

Exploiting Inaccurate A Priori Knowledge in Robot Exploration

Extended Abstract

Matteo Luperto
Università degli Studi di Milano
Milano, Italy
matteo.luperto@unimi.it

N. Alberto Borghese
Università degli Studi di Milano
Milano, Italy
alberto.borghese@unimi.it

Danilo Fusi
Politecnico di Milano
Milano, Italy
danilo.fusi@mail.polimi.it

Francesco Amigoni
Politecnico di Milano
Milano, Italy
francesco.amigoni@polimi.it

ABSTRACT

Exploration is a task in which autonomous mobile robots incrementally discover features of interest in initially unknown environments. Most of the current exploration approaches ignore prior knowledge about the environments that have to be explored. However, in some practical cases, such knowledge could be available. In this paper, we present a method that includes a priori knowledge in an exploration strategy that selects the next best locations the robot should reach in partially explored indoor environments by exploiting the (possibly inaccurate) knowledge of their floor plans.

KEYWORDS

robot exploration; exploration strategies; robot mapping

ACM Reference Format:

Matteo Luperto, Danilo Fusi, N. Alberto Borghese, and Francesco Amigoni. 2019. Exploiting Inaccurate A Priori Knowledge in Robot Exploration. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019*, IFAAMAS, 3 pages.

1 INTRODUCTION

Exploration is an important task for autonomous robotics, in which mobile robots have to incrementally discover features of interest by moving in initially unknown (or partially known) environments [3, 6]. We consider the problem of exploring for map building [8], in which the goal of a robot is to move in an initially unknown environment in order to build a map representing the locations of obstacles and the free space. The robot follows an *exploration strategy* to select the next best locations to reach in the partially explored environment [2, 4]. Most of the current exploration strategies ignore *prior knowledge* about the environment to explore that, in some cases, could be available. One of the few exceptions is [7], which shows that using accurate a priori knowledge has a positive impact on exploration performance. However, the question of whether also *inaccurate* a priori knowledge can improve exploration performance is still largely open.

In this paper, we address such question by presenting a method that includes a priori knowledge in an on-line exploration strategy

for a mobile robot that incrementally selects the next best locations the robot should reach by exploiting the knowledge of the floor plan of the indoor environment that is being explored. A floor plan is a two-dimensional representation of an environment composed of line segments (walls) that identify the spaces within the environment, like rooms and corridors. Such representation does not need to be fully accurate. For example, it usually does not include information about furniture, which can significantly limit the area that could be explored by a robot and affect path planning. Hence, although the floor plan is known, the map for safe navigation of a robot should be built and exploration is still required. We show that knowing a floor plan that is inaccurate can improve the exploration performance.

Our method can be practically applied to speed up the creation of maps of large environments exploiting (possibly inaccurate) prior knowledge, like in search and rescue, where the floor plan can be acquired from an evacuation map or from a blueprint, and in maintenance or cleaning tasks, that are repeated not very frequently, such that the environment is subject to some changes between different executions of the task (objects and furniture can change, while walls remain static). In this case, prior knowledge could be the map built in the previous execution of the task.

2 OUTLINE OF THE METHOD

We consider a robot, equipped with a laser range scanner with a given field of view and range, that explores an initially unknown planar indoor environment E , for which a floor plan E^{FP} is available. We do not assume that E^{FP} accurately represents E . The exploration process we consider is a typical frontier-based exploration composed of the following steps:

- (a) the robot perceives a portion of E from its current location p_R using the laser range scanner and integrates the new perception in the current map M_E of the environment,
- (b) the robot identifies the current set of frontiers in M_E , namely the boundaries between known and unknown space, and considers them as possible candidate locations,
- (c) the robot selects the most promising candidate location, according to an exploration strategy,
- (d) the robot reaches the selected location, updates p_R , and starts again from (a).

The above steps are repeated until no frontier is left and the map M_E represents all the free space of E .

The robot maintains a grid map M_E of the discovered environment using a SLAM (Simultaneous Localization And Mapping) algorithm. Each cell of M_E can be known or unknown and, in the former case, free or occupied. Given M_E , a *frontier* is a chain of free cells each one adjacent to at least an unknown cell. A *candidate location* is the cell that divides a frontier into two equal segments. Hence, given M_E , we have a set C of candidate locations.

Each candidate location $p \in C$ is evaluated in step (c) above according to an utility function $u(p)$ that combines distance and information gain [1]:

$$u(p) = \alpha \cdot d(p) + (1 - \alpha) \cdot i(p), \quad (1)$$

where parameter $\alpha \in [0, 1]$ weights the two components (we set $\alpha = 0.5$ in our experiments). In the above equation, $d(p)$ is the normalized *distance* utility value and is inversely proportional to the distance from p_R to p . Instead, $i(p)$ is the *information gain* utility value and is calculated as:

$$i(p) = \frac{I(p)}{I_{\max}}, \quad (2)$$

where $I(p)$ is the estimate of the amount of new (unexplored) area visible from p and I_{\max} is the maximum value of $I(p)$ over all the candidate locations $p \in C$. The next best candidate location p^* is thus selected from C as the one that maximises $u(p)$.

