Exploiting Inaccurate A Priori Knowledge in Robot Exploration*

Francesco Amigoni[1], Matteo Luperto[2], Danilo Fusi[1], and N. Alberto Borghese[2]

[1] Politecnico di Milano, Milano, Italy
[2] Università degli Studi di Milano, Milano, Italy

Abstract. Exploration is a task in which autonomous mobile robots incrementally discover features of interest in initially unknown environments. Usually, robots follow exploration strategies to select the next best locations to reach in partially explored environments. Most of the current exploration strategies ignore prior knowledge about the environments to explore that, in some practical cases, could be available. In this paper, we present a method that includes a priori knowledge in an exploration strategy for a mobile robot. In particular, our exploration strategy 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. Although the floor plan can be inaccurate (e.g., it typically does not include furniture and could represent a topology that does not fully match with that of the actual environment), we experimentally show, both in simulation and with real robots, that knowing the floor plan improves the exploration performance under a wide range of conditions.

Keywords: robot exploration · exploration strategies · robot mapping

1 Introduction

Exploration is an important task for autonomous mobile robots. It is employed when robots have to incrementally discover features of interest by moving in initially unknown (or partially known) environments [6, 10]. For example, exploration can discover the presence of occupied and free space, discover the concentration of substances in air or water, or discover thermal information in search and rescue operations [15]. In this paper, we consider the problem of exploring for map building [17], 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 of free space. The robot follows an *exploration strategy* to select the next best locations to reach in the partially explored environment [3, 8]. 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 [11], which exploits the knowledge of a topo-metric map of the environment

* This paper is the full version of an extended abstract accepted at the Robotics Track of AAMAS2019.

in which the robot is operating in order to plan an exploration path. While [11] shows that using accurate a priori knowledge has a positive impact on exploration performance, the question of whether also *inaccurate* a priori knowledge can improve exploration performance is still largely open.

In this paper, we address the above question by presenting a method that includes a priori knowledge in an on-line exploration strategy for a mobile robot. In particular, our exploration strategy 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* represents the static elements of an indoor environment, like walls and doorways, ignoring the dynamic elements, like furniture. A floor plan is a form of abstract prior knowledge that is usually easy to obtain from documents, blueprints, or evacuation plans. We show that knowing a floor plan that is inaccurate (e.g., that shows an incorrect topology) can nevertheless improve the exploration performance. Extensive experimental evaluation shows that our proposed exploration strategy outperforms exploration strategies that do not consider any a priori knowledge.

The original contributions of this paper are thus: (i) an on-line exploration strategy that originally exploits the knowledge of the floor plans of the environments being explored to select the most promising next locations for the robot and (ii) its extensive experimental evaluation, both in simulation and with real robots, under several conditions and several degrees of accuracy of the prior knowledge.

The method presented in this paper 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 of the building, and in maintenance or cleaning tasks, that are repeated not very frequently, like once a week or a month, 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.

This paper is structured as follows. The next section reviews related work and places the contributions of this paper against that background. Section 3 describes the proposed method, which is experimentally evaluated in Section 4. Finally, Section 5 concludes the paper.

2 Related Work

Exploration is the incremental process with which a robot (or a multirobot system) covers with its sensors an initially unknown environment. Two main families of approaches have been developed for exploration: *frontier-based* approaches, which move the robots to the geometrical boundaries between known and unknown portions of environments [21], and *information-based* approaches, which move the robots to the most informative locations, according to some information measure (e.g., [7, 13, 16]). In this paper, we focus on the first family of approaches, since they more naturally address the discovery of space for the problem of map building we consider.

Different exploration strategies have been proposed to select the next best frontier, all of them being greedy [19], due to the inherently on-line nature of the exploration problem. Usually, exploration strategies choose the next best frontier by evaluating candidate locations according to different criteria, but ignoring the prior knowledge about the environment that could be available. Although a complete survey is out of the scope of this paper (the reader can refer, e.g., to [10]), some examples of these exploration strategies follow. For instance, [8] evaluates each candidate location taking into account its distance from the robot's current position and the expected information gain (in terms of the maximum unexplored area that could be viewed from it). The two criteria are combined in an exponential utility function. In [20], the same two criteria are combined in a fractional utility function, where the information gain is the numerator and the distance is the denominator. Also [18] combines criteria related to distance and information gain in a complex utility function. In [2] and [3], more principled ways to aggregate criteria, based on multi-objective optimization, are proposed. In all the above cases, the combined criteria (related to distance and to information gain) are calculated only on the basis of the portion of the map that is already known.

Recently, some forms of prior knowledge have been exploited with the aim of improving the performance of exploration. In [14], predictions of the possible aspect of the unexplored parts of the environment are made by exploiting a database of previously mapped environments, in order to complete the partial maps obtained by the robot. A similar approach, but extended to multirobot settings, is that of [12], in which the performance of multirobot exploration is improved by using structural inference that completes the partially-explored portions of the environment by matching (parts of) maps contained in a library of previously explored structures. In both [14] and [12], differently from this paper, prior knowledge is relative to environments different from the one where the knowledge is used to inform robot's operations. In other words, they do not consider knowledge specific to the environment that is being explored.

