

Robot Exploration of Indoor Environments Using Incomplete and Inaccurate Prior Knowledge

Matteo Luperto^{a,*}, Michele Antonazzi^a, Francesco Amigoni^b, N. Alberto Borghese^a

^a*Dipartimento di Informatica, Università degli Studi di Milano*

^b*Dipartimento di Elettronica, Informatica e Bioingegneria, Politecnico di Milano*

Abstract

Exploration is a task in which autonomous mobile robots incrementally discover features of interest in initially unknown environments. We consider the problem of exploration for map building, in which a robot explores an indoor environment in order to build a metric map. Most of the current exploration strategies used to select the next best locations to visit ignore prior knowledge about the environments to explore that, in some practical cases, could be available. In this paper, we present an exploration strategy that evaluates the amount of new areas that can be perceived from a location according to a priori knowledge about the structure of the indoor environment being explored, like the floor plan or the contour of external walls. Although this knowledge can be incomplete and inaccurate (e.g., a floor plan typically does not represent furniture and objects and consequently may not fully mirror the structure of the real environment), we experimentally show, both in simulation and with real robots, that employing prior knowledge improves the exploration performance in a wide range of settings.

*Corresponding author

Email addresses: matteo.luperto@unimi.it (Matteo Luperto),
michele.antonazzi@studenti.unimi.it (Michele Antonazzi),
francesco.amigoni@polimi.it (Francesco Amigoni), alberto.borghese@unimi.it (N. Alberto Borghese)

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 [1, 2]. For example, exploration can discover the presence of occupied and free space, the concentration of substances in air or water [3], or the presence of victims in search and rescue operations [4]. In this paper, we consider the problem of exploring for map building [5], in which a robot moves in an initially unknown indoor environment in order to build a two-dimensional 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 [6, 7]. Most of the current exploration strategies ignore *prior knowledge* about the environment to explore even if, in some cases, such knowledge could be available. One of the few exceptions is [8], which exploits knowledge of a topo-metric map of the environment in which the robot is operating in order to improve the exploration performance. The question of whether also *incomplete* and *inaccurate* a priori knowledge can improve exploration performance is still largely open.

In this paper, we present a method that originally includes incomplete and inaccurate a priori knowledge in an on-line exploration strategy for a mobile robot operating in indoor environments. A priori knowledge is incomplete when it does not represent some elements of the real environment, like the furniture. A priori knowledge is inaccurate when it does not precisely represent the properties of some elements of the real environment, for instance, the locations of walls. Our exploration strategy incrementally selects the next best locations the robot should reach by exploiting prior knowledge related to the environment that is being explored. We consider different types of prior knowledge, with different levels of completeness and accuracy. *Floor plans* provide accurate knowledge about structure, namely about the rooms and the doorways of

the environments, but no knowledge about furniture, and can be obtained from
30 blueprints or evacuation maps. *Footprints* accurately represent the external
contour (shape) of environments, but have no information about the interior
arrangement of walls, and could be obtained from aerial views or satellite maps.
Bounding boxes are rectangular approximations enclosing the footprints of en-
vironments and can be easily estimated from the outside. In our experiments,
35 we consider also other types of prior knowledge: *hand-drawn maps*, which are
coarse approximations of floor plans drawn by people with some knowledge of
the corresponding environments, and *previous maps* built by robots in earlier
explorations of environments that change over time.

We show that exploiting prior knowledge that is incomplete (e.g., that does
40 not represent furniture) and inaccurate (e.g., that shows incorrect topology or
incorrect location of walls) can improve the exploration performance, even in
situations where the difference between the prior knowledge and the actual
environment is substantial. Extensive experimental evaluation shows that our
proposed exploration strategy outperforms exploration strategies that do not
45 consider any a priori knowledge.

The method presented in this paper can be practically applied to speed up
the creation of maps of large environments, for example in search and rescue,
where a floor plan can be acquired from an evacuation map or a footprint
can be obtained from aerial images, and in maintenance or cleaning tasks that
50 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 maps built in the previous execution of the task. In our
experimental evaluation, we show an example of this setting.

55 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.

A short preliminary version of some ideas presented here appears in [9]. This

60 paper is an extended version of [10]. Specifically, while in [10] we consider only
prior knowledge related to floor plans, here we investigate different types of prior
knowledge in order to assess the relationship between the accuracy of the prior
knowledge available to the robot and the speed-up of performance in exploration.
Accordingly, also the experimental results are significantly extended with respect
65 to [10].

2. Related Work

Exploration is the incremental process with which a robot (or a multi-robot
system) covers with its sensors an initially unknown environment. Two main
families of approaches have been developed for exploration: *frontier-based* ap-
70 proaches, which move the robots to the geometrical boundaries between known
and unknown portions of environments [11], and *information-based* approaches,
which move the robots to the most informative locations, according to some
information measure (e.g., [12, 13, 14]). In this paper, we focus on the first
family of approaches.

75 Different exploration strategies have been proposed to select the next best
frontier, all of them being greedy [15], 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.
80 Although a complete survey is out of the scope of this paper (the reader can
refer, e.g., to [1]), some examples of these exploration strategies follow. For
instance, [6] 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
85 are combined in an exponential utility function. [16] combines criteria related to
distance and information gain in a complex utility function. In [7] and [17], 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 based on the portion of the map that is
90 already known. Finally, in [18] authors propose a method that matches partially-
explored parts of a map with fully-explored parts of the same map to predict
the structure of unexplored regions in order to improve SLAM performance by
providing a better frontier evaluation.

Recently, some forms of prior knowledge have been exploited with the aim
95 of improving the performance of exploration. In [19], 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 [20]. In both [19] and [20], differently from this paper,
100 prior knowledge is relative to environments different from the one where the
knowledge is used to inform the robot’s operations. In other words, they do not
consider knowledge specific to the environment that is being explored. A recent
work [21] presents a method that uses deep learning to predict the missing
parts of partially explored rooms for speeding up the exploration. Although
105 the method of [21] shares some similarities with that of this paper, it does not
exploit prior knowledge of the same environment being explored. Moreover, it
is evaluated on clean and empty simulated environments, as the deep neural
network on which it is based needs to be trained on an extensive data set of
large-scale maps, while our method can be used on cluttered environments, as
110 we show later in this paper.

The authors of [8] propose an exploration approach that, knowing a repre-
sentation of the environment in terms of a topo-metric graph, whose nodes are
locations and edges are direct connections that roughly reflect the true metric
distances between locations, finds an exploration path. A global exploration
115 path is calculated off-line by solving a TSP (Travelling Salesperson Problem¹)
on the topo-metric graph and is then completed locally by performing on-line

¹In 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.

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, the two approaches differ in the nature of prior knowledge they use. In our case, it is an incomplete and inaccurate representation of the geometry of the environment, which can be obtained from various sources, including blueprints and aerial images. In the case of [8], the prior knowledge is a topo-metric graph that is manually built by the user. Further details are discussed in Section 4.5.

