

Guest Editorial: Special Issue on Active Perception

Robert Fitch¹ · Volkan Isler² · Pratap Tokekar³ · Davide Scaramuzza⁴

Received: date / Accepted: date

Active perception involves controlling sensor parameters to achieve a sensing task. Traditionally, active perception has been focused on tasks such as object inspection which take place in a confined space. Existing sensor planning approaches are typically reactive and are not informed by past experience. Recently there has been significant interest in extending the scope of active perception to (1) more sophisticated sensing tasks (e.g., object recognition and detection); (2) expansive environments such as farms; and (3) planning approaches which use prior knowledge, seek information and are non-myopic.

This issue of *Autonomous Robots* presents a collection of articles that illustrate the diversity and richness of current research in active perception and its particular relevance to modern robotics. Specific topics range from theory to algorithms and systems. Sensing modalities include monocular vision, stereo vision,

range sensors, and robot audition. Task formulations include object detection and recognition, target tracking and search, exploration and navigation in unknown environments, and adaptive sampling of environmental phenomena. Platforms range from fixed-based manipulators to mobile robots and multi-robot systems. Following an open call for papers, the 17 papers that comprise this special issue were selected from a total of 26 submissions.

We are pleased to open the special issue with the paper “Revisiting Active Perception,” a 30-year retrospective on the topic by pioneering researchers Ruzena Bajcsy, Yiannis Aloimonos, and John Tsotsos. Their paper provides a fitting introduction to the special issue and its themes by identifying the key challenges faced by any “active perceiver,” and discussing to what extent these challenges have been addressed by the research community.

The paper “A Comparison of Volumetric Information Gain Metrics for Active 3D Object Reconstruction” by Jeffrey Delmerico, Stefan Isler, Reza Sabzevari, and Davide Scaramuzza studies the question of how best to quantify the volumetric information gained from dense 3D input data with respect to a probabilistic volumetric map. The authors propose and evaluate several new metrics in the context of active dense 3D object reconstruction, and have released a publicly available ROS-based implementation of their work.

Authors Yash Satsangi, Shimon Whiteson, Frans Oliehoek, and Matthijs Spaan consider POMDP models for active perception problems in their paper, “Exploiting Submodular Value Functions for Scaling Up Active Perception.” The authors show how the assumptions of traditional methods for POMDP planning can be violated in active perception tasks, and propose new POMDP frameworks that, in certain cases, can exploit

Robert Fitch
rfitch@uts.edu.au

Volkan Isler
isler@cs.umn.edu

Pratap Tokekar
tokekar@vt.edu

Davide Scaramuzza
sdavide@ifi.uzh.ch

¹ University of Technology Sydney, Ultimo, NSW 2007, Australia

² University of Minnesota, Minneapolis, MN 55455, USA

³ Virginia Tech, Blacksburg, VA 24061, USA

⁴ University of Zurich, 8050 Zurich, Switzerland

the property of submodularity to drastically reduce computational cost.

The theme of exploring appropriate utility functions for active perception continues with the paper “Autonomous Robotic Exploration using a Utility Function based on Rényi’s General Theory of Entropy” by Henry Carrillo, Philip Dames, Vijay Kumar, and José Castellanos. The authors propose a new method for quantifying the value of potential actions in an exploration task that takes into account map uncertainty as well as uncertainty in robot pose.

The next three papers in the issue also deal with the task of exploration and mapping. In their paper “Environment Model Adaptation for Mobile Robot Exploration,” authors Erik Nelson, Micah Corah, and Nathan Michael propose an adaptive compressed map representation that relieves the computational burden of evaluating information-theoretic metrics in planning. The paper “Gaussian Processes Autonomous Mapping and Exploration for Range-Sensing Mobile Robots” by Maani Ghaffari Jadidi, Jaime Valls Miro, and Gamini Disanayake considers Gaussian process occupancy maps and proposes efficient methods to compute probabilistic frontiers and perform greedy exploration based on mutual information. The paper “Receding Horizon Path Planning for 3D Exploration and Surface Inspection” by Andreas Bircher, Mina Kamel, Kostas Alexis, Helen Oleynikova, and Roland Siegwart considers exploration in the context of inspecting 3D surfaces and proposes a planning algorithm that samples configurations from a geometric random tree.