The authors of [11] propose an exploration approach that, knowing a representation of the environment in terms of a topo-metric graph, finds an exploration path. The global exploration path is calculated solving a TSP (Travelling Salesperson Problem[3]) on the topo-metric graph and is completed locally by performing on-line explorations when the path is actually followed by the robot. Similarly to ours, this method exploits the knowledge of the same environment in which the robot operates. However, a difference between the two approaches is the nature of prior knowledge. In our case, it is a floor plan, which can be obtained from various sources, including blueprints and evacuation plans. In the case of [11], the prior knowledge is a topo-metric graph, whose nodes are lo-

[3] In a TSP, given a set of locations and a metric to calculate the distance between pairs of locations, the goal is to find the shortest tour that visits each location exactly once.

cations and edges are direct connections that roughly reflect the true metric distances between locations, that is manually built by the user.

3 Our Method

3.1 Overview of the Exploration Process

We consider a single robot, equipped with a laser range scanner with 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*, 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 updating 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 (reachable from the initial location of the robot).

The robot maintains a grid map M_E of the discovered environment using a SLAM algorithm. We use GMapping [9] in our experiments. Each cell of M_E can be known or unknown and, in the former case, free or occupied. Given M_E, we identify the chains of free cells that are adjacent to at least an unknown cell. Each of such chains is a frontier and the middle cell of each frontier is a candidate location. More precisely, a candidate location is the cell that divides a frontier into two equal segments. (Note that it is not safe for the robot to select a candidate location beyond a frontier.) 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 (e.g., as in [3, 8]). In particular ($\alpha \in [0, 1]$ is a parameter that weights the two components),

$$u(p) = \alpha \cdot d(p) + (1 - \alpha) \cdot i(p). \tag{1}$$

In the above equation, $d(p)$ is the *distance* utility value that is calculated as:

$$d(p) = \frac{D_{\max} - D(p, p_R)}{D_{\max}}, \tag{2}$$

where $D(p, p_R)$ is the Euclidean distance between the current location of the robot p_R and the candidate location p and $D_{\max}$ is the maximum $D(p, p_R)$ over all the candidate locations $p \in C$. In (1), $i(p)$ is the *information gain* utility value and is calculated as:

$$i(p) = \frac{I(p)}{I_{\max}}, \tag{3}$$

where $I(p)$ is the estimate of the amount of new (unexplored) area visible from p (calculated as described in Section 3.3) 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 as follows:

$$p^* = \operatorname*{argmax}_{p \in C} u(p). \tag{4}$$

According to the value of α, p^* represents the best balance between closeness and expected new area visible and, as such, is considered a good greedy choice for an efficient exploration of the environment [1]. (Note that with $\alpha = 1$ our method performs a closest-frontier exploration.) As shown in Section 4, the exploration performance is measured in terms of distance travelled and time employed to fully map the environment.

3.2 A Priori Knowledge

(a) (b) (c) (d)

Fig. 1: An example of floor plan (1a), a map built by the robot (1b), an overlay of the floor plan on the map (1c), and the environment simulated in Gazebo (1d).

In this paper, we focus on a specific type of a priori knowledge, which is the floor plan. For indoor environments, floor plans can be easily obtained from documents, blueprints, and even from evacuation plans [5]. A floor plan E^{FP} is a two-dimensional representation of the environment E composed of line segments (walls) that identify the spaces within the environment, like rooms and corridors. Note that E^{FP} does not need to be fully accurate, for example, it usually does not include information about furniture, which limits significantly the area of E that could be explored by a robot, and small objects, which affect path planning and whose number, type, and location cannot be known in advance. An example of floor plan is reported in Fig. 1a, while the corresponding (simulated) environment (with furniture) is in Fig. 1d. Moreover, E can have obstacles, as closed doors, that are not present in E^{FP} or can exhibit connections between locations that are not connected in E^{FP} (e.g., when E has been structurally modified and E^{FP} is outdated). In this sense, we say that E^{FP} can be topologically inaccurate. Hence, although the floor plan E^{FP} of E is known, the map M_E for safe navigation of a robot in E should be built and an exploration is still required. In the proposed approach, E^{FP} should be manually fed to the system. The required human effort required for this step involves: getting prior knowledge in the form of a floor plan (e.g., taking a picture of a paper blueprint), "cleaning" the floor plan image from unnecessary details (e.g., words indicating the name of the building and symbols like emergency exits), scaling and aligning the floor plan to the map built by the robot. The processing of E^{FP} is done once for each environment, and in our experience requires just few minutes.

In the proposed approach, E^{FP} is exploited to make informed decisions when evaluating and selecting the next best candidate location using Equations (3) and (4), respectively. More precisely, we compute an estimate of the area that the robot can perceive from a candidate location p (when it is oriented toward the unknown space, to maximize the new area perceived) by superimposing M_E on E^{FP}. We assume that E^{FP} and M_E are metrically consistent, namely that they are aligned and with the same scale, without deformations. In practice, this amounts to assume that the initial pose of the robot is known, that E^{FP} is represented as a grid map with the same resolution of M_E, where a cell of E^{FP} is either free or obstacle, and that E^{FP} and M_E are aligned in a global coordinate system. In our experiments, we perform the alignment manually, but an automatic method can be developed. For example, one can use the approach of [4] to localize the robot in E^{FP} and then calculate its alignment with M_E. Fig. 1 reports a map M_E (Fig. 1b) and the scaled and aligned floor plan E^{FP} overlapped to M_E (Fig. 1c).

