Information Gain-based Exploration Using Rao-Blackwellized Particle Filters

Wolfram Burgard[1] Cyrill Stachniss[1] Giorgio Grisetti[1,2]
[1]*Univ. of Freiburg, Dept. of Computer Science, D-79110 Freiburg, Germany*
[2]*Univ. La Sapienza, Dipart. Informatica e Sistemistica, I-00198 Rome, Italy*

Topic: estimation, control
Oral presentation

Robots that are able to acquire an accurate model of their environment are regarded as fulfilling a major precondition of truly autonomous mobile vehicles. To learn a map of the environment, three problems need to be addressed simultaneously, namely exploration, mapping, and localization. In this work, we present an integrated solution to these three problems. To solve the simultaneous localization and mapping (SLAM) problem, our approach uses a highly efficient variant of algorithm proposed by Murphy and colleagues [2, 4]. In this algorithm, a Rao-Blackwellized particle filter (RBPF) is used to efficiently represent the joint posterior about possible maps and trajectories taken by the robot. The key contribution of our approach is an efficient decision-theoretic algorithm for computing vantage points that reduce the expected uncertainty in the RBPF.

The approaches mostly related to our work have been presented by Makarenko *et al.* [3] and Bourgault *et al.* [1]. They use an Extended Kalman Filter (EKF) to solve the SLAM problem and introduce a utility function which trades-off the cost of exploring new terrain with the potential reduction of uncertainty by measuring at selected positions. A similar technique has been applied by Sim *et al.* [5], who consider actions to guide the robot back to a known place in order to reduce the pose uncertainty of the vehicle during exploration. In contrast to our work, these approaches assume that the environment contains landmarks that can be uniquely determined during mapping. Our approach, in contrast, learns occupancy grid maps and thus is not restricted to environments with pre-defined landmarks.

Compared to previous approaches, the novelty of the work reported here is that our algorithm simultaneously considers the uncertainty in the trajectory and in the map while building accurate occupancy grids. Based on an efficient scheme for computing the uncertainty of the joint posterior, we apply decision-theoretic framework for choosing appropriate actions. Thereby, we utilize the properties of the Rao-Blackwellization. In brief, the uncertainty of an RBPF is determined by its entropy, which consists of two components. Whereas the first one computes the uncertainty over the trajectory taken by the robot, the second calculates the uncertainty of the individual maps weighted by the likelihood of the corresponding particles. We also demonstrate how the filter can be used to efficiently simulate possible measurements to be obtained along a path through already mapped terrain. To determine potential actions, we consider different action types including actions that collect sensor information about unknown areas as well as actions designed to re-visit known places to reduce the pose uncertainty of the vehicle. To finally select an action, we take into account the potential measurements gathered along the path of the robot as well as the cost for carrying out this action.

Our approach has been implemented and tested in the real world and in simulation. Experimental results suggest that our approach leads to a robust exploration behavior that produces highly accurate grid maps. Furthermore, we illustrate the advantages of our action selection technique over previous approaches.

References

[1] F. Bourgoult, A.A. Makarenko, S.B. Williams, B. Grocholsky, and F. Durrant-Whyte. Information based adaptive robotic exploration. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2002.

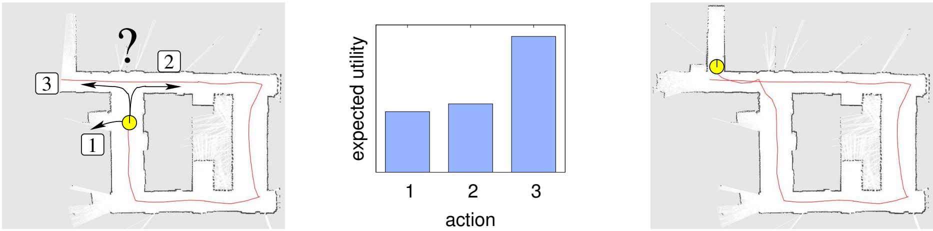

Figure 1: Suppose the robot has a high pose uncertainty and has to decide where to go next. Shown are three opportunities in the left image. By estimating the change of uncertainty which is introduced by executing an action, we can determine which action promises the highest uncertainty reduction (see middle image). Thus, our approach chooses action 3 as depicted in the right image.

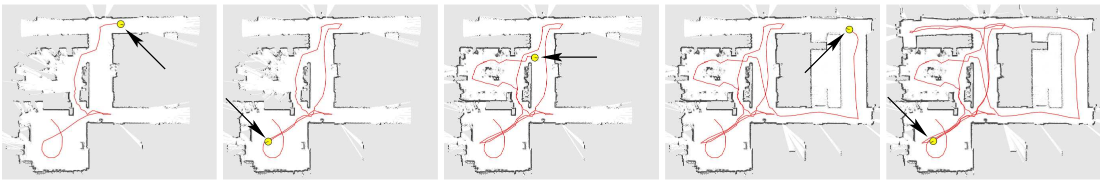

Figure 2: Five different stages of an autonomous exploration run on the second floor of building 106 at the University of Freiburg. The map was acquired fully autonomously by our integrated approach.

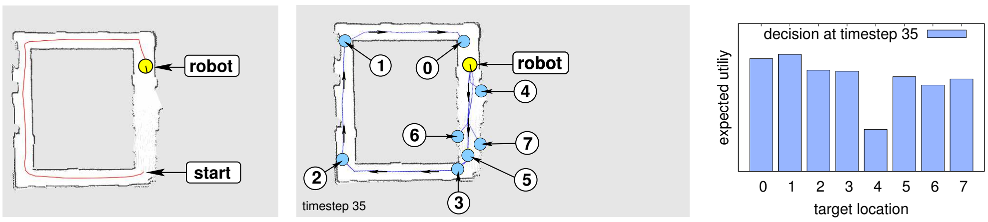

Figure 3: The robot had to decide where to move next during an exploration experiment. The left image depicts the trajectory of the robot so far and the middle image illustrates the best map in that situation as well as possible actions with corresponding target locations. The plot on the right hand side shows the expected utility of the individual goal locations. The robot selected location 1 and performed an active loop closure, since this action provided the highest expected utility.

[2] A. Doucet, J.F.G. de Freitas, K. Murphy, and S. Russel. Rao-blackwellized partcile filtering for dynamic bayesian networks. In *Proc. of the Conf. on Uncertainty in Artificial Intelligence (UAI)*, 2000.

[3] A.A. Makarenko, S.B. Williams, F. Bourgoult, and F. Durrant-Whyte. An experiment in integrated exploration. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2002.

[4] K. Murphy. Bayesian map learning in dynamic environments. In *Neural Info. Proc. Systems (NIPS)*, 1999.

[5] R. Sim, G. Dudek, and N. Roy. Online control policy optimization for minimizing map uncertainty during exploration. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2004.