


ROBOTICS

SCIENCE AND SYSTEMS III



EDITED BY

WOLFRAM BURGARD, OLIVER BROCK, AND CYRILL STACHNISS

TP24-53
R 666.2
2007

Robotics

Science and Systems III

edited by Wolfram Burgard, Oliver Brock, and Cyrill Stachniss



E2008001337

The MIT Press
Cambridge, Massachusetts
London, England

© 2008 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.

For information on quantity discounts, email special_sales@mitpress.mit.edu.

Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Robotics : science and systems III / edited by Wolfram Burgard, Oliver Brock, and Cyrill Stachniss.
p. cm.

Includes bibliographical references.

ISBN 978-0-262-52484-1 (pbk. : alk. paper)

1. Robotics. I. Burgard, Wolfram. II. Brock, Oliver. III. Stachniss, Cyrill.

TJ211.R57485 2008

629.8'92—dc22

2007041215

Preface

This volume contains the papers presented at the third *Robotics Science and Systems (RSS)* Conference, held in Atlanta, Georgia, from June 27 to June 30, 2007. Set in the intricately ornamented Egyptian Ballroom of the Fox Theatre, the conference brought together over 230 researchers from around the world. *RSS* strives to showcase the leading scientific research and the most advanced systems from all areas of robotics. This goal is reflected in the diversity of the topics covered in this volume. The papers collected here cover work in mechanism design, motion planning, control, computer vision, mobile robotics, mapping, navigation, machine learning, manufacturing, distributed systems, sensor networks, self-reconfigurable systems, dynamic simulation, robot safety, and field applications.

The *RSS* area chairs and program committee carefully reviewed 137 submissions and selected 41 papers for publication. The review process was double-blind, that is, program committee members were not told authors' identities in the reviewing process. Authors were given the opportunity to respond to the reviewers' comments in a rebuttal process. The final program was selected during an in-person meeting at MIT. This meeting was attended by the general chair, the program chair, and the area chairs. Decisions were based on the recommendations of three to seven program committee members (most papers received four reviews), discussions among the reviewers, the rebuttals of the authors, and intense discussions among the area chairs. Twenty-three of the papers were presented orally in a single conference track; eighteen additional papers were presented at a vibrant and stimulating poster session. Each poster was advertised in a plenary spotlight presentation session.

In addition to the contributed papers, the technical program also featured five invited talks. Shree Nayar from Columbia University gave an inspirational presentation on *Computational Cameras: Redefining the Image*; Arthur Horwich from Yale University took his audience on a breathtaking journey into the miraculous world of proteins, reporting on *GroEL: A Protein Folding Machine*; Atsushi Iriki from the Riken Brain Science Institute gave an insightful and fascinating account of *Latent Precursors of Human Intelligence in Monkey Tool Use Actions*; Daniel Wolpert from the University of Cambridge gave an exciting and thought-provoking presentation on *Probabilistic Models of Human Sensorimotor Control*; and Mathieu Desbrun from California Institute of Technology shared his deep insights in a superb presentation on *Calculus Ex Geometrica: Structure-Preserving Computational Foundations for Graphics and Simulation*.

This year's *RSS* introduced a new feature: the Early Career Spotlight Presentations. These presentations allow promising early-career researchers to present their research vision beyond the scope of a single conference or journal paper. Noah Cowan from Johns Hopkins University gave a fascinating account of his work on *Sensorimotor Integration in Robots and Animals: Signals, Geometry and Mechanics*; and Hod Lipson from Cornell University presented surprising and intriguing work on *Morphological and Behavioral Adaptation in Robotics*.

Ken Goldberg from the University of California at Berkeley gave a truly amazing banquet presentation about *Earth Art with Robots and Networks*. The aquariums surrounding the banquet room at the Georgia Aquarium provided a perfect setting for the symbiosis of art and science.

For the second year in a row, *RSS* featured a best student paper award, sponsored by Springer on behalf of the journal *Autonomous Robots*. At the area chair meeting, the area chairs had selected the following three finalists among the many nominees: *An Implicit Time-Stepping Method for Multibody Systems with Intermittent Contact* by Nilanjan Chakraborty, Stephen Berard, Srinivas Akella, and Jeff Trinkle; *Dimensionality Reduction Using Automatic Supervision for Vision-based Terrain Learning* by Anelia Angelova, Larry Matthies, Daniel Helmick, and Pietro Perona; and *A Fundamental Tradeoff Between Performance and Sensitivity Within Haptic Rendering* by Paul Griffiths, Brent Gillespie, and Jim Freudenber. Congratulations to these authors for their outstanding contribution! The best student paper award went to Nilanjan Chakraborty, Stephen Berard, Srinivas Akella, and Jeff Trinkle from the Rensselaer Polytechnic Institute.

Three days of technical sessions were followed by four one-day workshops: *Robot Manipulation: Sensing and Adapting to the Real World*, organized by Charles Kemp, Aaron Edsinger, Robert Platt, and Neo Ee Sian; *Robotic Sensor Networks: Principles and Practice*, organized by Gaurav Sukhatme and Wolfram Burgard; *Algorithmic Equivalences Between Biological and Robotic Swarms*, organized by James McLurkin and Paulina Varshavskaya; and *Research in Robots for Education*, organized by Doug Blank, Maria Hybinette, Keith O'Hara, and Daniela Rus. These workshops were complemented by the tutorial *Microsoft Robotics Studio (MSRS): A Technical Introduction*, organized by Stewart Tansley and Joseph Fernando. We would like to thank the organizers of these events.

The success of RSS is fueled by the dedication, enthusiasm, and effort of members of the robotics community. We express our gratitude to the area chairs, who spent endless hours reading papers, managing reviews, and debating acceptance decisions; some of them traveled between continents to attend the in-person area chair meeting: Nancy Amato (Texas A&M University), Darius Burschka (Technical University Munich), Jaydev P. Desai (University of Maryland, College Park), Dieter Fox (University of Washington), Hiroshi Ishiguro (University of Osaka), Yokoi Kazuhito (AIST), Yoky Matsuoka (University of Washington), Brad Nelson (ETH Zurich), Paul Newman (University of Oxford), Allison Okamura (Johns Hopkins University), Nicholas Roy (MIT), Roland Siegwart (ETH Zurich), Jeff Trinkle (Rensselaer Polytechnic Institute), Jing Xiao (University of North Carolina Charlotte), and Katsu Yamane (University of Tokyo). We would also like to thank the approximately 150 program committee members for providing extensive, detailed, and constructive reviews.