A recent work [22] uses prior maps taken from emergency maps to build graph-based SLAM representations. The method, which extends that of [23], merges the information from sensors and a prior map into a representation of the environment and corrects errors both in sensor readings and the prior map. Our focus is on exploration performance and not on the quality of the produced map (as in [22]), but the relationships between the two approaches could be further investigated.

3. The Proposed Method

3.1. Overview of the Exploration Process

In this work, we propose an on-line exploration strategy that exploits possibly incomplete and inaccurate a priori knowledge of the environment to determine the next best locations the robot should reach. We consider two-dimensional indoor environments and we assume to have a priori knowledge with a two-dimensional representation, like floor plans or footprints (see Section 3.2).

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 two-dimensional prior knowledge PK_E is available. We do not assume that PK_E completely and accurately represents E . Our approach follows the typical steps of frontier-based exploration [11, 24]:

- 145 (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,
- 150 (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).

155 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).

Our robot builds a grid map M_E of the discovered environment using the GMapping SLAM algorithm [25]. 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

160 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. Thus, a candidate location $p \in C$ is the cell that divides a frontier into two equal segments. Candidate locations are evaluated in step (c) above according to a utility function $u(p)$ that combines distance and information gain (e.g., as in

165 [6, 7]):

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

where $\alpha \in [0, 1]$ is a parameter that weights the two components and $d(p)$ is the *distance* utility value:

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

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

170 all the candidate locations $p \in C$. The Euclidean distance $D(p, p_R)$ provides an underestimate of the length of the actual path from p_R to p that is fast to

compute. In (1), $i(p)$ is the *information gain* utility value and is calculated as:

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

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 $p \in C$.

175 As illustrated in the following section, our main original contribution lies in calculating $I(p)$ exploiting prior knowledge. The next best candidate location p^* is selected as:

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

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
 180 for an efficient exploration of the environment [24]. 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. Exploiting A Priori Knowledge

In this paper, we consider different types of a priori knowledge PK_E , all of
 185 them related to the geometry and the shape of the building. More precisely, as anticipated in Section 1, we consider the floor plan FP_E , the footprint FO_E , and the bounding box BB_E . Moreover, we consider hand-drawn maps of E and maps of \bar{E} previously built by robots. These types of knowledge exhibit different degrees of completeness and accuracy and of difficulty in obtaining them. For
 190 the indoor environments we consider, prior knowledge about buildings can be easily obtained from documents, blueprints, evacuation maps [26], aerial images, external observations, human users, and records of earlier activities of robots. We now discuss in detail the types of prior knowledge we consider.

A *floor plan* FP_E is a two-dimensional representation of the environment E
 195 composed of line segments (walls) that identify the spaces within the environment, like rooms and corridors, and their connections, like doorways. Note that FP_E is usually accurate in representing the locations of external and internal walls, but is incomplete because it usually does not include information about




Figure 1: An example of different types of prior knowledge for the environment of Figure 1a. The floor plan (1b) provides more knowledge than the footprint (1c) and than the bounding box (1d). Moreover, a map built by a robot at an earlier time (1e) and a hand-drawn map (1f) could be available.

static furniture, which limits significantly the area of E that could be explored
 200 by a robot, and about small movable objects, which affects path planning and
 whose number, type, and location cannot be known in advance. An example of
 floor plan is reported in Figure 1b, while the corresponding (simulated) envi-
 ronment (with furniture) is in Figure 1a. Moreover, E can have obstacles, like
 closed doors, that are not present in FP_E or can exhibit connections between

205 locations that are not connected in FP_E (e.g., when E has been structurally modified and FP_E is outdated). In this sense, the topology of FP_E could misrepresent the topology of E .

A *footprint* FO_E is a more abstract form of representation of an environment that is composed only of the external walls, without considering the interior
210 divisions of the space into rooms and corridors. FO_E thus provides an accurate representation of external walls, but an incomplete representation of the interior of E .

The *bounding box* BB_E is the most abstract form of prior knowledge that we consider and is a box-like approximation that encloses the footprint, which
215 can be easily estimated even from outside a building. BB_E is thus usually incomplete and inaccurate in representing E . A visual comparison of the above types of prior knowledge for the environment of Figure 1a can be seen in Figures 1b-1d.

Moreover, in our experimental activity, we investigate other sources of prior
220 knowledge that could be available to the robot. In long-term deployments, a robot may build several maps of a changing environment at different times. Such *previous maps* represent prior knowledge that could be exploited by the robot at later times (Figure 1e). Another source of prior knowledge is represented by *maps hand-drawn* by human users (Figure 1f), which could be a valuable
225 and intuitive way of communication between humans and service robots [27]. Previous maps and hand-drawn maps could be complete, representing all the elements (e.g., furniture and walls) of the real environment, but inaccurate, because they could not exactly mirror the locations of these elements.

Note that, although prior knowledge PK_E is known, the map M_E for safe
230 navigation of a robot in E should be built and an exploration is still required (as also discussed in [22]).

In the proposed approach, PK_E should be manually fed to the system. The human effort needed for this step, which is done only once and requires few minutes, involves: getting prior knowledge represented as a grid, for example as
235 an image of an evacuation map, “cleaning” the image from unnecessary details

(e.g., words and symbols like emergency exits), scaling and aligning PK_E to the map M_E . We assume that PK_E is represented as a grid map, where a cell of PK_E is either free or obstacle. To ensure that PK_E and M_E are metrically consistent, namely that they are aligned and with the same scale, we first scale PK_E by assuming to know the width of the doors (which is regulated by the law in many countries). Then, alignment is performed by assuming that the initial pose of the robot is known in the coordinate systems of both PK_E and M_E . In our experiments, we perform the alignment manually, but an automatic method can be developed (as in [28]). In this way, we can find a correspondence between cells of the grid map M_E and cells of PK_E . Although no further adjustment is required in our experiments, methods like that of [29] can be used, if needed, to improve the alignment between PK_E and M_E .

In the proposed approach, PK_E is exploited to inform evaluation and selection of the next best candidate location using Equations (3) and (4), respectively.

In particular, we calculate $I(p)$, the estimate of the amount of the unexplored area visible from a candidate location p , by using PK_E .

The state-of-the-art approaches for estimating $I(p)$ measure the maximum visible area from p given the length of the frontier (as partially done, e.g., in [16]) or the perceivable area within the field of view and range of the robot’s laser range scanner (as done, e.g., in [6, 7], and shown in Figure 2b). Such 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 $p^{\text{PK}} \in \text{PK}_E$ (i.e., the cell with the closest center after scaling and alignment of M_E with PK_E). Then, for each unknown cell $c \in M_E$ that is within the perceivable area of the laser range scanner when the robot is in p , we find the corresponding $c^{\text{PK}} \in \text{PK}_E$. The expected area $I(p)$ visible from p is eventually computed by summing all cells c that satisfy all the following conditions:

- c^{PK} is free,




Figure 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), $I(p)$ (light blue area) computed according to the perceivable area of the robot’s laser range scanner (2b), and $I(p)$ calculated exploiting the knowledge of the floor plan FP_E (2e).