While it could be limiting to use the knowledge derived from floor plans only on-line (namely, during the exploration process), because it seems that the availability of floor plans is not fully exploited, this approach copes well with increasing inaccuracies in prior knowledge. Indeed, in these cases, plans calculated off-line (namely, before the exploration process starts) could become useless, because they are built on inaccurate knowledge, and their revision (e.g., replanning) can be costly. For example, in the method of [11], replanning requires to solve a TSP, which can be done efficiently using a solver (Concorde, in that case) for topo-metric maps with a limited number of nodes (less than 100 in [11]), but that does not efficiently scale to larger instances, being the TSP a NP-hard problem.

3.3 Expected Information Gain

We calculate $I(p)$, namely the estimate of the amount of the unexplored area visible from a candidate location p by using the a priori information obtainable from E^{FP}.

The state-of-the-art approaches for estimating $I(p)$ measure the maximum visible area from p given the footprint of the robot's laser range scanner (as done, e.g., in [3,8]) or the length of the frontier (as partially done, e.g., in [18]). These approaches are reasonable if no a priori knowledge about the environment

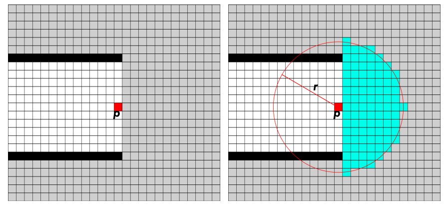

(a) (b)

Fig. 2: A candidate location p (red cell) in a grid map (free cells are white, obstacle cells are black, and unknown cells are gray) (2a) and $I(p)$ (light blue area) computed according to the footprint of the robot's laser range scanner (2b).

is available. Fig. 2 shows an example in which $I(p)$ is calculated as the maximum area of the unknown cells that can be perceived by the laser range scanner from p. This estimate is optimistic and implicitly assumes that the area beyond the frontier on which p is located is free of obstacles.

Instead, we calculate $I(p)$ as follows. Given $p \in M_E$, we find the corresponding cell $p^{\mathrm{FP}} \in E^{\mathrm{FP}}$ (i.e., the cell with the same center after scaling and alignment of M_E with E^{FP}). Then, for each unknown cell $c \in M_E$ that is within the footprint of the laser range scanner when the robot is in p, we find the corresponding $c^{\mathrm{FP}} \in E^{\mathrm{FP}}$. The cell c contributes to calculate the expected area $I(p)$ visible from p when all the following conditions are all satisfied:

- c^{FP} is free,
- c^{FP} is visible from p^{FP} in E^{FP}, namely the line segment connecting their centers does not touch any obstacle cell in E^{FP}, and
- c is visible from p in M_E.

Eventually, given the cells c that satisfy the above conditions, $I(p)$ is calculated by summing the areas of those cells. Fig. 3 shows an example, in which the method just described is used to calculate $I(p)$. It is interesting to contrast it with Fig. 2. Although the proposed approach appears to be a variant of classical frontier-based exploration approaches, it provides significant benefits to the performance of exploration also when E^{FP} is inaccurate, as we show in the next section.

4 Experimental Evaluation

To evaluate our approach and its ability to exploit inaccurate a priori knowledge for efficient exploration, we present several tests conducted both in simulation and with real robots. We measure, as exploration progresses, the *distance* travelled by the robot (as done, e.g., in [8, 18]) and the percentage of *covered area*, namely the percentage of free area of E mapped in M_E, as done, e.g., in [3]. To

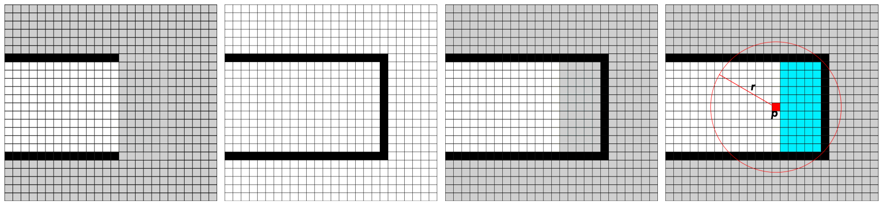

(a) Map M_E (b) Floor plan E^{FP} (c) Overlap (d) $I(p)$

Fig. 3: An example of how $I(p)$ is calculated exploiting the knowledge of the floor plan. See Fig. 2 for notation.

have a fair comparison, we present results up to 95 % of coverage, since some runs end without reaching full coverage, due to noise in localization and mapping.

4.1 Simulations

Simulations are performed with ROS Gazebo, using the ROS GMapping and Nav2D packages[4] for SLAM and robot navigation, respectively. We consider three indoor environments with different characteristics (Fig. 4): a *basic* environment (19 m × 10 m) that represents a small apartment, an *office* environment (90 m × 53 m) with several small rooms, and an *open* environment (57 m × 45 m) with few large rooms.