We thank the workshop chair, Udo Frese from the University of Bremen, for orchestrating the workshop proposal and selection process.

Frank Dellaert and Magnus Egerstedt from Georgia Tech were the local arrangement chairs this year. They combined Southern hospitality with excellent organization and fabulous venues. Their attention to detail ensured that the meeting and workshops went smoothly. We would like to especially thank Teri Russell, Jacque Berry, Vivian Chandler, Katrien Hemelsoet, and Jennifer Beattie for helping with the organization of poster sessions, the banquet, registration, and many other aspects of the conference, assisted by a wonderful group of student volunteers.

We thank the researchers from the following Georgia Tech robotics labs for live demonstrations: the Sting Racing Team (Dave Wooden, Matt Powers), the GRITS Lab (Jean-Pierre de la Croix, Magnus Egerstedt), the BORG Lab (Roozbeh Mottaghi, Frank Dellaert), and the HumAnSlab (Ayanna Howard). We also thank Henrik Christensen for the KUKA demo.

We are extremely grateful for the generosity of our sponsors. We would like to thank Stewart Tansley from Microsoft Corporation, Jeff Walz from Google, Rainer Bischoff from Kuka, Paolo Pirjanian from Evolution Robotics, and Melissa Fearon from Springer for their continued support of RSS. We would also like to thank the exhibitors from Microsoft Research (Stewart Tansley) and MIT Press (John Costello).

Finally, we would like to thank the members of the robotics community for their contributions. By submitting their outstanding research and by attending the conference, they provided the most essential ingredients for a successful scientific meeting. Through their active participation, RSS has become a forum for intellectual exchange as well as a high-quality robotics conference. The high ratio of attendees to accepted papers attests to this. We are looking forward to the next event in Zurich, Switzerland, in June 2008.

Wolfram Burgard, Albert-Ludwigs Universität Freiburg
Oliver Brock, University of Massachusetts Amherst
Cyrill Stachniss, Albert-Ludwigs Universität Freiburg
September 2007

Conference Organizers

General Chair	Wolfram Burgard, University of Freiburg
Program Chair	Oliver Brock, University of Massachusetts Amherst
Local Arrangement Chairs	Frank Dellaert, Georgia Tech Magnus Egerstedt, Georgia Tech
Publicity Chair	Cyrill Stachniss, University of Freiburg
Publication Chair	Cyrill Stachniss, University of Freiburg
Workshop Chair	Udo Frese, University of Bremen
Web Masters	Cyrill Stachniss, University of Freiburg Christian Plagemann, University of Freiburg
Conference Management System	Oliver Brock, University of Massachusetts Amherst
Childcare	Ayanna Howards, Georgia Tech Maren Bennewitz, University of Freiburg
Area Chairs	Nancy Amato, Texas A&M University Darius Burschka, TU Munich Jaydev P. Desai, University of Maryland, College Park Dieter Fox, University of Washington Hiroshi Ishiguro, University of Osaka Yokoi Kazuhito, AIST Yoky Matsuoka, University of Washington Brad Nelson, ETH Zurich Paul Newman, University of Oxford Allison Okamura, Johns Hopkins University Nicholas Roy, MIT Roland Siegwart, ETH Zurich Jeff Trinkle, Rensselaer Polytechnic Institute Jing Xiao, University of North Carolina Charlotte Katsu Yamane, University of Tokyo
Conference Board	Oliver Brock, University of Massachusetts Amherst Wolfram Burgard, University of Freiburg Greg Dudek, McGill University Dieter Fox, University of Washington Lydia Kavraki, Rice University Eric Klavins, University of Washington Nicholas Roy, Massachusetts Institute of Technology Daniela Rus, Massachusetts Institute of Technology Sven Koenig, University of Southern California John Leonard, Massachusetts Institute of Technology Stefan Schaal, University of Southern California Gaurav Sukhatme, University of Southern California Sebastian Thrun, Stanford University

Advisory Board

George Bekey, University of Southern California
Raja Chatila, Centre National de la Recherche Scientifique
Henrik Christensen, Royal Institute of Technology
Hugh Durrant-Whyte, Australian Centre for Field Robotics
Ken Goldberg, UC Berkeley
John Hollerbach, University of Utah
Hiroshi Ishiguro, University of Osaka
Vijay Kumar, University of Pennsylvania
Matt Mason, Carnegie Mellon University
Yoshihiko Nakamura, University of Tokyo
Brad Nelson, ETH Zurich
Bruno Siciliano, Universit degli Studi di Napoli Federico II
Roland Siegwart, ETH Zurich

Program Committee

Farhad Aghili	Francois Chaumette	Hideki Hashimoto
Peter Allen	Gordon Cheng	Shinichi Hirai
Sean Andersson	Woojin Chung	Yasuhisa Hirata
Kai Arras	Jacob Crandall	Koh Hosoda
James Bagnell	Andrew Davison	David Hsu
Tim Bailey	Matt Deans	Tetsunari Inamura
Ravi Balasubramanian	Rdiger Dillmann	Pekka Isto
Devin Balkcom	Zachary Dodds	Patric Jensfelt
Stacy Bamberg	Lixin Dong	Yan-Bin Jia
Christoph Bartneck	Gregory Dudek	Chris Jones
Burchan Bayazit	Ryan Eustice	Shuuji Kajita
Patrick Beeson	Ronald Fearing	Takayuki Kanda
Michael Beetz	John Folkesson	Imin Kao
Mike Benjamin	Emilio Frazzoli	Lydia Kavraki
Peter Berkelman	Antonio Frisoli	Alonzo Kelly
Antonio Bicchi	Edwardo Fukushima	Abderrahmane Kheddar
Philippe Bidaud	Darren Gawley	Georg Klein
Michael Bosse	Maria Gini	Atsushi Konno
Michael Bowling	Ken Goldberg	Danica Kragic
Jean-Marc Breguet	Geoff Gordon	Venkat Krovi
Bambi Brewer	Giorgio Grisetti	Ben Krse
Herman Bruyninckx	Horst-Michael Gross	James Kuffner
Karl Bhringer	Li Han	Steve LaValle
Cenk Cavusoglu	Kensuke Harada	Cecilia Laschi
Raja Chatila	Matthias Harders	John Leonard

Program Committee (cont.)