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

270 Figure 2e shows an example, in which the method just described is used to calculate $I(p)$ in the case the prior knowledge PK_E is the floor plan FP_E . It is interesting to contrast it with Figure 2b. 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 PK_E 275 is highly incomplete and inaccurate, as we show in the next section.

Using prior knowledge only on-line (namely, during the exploration process, as we do in this paper) appears not to fully exploit the availability of such knowledge, which could be used, for instance, to plan an exploration path off-line, before the exploration process starts. However, one should consider that 280 off-line planning of an exploration path can be performed only when the prior knowledge is sufficiently complete and accurate, like in the case of a floor plan. There is no way to plan a meaningful path given the footprint or the bounding box of an environment (see Figures 1c-1d). Moreover, also in the case a floor plan is available, differences between FP_E and E may be substantial, as the presence

285 of furniture and objects significantly reduces the area of the environment that
can be explored by the robot (see Section 4.6 for an example). Hence, also when
they can be planned, off-line exploration paths could become useless, because
they are built on incomplete knowledge, and their revision (e.g., replanning)
can be costly. For example, in the method of [8], replanning requires to solve a
290 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 [8]), but
that is not expected to efficiently scale to larger instances, as TSP is a NP-hard
problem. In this sense, we claim that our on-line approach copes well with
increasing incompleteness and inaccuracies in prior knowledge.

295 4. Experimental Evaluation

To evaluate our approach and its ability to exploit incomplete and inaccurate
a priori knowledge for efficient exploration, we present several tests conducted
both in simulation and with real robots. We measure, as exploration time
progresses, the *distance* travelled by the robot (as done, e.g., in [6, 16]) and the
300 percentage of *covered area*, namely the percentage of free area of E mapped in
 M_E , as done, e.g., in [7]. To have a fair comparison, we present results up to
95% of coverage, since some runs end without reaching full coverage, because
of noise in localization and mapping².

4.1. Simulations Using Different Types of Prior Knowledge

305 Simulations are performed with ROS Gazebo, using the ROS GMapping
and Nav2D packages³ for SLAM and robot navigation, respectively. We con-
sider three indoor environments with different characteristics (Figure 3): a *basic*

²The code of our method and the simulated environments (with the corresponding ex-
perimental settings) used in this section are available at; [https://github.com/goldleaf3i/
prior-maps-exploration](https://github.com/goldleaf3i/prior-maps-exploration)

³<http://wiki.ros.org/{gazebo,gmapping,nav2d}>. Gazebo is a 3D dynamic robot simu-
lator.




Figure 3: The floor plans of environments used in simulations.

environment ($19\text{ m} \times 10\text{ m}$) that represents a small apartment, an *office* environment ($90\text{ m} \times 53\text{ m}$) with several small rooms, and an *open* environment ($57\text{ m} \times 45\text{ m}$) with few large rooms. Two settings are considered for each environment, namely with and without furniture. The furniture is consistent with the type of the environment. For the basic environment, the furniture includes couches, chairs, tables, beds, nightstands, and so on, which cover 11 % of the total area. Furniture includes desks and chairs for the office environment (covering the 2 % of the total area) and shelves for the open environment (covering 3 % of the total area). Environments are static and differences between them and the floor plans FP_E are due to the presence of furniture. Differences between the environments and the corresponding footprints FO_E and bounding boxes BB_E are due to the fact that the inner space is not represented.

Exploration is performed by a simulated robot equipped with a laser range scanner with a field of view of 180° and a range of 10 m for the basic environment and of 25 m for the other two environments. These different settings are used to make explorations more challenging and simulate the equipment of robots used in real applications. Longer ranges for the laser range scanner generally result in less exploration steps and, consequently, in less decisions. Small-size robots with limited payload and short-range sensors are usually employed in small house-like environments (as the basic environment). Middle-sized robots with heavier payload and long-range sensors are instead employed in office-like and warehouse-like large-scale environments (as the office and open environments). We use robots with similar features in experiments performed with real robots

(Section 4.6). Linear and angular speeds of the robot are fixed, but different from each other, with the robot moving slower when rotating. Thus the robot follows a path with few turns faster than a path with the same length but with more turns. This motivates the decision of reporting both distance \mathcal{D} and time
335 \mathcal{T} as performance measures.

After several preliminary tests, we set $\alpha = 0.5$ in (1) to equally balance distance and information gain and guarantee a good overall performance in exploration. A value of $\alpha \approx 1$ causes the robot to select the closest frontier, without considering the information gain and, ultimately, slowing down the exploration
340 process by choosing frontiers with a limited information gain. A value $\alpha \approx 0$ causes the robot to select the most promising frontier in terms of information gain even if it is far from the current robot position. This causes an exploration pattern where the robot could travel to the extremities of the environment to reach frontiers with high $i(p)$, ultimately slowing down the exploration process
345 and increasing the travel cost. A balance between the two components of (1), obtained with values of α around 0.5, results in more stable and more efficient performance across different environments.

For each environment, we perform 10 exploration runs (starting from the same pose) and average results over the runs. Gaussian noise with zero mean
350 and 0.1 m standard deviation is added to laser scans. 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, optimistically, as in [6, 7]
355 (see Figure 2b). Note that the computing effort of our approach is negligible (selection of frontiers takes a time in the order of milliseconds). In this, it is fully comparable with the exploration strategies that compute $I(p)$ only on the basis of the map M_E .

360 Tables 1 and 2 show the results obtained in the basic environment. The use of the floor plan FP_E allows to explore the environment travelling a significantly

shorter distance and employing a significantly shorter time, especially when the covered area is less than 80%. For example, the difference at 80% for the unfurnished basic environment has p -value $\ll 0.0001$ in one-way ANOVA and
365 the difference at 50% and 60% for the furnished basic environment has a p -value $\ll 0.0001$. The speed-up is substantial even considering the small area of the environment, which results in a fast exploration and in a limited number of frontiers to be reached. This is because frontiers that are close to walls are correctly evaluated to have small $I(p)$ using our method, while the same
370 frontiers can have large $I(p)$ when no a priori knowledge is used. The gain reduces towards the end of the exploration, as the impact of a wrong decision is less critical (see Figure 4).

Similar results are found both for unfurnished and furnished basic environments. In the second case, the trends of the two strategies become very
375 different after reaching a coverage of 70%. 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 70% and 80%, with p -value $\ll 0.0001$ and p -value = 0.0024, respectively for a one-way ANOVA test.
380 Note that the prior knowledge of FP_E is incomplete in the case of the furnished environment, while it is complete in the case of the unfurnished one.