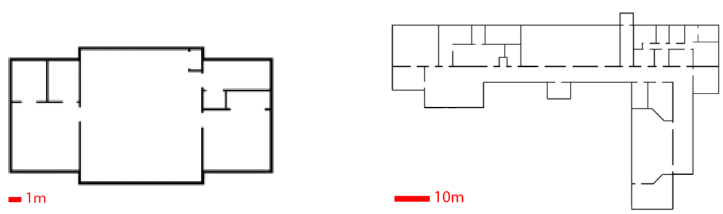

(a) basic (19 × 10 m) (b) office (90 × 53 m)

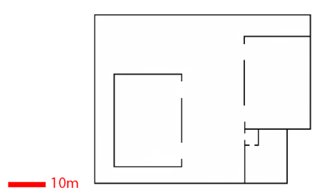

(c) open (57 × 45 m)

Fig. 4: The floor plans of environments used in simulations.

[4] `http://wiki.ros.org/{gazebo,gmapping,nav2d}`. Gazebo is a 3D dynamic robot simulator.

Two settings are considered for each environment, namely with and without furniture. The furniture is coherent with the type of the environment. For the basic environment, the furniture includes couches, chairs, tables, beds, night stands, and so on (Fig. 1d). Furniture include desks and chairs for the office environment and shelves for the open environment. Environments are static and differences between them and the floor plans are due to the presence of furniture.

Exploration is performed by a simulated robot equipped with a laser range scanner with a field of view of 180° and range of 10 m for the basic environment and of 25 m for the other two environments. In our experiments, longer ranges for the laser range scanner result in less frontiers and in less decisions. Linear and angular speeds of the robot are fixed, so the robot moves at a constant velocity. After several preliminary tests, we set $\alpha = 0.5$ in (1) to equally balance distance and information gain. For each environment, we perform 10 exploration runs (starting from the same pose) and average results over the runs. Gaussian noise with zero mean and 0.1 m standard deviation is added to scans acquired by the laser range scanner. The maps obtained in different runs are, as a consequence, slightly different from each other, thus resulting in different frontiers being detected and, ultimately, in different choices being made by the robot.

We compare our approach to a state-of-the-art approach where the information gain $I(p)$ is evaluated without prior knowledge, taking an optimistic stance, as in [3, 8] (see Fig. 2).

We start from the open environment without furniture. Table 1 shows that our approach is able to explore a given amount of area in a significantly shorter time $\mathcal{T}$ than the approach without prior knowledge, especially when the covered area is less than 80 %. As an example, the difference at 60 % has p-value= 0.00034 in one-way ANOVA. Similarly, the robot reaches the same percentage of covered area travelling a shorter distance $\mathcal{D}$ when using our approach. This is because frontiers that are close to walls are correctly evaluated to have small $I(p)$ using our method, while the same frontiers can have large $I(p)$ when no a priori knowledge is used. The gain reduces towards the end of the exploration because the selection of the few frontiers left is less critical, thus reducing the impact of a wrong decision.

Similar results are found for the furnished open environment, as shown by Table 2, where the trends of the two strategies become very different after reaching a coverage of 90%. From that point, our exploration strategy has a consistently better performance with respect to the strategy without prior knowledge. For instance, the difference between the time required by the two strategies is significant for coverage values of 90 % and 95 %, with p-value= 0.0012 and p-value= 0.00035, respectively. Note that the prior knowledge represented by E^{FP} is less accurate in the case of the furnished environment than in the case of the unfurnished environment.

The same results are qualitatively found also for the other two environments (full data are not reported). Finally, Table 3 reports the overall results obtained in the three simulated environments for 95 % of coverage. Using a priori knowledge improves considerably the performance in all the three cases, allowing the

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%	71.15	14.43	101.39	21.41	43.99	3.51	66.07	4.97	-38%	-35%
80%	83.53	13.75	116.29	21.04	73.95	5.30	104.08	10.64	-11%	-10%
90%	140.30	22.30	188.79	27.62	127.20	11.68	176.70	16.48	-9%	-6%
95%	186.87	30.14	244.82	39.61	161.40	17.52	218.12	24.27	-14%	-11%

Table 1: Results for the unfurnished open environment (Fig. 4c). $\mathcal{D}$ 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.

robot to cover large portions of the environments travelling a shorter distance and spending less time. Note that, as the robot speed is fixed, differences in average speeds are due to the different exploration paths. Using prior knowledge results in more direct and straight paths, with less rotations of the robot.

We now consider two environments from [11] (reported in Figs. 5a and 5b) and we use the same range of the robot's laser range scanner (6 m) and the same initial positions (in red in the figures) used in [11]. The first environment has approximately a size of 34 m × 34 m, while the second one has a size of 38 m × 25 m. The evaluation is done by performing one exploration run for each initial position. Results are then averaged over runs.