Zexiang Li
Jyh-Ming Lien
Ming Lin
Guanfeng Liu
Yunhui Liu
Max Lungarella
Yusuke Maeda
Ezio Malis
Dinesh Manocha
Eric Marchand
Agostino Martinelli
Stephen Mascaro
Yoshio Matsumoto
Takashi Minato
Bud Mishra
Mark Moll
Michael Montemerlo
Taketoshi Mori
Jun Nakanishi
Jos Neira
Andrew Ng
Gunter Niemeyer
Jenkins Odest
Tetsuya Ogata

Paul Oh
Edwin Olson
Shinsuk Park
William Peine
Roland Philippsen
Joelle Pineau
Ingmar Posner
Domenico Prattichizzo
Thomas Preuss
Fabio Ramos
Rajesh Rao
Christopher Rasmussen
Iead Rezek
Chris Roman
Stergios Roumeliotis
Jee-Hwan Ryu
Jesus Savage
Derik Schroeter
Luis Sentis
Tomohiro Shibata
Gabe Sibley
Rob Sim
Thierry Simeon
Metin Sitti

Dezhen Song
Guang Song
Mark Spong
Cyrill Stachniss
Raul Suarez
Thomas Sugar
Salah Sukkarieh
Yu Sun
Yasutake Takahashi
Hong Tan
Juan Tardos
Russ Tedrake
Frank Tendick
Sebastian Thrun
Rudolph Triebel
Markus Vincze
Michael Wang
Alfredo Weitzenfeld
Stefan Williams
Dirk Wollherr
Robert Wood
Jizhong Xiao
Yuichiro Yoshikawa
Nakamura Yutaka

Contents

Preface.....	vii
Semantic Modeling of Places using Objects.....	1
<i>Ananth Ranganathan and Frank Dellaert</i>	
Design of a Bio-inspired Dynamical Vertical Climbing Robot	9
<i>Jonathan E. Clark, Daniel I. Goldman, Pei-Chun Lin, Goran Lynch, Tao S. Chen, Haldun Komsuoglu, Robert J. Full, and Daniel Koditschek</i>	
Online Learning for Offroad Robots: Using Spatial Label Propagation to Learn Long-Range Traversability	17
<i>Raia Hadsell, Pierre Sermanet, Ayse Naz Erkan, Jan Ben, Jefferson Han, Beat Flepp, Urs Muller, and Yann LeCun</i>	
Composition of Vector Fields for Multi-Robot Manipulation via Caging	25
<i>Jonathan Fink, Nathan Michael, and Vijay Kumar</i>	
Closed Loop Control of a Gravity-assisted Underactuated Snake Robot with Application to Aircraft Wing-Box Assembly.....	33
<i>Binayak Roy and H. Harry Asada</i>	
Predicting Partial Paths from Planning Problem Parameters.....	41
<i>Sarah Finney, Leslie Kaelbling, and Tomás Lozano-Pérez</i>	
Emergent Task Allocation for Mobile Robots	49
<i>Nuzhet Atay and Burchan Bayazit</i>	
Passivity-Based Switching Control for Stabilization of Wheeled Mobile Robots	57
<i>Dongjun Lee</i>	
A Tree Parameterization for Efficiently Computing Maximum Likelihood Maps using Gradient Descent	65
<i>Giorgio Grisetti, Cyrill Stachniss, Slawomir Grzonka, and Wolfram Burgard</i>	
Spatially-Adaptive Learning Rates for Online Incremental SLAM	73
<i>Edwin Olson, John Leonard, and Seth Teller</i>	
Adaptive Non-Stationary Kernel Regression for Terrain Modeling	81
<i>Tobias Lang, Christian Plagemann, and Wolfram Burgard</i>	
Fishbone Model for Belt Object Deformation	89
<i>Hidefumi Wakamatsu, Eiji Arai, and Shinichi Hirai</i>	
Context and Feature Sensitive Re-sampling from Discrete Surface Measurements	97
<i>Dave M. Cole and Paul M. Newman</i>	
Simultaneous Localisation and Mapping in Dynamic Environments (SLAMIDE) with Reversible Data Association.....	105
<i>Charles Bibby and Ian Reid</i>	
Sliding Mode Formation Tracking Control of a Tractor and Trailer-Car System.....	113
<i>Fabio Morbidi and Domenico Prattichizzo</i>	
Map-Based Precision Vehicle Localization in Urban Environments.....	121
<i>Jesse Levinson, Michael Montemerlo, and Sebastian Thrun</i>	
Dense Mapping for Range Sensors: Efficient Algorithms and Sparse Representations	129
<i>Manuel Yguel, Christopher Tay Meng Keat, Christophe Braillon, Christian Laugier, and Olivier Aycard</i>	
Gaussian Beam Processes: A Nonparametric Bayesian Measurement Model for Range Finders.....	137
<i>Christian Plagemann, Kristian Kersting, Patrick Pfaff, and Wolfram Burgard</i>	
Vision-Aided Inertial Navigation for Precise Planetary Landing: Analysis and Experiments	145
<i>Anastasios I. Mourikis, Nikolas Trawny, Stergios I. Roumeliotis, Andrew Johnson, and Larry Matthies</i>	
Optimal Kinodynamic Motion Planning for 2D Reconfiguration of Self-Reconfigurable Robots	153
<i>John Reif and Sam Slee</i>	
A Discrete Geometric Optimal Control Framework for Systems with Symmetries.....	161
<i>Marin Kobilarov, Mathieu Desbrun, Jerrold E. Marsden, and Gaurav S. Sukhatme</i>	

