Robotics and Autonomous Systems ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Robotics and Autonomous Systems

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

Robust exploration and homing for autonomous robots

Daniel Perea Ström [a], Igor Bogoslavskyi [b], Cyrill Stachniss [b,*]

[a] Universidad de La Laguna, Departamento de Ingeniería Informática, Av. Francisco Sánchez, 38206 Tenerife, Spain
[b] University of Bonn, Institute for Geodesy and Geoinformation, Nussallee 15, 53115 Bonn, Germany

ARTICLE INFO

Article history:
Available online xxxx

Keywords:
Exploration
Background knowledge
Homing
Navigation
SLAM

ABSTRACT

The ability to explore an unknown environment is an important prerequisite for building truly autonomous robots. Two central capabilities for autonomous exploration are the selection of the next view point(s) for gathering new observations and robust navigation. In this paper, we propose a novel exploration strategy that exploits background knowledge by considering previously seen environments to make better exploration decisions. We furthermore combine this approach with robust homing so that the robot can navigate back to its starting location even if the mapping system fails and does not produce a consistent map. We implemented the proposed approach in ROS and thoroughly evaluated it. The experiments indicate that our method improves the ability of a robot to explore challenging environments as well as the quality of the resulting maps. Furthermore, the robot is able to navigate back home, even if it cannot rely on its map.

1. Introduction

Exploration is the task of selecting view points so that a robot can cover the environment with its sensors to build a map. The ability to robustly operate without user intervention is an important capability for exploration robots, especially if there is no means for communication between the robot and an operator. Most exploration robots always start assuming zero knowledge and do not exploit any background knowledge about the environment or typical environments. They build a map of the environment online and make all navigation decisions based on this map. As long as this map is consistent, the robot can perform autonomous navigation by planning the shortest path – for example using A^* – from its current location to its next vantage point using the map. Although recent SLAM systems are fairly robust, there is a chance that they fail, for example, due to wrong data associations generated by the front-end. Even current state-of-the-art SLAM approaches cannot guarantee the consistency of the resulting map. Computing a path based on an inconsistent map, however, is likely to lead to a failure and possibly to losing the robot if operating in a hazardous environment. Thus, exploring robots should always decide where to go next and at the same time verify if their map is still consistent (see sketch in Fig. 1). Considering existing approaches, however, it is fair to say that most exploration systems follow the paradigm that they (a) make their navigation and exploration decisions using the current map only and (b) assume that the map is consistent and thus can be used as the basis for path planning and navigation.

In this paper, we aim at relaxing these assumptions. The key idea is to consider the knowledge gained from previously conducted exploration missions to support the navigation system of the robot. This is motivated by the fact that selecting appropriate target locations during exploration supports the mapping process, and can increase the probability of building a consistent map. Furthermore, we want to be able to safely navigate our robot back to its starting location, even if the mapping process failed.

The first contribution of this paper is a novel approach to exploiting background knowledge while generating exploration behaviors to support mapping. The key idea is to use previously experienced environments to reason about what to find in the unknown parts of the world. To achieve this, we equip our robot with a database to store all acquired (local) maps and exploit this knowledge when selecting target locations. Our research is motivated by an exploration project for autonomously digitizing the Roman catacombs, which are complex underground environments with repetitive structures. To predict possible geometries of the environment the robot may experience during exploration, we exploit previously visited areas and consider the similarities with the area around the currently planned next view point. This allows the robot to actively seek for loop-closures and in this way actively reduce its pose uncertainty. Our experiments indicate that this approach is beneficiary for robots when comparing it to a standard frontier-based exploration method.

The second contribution is a robot homing approach with the goal of retrieving our robot even if the SLAM system failed to build

* Corresponding author.
E-mail addresses: dani@isaatc.ull.es (D.P. Ström), igor.bogoslavskyi@uni-bonn.de (I. Bogoslavskyi), cyrill.stachniss@igg.uni-bonn.de (C. Stachniss).

http://dx.doi.org/10.1016/j.robot.2016.08.015
0921-8890/

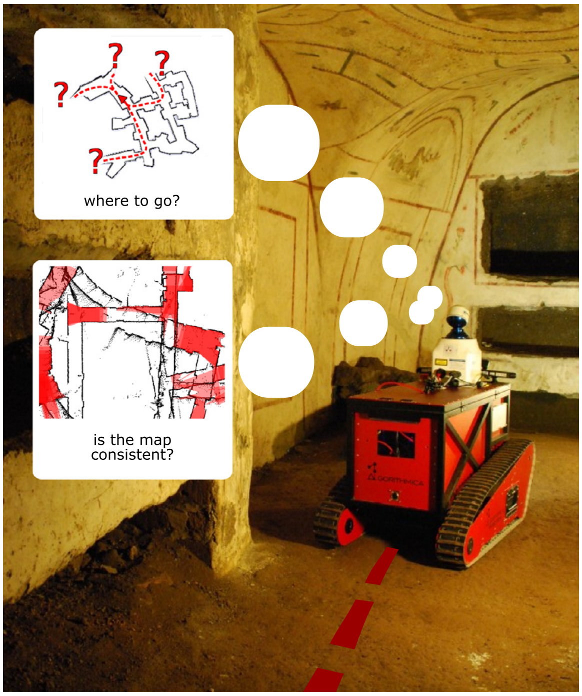

Fig. 1. Mobile robot exploration has to answer the question: "Where to go next?". Our approach exploits previously mapped environments to predict potential future loop closures and thus to select better target locations. When the statistical map consistency tester provides the robot with the information that the map is not consistent anymore the robot starts rewinding the trajectory using our robust homing method.