Authors of [11] report that, in case of fully accurate a priori knowledge ($E^{\mathrm{FP}} = E$ in our notation), the exploration path built with their method reaches a coverage of 100 % by travelling a distance of about 239 m in the first environment and of about 171 m in the second one. Our approach fully explores the environments travelling a longer distance, about 433 m and 429 m, respectively. This is expected because the method of [11] plans off-line a global exploration path that is completed on-line, while our approach is fully on-line. However, after 239 m and 171 m, our approach remarkably covers $\simeq$ 87 % and $\simeq$ 91 % of

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%	78.86	12.84	109.87	17.89	67.51	9.79	92.91	15.54	-14%	-15%
80%	117.23	13.27	159.02	16.37	111.87	6.95	149.95	149.95	-5%	-6%
90%	237.26	23.03	315.90	31.44	185.24	23.26	245.85	42.11	-22%	-22%
95%	305.83	28.43	407.35	42,44	223.56	18.54	293.88	36.20	-27%	-28%

Table 2: Results for the furnished open environment (Fig.4c). See Table 1 for notation.

		basic environment			office environment			open environment		
		without prior knowledge	with prior knowledge	difference	without prior knowledge	with prior knowledge	difference	without prior knowledge	with prior knowledge	difference
without	$\mathcal{T}$	100.69	70.09	$-30\,\%$	570.81	500.22	$-12\,\%$	244.82	218.12	$-11\,\%$
furniture	$\mathcal{D}$	36.14	28.99	$-20\,\%$	407.55	370.83	$-9\,\%$	186.87	161.40	$-14\,\%$
with	$\mathcal{T}$	162.34	137.99	$-15\,\%$	645.80	568.40	$-12\,\%$	407.35	293.89	$-28\,\%$
furniture	$\mathcal{D}$	57.89	56.95	$-1\,\%$	460.04	414.29	$-9\,\%$	305.83	223.56	$-27\,\%$

Table 3: Performance evaluation at 95 % of coverage for the simulated environments of Fig. 4. See Table 1 for notation.

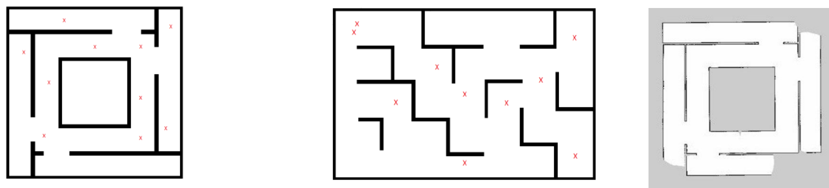

(a) 34 m × 34 m (b) 38 m × 25 m (c) map of (5a)

Fig. 5: Two environments from [11] (5a and 5b) and a map covering the 90 % of the first environment (5c).

the area of the two environments, respectively (with a standard deviation of 2 % and 3 %, respectively). This suggests that our approach makes very good initial decisions and quickly covers almost completely the environments but then it needs to travel back to cover small portions that have been left behind (like the top right and the bottom left corners in Fig. 5c).

Results from [11] show that, when Gaussian noise with standard deviation larger than 2.5 m is added to the node positions of the topo-metric map, their method performance is on par with that of a method that does not consider prior knowledge. We cannot compare our method directly to these results, because the noise added in [11] to the user-provided topo-metric map cannot be transferred to the floor plan we consider. However, in the following, we show that, also when available knowledge of E^{FP} is inaccurate, our method performs consistently better than an exploration strategy that does not consider prior knowledge.

The computing effort of our approach is negligible (selection of frontiers takes a time in the order of milliseconds).

4.2 Hand-Drawn Floor Plans

In this section, we test our approach using inaccurate floor plans in order to evaluate its robustness. More specifically, we use three hand-drawn floor plans as E^{FP} and we compare, as before, the performance of our exploration strategy with that of the exploration strategy that does not use any prior knowledge.

The floor plans represent a building and are hand-drawn by three different people who work in the same building, and are later digitalized and scaled. Draws are based only on their memory, without any support for recollection. The environment has approximately a size of 50 m × 43 m. The correct floor plan and the three digitalized hand-drawn floor plans are reported in Fig. 6, along with one of the original drawings. While the corridors are drawn almost correctly and consistently in the three floor plans, the number and size of rooms are different and several small rooms are missing or are merged to adjacent bigger rooms. In practice, while the hand-drawn floor plans might appear visually similar to the real floor plan, a point in a room of a hand-drawn floor plan could correspond to a point in a completely different room of the real floor plan.

Experiments are made in a simulated version of the building, performing 5 exploration runs starting from the same initial position (results are averaged over the runs). The range of the robot's laser range scanner is set to 25 m.

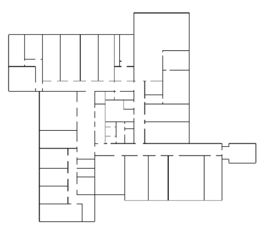

(a) floor plan (50 m x 43 m)

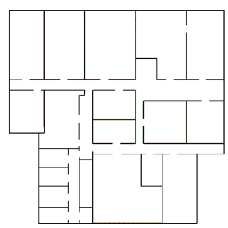

(b) hand-drawn 1

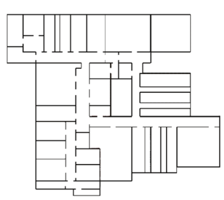

(c) hand-drawn 2

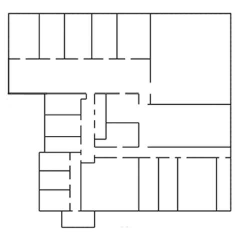

(d) hand-drawn 3

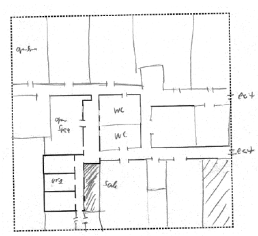

(e) original draw of (6b)

Fig. 6: The correct floor plan (6a), three hand-drawn floor plans (6b-6d), and an original draw (6e).