BS-SLAM: Shaping the World	169
<i>Luis Pedraza, Gamini Disisanayake, Jaime Valls Miro, Diego Rodriguez-Losada, and Fernando Matia</i>	
An Implicit Time-Stepping Method for Multibody Systems with Intermittent Contact	177
<i>Nilanjan Chakraborty, Stephen Berard, Srinivas Akella, and Jeff Trinkle</i>	
Synthesis of Constrained nR Planar Robots to Reach Five Task Positions	185
<i>Gim Song Soh and J. Michael McCarthy</i>	
Automatic Scheduling for Parallel Forward Dynamics Computation of Open Kinematic Chains	193
<i>Katsu Yamane and Yoshihiko Nakamura</i>	
CRF-Matching: Conditional Random Fields for Feature-Based Scan Matching	201
<i>Fabio Ramos, Dieter Fox, and Hugh Durrant-Whyte</i>	
Control of Many Agents Using Few Instructions	209
<i>Timothy Bretl</i>	
Safety Evaluation of Physical Human-Robot Interaction via Crash-Testing.....	217
<i>Sami Haddadin, Alin Albu-Schäffer, and Gerd Hirzinger</i>	
Dimensionality Reduction Using Automatic Supervision for Vision-Based Terrain Learning	225
<i>Anelia Angelova, Larry Matthies, Daniel Helmick, and Pietro Perona</i>	
The Stochastic Motion Roadmap: A Sampling Framework for Planning with Markov Motion Uncertainty	233
<i>Ron Alterovitz, Thierry Siméon, and Ken Goldberg</i>	
A Fundamental Tradeoff between Performance and Sensitivity within Haptic Rendering	241
<i>Paul G. Griffiths, R. Brent Gillespie, and Jim S. Freudenberg</i>	
Motion Strategies for Surveillance	249
<i>Sourabh Bhattacharya, Salvatore Candido, and Seth Hutchinson</i>	
Learning Omnidirectional Path Following Using Dimensionality Reduction.....	257
<i>J. Zico Kolter and Andrew Y. Ng</i>	
A Fast and Practical Algorithm for Generalized Penetration Depth Computation	265
<i>Liangjun Zhang, Young J. Kim, and Dinesh Manocha</i>	
Planning and Control of Meso-scale Manipulation Tasks with Uncertainties	273
<i>Peng Cheng, David J. Cappelleri, Bogdan Gavrea, and Vijay Kumar</i>	
Data Association in $O(n)$ for Divide and Conquer SLAM	281
<i>Lina M. Paz, José Guivant, Juan D. Tardós, and José Neira</i>	
An Experimental Study of Exploiting Multipath Fading for Robot Communications.....	289
<i>Magnus Lindhé, Karl Henrik Johansson, and Antonio Bicchi</i>	
Mapping Large Loops with a Single Hand-Held Camera.....	297
<i>Laura A. Clemente, Andrew J. Davison, Ian D. Reid, José Neira, and Juan D. Tardós</i>	
Dynamic Coverage Verification in Mobile Sensor Networks via Switched Higher Order Laplacians	305
<i>Abubakr Muhammad and Ali Jadbabaie</i>	
Discrete Search Leading Continuous Exploration for Kinodynamic Motion Planning	313
<i>Erion Plaku, Lydia E. Kavvaki, and Moshe Y. Vardi</i>	
Active Policy Learning for Robot Planning and Exploration under Uncertainty	321
<i>Ruben Martinez-Cantin, Nando de Freitas, Arnaud Doucet, and José A. Castellanos</i>	
Robot Manipulation: Sensing and Adapting to the Real World.....	329
<i>Charles C. Kemp, Aaron Edsinger, Robert Platt, Neo Ee Sian</i>	
Robotic Sensor Networks: Principles and Practice.....	331
<i>Gaurav S. Sukhatme and Wolfram Burgard</i>	
Workshop on Algorithmic Equivalences Between Biological and Robotic Swarms.....	333
<i>Paulina Varshavskaya and James McLurkin</i>	
Workshop on Research in Robots for Education	335
<i>Keith O'Hara, Doug Blank, Maria Hybinette, and Daniela Rus</i>	

Semantic Modeling of Places using Objects

Ananth Ranganathan and Frank Dellaert
 College of Computing, Georgia Institute of Technology
 {ananth, dellaert}@cc.gatech.edu

Abstract—While robot mapping has seen massive strides recently, higher level abstractions in map representation are still not widespread. Maps containing semantic concepts such as objects and labels are essential for many tasks in manmade environments as well as for human-robot interaction and map communication. In keeping with this aim, we present a model for places using objects as the basic unit of representation. Our model is a 3D extension of the constellation object model, popular in computer vision, in which the objects are modeled by their appearance and shape. The 3D location of each object is maintained in a coordinate frame local to the place. The individual object models are learned in a supervised manner using roughly segmented and labeled training images. Stereo range data is used to compute 3D locations of the objects. We use the Swendsen-Wang algorithm, a cluster MCMC method, to solve the correspondence problem between image features and objects during inference. We provide a technique for building panoramic place models from multiple views of a location. An algorithm for place recognition by comparing models is also provided. Results are presented in the form of place models inferred in an indoor environment. We envision the use of our place model as a building block towards a complete object-based semantic mapping system.

I. INTRODUCTION

Robot mapping has in recent years reached a significant level of maturity, yet the level of abstraction used in robot-constructed maps has not changed significantly. Simultaneous Localization and Mapping (SLAM) algorithms now have the capability to accurately map relatively large environments [8], [22]. However, grid-based and feature-based maps constructed using lasers or cameras remain the most common form of representation. Yet higher level abstractions and advanced spatial concepts are crucial if robots are to successfully integrate into human environments.

We have chosen objects and their location to be the semantic information in our maps based on the object-centricity of most man-made environments. People tend to associate places with their use or, especially in workspaces, by the functionality provided by objects present there. Common examples are the use of terms such as “printer room”, “room with the coffee machine”, and “computer lab”. Even in outdoor spaces, people often remember locations by distinguishing features that most often turn out to be objects such as store signs and billboards [12]. Thus, objects form a natural unit of representative abstraction for man made spaces. While we do not claim that representing objects captures all the salient information in a place of interest, it is an important dimension that is useful in a wide variety of tasks.