Also, the other two types of prior knowledge result in a substantial gain in term of exploration time, for both the furnished and unfurnished versions. The floor plan FP_E provides an initially higher speed-up by making better
385 frontier choices during the first stages of exploration, also resulting in a lower standard deviation. When more portions of the environment are explored, with a coverage above 80%, however, the use of more incomplete and inaccurate sources of knowledge like FO_E and BB_E provides a faster exploration, because one of the side effects of making better frontier selections at the beginning of
390 exploration (as done by FP_E) is to leave several small unexplored areas scattered around the environment. Overall, the gain in terms of exploration time wrt the strategy that does not exploit any prior knowledge is substantial for all the types

coverage	without prior knowledge		floor plan			footprint			bounding box		
with furniture	\mathcal{D}	σ	\mathcal{D}	σ	difference	\mathcal{D}	σ	difference	\mathcal{D}	σ	difference
70%	22.97	3.46	12.27	1.84	-47%	16.31	3.57	-29%	19.06	4.53	-17%
80%	33.23	6.9	16.15	1.33	-51%	23.43	3.55	-29%	25.16	3	-24%
90%	38.86	6.15	32.29	2.96	-17%	39.43	6.89	+1%	35.68	4.38	-8%
95%	50.77	6.62	47.82	9.44	-6%	52.39	5.08	+3%	52.37	4.56	+3%
without furniture	\mathcal{D}	σ	\mathcal{D}	σ	difference	\mathcal{D}	σ	difference	\mathcal{D}	σ	difference
70%	7.12	1.3	7.32	1.37	+3%	11.3	1.04	+59%	11.91	1.33	+67%
80%	29.39	4.73	15.7	4.92	-47%	22.29	1.5	-24%	24.21	1.55	-18%
90%	32.55	4.43	24.17	4.11	-26%	24.69	1.79	-24%	26.18	1.64	-20%
95%	35.55	3.64	27.7	2.38	-22%	28.53	1.83	-20%	29.78	1.54	-16%

Table 1: Performance according to distance using different types of prior knowledge for the basic environment. \mathcal{D} is the distance in m and σ is the corresponding standard deviation. The column difference shows the percentage difference in performance of the strategy with prior knowledge over that without prior knowledge: negative numbers mean that the former performs better than the latter. Bold entries are the best for each row.

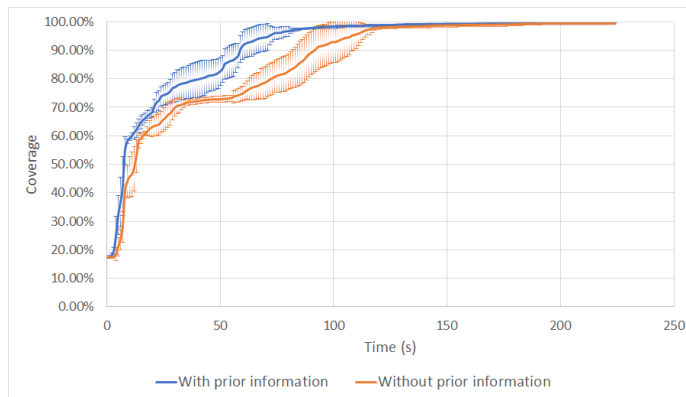


Figure 4: Coverage vs. time for the unfurnished basic environment.

of prior knowledge we consider.

Similar results can be seen in Tables 3 and 4 for the office environment. Remarkably, the use of FP_E results into a speed-up for covering 70 % of the free area of 25 % wrt the use of no prior knowledge for the furnished version and of almost 40 % for the unfurnished one. It is interesting to see how the exploration time is similar for the furnished and unfurnished versions of the environment when no prior knowledge is considered, while the use of prior knowledge provides

coverage	without prior knowledge		floor plan			footprint			bounding box		
with furniture	\mathcal{T}	σ	\mathcal{T}	σ	difference	\mathcal{T}	σ	difference	\mathcal{T}	σ	difference
70%	74.87	16.39	34.01	5.5	-55%	36.37	9.25	-51%	40.04	8.57	-47%
80%	96.67	21.49	44.3	4.65	-54%	48.89	9.33	-49%	50.44	6.69	-48%
90%	115.29	28.05	76.33	9.56	-33%	80.75	13.38	-30%	71.19	6.42	-38%
95%	143.99	25.71	116.84	24.92	-18%	105.51	14.22	-27%	101.42	11.23	-30%
without furniture	\mathcal{T}	σ	\mathcal{T}	σ	difference	\mathcal{T}	σ	difference	\mathcal{T}	σ	difference
70%	28.42	3.63	21.97	2.31	-10%	25.28	3.08	+4%	24.78	1.91	+1%
80%	81.47	14.28	41.24	13.3	-49%	42.3	3.88	-48%	44.06	2.19	-46%
90%	90.72	12.99	59.05	8.34	-35%	45.4	3.95	-50%	47.63	2.32	-48%
95%	98.6	11.17	67.1	6.75	-32%	55.68	4.75	-44%	55.4	3.05	-44%

Table 2: Performance according to time using different types of prior knowledge for the basic environment. \mathcal{T} is time in s. See Table 1 for the rest of notation.

coverage	without prior knowledge		floor plan			footprint			bounding box		
with furniture	\mathcal{D}	σ	\mathcal{D}	σ	difference	\mathcal{D}	σ	difference	\mathcal{D}	σ	difference
70%	166.66	37.42	133.91	35.9	-20%	141.8	11.72	-15%	178.63	8.67	+7%
80%	291.25	12.8	217.03	27.42	-25%	238.03	12.74	-18%	208.93	13.56	-28%
90%	317.81	14.1	321.08	19.3	+1%	273.75	22.4	-14%	291.82	18.45	-8%
95%	431.84	55.11	404.77	19.75	-6%	360.31	33.42	-17%	433.67	38.94	0%
without furniture	\mathcal{D}	σ	\mathcal{D}	σ	difference	\mathcal{D}	σ	difference	\mathcal{D}	σ	difference
70%	164.83	31.28	111.51	4.42	-32%	146.23	12.76	-11%	152.56	14.56	-7%
80%	236.13	22.88	190.91	24.26	-19%	241.19	13.3	+2%	196.53	6.5	-17%
90%	308.64	35.39	266.04	15.76	-14%	282.45	22.87	-8%	270.8	15.58	-12%
95%	390.43	42.88	359.39	18.57	-8%	366.38	19.31	-6%	376.21	17.55	-4%


Table 3: Performance according to distance using different types of prior knowledge for the office environment.

400 different results. Interestingly, the use of BB_E as prior knowledge results into a relatively limited degradation of the performance wrt other forms of prior knowledge, despite the fact that BB_E greatly overestimates the true area of the L-shaped environment, as shown in Figure 5.