In Table 4 we see how the accuracy of a priori knowledge affects the performance of the exploration. The use of the more accurate prior knowledge provided by the correct floor plan leads to a significant improvement in performance. The use of hand-drawn floor plans as a priori knowledge sometimes leads to wrong estimates of the information gain, delaying the exploration of frontiers which could bring the robot to perceive larger amounts of new area of the environment. In fact, the presence of an obstacle in a hand-drawn floor plan could cause an underestimation of the actual information gain from a frontier, leading the robot to explore other frontiers with an actual lower information gain and, therefore, worsening the system performance. However, it is remarkable that, despite the inaccurate knowledge, in all cases, the exploration strategies with prior knowledge make the robot travel a shorter distance than the strategy without prior knowledge, covering the 95 % of the area with a gain that ranges from 3 % to 9 %.

coverage	without prior knowledge		floor plan			hand-drawn 1			hand-drawn 2			hand-drawn 3		
	$\mathcal{D}$	σ	$\mathcal{D}$	σ	difference	$\mathcal{D}$	σ	difference	$\mathcal{D}$	σ	difference	$\mathcal{D}$	σ	difference
70%	129.22	16.42	101.84	6.83	-21%	118.60	7.16	-8%	143.29	15.50	+10%	152.91	17.32	+18%
80%	187.87	36.22	151.35	25.76	-19%	189.34	13.19	0%	179.78	12.72	-4%	195.59	18.92	-4%
90%	253.24	42.74	235.18	26.06	-7%	239.92	20.83	-5%	225.13	10.79	-11%	236.60	27.39	-7%
95%	320.91	31.33	302.33	17.79	-6%	298.11	45.42	-7%	290.45	33.36	-9%	310.49	19.97	-3%

Table 4: Performance comparison between the exploration strategy without prior knowledge and the exploration strategies with prior knowledge with correct floor plan and with the approximate hand-drawn floor plans in the environment of Fig. 6. See Table 1 for notation. Results are over 5 runs.

4.3 Experiments with Real Robots

In this section we describe the results of the experiments in different environments conducted with the implementation of our approach on two autonomous mobile robots, running the same ROS configuration used for simulations. Also in this case, we compare our approach with an exploration strategy that does not use a priori knowledge.

The first set of experiments is performed on a three-wheeled differential drive robot, called Robocom, equipped with a SICK LMS100 laser range scanner with a field of view of 270° and a range of 20 m (Fig. 7a). The runs are performed in a portion of the environment of Fig. 6a, with a size of 36 m × 27 m and shown in Fig. 7b, performing 3 exploration runs from the same initial position. Results are averaged over the runs. Note that discrepancies between the actual map and the floor plans can change due to the changes of furniture in different runs.

The results in Table 5 confirm that our exploration strategy outperforms the exploration strategy without a priori knowledge. Also in this case, the main reason why our approach has better performance lies in a better information gain estimate. In particular, our approach gives low $I(p)$ to frontiers that are close to walls and large $I(p)$ to frontiers that are in cluttered areas but, according to the floor plan, are far from walls. This leads the robot to first explore frontiers with a higher information gain, reaching large percentages of explored area in a shorter time. This behavior is more evident when clutter and occlusions increase, e.g., when the number of obstacles (like furniture) increases. Because of this, the

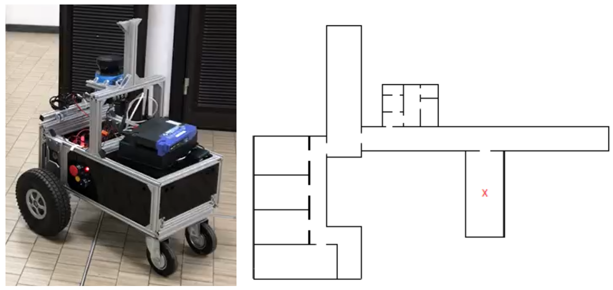

(a) (b)

Fig. 7: Robocom (7a) used in experiments in the environment with floor plan in 7b. In red, the initial position of the robot.

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 5: Results (over 3 runs) of the experiments with the Robocom robot. See Table 1 for notation.

coverage	without prior knowledge				with plain floor plan				difference		with modified floor plan				difference	
	$\mathcal{D}$	σ	$\mathcal{T}$	σ	$\mathcal{D}$	σ	$\mathcal{T}$	σ	$\mathcal{D}$	$\mathcal{T}$	$\mathcal{D}$	σ	$\mathcal{T}$	σ	$\mathcal{D}$	$\mathcal{T}$
70%	3.96	0.45	25.86	2.00	4.54	0.58	24.00	0.99	+15%	-7%	4.43	0.22	23.36	1.93	+11%	-9%
80%	7.56	2.92	56.06	17.81	5.91	0.09	33.70	2.57	-21%	-40%	4.80	0.43	25.52	1.82	-37%	-55%
90%	10.56	3.58	75.60	26.64	9.98	1.21	58.99	6.74	-6%	-22%	9.04	1.93	53.11	13.00	-15%	-30%
95%	24.68	5.04	171.05	31.15	17.70	6.68	116.87	26.90	-28%	-32%	-	-	-	-	-	-

Table 6: Results (over 3 runs) of the experiments with the TurtleBot3 Burger robot, using a plain floor plan and a modified one as a priori knowledge. See Table 1 for notation.

application of our method to real world settings provides even more advantages than in the more controlled environments used in simulations. It can also be noted that the strategy without prior knowledge has a higher standard deviation than our approach, due to more variable decisions based on an overestimated information gain. This is also confirmed by looking at the average speed along the different exploration paths (remember that linear and angular speeds are fixed). Paths obtained using prior knowledge are more direct and straight (followed at about 10 cm/s) than the paths obtained without prior knowledge (followed at less than 9 cm/s).

