Auton Robot (2011) 30: 1–2
DOI 10.1007/s10514-010-9214-z

Special issue on selected papers from Robotics: Science and Systems 2009

Russ Tedrake · Yoky Matsuoka

Received: 9 November 2010 / Accepted: 14 November 2010 / Published online: 24 November 2010

This volume is one of two special journal issues compiled from the best papers presented at the fifth Robotics: Science and Systems (RSS) conference, held in Seattle from June 28 to July 1, 2009. RSS is a meeting that values research papers focusing on the development of, and experimentation with, new complex robotic systems in real-world environments as much as it values new theoretical and algorithmic advances. The other special issue was published by the International Journal of Robotics Research.

The Seattle meeting of RSS brought together more than 300 researchers from Europe, Asia, North America, and Australia. The meeting received 154 submitted papers. After a rigorous review process, 39 papers were accepted covering a wide range of topics in robotics spanning manipulation, locomotion, machine learning, localization, visual SLAM, haptics, biological inspired design and control, etc. Following the tradition of RSS, there were a number of invited speakers who are global leaders in research areas on the frontiers of robotics.

Papers selected for a special issue were selected first, and then special issues editors and editor-in-chiefs from Autonomous Robots and International Journal of Robotics Research carefully partitioned into two groups for the two special issues.

The first paper by *Ko et al.* presents a novel approach to non-parametric system identification, using Gaussian Process Latent Variable Models for subspace identification.

R. Tedrake (✉)
MIT 32-380 32 Vassar St., Cambridge, MA, 02139 USA
e-mail: russt@mit.edu

Y. Matsuoka
University of Washington, BOX 352350, Seattle, Washington, 98195 USA
e-mail: yoky@cs.washington.edu

The technique is able to learn a model for GP-BayesFilters, a general framework for integrating Gaussian Process models into Bayesian filtering techniques, without the need for ground truth training data. The presented approach extends Gaussian Process Latent Variable Models to the setting of dynamical robotics systems. It is capable of incorporating sparse and noisy labels into the training procedure. In addition, the authors show how the framework can be used for time alignment of multiple training trajectories and for controlling dynamical systems based on demonstrations. The algorithms are evaluated on a challenging toy slotcar system and a Barrett WAM robotic arm.

The paper by *Kümmerle et al.* presents an approach to utilize aerial images as a prior for graph-based SLAM. The method obtains relative constraints between nearby robot poses by combining measurements from 3D lasers, stereo-cameras, and wheel encoders. It furthermore localizes the robot in the aerial images to determine a prior over the robot positions in a global reference frame. These priors are inserted into the underlying pose-graph as single-node constraints. The approach then computes the most likely map using a state-of-the-art optimization method. Experiments carried out in mixed indoor and outdoor environments demonstrate that exploiting such prior information results in a more accurate map than with a standard SLAM approach that does not utilize the prior.

The paper by *Jones et al.* focuses on a multi-agent domain where fire trucks are sent out to extinguish fires caused by a large-scale disaster. Because of the disaster, roads are blocked by debris that can only be cleared by bulldozer robots. Coordination in this scenario amounts to determining which routes the fire trucks should take to extinguish which fires and how bulldozers should be used to clear the way. Allocating the tasks so that fires are extinguished quickly and efficiently is challenging because of the explosion in possible combinations of agents, tasks and routes. To address this challenge, Jones et al. propose two approaches. In the first, agents bid on groups of tasks to be accomplished over time and auctions are then held to distribute the tasks. The second approach searches over all possible solutions by using a genetic algorithm. Experiments in simulation show that the genetic algorithm, if given enough time, can improve on the performance of the auction-based systems; however, obtaining this higher performance requires orders of magnitude additional processing time.

The paper by *Napp et al.* presents a feedback regulator for controlling the mean output of discrete state Markov processes. The system input is tunable parameter in the transition rates between some of the states. This control scheme is applied to a set of simple stochastically interacting robots that self-assemble on an air-table to regulate the mean concentration of a target assembly. Experiments demonstrate tracking at different setpoints. The goal is to facilitate engineering with stochastic self-assembly by providing a method to create predictable, robust operating conditions in such systems.

The paper by *Michael et al.* focuses on cooperative manipulation using multiple quadrotor helicopters attached to a object via a cable suspension. The paper first analyzes the static equilibrium conditions of the 3D object, then carefully examines the set of candidate equilibria achievable for three robots with cable tension constraints. The paper combines these conditions with a simplified model of the quadrotor dynamics to design a controller, and then demonstrates experimental results for three quadrotors stabilizing and translating a cable-suspended payload.

The paper by *Yamane et al.* presents an approach for building a large human motion database. They developed an efficient clustering algorithm for organizing motion data into a binary tree data structure that can be used for quickly searching similar poses. The database also includes a graph of possible transitions among the poses to perform search and synthesis of motion sequences. The database is successfully applied to human motion prediction and recognition, as well as humanoid motion planning.

The paper by *Huang et al.* studies the consistency problem of extended Kalman filter (EKF)-based cooperative localization (CL) from the perspective of system observability. It is analytically shown that the error-state system model employed in the standard EKF-CL always has an observable subspace of higher dimension than that of the actual nonlinear CL system. This results in unjustified reduction of the EKF covariance estimates in directions of the state space where no information is available, and thus leads to inconsistency. To address this problem, two new observability-constrained (OC)-EKF estimators are proposed, which select the linearization points so as to ensure that the EKF linearized system model has an observable subspace of correct dimension. Both simulation and experimental results show that the proposed OC-EKFs significantly outperform the standard EKF in terms of both accuracy and consistency.

Enjoy the papers!