The task of tracking a moving target using a mobile robot is studied by Yoonseon Oh, Sungjoon Choi, and Songhwa Oh in their paper “Chance-Constrained Target Tracking using Sensors with Bounded Fan-Shaped Sensing Regions.” The authors assume a realistic, non-omnidirectional sensor model and propose a robust tracking algorithm that provides probabilistic performance guarantees. The related problem of searching for a target is considered in the paper “Near-Optimal Probabilistic Search Using Spatial Fourier Sparse Set” by Kuo-Shih Tseng and Bérénice Mettler. The authors use reinforcement learning to compute a set of subgoal positions that allow the robot to search the environment efficiently.

Two papers consider active perception in multi-robot systems. The paper “Cooperative Multi-Robot Belief Space Planning for Autonomous Navigation in Unknown Environments” by Vadim Indelman studies multi-robot navigation in unknown, GPS-denied environments and proposes a new multi-robot belief space planning framework where robots act cooperatively to reduce map uncertainty. The paper “Distributed Information-Based

Guidance of Multiple Mobile Sensors for Urban Target Search” by Juan-Pablo Ramirez-Paredes, Emily Doucette, Jess Curtis, and Nicholas Gans studies the problem of searching for a moving target with a team of mobile robots, and proposes new data fusion and planning algorithms that allow robots to act independently when communication becomes unavailable.

Timothy Patten, Wolfram Martens, and Robert Fitch study object classification from 3D laser scans in unknown outdoor environments in their paper “Monte Carlo Planning for Active Object Classification.” The authors propose a new non-myopic planning framework based on Monte Carlo methods and validate it with a mobile robot in a farm environment. The planner chooses a sequence of viewpoints to improve object class and pose estimates and provides analytical performance guarantees. Object pose estimation is also studied in the paper “Improving Object Orientation Estimates by Considering Multiple Viewpoints” by Zoltán Csaba Márton, Serkan Türker, Christian Rink, Manuel Brucker, Simon Kriegel, Tim Bodenmüller, and Sebastian Riedel. The authors present a new probabilistic method based on the histogram filter and integrate it with a next-best-view algorithm to improve the accuracy of pose estimates by choosing a sequence of informative views.

The paper “Contour-Based Next-Best View Planning from Point Cloud Segmentation of Unknown Objects” by Riccardo Monica and Jacopo Aleotti considers the problem of choosing the next-best-view for scene exploration in a table-top scenario with a robot manipulator and an eye-in-hand range sensor. The authors propose an algorithm that exploits point cloud segmentation to evaluate viewpoints, and they implement their algorithm using a GPU.

A bio-inspired approach for artificial selective visual attention is studied in the paper “On the Advantages of Foveal Mechanisms for Active Stereo Systems in Visual Search Tasks” by Rui Pimentel de Figueiredo, Alexandre Bernardino, José Santos-Victor, and Helder Araújo. In the authors’ foveated stereo vision model, depth resolution is nonuniform and is highest (most acute) at a given fixation point. The authors propose algorithms for actively controlling the position of the sensor’s fixation point and show how such a system is computationally advantageous in visual search and exploration tasks.

Auditory attention is studied in the paper “An Information Based Feedback Control for Audio-Motor Binaural Localization” by Gabriel Bustamante, Patrick Danès, Thomas Fergue, Ariel Podlubne, and Jérôme Manhès. The authors consider the problem of stereo sound source localization and propose an algorithm for choosing the most informative motor action for an actuated binaural

sensor. The algorithm is validated using a real binaural robot head.

The final paper in the issue is “Adaptive Sampling of Cumulus Clouds with UAVs” by Christophe Reymann, Alessandro Renzaglia, Fayçal Lamraoui, Murat Bronz, and Simon Lacroix. The authors consider an environmental monitoring scenario where a fleet of UAVs must map atmospheric variables by collecting data in low-altitude clouds. The atmospheric maps are represented with a Gaussian process regression model where the hyperparameters are learned online. Trajectories for the UAVs are optimised with consideration to information gain and realistic flight dynamics, and are evaluated in simulation examples.

The special issue was motivated by the workshop “Scaling Up Active Perception,” organized by Yianis Aloimonos, Andrea Censi, Kostas Daniilidis, Volkan Isler, and Stefano Soatto (in alphabetical order) and held at the 2015 IEEE International Conference on Robotics and Automation (May 2015 in Seattle, Washington). The articles have undergone rigorous peer-review according to the journal’s high standards.

We would like to thank the authors and reviewers whose contributions made this special issue possible. Special thanks go to Editor-in-Chief Gaurav Sukhatme for his unwavering support. We are also grateful to the Springer staff who worked tirelessly in helping to bring this project to fruition.

Guest editors:

Robert Fitch
Volkan Isler
Pratap Tokekar
Davide Scaramuzza

December, 2017