A major concern in constructing maps with objects is object detection, which has been a major area of research in




Fig. 1. A generative model for representing places in a mapped environment. The place label L generates a set of N objects O , each having a shape S , an appearance A , and a 3D location T . The objects, transformed to location T , give rise to 3D feature points u_{3d} . These features are observed in an image as n features, each with pixel location u , appearance v , and stereo depth d (shaded variables). The schematic on the right illustrates this process.

computer vision for a significant period of time. Due to the difficulty in general purpose object detection, many semantic mapping algorithms assume object detection as a black box [5], and so, sidestep a major component of the mapping problem. However, recent advances in stable feature detection algorithms have enabled featurebased object detection methods that have revolutionized the field. In particular, use of SIFT descriptors [11] along with affine invariant features [13] has been particularly popular since binaries for detecting these features are publicly available.

In this paper, we present one of the first instances of semantic mapping in robotics that integrates state-of-the-art object detection techniques from computer vision. We present a 3D generative model for representing places using objects and develop learning and inference algorithms for the construction of these models. The model for a place is constructed using images and depth information obtained from a stereo camera. Our model is a 3D extension of the constellation models popular in the computer vision literature [4]. In particular, as illustrated in Figure 1, a place is represented as a set of objects O with 3D locations T specified in a local coordinate frame.

In turn, an object is modeled as having a particular shape and appearance, and gives rise to features in the image. This generative model is discussed in detail in Section III below.

The models for the objects are learned in a supervised manner, as will be developed in Section IV. The shape models for the objects are learned using stereo range data, while corresponding appearance models are learned from features detected on the images. Training data is provided to the learning algorithm in the form of roughly segmented images from the robot camera in which objects have been labeled. While unsupervised learning is preferable from the standpoint of automation, it requires that the objects to be learned be prominent in the training images, a condition that is not satisfied by our training images. Supervised learning results in more reliable and robust models.

Once the object models have been learned, inference for the place model is performed at runtime using the Swendsen-Wang algorithm, a cluster Markov Chain Monte Carlo (MCMC) technique [1]. Approximate inference through MCMC is required since finding the correspondence between image features and objects is a combinatorial problem. Features in the test image are connected in a Markov Random Field (MRF) which promotes clustering based on appearance and locality in 3D space. Subsequently, we employ the Swendsen-Wang algorithm to perform sampling-based inference on the MRF. Range data and feature appearance provide strong constraints on the distributions resulting in rapid convergence. The details are given in Section V.

Finally, the place models can be used to perform place recognition, for which purpose we provide an algorithm in Section VI. We also describe a technique to build 360° panoramic models from multiple views of a place using relative robot pose.

Experiments are presented on a robot mounted with a Triclops stereo camera system. We present place models constructed using a fixed object vocabulary in an indoor environment to validate our learning and inference algorithms, in Section VII. We start with related work in the next section.

II. RELATED WORK

The need to include semantic information in robot representations of the environment has long been recognized [10]. A number of instances of functional and semantic representations of space exist in the literature, of which we mention a few recent examples. [18] builds a map of the environment consisting of regions and gateways and augments it to include rectangular objects [17]. Supervised learning of labels on regions is performed in [16] using cascade detectors to detect objects of interest. [5] describes a hierarchical representation that includes objects and semantic labeling of places in a metric map but assumes the identities of objects to be known. [14] performs 3D scene interpretation on range data and computes a map consisting of semantically labeled surfaces. [24] lays out a detailed program for creating cognitive maps with objects as the basic unit of representation. However, their object detection technique uses simple SIFT feature matching




Fig. 2. Features detected on a typical training image using (a) Harris-affine corners (b) Canny edges (c) MSER detector.

which does not scale to larger objects. Our method is more comprehensive than the above-mentioned techniques since it incorporates inference for the 3D location, type, and number of objects in the scene.

More computer vision oriented approaches to the problem of scene modeling also exist. [7] gives a technique for 3D scene analysis and represents the analyzed 3D scene as a semantic net. [23] implements place recognition using features extracted from an image and uses place context to detect objects without modeling their location. More similar to our approach is the one in [20] that creates 3D models of scenes using the Transformed Dirichlet Process. In contrast to this, our approach is simpler and more robust as it uses supervised learning and Swendsen-Wang sampling.

III. CONSTELLATION MODELS FOR PLACES

We model places using a 3D extension of the popular constellation model in computer vision. More precisely, we use the “star” modification of the constellation model [4]. The model represents each place as a set of objects with 3D locations in a local coordinate frame. We assume that given the origin of this coordinate frame, hereafter called the *base location*, the objects are conditionally independent of each other. While full constellation models can also model relative locations between objects, the associated increase in complexity is substantial. More discussion on this subject can be found in [4].

A graphical representation of the generative place model is given in Figure 1. The place label L generates a set of objects O , where the number of objects N is variable. Each object gives rise to a set of 3D feature points u_{3d} according to a shape distribution S . Further, the 3D location of the object in local coordinates is represented using the translation variable T , which is unique for each object and also depends on the place label. Finally, the 3D points, transformed in space according to T , give rise to image features at locations u with appearance v distributed according to an object specific distribution A . The 3D points also produce range measurements d , obtained using stereo, corresponding to the observed features.

The shape distribution S and the appearance distribution A taken together model an object class. In our object models, we represent the shape of an object using a Gaussian distribution in 3D, while its appearance is modeled as a multinomial distribution on vector quantized appearance words.

A. Feature detection and representation

For appearance measurements, we use three complementary types of features in our work that are subsequently




Fig. 3. An example training image with roughly segmented and labeled objects, a monitor in cyan and a drawer in brown.

discretized to facilitate learning. Following previous work [20], we use Harris-affine corners [13], maximally stable extremal regions [13], and clustered edges obtained from a Canny edge detector [2]. We used the public implementations of the first two feature detectors available at <http://www.robots.ox.ac.uk/~vgg/research/affine>. Figure 2 highlights the complementary characteristics of the detected features of each type for a sample image. As noted in [20], edge features are crucial for modeling objects with textureless surfaces such as monitors. The features are represented using SIFT descriptors in a 128 dimensional vector space. We vector quantize the SIFT descriptors using K-means clustering to produce a set of appearance “words”. Each feature is subsequently described by the bin corresponding to its closest appearance word, and its 2D pixel location in the image.