Tables 5 and 6 finally show the results obtained for the open environment.
405 As that environment is a large-scale, box-like environment, its bounding box and footprint are very similar and, consequently, results are almost the same. We thus report only those obtained with FO_E . In this environment, despite the fact that the use of incomplete prior knowledge provided by FO_E still results into a substantial speed-up in terms of exploration time, the use of the more complete

coverage	without prior knowledge		floor plan			footprint			bounding box		
with furniture	\mathcal{T}	σ	\mathcal{T}	σ	difference	\mathcal{T}	σ	difference	\mathcal{T}	σ	difference
70%	245.57	58.6	184.06	50.95	-25%	218.14	17.9	-11%	247.09	14.3	+1%
80%	407.94	23.15	300.84	34.95	-26%	338.97	16.25	-17%	288.46	21.15	-29%
90%	443.24	24.3	436.77	24.7	-2%	394.1	30.05	-11%	407.56	29.28	-8%
95%	603.11	76.05	552.75	25.78	-8%	523.12	51.75	-13%	625.47	62.71	-3%
without furniture	\mathcal{T}	σ	\mathcal{T}	σ	difference	\mathcal{T}	σ	difference	\mathcal{T}	σ	difference
70%	247.66	47.77	150.11	6.66	-39%	219.07	21.03	-12%	216.33	18.02	-12%
80%	345.11	24.95	261.48	29.73	-24%	339.86	24.4	-2%	269.41	13.93	-22%
90%	440.39	54.17	359.19	18.2	-18%	399.99	39.87	-9%	373.21	25.08	-15%
95%	546.46	51	483.09	23.45	-11%	520.04	35.65	-5%	529.21	25.21	-3%

Table 4: Performance according to time using different types of prior knowledge for the office environment.



(a)

Figure 5: The floor plan of the environment of Figure 3b overlapped with the bounding box BB_E (in orange).

410 knowledge provided by FP_E obtains better results, as it allows the robot to make better choices when exploring the two corridors in the bottom-left part of the building.

Using a priori knowledge improves considerably the performance in all the three cases, allowing the robot to cover large portions of the environments travelling a shorter distance and spending less time. Note that, as the robot moves 415 faster when travelling straight than when rotating, differences in average speeds for completing explorations are due to the different exploration paths. Using prior knowledge results in more direct and straight paths, which are faster for the robot to follow.

coverage	without prior knowledge		floor plan			footprint		
with furniture	\mathcal{D}	σ	\mathcal{D}	σ	difference	\mathcal{D}	σ	difference
70%	78.96	12.84	67.51	9.79	-15%	71.34	6.62	-10%
80%	117.23	13.37	111.87	6.95	-5%	131.42	14.94	+12%
90%	237.26	23.03	185.23	23.26	-22%	175.61	8.62	-26%
95%	305.83	28.43	223.56	18.54	-27%	223.9	24.5	-27%
without furniture	\mathcal{D}	σ	\mathcal{D}	σ	difference	\mathcal{D}	σ	difference
70%	71.15	14.42	43.98	3.51	-38%	47.2	11.2	-34%
80%	83.53	13.75	73.95	5.3	-11%	65.39	10.08	-22%
90%	140.3	23.3	127.2	11.68	-9%	138.85	13.33	-1%
95%	186.87	30.14	161.4	17.52	-14%	17356	17.98	-8%

Table 5: Performance according to distance using different types of prior knowledge for the open environment.

coverage	without prior knowledge		floor plan			footprint		
with furniture	\mathcal{T}	σ	\mathcal{T}	σ	difference	\mathcal{T}	σ	difference
70%	109.86	17.89	92.91	15.54	-15%	103.88	12.02	-5%
80%	159.02	16.37	149.95	14.28	-6%	186.55	21.21	+17%
90%	315.91	31.44	245.85	42.11	-22%	247.91	12.23	-22%
95%	407.35	42.44	293.88	36.2	-28%	308.63	30.9	-24%
without furniture	\mathcal{T}	σ	\mathcal{T}	σ	difference	\mathcal{T}	σ	difference
70%	101.4	47.77	66.07	4.98	-35%	80.9	21.03	-20%
80%	116.29	24.95	104.08	10.64	-10%	104.51	24.4	-10%
90%	188.8	54.17	176.7	16.48	-6%	202.68	39.87	+7%
95%	244.82	51	218.12	24.28	-11%	247.98	35.65	+1%

Table 6: Performance according to time using different types of prior knowledge for the open environment.

coverage	without prior knowledge		floor plan			previous map		
M_{t_1}	\mathcal{D}	σ	\mathcal{D}	σ	difference	\mathcal{D}	σ	difference
70%	19.25	2.77	11.67	1.72	-39%	11.09	1.14	-42%
80%	29.9	11.82	19.73	4.15	-34%	23.52	4.06	-21%
90%	47.4	8.19	47.73	3.94	+1%	42.34	2.6	-11%
95%	74.33	11.74	76.44	9.1	+3%	75.08	7.4	+1%
M_{t_2}	\mathcal{D}	σ	\mathcal{D}	σ	difference	\mathcal{D}	σ	difference
70%	24.06	3.38	17.17	2.95	-29%	14.72	1.11	-39%
80%	33.74	7.2	25.01	2.64	-26%	24.8	3.59	-26%
90%	58.25	19.2	39.13	4.52	-33%	45.9	5.7	-21%
95%	67.35	15.35	57.91	6.41	-14%	53.69	6.69	-20%

Table 7: Performance according to distance using previous maps as prior knowledge while performing repeated exploration runs in a changing environment.

coverage	without prior knowledge		floor plan			previous map		
M_{t_1}	\mathcal{T}	σ	\mathcal{T}	σ	difference	\mathcal{T}	σ	difference
70%	43.78	6.44	24.16	15.54	-44%	21.71	3.97	-50%
80%	63.42	21.69	44.18	14.28	-30%	47.62	6.83	-25%
90%	101.57	14.72	98.6	42.11	-3%	81.73	6.56	-20%
95%	160.27	27.68	159.17	36.2	-1%	149.1	19.08	-7%
M_{t_2}	\mathcal{T}	σ	\mathcal{T}	σ	difference	\mathcal{T}	σ	difference
70%	49.19	7.71	39.23	9.18	-35%	31.76	5.08	-35%
80%	66.83	14.61	51.31	5.85	-10%	53.57	8.78	-20%
90%	119.7	44.53	78.44	8.92	-6%	92.23	16.08	-23%
95%	139.33	34.17	115.21	11.7	-11%	110.17	16.68	-21%

Table 8: Performance according to time using previous maps as prior knowledge while performing repeated exploration runs in a changing environment.

420 *4.2. Simulations with Previous Maps as Prior Knowledge*

In this section we evaluate the possible application of our method to a scenario where, from time to time, a robot has to explore an indoor environment, which changes over time. In such a setting, the map of the environment obtained in the previous exploration may become outdated [30], for example due
425 to changes in the furniture, and consequently it needs to be updated by per-



Figure 6: Three versions of the basic environment simulating the evolution of the environment over time as it can happen in a long-term robot deployment.

forming a new exploration task. For this experiment, we consider the basic environment of Figure 3a in its furnished version and the same simulation and experimental settings of Section 4.1.

We consider three different versions of the same environment, E_{t_0} , E_{t_1} , and
 430 E_{t_2} , where the furniture is arranged in different ways (Figure 6). Arranging the furniture differently causes significant changes in the map acquired by the robot, while the structure of the walls remains static.