a consistent map. To avoid that our robot gets lost, we propose a robust homing system consisting of two distinct parts. Part A performs a statistical analysis of the map and thus provides the information about its consistency. We build upon our previous work [1] for performing a cascade of pair-wise consistency checks using the observations perceiving the same areas. To avoid performing such checks on the overall map, we reduce the area to analyze by planning the shortest homing route for the robot assuming a consistent map. We then analyze the map consistency only along that path and can estimate on the fly if the map around this path is consistent or not with a given confidence level. If it is consistent, we navigate back on the verified homing path. Part B of our approach is responsible for driving the robot back to its starting location without a map. We achieve this by rewinding the trajectory that the robot took to reach its current pose. If the motions of the robot were perfect, i.e. would lead to the desired robot pose in the world frame, we would be able to simply invert the motion commands performed by the robot and could safely reach the starting location. Motion execution and odometry, however, are often noisy. As a result, simply following inverse motion commands will not bring the robot to the starting location in the real world in most cases. Therefore, we take into account the sensor information to guide the robot back by matching the observations with the past.

2. Related work

The majority of techniques for mobile robot exploration focus on generating motion commands that minimize the time needed to cover the whole terrain. Several techniques also assume that an accurate position estimate is available during exploration [2,3]. Whaite and Ferrie [4] present an approach that uses the entropy to measure the uncertainty in the geometry of objects that are scanned with a laser range sensor. Similar techniques have been applied to mobile robots [5,6], but such approaches still assume to know the correct pose of the vehicle. Such approaches take the map but not the pose uncertainty into account when selecting the next vantage point. There are, however, exploration approaches that have been shown to be robust against uncertainties in the pose estimates [7,8].

Besides the idea of navigating to the next frontier [3], techniques based on stochastic differential equations for goal-directed exploration have been proposed by Shen et al. [9]. Similar to that, constrained partial differential equations that provide a scalar field into unknown areas have been presented by Shade et al. [10]. An information-theoretic formulation that seeks to minimize the uncertainty in the belief about the map and the trajectory of the robot has been proposed by Stachniss et al. [11]. This approach builds upon the works of Makarenko et al. [12] and Bourgault et al. [13]. Both extract landmarks out of laser range scans and use an Extended Kalman Filter to solve the underlying SLAM problem. They furthermore 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 presented by Sim et al. [14], who consider actions to guide the robot back to a known place in order to reduce the pose uncertainty of the vehicle. Such information-driven techniques have also been used for perception selection to limit the complexity of the underlying optimization problems in SLAM [15].

In general, the computation of the expected entropy reductions is a complex problem, see Krause and Guestrin [16], and in all real world systems, approximations are needed. Suitable approximations often depend on the environment model, the sensor data, and the application. In some cases, efficient approximations can be found, for example in the context of monitoring lakes using autonomous boats [17].

Other approaches, especially in the context of autonomous micro aerial vehicles (MAVs), seek to estimate the expected feature density in the environment in order to plan a path through areas that support the helicopter localization [18]. This can be seen as related to information-theoretic approaches, although Sadat et al. [18] do not formulate their approach in this framework. A related approach to MAV exploration seeks to select new vantage points during exploration, so that the expected number of visible features is maximized, see Mostegel et al. [19].

An interesting approach by Fox et al. [20] aims at incorporating knowledge about *other* environments into a cooperative mapping and exploration system for multiple robots. This allows for predicting simplified laser scans of an unknown environment. This idea was an inspiration for our paper for predicting possible loop closures given the environment structure explored so far. We use this approach for exploring ancient catacombs, which are repetitive underground environments, with a mobile platform, see Fig. 1. Chang et al. [21] propose an approach for predicting the environment using repetitive structures for SLAM. Other background knowledge about the environment, for example semantic information [22], can support the exploration process as shown by Wurm et al. [23], Stachniss et al. [24] as well as Holz et al. [25].

A central problem in robust exploration, however, is that in case of a SLAM failure, the map becomes inconsistent. This can prevent the robot from continuing its exploration mission and – even worse – from being able to navigate back. It is therefore important to be able to perform reliable navigation without relying on a map.

Sprunk et al. [26] present a lidar-based teach-and-repeat method to follow a route given by the user. The approach relies on precise localization of the robot based on the lidar measurements with respect to a taught-in trajectory. Similarly, Furgale et al. [27] perform the ICP-based teach-and-repeat approach on an autonomous robot equipped with a high precision 3D spinning lidar. They extend the standard teach-and-repeat approach by adding a local motion planner to account for dynamic changes in the

Please cite this article in press as: D.P. Ström, et al., Robust exploration and homing for autonomous robots, Robotics and Autonomous Systems (2016), http://dx.doi.org/10.1016/j.robot.2016.08.015