The state-of-the-art approaches for estimating $I(p)$ measures the maximum visible area from p given the footprint of the robot's laser range scanner (as done, e.g., in [2, 4]) or the length of the frontier (as partially done, e.g., in [9]). These approaches are reasonable if no a priori knowledge about the environment is available. This estimate is optimistic and implicitly assumes that the area beyond the frontier on which p is located is free of obstacles. In our approach, we calculate $I(p)$ by using the a priori information obtainable from E^{FP} . Given M_E and the floor plan E^{FP} (Figures 1a and 1b), we overlap them (Figure 1c) and calculate the amount of new area visible from p (Figure 1d).

3 EXPERIMENTS AND DISCUSSION

We implemented our approach using the ROS navigation stack and GMapping [5] for SLAM¹. Experiments are performed on a three-wheeled differential drive robot, called Robocom, equipped with a SICK LMS100 laser range scanner with a field of view of

¹http://wiki.ros.org/navigation_gmapping.




Figure 1: An example of how $I(p)$ (light blue area) is calculated exploiting the knowledge of the floor plan. Candidate location p is the red cell, free cells are white, obstacle cells are black, and unknown cells are gray.

coverage	without prior knowledge				with prior knowledge				difference	
	\mathcal{D}	σ	\mathcal{T}	σ	\mathcal{D}	σ	\mathcal{T}	σ	\mathcal{D}	\mathcal{T}
70%	33.49	8	386.03	94.86	26.76	2.36	281.20	25.80	-20%	-27%
80%	37.96	8	425.11	99.15	30.64	2.24	317.13	24.46	-19%	-25%
90%	44.17	7.87	488.46	120.29	37.33	1.82	368.33	27.46	-15%	-25%
95%	47.10	7.8	528.31	96.75	41.77	2.96	411.11	14.60	-11%	-22%

Table 1: Results (over 3 runs) of the experiments with the Robocom robot. \mathcal{D} is distance in m, \mathcal{T} is time in s, and σ is the corresponding standard deviation. The last two columns show the percentage difference in performance, according to \mathcal{D} and \mathcal{T} , of the strategy with prior knowledge over that without prior knowledge: negative numbers mean that the former performs better than the latter.

270° and a range of 20 m (Figure 2a). We measure, as exploration progresses, the *distance* \mathcal{D} travelled and *time* \mathcal{T} required by the robot and the percentage of *covered area*, namely the percentage of free area of E mapped in M_E . The runs are performed in an environment of size 36 m \times 27 m, with 3 exploration runs from the same initial position (Figure 2b). Results are averaged over the runs. Note that the discrepancies between the actual map and the floor plan can change due to the changes of furniture in different runs. We compare our approach to a state-of-the-art approach where the information gain $I(p)$ is evaluated without prior knowledge, measuring the maximum visible area from p , as in [2, 4].

Table 1 shows that our exploration strategy outperforms the exploration strategy without a priori knowledge. Our approach leads the robot to first explore frontiers with a higher information gain, reaching large percentages of explored area in a shorter time. The strategy without prior knowledge has a higher standard deviation than our approach, due to more variable decisions based on an overestimated information gain.

Overall, experiments suggest that the use of a priori knowledge can be particularly useful in human-inhabited settings where objects, furniture, people, and obstacles (as partially open doors) can negatively affect the perception of the robot. In these settings, the use of a floor plan, even if it does not faithfully represent the environment, provides an effective mean to drive the robot to select the next best locations for exploration.




Figure 2: Robocom (2a) and the floor plan of the test environment (2b). In red, the initial position of the robot.

REFERENCES

- [1] F. Amigoni. 2008. Experimental evaluation of some exploration strategies for mobile robots. In *Proc. ICRA*. 2818–2823.
- [2] N. Basilico and F. Amigoni. 2011. Exploration strategies based on multi-criteria decision making for searching environments in rescue operations. *Auton Robot* 31, 4 (2011), 401–417.
- [3] E. Galceran and M. Carreras. 2013. A survey on coverage path planning for robotics. *Robot Auton Syst* 61, 12 (2013), 1258 – 1276.
- [4] H. González-Baños and J. Latombe. 2002. Navigation Strategies for Exploring Indoor Environments. *Int J Robot Res* 21, 10-11 (2002), 829–848.
- [5] G. Grisetti, C. Stachniss, and W. Burgard. 2007. Improved Techniques for Grid Mapping with Rao-Blackwellized Particle Filters. *IEEE T Robot* 23 (2007), 34–46.
- [6] M. Julia, A. Gil, and O. Reinoso. 2012. A comparison of path planning strategies for autonomous exploration and mapping of unknown environments. *Auton Robot* 33, 4 (2012), 427–444.
- [7] S. Obšwald, M. Bennewitz, W. Burgard, and C. Stachniss. 2016. Speeding-Up Robot Exploration by Exploiting Background Information. *IEEE RA-L* 1, 2 (2016), 716–723.
- [8] S. Thrun. 2003. Robotic Mapping: A Survey. In *Exploring Artificial Intelligence in the New Millenium*, G. Lakemeyer and B. Nebel (Eds.). Morgan Kaufmann, 1–35.
- [9] B. Tovar, L. Munoz, R. Murrieta-Cid, M. Alencastre, R. Monroy, and S. Hutchinson. 2006. Planning exploration strategies for simultaneous localization and mapping. *Robot Auton Syst* 54, 4 (2006), 314–331.