robot is then able to reason about the objects; its knowledge becomes inspectable and communicable. These arguments lead to the necessity of automatic and fast semantic environment modelling in robotics. A revolutionary method for gaging environments are 3D scanners, which enable robots to scan objects in a non-contact way in three dimensions. The presented work examines and evaluates the algorithms needed for automatic semantic 3D map building using a 3D laser range finder and the mobile robot Kurt3D.

The first part of this book deals with the task of registering 3D scans in a common coordinate system. Correct, globally consistent models result from a 6D SLAM algorithm. Hereby 6 degrees of freedom of the robot pose are considered, closed-loops are detected and the global error is minimized. 6D SLAM is based on a very fast ICP algorithm. While initially written for the mobile robot Kurt3D, the algorithms have also been applied in many other contexts. The second part presents a wide variety of working experiments and applications, including the results of presentations, like the Kurt3D presentation at RoboCup.

In the last part semantic descriptions are derived from the 3D point model. For that purpose 3D planes are detected and interpreted in the digitalized 3D scene. After that an efficient algorithm detects objects and estimates their pose with 6 degrees of freedom.

Of course, this book is the result of my active cooperation with many researchers. I would like to express my thanks to my supervisor Joachim Hertzberg who has always supported me. I am also thankful for the joint preceding research on 3D robotic mapping with: Dorit Borrmann, Thomas Christaller, Jan Elseberg, Simone Frintrop, Dominik Giel, Matthias Hennig, Joachim Hertzberg, Achim Lilienthal, Kai Lingemann, Martin Magnusson,

Stefan Stiene, Hartmut Surmann, Sebastian Thrun, Bernardo Wagner, Oliver Wulf. Special thanks go to Alexander Loktyuhsin for proofreading the book.

My intention is that the reader of this book will use this compilation of algorithms and results to build smarter robotic systems that can cope with real world environments using 3D sensor data.

Osnabrück, Germany
Juli 2008

Andreas Nüchter

Contents


| | | |
|----------|--|----|
| 1 | Introduction | 1 |
| 1.1 | A Brief Introduction to Robotic Mapping | 2 |
| 1.2 | The Big Picture | 7 |
| 2 | Perceiving 3D Range Data | 9 |
| 2.1 | Basic Principles | 9 |
| 2.1.1 | Time-of-Flight Measurement | 9 |
| 2.1.2 | Phase-Shift Measurement | 10 |
| 2.1.3 | Triangulation | 11 |
| 2.2 | Laser Scanner | 11 |
| 2.2.1 | 1D Laser Measurement System | 11 |
| 2.2.2 | 2D Laser Scanner | 11 |
| 2.2.3 | 3D Laser Scanner | 12 |
| 2.2.4 | Projection 3D Scanner | 14 |
| 2.3 | Cameras and Camera Models | 15 |
| 2.3.1 | The Pinhole Camera Model and Perspective Projection | 16 |
| 2.3.2 | Stereo Cameras | 19 |
| 2.4 | 3D Cameras | 26 |
| 3 | State of the Art | 29 |
| 3.1 | Categorization of Robotic Mapping Algorithms | 29 |
| 3.1.1 | Planar 2D Mapping | 29 |
| 3.1.2 | Planar 3D Mapping | 31 |
| 3.1.3 | Slice-Wise 6D SLAM | 32 |
| 3.1.4 | Full 6D SLAM | 32 |
| 3.1.5 | Recent Trends | 32 |
| 3.2 | Globally Consistent Range Image Alignment | 33 |