IV. SUPERVISED LEARNING OF 3D OBJECT MODELS

The individual object models in the object vocabulary of the robot are learned in a supervised manner. While unsupervised learning is the sought-after goal in most applications, learning objects in this manner poses many problems. Firstly, the objects of interest have to be displayed prominently in the images with almost no background for reliable unsupervised object discovery to be possible. Further, currently most unsupervised object recognition methods rely on some form of “topic” discovery on the set of training images to learn object models [19]. In images with varied background, the topic discovered by the algorithm may not correspond to objects of interest. Supervised learning sidesteps these issues and is accordingly more suitable for our application.

Training data is presented to the algorithm in the form of roughly segmented stereo images along with the associated stereo range data. Objects in the image are labeled while all unlabeled regions are assumed to be background. An example training image is given in Figure 3.

The shape Gaussian for an object is estimated from the range data labeled as belonging to the object, and the appearance is similarly learned from the features corresponding to the object. Learning the object models is thus straight-forward.

V. INFERENCE FOR CONSTRUCTING PLACE MODELS

During the testing phase, the robot observes an image along with associated stereo range information, and has to infer the label and model for the place, i.e the types of objects and

their 3D locations. We denote the set of appearance and shape models learned during the training phase as $\mathcal{A} = \{A_{1:m}\}$ and $\mathcal{S} = \{S_{1:m}\}$ respectively, where m is the number of objects in the robot’s vocabulary. The pixel locations of the features observed in the image are denoted by the set $U = \{u_{1:n_f}\}$, while the corresponding appearance descriptors are written as $V = \{v_{1:n_f}\}$, n_f being the number of features. The depth from stereo corresponding to each of these features is represented as $D = \{d_{1:n_f}\}$. In the interest of brevity, in the following exposition we will compress the set of measurements to $Z = \{U, V, D\}$ whenever possible.

We infer the place model and label in a Bayesian manner by computing the joint posterior distribution on the place label, the types of objects, and their locations. This posterior can be factorized as

$$p(L, O, T | \mathcal{A}, \mathcal{S}, Z) = p(O, T | \mathcal{A}, \mathcal{S}, Z) p(L | O, T, \mathcal{A}, \mathcal{S}, Z) \quad (1)$$

where L is the place label, O is the set of object types, and T is the corresponding 3D locations. Note that the number of objects at a place, i.e the cardinality of set O , is unknown.

The inference problem can be divided into two parts, namely place *modeling* and place *recognition*, which correspond to the two terms on the right side of (1) respectively. The modeling problem consists of inferring the objects and their locations, while the recognition problem involves finding the label of the place given the objects and their locations. In this section we focus on the modeling problem and return to the recognition problem in the next section.

Since the measurements are in the form of image features, inference cannot proceed without the correspondence between features and the object types they are generated from. We incorporate correspondence by marginalizing over it so that the model posterior of interest can be written as

$$p(O, T | \mathcal{A}, \mathcal{S}, Z) = \sum_{\Pi} p(O, T | \mathcal{A}, \mathcal{S}, Z, \Pi) p(\Pi | \mathcal{A}, \mathcal{S}, Z) \quad (2)$$

where Π is a discrete correspondence vector that assigns each image feature to an object type. We call (2) the *place posterior*.

Since computing the place posterior analytically is intractable, we employ a sampling-based approximation. The intractability arises from the need to compute the distribution over correspondences, which is a combinatorial problem. One technique for overcoming this intractability is using Monte Carlo EM (MCEM) [21], in which a Monte Carlo estimate of the distribution over correspondences is used to maximize the posterior over the other hidden variables iteratively. In our case, this would involve a maximization over a possibly large discrete space of object types. Further, multi-modal distributions in this space cannot be discovered using MCEM, which only computes the MAP solution. These reasons motivate our use of Markov Chain Monte Carlo (MCMC) methods for computing a sampling-based approximation to the posterior.

To compute the distribution over correspondences, we note that features corresponding to a particular object type occur in clusters in the image. Hence, appearance and stereo depth provide important clues to the correspondence of a feature in




Fig. 4. A test image with detected features (top) and the corresponding MRF of features on which sampling is performed (bottom). Note that the MRF may be disconnected as shown here.

the sense that if a feature is similar to its neighboring features with respect to its appearance and 3D location, it is highly likely that it belongs to the same object type as its neighbors.

We take advantage of this spatially clustered nature of the correspondences by placing the image features in a Markov Random Field (MRF). Each feature is connected to its k closest neighbors in the image, where the neighborhood k is a parameter. Larger values of k make large scale correlations visible while increasing complexity ($k = n_f$ gives a fully connected graph). The value of k depends on the expected size of the projected objects in the images. Figure 4 shows the MRF corresponding to features in an image for $k = 10$.

We define discriminative probabilities, also called edge compatibilities, on the edges of the MRF. These are defined as functions of the depth and appearance of the features involved in the edge, where both functions are the Mahalanobis distance between the respective feature values. Denoting the functions on depth and appearance as f_d and f_a , the discriminative probability is

$$p_e \propto f_d(d_i, d_j) \times f_a(v_i, v_j) \quad (3)$$

where

$$-\log f_d = \left(\frac{d_i - d_j}{\sigma_d} \right)^2 \text{ and } -\log f_a = (v_i - v_j)^T \Sigma_a^{-1} (v_i - v_j)$$

and σ_d and Σ_a are depth variance and appearance covariance respectively, that encode the size of objects and their uniformity of appearance.

The overall sampling scheme to compute the posterior can now be seen to have the following form. We sample clusters in the MRF according to the edge compatibilities and subsequently assign an object type to each cluster according to some prior distribution. The sample configuration is then evaluated using the measurement likelihoods based on the learned object models.

We employ a cluster MCMC sampling algorithm to efficiently implement the above scheme. Common MCMC sampling techniques such as Gibbs sampling and Metropolis-Hastings change the value of only a single node in a sampling step, so that mixing time for the chain, i.e the expected time to move between widely differing states, is exponential. Cluster MCMC methods change the value of multiple nodes at each sampling step, leading to fast convergence.

