Robotics and Autonomous Systems 69 (2015) 68–79

Contents lists available at ScienceDirect

Robotics and Autonomous Systems

journal homepage: www.elsevier.com/locate/robot

Potential information fields for mobile robot exploration[☆,☆☆]

Joan Vallvé, Juan Andrade-Cetto *

Institut de Robòtica i Informàtica Industrial, CSIC-UPC, Llorens Artigas 4-6, 08028 Barcelona, Spain

HIGHLIGHTS

- Novel exploration method based on joint path and map entropy minimization.
- Minimizing both map and path entropies produces more reliable maps.
- Exploration is pursued by gradient descent over the potential information field.

ARTICLE INFO

Article history:
Available online 8 September 2014

Keywords:
Mobile robotics
SLAM
Exploration

ABSTRACT

We present a decision theoretic approach to mobile robot exploration. The method evaluates the reduction of joint path and map entropy and computes a potential information field in robot configuration space using these joint entropy reduction estimates. The exploration trajectory is computed descending on the gradient of this field. The technique uses Pose SLAM as its estimation backbone. Very efficient kernel convolution mechanisms are used to evaluate entropy reduction for each sensor ray, and for each possible robot orientation, taking frontiers and obstacles into account. In the end, the computation of this field on the entire configuration space is shown to be very efficient. The approach is tested in simulations in a pair of publicly available datasets comparing favorably both in quality of estimates and in execution time against an RRT*-based search for the nearest frontier and also against a locally optimal exploration strategy.

1. Introduction

We consider the problem of autonomous mobile robot exploration, and frame it as that of reducing both localization and map uncertainties. Exploration strategies driven by uncertainty reduction date back to the seminal work of Whaite [1] for the acquisition of 3-D models of objects from range data. Within the context of SLAM, it is the work of Feder et al. [2], who first proposed a metric to evaluate uncertainty reduction as the sum of the independent robot and landmark entropies with an exploration horizon of one step to autonomously produce occupancy maps. Bourgault et al. [3] alternatively proposed a utility function for exploration that trades off the potential reduction of vehicle localization uncertainty, measured as entropy over a feature-based map, and the information gained over an occupancy grid. In contrast to these approaches, which consider independently the reduction of vehicle and map entropies [4]. Vidal-Calleja et al., [5] tackled the issue of joint robot and map entropy reduction, taking into account robot and map cross correlations for the Visual SLAM EKF case.

Action selection in SLAM can also be approached as an optimization problem using receding horizon strategies [6–8]. Multistep look ahead exploration in the context of SLAM makes sense only for situations in which the concatenation of prior estimates without measurement evidence remains accurate for large motion sequences. For highly unstructured scenarios and poor odometry models, this is hardly the case. So, we stick to the one step look ahead case.

One technique that tackles the problem of exploration in SLAM as a one step look ahead entropy minimization problem makes use of Rao-Blackwellized particle filters [9]. The technique extends the classical frontier-based exploration method [10] to the full SLAM case. When using particle filters for exploration, only a very narrow action space can be evaluated due to the complexity in computing the expected information gain. The main bottleneck is the generation of the expected measurements that each action

☆ This work has been supported by the Spanish Ministry of Economy and Competitiveness under Project DPI-2011-27510 and by the EU Project Cargo-ANTs FP7-605598.

☆☆ This paper is an extended and revised version of the work of Vallvé and Andrade-Cetto (2013) [4].

* Corresponding author. Tel.: +34 934010775.
E-mail addresses: jvallve@iri.upc.edu (J. Vallvé), cetto@iri.upc.edu (J. Andrade-Cetto).

http://dx.doi.org/10.1016/j.robot.2014.08.009
0921-8890/© 2014 Elsevier B.V. All rights reserved.

sequence would produce, which is generated by a ray-casting operation in the map of each particle. In contrast, measurement predictions in a Pose SLAM implementation, such as ours, can be computed much faster, having only one map posterior per action to evaluate, instead of the many that a particle filter requires. Moreover, in [9], the cost of choosing a given action is subtracted from the expected information gain with a user selected weighting factor. In our approach, the cost of long action sequences is taken into consideration during the selection of goal candidates, using the same information metrics that help us keep the robot localized during path execution.

In [11] our group proposed a solution to the exploration problem that maximizes information gain in both the map and path estimates. The method evaluates both exploratory and loop closure candidate trajectories, computing entropy reduction estimates from a coarse resolution realization of occupancy maps. The final trajectory is computed using A* in the occupancy grid, just as [12] does so over an initial reference trajectory. The computational bottleneck of [11] was in the estimation of the occupancy map. In this paper we present an alternative method, in which we compute directly the global entropy reduction estimate for each possible robot configuration. The use of very efficient kernel convolutions allows us to compute this estimate very fast and without the need to reduce the grid resolution. In [11], exploratory actions considered omnidirectional sensing and evaluated paths toward positions near frontiers, disregarding orientation. In a more general setting, a sensor, such as a laser range finder or a camera, would have a narrow field of view, and hence, we need to deal with full poses not just positions. In this paper we take this issue into account and compute instead entropy reduction estimates for the whole configuration space (C-space).