| | | |
|----------|--|-----|
| 4 | 3D Range Image Registration | 35 |
| 4.1 | The ICP Algorithm | 35 |
| 4.1.1 | Direct Solutions of the ICP Error Function | 36 |
| 4.1.2 | Approximate Solution of the ICP Error Function by a Helical Motion | 48 |
| 4.1.3 | Linearized Solution of the ICP Error Function | 50 |
| 4.1.4 | Computing Closest Points | 52 |
| 4.1.5 | Implementation Issues | 64 |
| 4.1.6 | The Parallel ICP Algorithm | 67 |
| 4.2 | Evaluation of the ICP Algorithm | 69 |
| 5 | Globally Consistent Range Image Registration | 77 |
| 5.1 | Explicit Loop Closing | 77 |
| 5.1.1 | Extrapolating the Odometry to Six Degrees of Freedom | 77 |
| 5.1.2 | Closed Loop Detection | 78 |
| 5.1.3 | Error Diffusion | 79 |
| 5.2 | GraphSLAM as Generalization of Loop Closing | 81 |
| 5.2.1 | Algorithm Overview | 81 |
| 5.2.2 | Probabilistic Global Relaxation Based on Scan Matching | 81 |
| 5.2.3 | Transforming the Solution | 87 |
| 5.3 | Other Modeling Approaches | 93 |
| 5.3.1 | Linearization Using Quaternions | 93 |
| 5.3.2 | GraphSLAM Using Helical Motion | 99 |
| 5.3.3 | GraphSLAM Using Linearization of a Rotation | 104 |
| 6 | Experiments and Results | 109 |
| 6.1 | 3D Mapping with Kurt3D | 109 |
| 6.1.1 | The 3D Laser Range Finder | 109 |
| 6.1.2 | The Mobile Robot Kurt3D | 110 |
| 6.1.3 | Full 6D SLAM in a Planar Environment | 111 |
| 6.1.4 | Full 6D SLAM in an Indoor/Outdoor Environment .. | 112 |
| 6.1.5 | Full 6D SLAM in an Outdoor Environment | 114 |
| 6.1.6 | Kurt3D @ Competitions | 121 |
| 6.2 | Globally Consistent Registration of High Resolution Scans .. | 123 |
| 6.3 | Mapping Urban Environments – Registration with Dynamic Network Construction | 125 |
| 6.4 | Benchmarking 6D SLAM | 130 |
| 6.4.1 | Ground Truth Experiments | 133 |
| 6.4.2 | The Benchmarking Technique | 134 |
| 6.4.3 | Experimental Results | 138 |
| 6.4.4 | Justification of the Results | 141 |
| 6.5 | Further Applications | 145 |
| 6.5.1 | Mapping Abandoned Mines | 145 |
| 6.5.2 | Applications in Medicine | 151 |

| | | |
|----------|---|-----|
| 7 | 3D Maps with Semantics | 155 |
| 7.1 | Scene Interpretation | 155 |
| 7.1.1 | Feature Extraction | 155 |
| 7.1.2 | Interpretation | 156 |
| 7.1.3 | Application: Model Refinement..... | 160 |
| 7.2 | Object Localization in 3D Data | 163 |
| 7.2.1 | Object Detection | 163 |
| 7.2.2 | Object Localization | 169 |
| 7.3 | Semantic 3D Maps | 170 |
| 8 | Conclusions and Outlook | 173 |
| A | Math Appendix | 177 |
| A.1 | Representing Rotations..... | 177 |
| A.1.1 | Euler Angles | 177 |
| A.1.2 | Axis Angle | 178 |
| A.1.3 | Unit Quaternions..... | 179 |
| A.1.4 | Converting Rotation Representations | 181 |
| A.2 | Plücker Coordinates | 182 |
| A.3 | Additional Theorems from Linear Algebra | 183 |
| A.4 | Solving Linear Systems of Equations | 186 |
| A.5 | Computing Constrained Extrema | 187 |
| A.5.1 | Computing the Minimum of a Quadratic Form with Subject to Conditions | 188 |
| A.6 | Numerical Function Minimization | 190 |
| | References | 193 |

Notation

| | |
|---|--|
| \mathbb{R} | set of reel numbers |
| $\mathbb{1}$ | identity matrix |
| M, D | set of measured data, subset of \mathbb{R}^3 |
| x, y, z, \dots | scalars |
| $\mathbf{x}, \mathbf{y}, \mathbf{z}, \dots$ | vectors |
| $\hat{\mathbf{q}}, \hat{\mathbf{s}}, \dots$ | quaternions |
| $\mathbf{R}, \mathbf{H}, \dots$ | matrices, including rotation matrices |
| $\mathbf{M}, \mathbf{G}, \dots$ | matrices and vectors that have been concatenated |
| $\text{tr}(\mathbf{M})$ | trace of the matrix \mathbf{M} |



Definition of the coordinate system.