Results are reported in Tables 7 and 8 and show that using previous maps as prior knowledge (i.e., using M_{t_0} when exploring for building M_{t_1} and M_{t_1}
 435 when exploring for building M_{t_2}) provides an advantage over using the floor plan. Note that, in this case, previous maps are a more complete type of prior knowledge than the floor plan, because they include the presence of furniture, although not necessarily in the right locations. We stress that our method can use either form of prior knowledge, providing a significant increase in exploration
 440 performance.

4.3. Simulations with Topologically Inaccurate Floor Plans as Prior Knowledge

In this section, we evaluate the effects of considering prior knowledge that is topologically inaccurate wrt the actual environment. We consider a typical scenario that may be encountered during the daily use of a service robot. In
 445 indoor environments, it is often the case that not all areas are reachable by the robot, for example due to closed doors or blocked doorways. Since floor plans represent all doors as open, they can incorrectly represent the topology of the

coverage	without prior knowledge		floor plan		
	\mathcal{T}	σ	\mathcal{T}	σ	difference
70%	157	12.55	121.68	2.54	-22%
80%	212.96	25.9	152.31	8.41	-28%
90%	315.44	16.69	285.35	16.35	-10%
95%	320.05	15.98	301.01	17.43	-6%

Table 9: Performance according to time in an environment in which some areas are not accessible.




Figure 7: The office environment of Figure 3b where some areas are unreachable due to closed doors (red crosses).

actual environment. As a consequence, the robot can make decisions that, for example, select paths along corridors ending with closed doors.

450 We consider the office environment of Figure 3b, where 12 of the 32 doors are closed. The setting is displayed in Figure 7, where closed doors are marked by red crosses. Some rooms are not accessible by the robot, while in some other rooms that have multiple doors, the choices of the robot are constrained as some paths are blocked by closed doors. Overall, 25 % of the free area of the
455 environment cannot be reached by the robot.

Table 9 shows that, although the floor plan greatly overestimates the amount of the area that the robot can explore, the use of such prior knowledge still

results in a significant speed-up wrt not using prior knowledge. This confirms the fact that the use of prior knowledge provides benefits even if it inaccurately
460 represents the topology of the actual environment.

4.4. Simulations with Hand-Drawn Floor Plans as Prior Knowledge

In this section, we test our approach using highly inaccurate prior knowledge in order to evaluate its robustness. We use three hand-drawn floor plans as PK_E and we compare, as before, the performance of our exploration strategy with
465 that of the exploration strategy that does not use any prior knowledge.

The prior knowledge consists of floor plans of a building that are hand-drawn by three different people who work in the same building. The environment has approximately a size of $50\text{ m} \times 43\text{ m}$. The correct floor plan and the three digitalized hand-drawn floor plans are reported in Figure 8, along with one
470 of the original drawings. People are asked to draw on paper the floor plans using only their memory, without any support for recollection. Moreover, they are given a pre-defined area in which to draw, in order to preserve scale (the drawing area is represented by a dashed line in Figure 8e). The original drawings made on paper are digitalized and scaled. Digitalization is performed manually
475 by straightening walls to vertical and horizontal lines following the Manhattan assumption and by removing canceled or blackened parts, letters, words, and inaccessible areas (e.g., the area at the bottom right of Figure 8e). (A method similar to that in [31] could be used for automating digitalization.) In the three hand-drawn floor plans, the number and size of rooms are different and several
480 small rooms are missing or are merged to adjacent bigger rooms. In practice, although they 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. (See [32] for a discussion on similar issues arising when matching sketch maps to ground truth.)

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




Figure 8: The correct floor plan (8a), three hand-drawn floor plans (8b-8d), and the original draw of 8b (8e).

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%	118.6	7.16	-8%	143.29	15.5	+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%	236.6	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 10: Performance evaluation using the hand-drawn floor plans of Figure 8.

In Table 10 we see how the completeness and accuracy of a priori knowledge affect 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 incomplete and 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%.

4.5. Comparison with the Method of [8]

We compare our on-line exploration strategy and the method of [8], which uses another kind of prior knowledge, a topo-metric map, to build an exploration path off-line that is expanded on-line, as discussed in Section 2. Recall also the discussion of Section 3.2 about the applicability of off-line approaches, as that of [8], only when prior knowledge is sufficiently complete and accurate, like a floor plan. We consider two environments, Map 2 and Map 4, reported in Figures 9a and 9b, respectively, from [8] and we use the same range of the robot's laser range scanner (6 m) and the same initial positions (10 per environment) used in [8]. The first environment has approximately a size of $34\text{ m} \times 34\text{ m}$, while the second one has a size of $38\text{ m} \times 25\text{ m}$. We perform one exploration run for each initial position. Results are then averaged over the runs.




Figure 9: Two environments from [8] (9a and 9b) and a map covering the 90% of the first environment (9c).

Authors of [8] report that, when the topo-metric graph of the environment is reliable, the best 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 two environments travelling a longer distance, about 433 m and 429 m, respectively. This is expected because the method of [8] 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 the area of the two environments (Figure 9c), 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
525 it needs to travel back to cover small portions that have been left behind (as we
have already observed). Moreover, results from [8] show that, when Gaussian
noise with standard deviation larger than 2.5 m is added to the node positions of
the topo-metric map, making it inaccurate, their method performance is on par
with that of a method that does not consider prior knowledge. In the following,
530 we show with real robot experiments that our method performs consistently
better than an exploration strategy that does not consider prior knowledge.

4.6. Experiments with Real Robots

In this section we describe the results of the experiments performed in dif-
ferent environments with the implementation of our approach on real robots,
535 running the same ROS configuration used for simulations. We consider floor
plans as prior 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. The runs are performed in
540 a portion of the environment of Figure 8a, with a size of $36\text{ m} \times 27\text{ m}$ and
shown in Figure 10b, performing 3 exploration runs from the same initial pose.
Results are averaged over the runs. Note that discrepancies between the actual
environment and the floor plan can change due to the changes of furniture in
different runs.

545 The results, reported in Table 11, confirm that our exploration strategy
outperforms the exploration strategy without a priori knowledge by providing
a better information gain estimate. This is more evident when clutter and
occlusions increase, e.g., when the number of obstacles (like furniture) increases.
Because of this, the application of our method to real-world settings seems to
550 provide even more advantages than in the more controlled environments used
in simulations. It is interesting to look at the average speed along the different
exploration paths (linear and angular speeds are fixed, as in simulation). Using
prior knowledge, the robot selects frontiers that, when visited, make 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	281.2	25.8	-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%	47.1	7.8	528.31	96.75	41.77	2.96	411.11	14.6	-11%	-22%

Table 11: Results (over 3 runs) of the experiments with the Robocom robot.

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	4.54	0.58	24	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	33.7	2.57	-21%	-40%	4.8	0.43	25.52	1.82	-37%	-55%
90%	10.56	3.58	75.6	26.64	9.98	1.21	58.99	6.74	-6%	-22%	9.04	1.93	53.11	13	-15%	-30%
95%	24.68	5.04	171.05	31.15	17.7	6.68	116.87	26.9	-28%	-32%	-	-	-	-	-	-

Table 12: Results of experiments with TurtleBot3 Burger using plain and modified floor plans.