To find candidate exploration paths, the entropy reduction grid in C-space is transformed into a potential field, taking into account frontiers and obstacles. The path is obtained by gradient descent on this field. Potential field methods have been previously used for exploration [13,14], but different from our approach, these methods directly evaluate boundary conditions on deterministic maps of obstacles and frontiers, without taking uncertainty into account. Our method follows the idea of gradient descent to a desired exploratory or loop closing location, due to the minimization of joint map and path entropies.

In summary, the proposed method iteratively proceeds as follows. First, from the current Pose SLAM estimate (Section 2), a log odds occupancy map is synthesized from raw sensor data as shown in Section 3. We use this map to compute a potential information field (Section 4), and plan exploration trajectories as gradient descent along this field. Once the current exploration goal is reached, or a loop closure is obtained, a new exploration candidate is computed in the next iteration. The method is compared against frontier-based exploration and locally optimal planning in Section 5, and conclusions are drawn in Section 6.

2. Pose SLAM

The proposed exploration strategy uses Pose SLAM as its estimation backbone. In Pose SLAM [15], a probabilistic estimate of the robot pose history is maintained as a sparse graph with a canonical parametrization $p(\mathbf{x}) = \mathcal{N}^{-1}(\boldsymbol{\eta}, \boldsymbol{\Lambda})$, using an information filter, with information vector $\boldsymbol{\eta} = \boldsymbol{\Lambda}\boldsymbol{\mu}$, and an information matrix $\boldsymbol{\Lambda} = \boldsymbol{\Sigma}^{-1}$. This parametrization has the advantage of being exactly sparse [16]. State transitions result from the composition of motion commands to previous poses,

$$x_k = f(x_{k-1}, u_k) = x_{k-1} \oplus u_k, \tag{1}$$

and the registration of sensory data also produces relative motion constraints, but now between non-consecutive poses,

$$z_{ik} = h(x_i, x_k) = \ominus x_i \oplus x_k, \tag{2}$$

where the operators $\oplus$ and $\ominus$ are used to indicate the forward and backward compositions of one coordinate frame relative to another [17].

When establishing a link between the current robot pose, say k, and any other previous pose, say i, the update operation only modifies the corresponding diagonal blocks of the information matrix $\boldsymbol{\Lambda}$ and introduces new off-diagonal blocks at locations ik, and ki. These links enforce graph connectivity, or loop closure in SLAM parlance, and revise the entire path state estimate, reducing overall uncertainty, hence entropy.

To enforce sparseness in Pose SLAM, only the non redundant poses and the highly informative links are added to the graph. A new pose is considered redundant when it is too close to another pose already in the trajectory and not much information is gained by linking this new pose to the map. However, if the new pose allows to establish an informative link, both the link and the pose are added to the map. In other words, in Pose SLAM, all decisions to update the graph, either by adding more nodes or by closing loops, are taken in terms of overall information gain.

The information gain of a link, i.e., the difference in entropies on the entire map before and after the link is established, takes the form

$$\mathcal{I}_{ik} = \frac{1}{2} \ln \frac{|\mathbf{S}_{ik}|}{|\boldsymbol{\Sigma}_y|}, \tag{3}$$

where $\boldsymbol{\Sigma}_y$ is the sensor registration covariance, $\mathbf{S}_{ik}$ is the innovation covariance

$$\mathbf{S}_{ik} = \boldsymbol{\Sigma}_y + [\mathbf{H}_i \mathbf{H}_k] \begin{bmatrix} \boldsymbol{\Sigma}_{ii} & \boldsymbol{\Sigma}_{ik} \\ \boldsymbol{\Sigma}_{ik}^\top & \boldsymbol{\Sigma}_{kk} \end{bmatrix} [\mathbf{H}_i \mathbf{H}_k]^\top, \tag{4}$$

$\mathbf{H}_i$, $\mathbf{H}_k$ are the Jacobians of h with respect to poses i and k evaluated at the state means μ_i and μ_k, $\boldsymbol{\Sigma}_{ii}$ is the marginal covariance of pose i, and $\boldsymbol{\Sigma}_{ik}$ is the cross correlation between poses i and k. Links that provide information above a threshold γ are added to the graph, either to link previous states, or to connect a new pose with the map prior.

3. Log odds occupancy grid

Pose SLAM does not maintain a grid representation of the environment. It only encodes relations about robot poses. The environment however, can be synthesized at any instance in time using the pose means in the graph and the raw sensor data. The resolution at which the map is synthesized depends on the foreseen use of this map. For instance, in [11] occupancy grid maps at very coarse resolution are produced to evaluate the effect of candidate trajectories in entropy reduction. In some cases one might not even need to render a map. Such is the case in [18], where only the graph is needed to plan optimal trajectories in a belief roadmap.

In this paper, we use the Pose SLAM estimate and raw sensor data to synthesize an occupancy map, and from this map, we then build an entropy reduction field in configuration space. The quality of the occupancy grid produced is a key element of our exploration strategy. The mapping of frontiers near obstacles in the presence of uncertainty might drive the robot to areas near collision, a situation we need to avoid. Moreover, there is a compromise between tractability and accuracy in choosing the resolution at which the occupancy cells are discretized.

To provide an accurate computation of the occupancy map, which is necessary for the proper computation of the potential information field, we render the map from all poses in the Pose SLAM graph, and not only a limited number of them. Moreover, the resolution at which the occupancy grid map is computed is finer than what we were able to compute in [11]. Instead of repeating the ray-casting operation at each iteration, we store local log odds