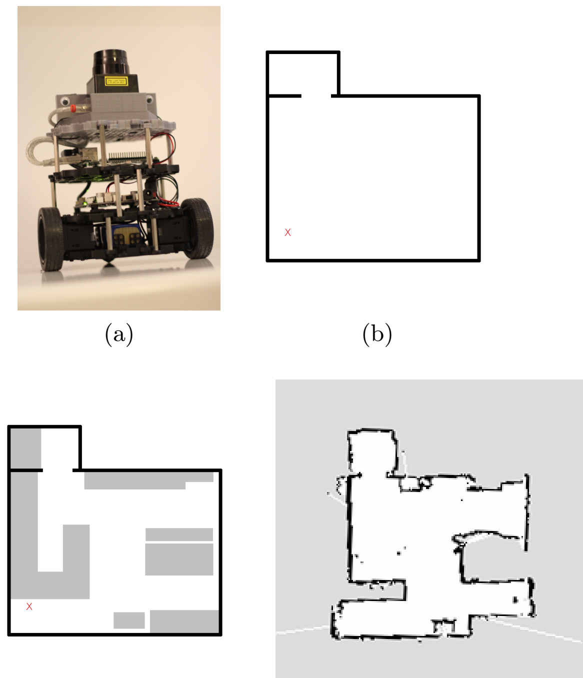

(a) (b) (c) (d)

Fig. 8: The TurtleBot3 Burger (8a) used in experiments in the environment with plain floor plan shown in (8b). In red, the initial position of the robot. (8c) shows the modified floor plan including furniture, in gray, while (8d) shows the map created by the robot in one of the runs.

The second set of experiments is performed using a TurtleBot3 Burger[5] (Fig. 8a), a small-size two-wheeled differential drive robot, equipped with a Hokuyo URG-04LX-UG laser range scanner with a field of view of 180° and a range of 5.6 m.

The runs are performed in a house environment with a size of 6 m × 6 m with an area of 30 m^2 (Fig. 8b), performing 3 exploration runs from the same initial position and averaging over the runs. The total area that could be explored by the robot is limited by the furniture, as 9 m^2 of the total 30 m^2 are occupied by static furniture, like wardrobes, shelves, and a kitchen, and cannot be accessed by the robot. Others 2 m^2 are occupied by dynamic furniture like chairs and small objects lying on the floor. Overall, 30% of the area is covered by (static) furniture and about 5% is covered by clutter, resulting in the floor plan E^{FP} overestimating the amount of explorable area by 35%. This difference between the floor plan E^{FP} and the actual environment E where the robot operates allows

[5] http://www.robotis.us/turtlebot-3/

us to test the robustness of our approach, and in particular to investigate how the accuracy of the prior knowledge affects its performance in the real world.

To this end, three configurations are used in this setting: an exploration strategy that does not use a priori knowledge, an exploration strategy that uses the floor plan of the house as a priori knowledge E^{FP}, and an exploration strategy that uses E^{FP} modified in order to include the static furniture as a priori knowledge. The modified E^{FP} is a more accurate representation of the environment than the plain E^{FP}, but it still does not include small objects such as chairs.

Results are shown in Table 6 and confirm that our exploration strategy is more efficient than that without a priori knowledge. Despite the fact that the plain E^{FP} is a rough estimate of the actual environment, the gain in terms of time and distance is consistent. A more significative improvement is obtained with the use of the modified floor plan that includes knowledge of static furniture. In general, with both types of prior knowledge, the advantages of our approach over the one not using prior knowledge are clear, especially after the robot has explored some portions of the environment (when the covered area is more than 75% in Table 6). Using or not prior knowledge could make little difference in the early stages of exploration, since the robot typically reaches frontiers close to the starting point. After this initial phase, the use of prior knowledge drives the robot directly to the most interesting frontiers, thus reducing the distance and the time required to map the environment.

Interestingly, the use of the more accurate floor plan (modified E^{FP}, which includes static furniture) stops the exploration process at about 92% of the explorable area, since no further frontier is detected (see '-' in Table 6). This is because the robot, when not considering any prior knowledge or when considering the empty E^{FP}, can map thoroughly also the small gaps between pieces of furniture or between furniture and walls. For example, in our runs without prior knowledge, a small (approximatively 10 cm) gap between a sofa and a wall is detected as an interesting and explorable area, and the robot tries to reach a location from which it can observe such small gap (although it cannot enter in that narrow space). When the modified floor plan is used, the robot "knows" that the gap area is uninteresting and does not try to observe it, selecting more interesting frontiers.

Overall, experiments performed with real robots in real environments 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 is far from faithfully representing the environment (as in the case of the plain E^{FP}), provides an effective mean to drive the robot to select the next best locations for exploration. In conclusion, the performance improvement of our proposed exploration strategy over the strategy that does not consider prior knowledge is more evident in the real world, which is inherently more complex and noisy, than in simulations.