555 travel along paths that are more straight than those travelled in the case no prior knowledge is used. As a consequence, the robot completes the exploration at an average speed of about 10 cm/s in former case and at an average speed of less than 9 cm/s in the latter case.




Figure 10: Robocom (10a) used in experiments in the environment with floor plan in 10b. In red, the initial position of the robot.

The second set of experiments is performed using a TurtleBot3 Burger (Figure 11a), a small-size two-wheeled differential drive robot, equipped with a Hokuyo URG-04LX-UG laser range scanner (field of view of 180° and range of 560

5.6 m). The runs are performed in a house environment with a size of $6\text{ m} \times 6\text{ m}$ with an area of 30 m^2 (Figure 11b), performing 3 exploration runs from the same initial pose 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 a kitchen, and cannot be accessed by the robot. Other 2 m^2 are occupied by more dynamic furniture like chairs. Overall, 30% of the area is covered by (static) furniture and about 5% is covered by clutter, resulting in the floor plan FP_E overestimating the amount of explorable area by 35%. This difference between the floor plan FP_E and the actual environment E where the robot operates allows us to test the robustness of our approach, and in particular to investigate the effects of the completeness and accuracy of the prior knowledge. 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 FP_E (Figure 11b), and an exploration strategy that uses FP_E modified in order to include the static furniture (Figure 11c), which is a more complete representation than the plain FP_E , but that still does not include small objects like chairs.

Results are shown in Table 12 and confirm that our exploration strategy is more efficient than that without a priori knowledge. Despite the fact that the




Figure 11: The TurtleBot3 Burger (11a) used in experiments in the environment with plain floor plan shown in (11b). In red, the initial position of the robot. (11c) shows the modified floor plan including furniture, in gray, while (11d) shows the map created by the robot in one of the runs.

580 plain FP_E is a rough estimate of the actual environment, the gain in terms of
time and distance is substantial. A larger 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
585 some portions of the environment (when the covered area is more than 75%
in Table 12). 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 position. After this initial phase, the use of prior knowledge drives
the robot directly to the most interesting frontiers, thus reducing the distance
590 and the time required to map the environment.

Interestingly, the use of the more complete floor plan (modified FP_E , which
includes static furniture) stops the exploration process at about 92% of the ex-
plorable area, since no further frontier is detected (see ‘-’ in Table 12). This is
because the robot, when not considering any prior knowledge or when consid-
595 ering the plain FP_E , can map thoroughly also the small gaps between furniture
and walls. For example, in our runs without prior knowledge, the robot tries
to reach a location to observe a small (approximately 10 cm) gap between a
sofa and a wall, although it cannot enter in it. When the modified floor plan
is used, the robot “knows” that the gap area is uninteresting and selects more
600 interesting frontiers.

Overall, experiments performed with real robots suggest that the use of a
priori knowledge can be particularly useful in human-inhabited settings where
objects, furniture, people, and obstacles can negatively affect the perception of
the robot. In these settings, the use of prior knowledge, even if it is far from
605 faithfully representing the environment (as in the case of the plain FP_E), pro-
vides an effective mean to drive the robot to select the next best locations for
exploration. The performance improvement of our proposed exploration strat-
egy 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.

610 4.7. Discussion

In our experiments, we have investigated the impact of several types of a priori knowledge on the exploration of indoor environments performed by a mobile robot. We have considered two aspects that characterize prior knowledge: completeness, namely the extent to which prior knowledge represents all the
615 elements of the real environment, and accuracy, namely the fidelity with which it represents properties of the elements of a real environment. For instance, floor plans are accurate in representing the locations of walls, but they are incomplete because they do not represent furniture. Conversely, hand-drawn maps could be complete in representing furniture, but they are often largely
620 inaccurate in representing the locations of walls. Experimental results suggest that using accurate but incomplete knowledge, like those provided by footprints, provides more advantages than using less accurate but more complete knowledge represented by hand-drawn maps. For example, consider the environment of Figure 8a. Using the footprint results in an improvement of 5–10% for coverage
625 of 70 to 80% (full data are not reported here), while the hand-drawn maps have a worse performance (see Table 10).

In order to push forward such considerations, we need a more detailed characterization of prior knowledge. This is beyond the scope of this paper, but we provide some hints that could be useful towards the definition of a general
630 framework. Completeness refers to the *information content* of prior knowledge and could be specialized by referring it to single elements of the environment, from those general (walls and rooms) to those particular (furniture, objects). Accuracy, instead, refers to the *quality* of the prior knowledge and could be specialized by referring to specific properties of elements, like locations and extents
635 of walls and size and shape of rooms and of furniture. In this way, we could describe prior knowledge as complete in representing the rooms of an environment but incomplete in representing its furniture, and as accurate in representing the length of corridors but inaccurate in representing the shape of the rooms. The accuracy could also be expressed by metric bounds. In this way, we could quantify
640 how much the lengths of the walls represented in an architectural drawing

are more accurate than those represented in a previous map built by a SLAM system.

Topological differences between prior knowledge and real environments could have a large impact on exploration and should be represented properly. For indoor environments, differences in topology are mainly (but not exclusively) related to doors, which can give rise to a number of combinations, like a door can be open in prior knowledge but closed in real environment (as in Section 4.3) or vice versa.

Finally, the sources and the availability of prior knowledge for different applications should be analyzed. For example, previous maps acquired at some earlier time (Section 4.2) could be a source of prior knowledge for service robots operating in houses, but may be not available for search and rescue applications for which the availability of floor plans and footprints is more likely, since they can be easily acquired from evacuation maps and aerial images.

5. Conclusions

In this paper, we introduced an on-line exploration strategy that exploits a priori knowledge, in the form of floor plans, footprints, bounding boxes, hand-drawn maps, and previous maps to select the next best locations for a robot exploring indoor environments. Experiments assessed the effectiveness of the proposed method, also when the prior knowledge is incomplete and inaccurate. Our results show that, while it is intuitive that complete and accurate prior knowledge can improve the performance of the exploration process, also incomplete but accurate prior knowledge provides relevant advantages and even incomplete and inaccurate prior knowledge can provide some benefits, which is far less intuitive. In this sense, our approach could be beneficial to several real-world application domains, where it is usually easy to obtain incomplete and inaccurate prior knowledge about the environments in which robots operate.

The question of how to best balance the effort to get complete and accurate prior knowledge and the improvements on the exploration performance is still

670 open and is a direction for future work. We will also investigate the use of other
forms of a priori knowledge, like pictures of evacuation maps (as in [22]), that can
be easily obtained in large buildings, and topological maps obtained from hand-
drawn maps that, although metrically inconsistent, could nevertheless provide
a good estimate of the structure of the environment. Better characterization of
675 prior knowledge available in different application contexts will be investigated,
as discussed in Section 4.7. For example, a priori knowledge could be used to
inform the adaptive tuning of exploration strategies (e.g., the parameter α in
our method) as the exploration progresses. Moreover, means to represent, and
include in the evaluation of the information gain, the uncertainty of the prior
680 knowledge 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 prior knowledge and of combination of different types
of prior knowledge could be studied, using an approach similar to that of [33].