A. The Swendsen-Wang algorithm

We now describe the Swendsen-Wang (SW) algorithm, which we use to compute the place posterior with fast convergence. The SW algorithm has been interpreted in many ways - as a random cluster model, as an auxiliary sampling method, and as a graph clustering model using discriminative edge probabilities [1]. It is in the latter manner that we use the algorithm.

A sample is produced using the SW algorithm by independently sampling the edges of the MRF to obtain connected components. Consider the graph $G = (V, E)$ of image features, as defined above, with discriminative probabilities p_e , $e \in E$ defined in (3) on the edges. We sample each edge of the graph independently and turn "on" each edge with probability p_e . Now only considering the edges that are on, we get a second graph which consists of a number of disjoint connected components. If the discriminative edge probabilities encode the local characteristics of the objects effectively, the connected components will closely correspond to a partition Π of the graph into various objects and the background. The distribution over partitions Π is given as

$$p(\Pi | \mathcal{A}, \mathcal{X}, Z) = \prod_{e \in E_0} p_e \prod_{e \in E \setminus E_0} (1 - p_e) \quad (4)$$

Samples obtained from a typical image feature graph are shown in Figure 5.

To sample over correspondence between image features and object types, we assign an object type to each component of the partition according to the posterior on object types $p(O, T | \mathcal{A}, \mathcal{S}, Z, \Pi)$. Computation of the posterior on object types involves only the appearance measurements since the other measurements also depend on the 3D location of the object T . Applying Bayes Law on the posterior, we get the distribution on object types as

$$p(O_c | \mathcal{A}, \mathcal{S}, Z, \Pi) \propto p(Z | O_c, \mathcal{A}, \mathcal{S}, \Pi) p(O_c | \mathcal{A}, \mathcal{S}, \Pi) \quad (5)$$

where the second term on the right is a prior on object types that can potentially incorporate information regarding the size and appearance of the component, and the frequency with which the object type has been observed in the past. We employ a uniform prior on object types for simplicity since the prior is largely overshadowed by the data in this case. The object type affects only appearance measurements and so, the measurement likelihood in (5) collapses to just the conditionally independent appearance likelihoods on the




Fig. 5. Samples from the SW algorithm for the MRF corresponding to Figure 4, obtained by independently sampling the MRF edges according to edge compatibilities. Each connected component is colored differently and the edges that are turned “off” are shown in gray. The SW algorithm works by assigning an object label to each connected component and subsequently scoring the sample based on the learned object models. In practice, most components get labeled as background.

Algorithm 1 The Swendsen-Wang algorithm for sampling from the place posterior

- 1) Start with a valid initial configuration (Π, O) and repeat for each sample
- 2) Sample the graph G according to the discriminative edge probabilities (3) to obtain a new partition Π'
- 3) For each set in the partition $c \in \Pi'$, assign the object type by sampling from $p(O_c|Z, \Pi')$ as given by (5)
- 4) Accept the sample according to the acceptance ratio computed using (7)

features that are evaluated using the multinomial appearance distribution for the object type

$$p(Z|O_c, \mathcal{A}, \mathcal{S}, \Pi) = \prod_{i=1}^{n_f} p(v_i|O_c, \mathcal{A}, \mathcal{S}, \Pi) \quad (6)$$

The sample thus generated, consisting of a graph partition and object type assignments to each component in the partition, is accepted based on the Metropolis-Hastings acceptance ratio [6], given as

$$\alpha(\Pi' \rightarrow \Pi) = \min \left(1, \frac{\prod_{c' \in \Pi'} p(T_{c'}|O_{c'}, \mathcal{A}, \mathcal{S}, Z, \Pi')}{\prod_{c \in \Pi} p(T_c|O_c, \mathcal{A}, \mathcal{S}, Z, \Pi)} \right) \quad (7)$$

The acceptance ratio (7) can be understood by considering the factorization of the place posterior (2) as

$$p(O_c, T_c|\mathcal{A}, \mathcal{S}, Z, \Pi) = p(O_c|\mathcal{A}, \mathcal{S}, Z, \Pi) p(T_c|\mathcal{A}, \mathcal{S}, Z, \Pi, O_c) \quad (8)$$

and noting that only the second term is involved in the acceptance ratio since the first term, the posterior on object types, has been used in the proposal distribution for the sample above. The acceptance ratio also makes the assumption that the partition components are independent of each other. Note that common factors in the calculation of the acceptance ratio can be omitted to improve efficiency. A summary of the SW algorithm for sampling from the place posterior is given in Algorithm 1.

B. Computing the target distribution

Computing the acceptance ratio (7), involves evaluating the posterior on object locations given a partition of the feature graph. The location posterior, which is the second term on the




Fig. 6. An illustration of the use of scale to obtain a prior on object location. (a) The scale of the projection in the image gives a good estimate of location if the size of the object in 3D is known. (b) The projection of a 3D object model (shown as a Gaussian ellipsoid) is a 2D covariance, the size of which is used to estimate the object's location. Images taken from [9].

right in (8), can be factorized into a prior on object locations and the stereo depth likelihood using Bayes law

$$p(T_c|\mathcal{A}, \mathcal{S}, Z, \Pi, O_c) \propto p(D|\mathcal{S}, \Pi, O_c, T_c) p(T_c|\mathcal{S}, U, \Pi, O_c) \quad (9)$$

where appearance measurements have been neglected since they are assumed independent of the location of the object.