5 Conclusions

In this paper, we introduced an on-line exploration strategy that exploits a priori knowledge, in the form of floor plans, to select the next best locations for a robot exploring indoor environments. Experiments assessed the effectiveness of the proposed method, also when the floor plans are inaccurate. In a nutshell, our results show that, while it is intuitive that accurate prior knowledge can improve the performance of the exploration process, also inaccurate prior knowledge can provide some benefits, which is far less intuitive.

The question of how to best balance the effort to get accurate prior knowledge and the improvements on the exploration performance is still open and is a direction for future work. Future work will also investigate the use of other forms of a priori knowledge, like pictures of evacuation maps that can be easily obtained in large buildings, and the quantitative relationship between the quality of a priori knowledge and the exploration performance. Moreover, means to represent, and include in the evaluation of the information gain, the uncertainty of the prior floor plan will be considered. Inspirations could come from methods to update maps when robots discover new features that do not match current expectations. Finally, the use of partial or empty floor plans could be studied. With empty floor plans, our method basically becomes a closest-frontier exploration, with information gain equal for all frontiers. With partial floor plans, our method works as shown in this paper for the known parts and as a closest-frontier exploration for the unknown parts of the floor plan.

Acknowledgements

We are glad to thank Stefan Oßwald and authors of [11] for kindly providing datasets and details of their method.

References

1. Amigoni, F.: Experimental evaluation of some exploration strategies for mobile robots. In: Proc. ICRA. pp. 2818–2823 (2008)
2. Amigoni, F., Gallo, A.: A multi-objective exploration strategy for mobile robots. In: Proc. ICRA. pp. 3850–3855 (2005)
3. Basilico, N., Amigoni, F.: Exploration strategies based on multi-criteria decision making for searching environments in rescue operations. Auton Robot **31**(4), 401–417 (2011)
4. Boniardi, F., Behzadian, B., Burgard, W., Tipaldi, G.: Robot navigation in hand-drawn sketched maps. In: Proc. ECMR. pp. 1–6 (2015)
5. Boniardi, F., Caselitz, T., Kümmerle, R., Burgard, W.: Robust lidar-based localization in architectural floor plans. Adv Robotics **26**(17), 2021–2041 (2012)
6. Galceran, E., Carreras, M.: A survey on coverage path planning for robotics. Robot Auton Syst **61**(12), 1258 – 1276 (2013)
7. Girdhar, Y., Whitney, D., Dudek, G.: Curiosity based exploration for learning terrain models. In: Proc. ICRA. pp. 578–584 (2014)

8. González-Baños, H., Latombe, J.: Navigation strategies for exploring indoor environments. Int J Robot Res **21**(10-11), 829–848 (2002)
9. Grisetti, G., Stachniss, C., Burgard, W.: Improved techniques for grid mapping with Rao-Blackwellized particle filters. IEEE T Robot **23**, 34–46 (2007)
10. Julia, M., Gil, A., Reinoso, O.: A comparison of path planning strategies for autonomous exploration and mapping of unknown environments. Auton Robot **33**(4), 427–444 (2012)
11. Oßwald, S., Bennewitz, M., Burgard, W., Stachniss, C.: Speeding-up robot exploration by exploiting background information. IEEE RA-L **1**(2), 716–723 (2016)
12. Smith, A., Hollinger, G.: Distributed inference-based multi-robot exploration. Auton Robot **42**(8), 1651–1668 (2018)
13. Stachniss, C., Grisetti, G., Burgard, W.: Information gain-based exploration using Rao-Blackwellized particle filters. In: Proc. RSS. pp. 65–72 (2005)
14. Ström, D., Bogoslavskyi, I., Stachniss, C.: Robust exploration and homing for autonomous robots. Robot Auton Syst **90**, 125–135 (2017)
15. Tabib, W., Whittaker, R., Michael, N.: Efficient multi-sensor exploration using dependent observations and conditional mutual information. In: Proc. SSRR. pp. 42–47 (2016)
16. Thompson, D., Wettergreen, D., Foil, G., Furlong, M., Kiran, A.R.: Spatio-spectral exploration combining in situ and remote measurements. In: Proc. AAAI. pp. 3679–3685 (2015)
17. Thrun, S.: Robotic mapping: A survey. In: Lakemeyer, G., Nebel, B. (eds.) Exploring Artificial Intelligence in the New Millenium, pp. 1–35. Morgan Kaufmann (2003)
18. Tovar, B., Munoz, L., Murrieta-Cid, R., Alencastre, M., Monroy, R., Hutchinson, S.: Planning exploration strategies for simultaneous localization and mapping. Robot Auton Syst **54**(4), 314–331 (2006)
19. Tovey, C., Koenig, S.: Improved analysis of greedy mapping. In: Proc. IROS. pp. 3251–3257 (2003)
20. Visser, A., Slamet, B.A.: Balancing the information gain against the movement cost for multi-robot frontier exploration. In: Proc. EUROS. pp. 43–52 (2008)
21. Yamauchi, B.: A frontier-based approach for autonomous exploration. In: Proc. CIRA. pp. 146–151 (1997)