Acknowledgements

685 We are glad to thank Stefan Oskwald and authors of [8] for kindly providing
data sets and details of their method. We also thank Danilo Fusi for his help in
developing the previous version of our proposed method.

We gratefully acknowledge the European Commission H2020 project Move-
Care grant ICT-26-2016b – GA 732158 for financial partial support to some of
690 the authors of this work.

References

- [1] M. Julia, A. Gil, O. Reinoso, A comparison of path planning strategies for
autonomous exploration and mapping of unknown environments, *Auton
Robot* 33 (4) (2012) 427–444.
- 695 [2] E. Galceran, M. Carreras, A survey on coverage path planning for robotics,
Robot Auton Syst 61 (12) (2013) 1258 – 1276.

- [3] P. Neumann, S. Asadi, J. Schiller, A. Lilienthal, M. Bartholmai, An artificial potential field based sampling strategy for a gas-sensitive micro-drone, in: Proc. IROS Workshop on Robotics for Environmental Monitoring, 2011, pp. 34–38.
- [4] W. Tabib, R. Whittaker, N. Michael, Efficient multi-sensor exploration using dependent observations and conditional mutual information, in: Proc. SSRR, 2016, pp. 42–47.
- [5] S. Thrun, Robotic mapping: A survey, in: G. Lakemeyer, B. Nebel (Eds.), Exploring Artificial Intelligence in the New Millenium, Morgan Kaufmann, 2003, pp. 1–35.
- [6] H. González-Baños, J. Latombe, Navigation strategies for exploring indoor environments, *Int J Robot Res* 21 (10-11) (2002) 829–848.
- [7] N. Basilico, F. Amigoni, Exploration strategies based on multi-criteria decision making for searching environments in rescue operations, *Auton Robot* 31 (4) (2011) 401–417.
- [8] S. Ośwald, M. Bennewitz, W. Burgard, C. Stachniss, Speeding-up robot exploration by exploiting background information, *IEEE RA-L* 1 (2) (2016) 716–723.
- [9] F. Amigoni, D. Fusi, M. Luperto, A. Borghese, Exploiting inaccurate a priori knowledge in robot exploration (extended abstract), in: Proc. AAMAS, 2019, pp. 2102–2104.
- [10] F. Amigoni, D. Fusi, M. Luperto, A. Borghese, Robot exploration using knowledge of inaccurate floor plans, in: Proc. ECMR, 2019.
- [11] B. Yamauchi, A frontier-based approach for autonomous exploration, in: Proc. CIRA, 1997, pp. 146–151.
- [12] C. Stachniss, G. Grisetti, W. Burgard, Information gain-based exploration using Rao-Blackwellized particle filters, in: Proc. RSS, 2005, pp. 65–72.

- [13] D. Thompson, D. Wettergreen, G. Foil, M. Furlong, A. Kiran, Spatio-
725 spectral exploration combining in situ and remote measurements, in: Proc.
AAAI, 2015, pp. 3679–3685.
- [14] Y. Girdhar, D. Whitney, G. Dudek, Curiosity based exploration for learning
terrain models, in: Proc. ICRA, 2014, pp. 578–584.
- [15] C. Tovey, S. Koenig, Improved analysis of greedy mapping, in: Proc. IROS,
730 2003, pp. 3251–3257.
- [16] B. Tovar, L. Munoz, R. Murrieta-Cid, M. Alencastre, R. Monroy,
S. Hutchinson, Planning exploration strategies for simultaneous localiza-
tion and mapping, *Robot Auton Syst* 54 (4) (2006) 314–331.
- [17] F. Amigoni, A. Gallo, A multi-objective exploration strategy for mobile
735 robots, in: Proc. ICRA, 2005, pp. 3850–3855.
- [18] J. Chang, G. Lee, Y.-H. Lu, C. Hu, P-SLAM: Simultaneous localization and
mapping with environmental-structure prediction, *IEEE T Robot* 23 (2)
(2007) 281–293.
- [19] D. Ström, I. Bogoslavskyi, C. Stachniss, Robust exploration and homing
740 for autonomous robots, *Robot Auton Syst* 90 (2017) 125–135.
- [20] A. Smith, G. Hollinger, Distributed inference-based multi-robot explo-
ration, *Auton Robot* 42 (8) (2018) 1651–1668.
- [21] R. Shrestha, F. Tian, W. Feng, P. Tan, R. Vaughan, Learned map pre-
745 diction for enhanced mobile robot exploration, in: Proc. ICRA, 2019, pp.
1197–1204.
- [22] M. Mielle, M. Magnusson, A. Lilienthal, The auto-complete graph: Merging
and mutual correction of sensor and prior maps for SLAM, *Robotics* 8 (2)
(2019) 40.

- [23] M. Mielle, M. Magnusson, H. Andreasson, A. Lilienthal, SLAM auto-
750 complete: Completing a robot map using an emergency map, in: Proc.
SSRR, 2017, pp. 35–40.
- [24] F. Amigoni, Experimental evaluation of some exploration strategies for
mobile robots, in: Proc. ICRA, 2008, pp. 2818–2823.
- [25] G. Grisetti, C. Stachniss, W. Burgard, Improved techniques for grid map-
755 ping with Rao-Blackwellized particle filters, IEEE T Robot 23 (2007) 34–46.
- [26] F. Boniardi, T. Caselitz, R. Kümmerle, W. Burgard, Robust lidar-based
localization in architectural floor plans., Adv Robotics 26 (17) (2012) 2021–
2041.
- [27] M. Blades, The reliability of data collected from sketch maps, J Environ
760 Psychol 10 (4) (1990) 327 – 339.
- [28] F. Boniardi, B. Behzadian, W. Burgard, G. Tipaldi, Robot navigation in
hand-drawn sketched maps, in: Proc. ECMR, 2015, pp. 1–6.
- [29] S. Shahbandi, M. Magnusson, 2D map alignment with region decomposi-
tion, Auton Robot 43 (5) (2019) 1117–1136.
- 765 [30] L. Kunze, N. Hawes, T. Duckett, M. Hanheide, T. Krajník, Artificial In-
telligence for Long-Term Robot Autonomy: A Survey, IEEE RA-L 3 (4)
(2018) 4023–4030.
- [31] M. Luperto, F. Amigoni, Extracting structure of buildings using layout
reconstruction, in: Proc. IAS, 2018, pp. 652–667.
- 770 [32] M. Mielle, M. Magnusson, A. Lilienthal, Using sketch-maps for robot nav-
igation: Interpretation and matching, in: Proc. SSRR, 2016, pp. 252–257.
- [33] M. Luperto, F. Amigoni, Predicting the layout of partially observed rooms
from grid maps, in: Proc. ICRA, 2019, pp. 6898–6904.