The prior on object locations $p(T_c|\mathcal{S}, U, \Pi, O_c)$ incorporates the 2D pixel locations of the features in the graph component and is evaluated using projection geometry. The graph component is a projection of the object ellipsoid in space and the scale of the projection thus gives an estimate of the 3D object location. Assuming a linear camera projection model, the mean location of the object is projected onto the mean of the feature locations. Analogously, the covariance of the feature locations is a projection of the 3D covariance of the object shape model, which is known. Hence, a proposed 3D location for the object can be evaluated through the norm-1 error between the predicted 2D feature location covariance and the actual covariance. Our use of scale to estimate location in this manner is illustrated in Figure 6. If the observed 2D covariance is Σ_f , the 3D object shape covariance is Σ_O and the Jacobian of the camera projection matrix at T_c is represented as P , the object location prior can be written as

$$\log p(T_c|\mathcal{S}, U, \Pi, O_c) \propto - \left\| \Sigma_f - P \Sigma_O P^T \right\| \quad (10)$$

In practice, the location distribution (9) displays a peaked nature, so that its evaluation is approximated by a maximum

Algorithm 2 Inference algorithm for constructing the place model

For each image I and stereo depth map D obtained at a landmark location, do

- 1) Detect features using the 3 types of feature detectors on the image I .
- 2) Create an MRF using the edge potentials as defined in (3)
- 3) Sample over partitions of the graph using the Swendsen-Wang algorithm 1 and obtain the posterior (8) over object types and their locations
- 4) The posterior on object types and locations is the required place model

a priori (MAP) estimate. This also saves us from sampling over a large continuous space. The maximum of the posterior defined in (9) is obtained by performing a line search along the projection ray of the mean of the feature locations of the graph component. The map value for the object location, T_c^* , is given as

$$T_c^* = \underset{T_c}{\operatorname{argmax}} p(D|S, U, \Pi, O_c, T_c) p(T_c|S, U, \Pi, O_c) \quad (11)$$

We compute the individual stereo depth likelihoods by marginalizing over the “true” depth of the feature as measured along the projection ray

$$p(D|S, \Pi, O_c, T_c^*) = \prod_{i=1}^{n_f} \int_{u_{3d}^i} p(d_i|u_{3d}^i) p(u_{3d}^i|S, U, O_c, T_c^*) \quad (12)$$

where u_{3d}^i is the true depth measured as d_i by the stereo. The stereo measurement model $p(d_i|u_{3d}^i)$ is modeled as a Gaussian distribution that is learned from measurements obtained in a scenario where ground-truth is known.

The prior on true depth $p(u_{3d}^i|S, U, O_c, T_c^*)$ is obtained as a 1D marginal along the projection ray of the object shape model, i.e it is the Gaussian on the projection ray induced by its intersection with the object. This is necessitated by the fact that each feature in the graph component picks its own true depth based on its projection ray. Since both the depth likelihood and true depth prior are Gaussian, the integral in (12) can be computed analytically.

We now have all the pieces to compute the target distribution (8) and the acceptance ratio (7) in the Swendsen-Wang Algorithm 1. Appearance and stereo depth likelihoods are computed using equations (6) and (12) respectively, while a MAP estimate of the object locations is given by (11). A summary of the inference algorithm is given in Algorithm 2.

C. Extension to panoramic models

The algorithm discussed thus far computes the place model for a single image but cannot find a 360° model for a place unless a panoramic image is used. To overcome this limitation, we propose a simple extension of the algorithm.

We compute the panoramic model of a place from multiple images by “chaining” the models from the individual images using odometry information. For example, the robot is made to spin around at the place of interest to capture multiple images. We designate robot pose corresponding to the first of these

images as the base location for the model and marginalize out the poses corresponding to all the other images to create a combined model from all the images.

Denoting the measurements from each of the n images of a place as Z_1, Z_2, \dots, Z_n , and the corresponding detected objects and locations by O_1, \dots, O_n and T_1, \dots, T_n respectively, the panoramic model of the place is computed as

$$\begin{aligned} p(O, T|Z_{1:n}, o^n) &= \int_{x_{1:n}} p(O_{1:n}, T_{1:n}|Z_{1:n}, x_{1:n}) p(x_{1:n}|o^n) \\ &= p(O_1, T_1|Z_1) \\ &\quad \prod_{i=2}^n \int_{x_i} p(O_i, T_i|Z_i, x_i) p(x_i|o_{i-1}) \quad (13) \end{aligned}$$

where x_i is the pose corresponding to the i th image, o_{i-1} is the odometry between poses x_{i-1} and x_i , and $o^n = o_{1:n-1}$ is the set of all odometry. x_1 is assumed to be the origin as it is the base location and the pose distribution $p(x_i|o_{i-1})$ is evaluated using the odometry model. Note that (13) uses the fact that the (O_i, T_i) are conditionally independent of each other given the robot pose.

VI. PLACE RECOGNITION

Place recognition involves finding the distribution on place labels given the detected objects and their locations, i.e. finding the *recognition posterior* $p(L|O, T, \mathcal{A}, S, Z)$ from (1). While robust techniques for place recognition using feature matching are well-known [15], [3], the detected objects can be effectively used to localize the robot, and can be expected to improve place recognition as they provide higher-level distinguishing information. We now give a technique to accomplish this. Applying Bayes law to the recognition posterior from (1)

$$p(L|O, T, \mathcal{A}, S, Z) \propto p(O|L, \mathcal{A}, S, Z) p(T|L, O, \mathcal{A}, S, Z) p(L) \quad (14)$$

If the sequence of labels observed in the past is available, a Dirichlet label prior $p(L)$ is suitable. We assume that such history is unavailable and so use a uniform label distribution.

The object likelihood $p(O|L, \mathcal{A}, S, Z)$ is evaluated as the distance between the observed discrete distribution on object types and prediction assuming the label L . Denoting the object type distribution corresponding to L as O_L , the likelihood is

$$\log p(O|L, \mathcal{A}, S, Z) = -||O - O_L||_2$$

The location likelihood $p(T|L, O, \mathcal{A}, S, Z)$ is computed by minimizing the distance between corresponding objects in T and the model for the place label L . Nearest neighbor (NN) correspondence is used for this purpose. However, since the robot is unlikely to be in exactly the same location even if it visits the same place, we also optimize over a 2D rigid transformation that determines the current pose of the robot in the local coordinates of the place model for the label L

$$p(T|L, O, \mathcal{A}, S, Z) = \int_{X_r} p(T|L, O, \mathcal{A}, S, Z, X_r) p(X_r) \quad (15)$$

where X_r , for which we use a flat Gaussian prior, is the location of the robot in the coordinates of the base location of the place L . In practice, we iterate between optimizing for X_r given object correspondence and finding NN correspondence