Mapping Beyond What You Can See: Predicting the Layout of Rooms Behind Closed Doors

Matteo Luperto[a], Federico Amadelli[b], Moreno Di Berardino[b], Francesco Amigoni[b]

[a]Dipartimento di Informatica, Università degli Studi di Milano
[b]Dipartimento di Elettronica, Informatica e Bioingegneria, Politecnico di Milano

Abstract

The availability of maps of indoor environments is often fundamental for autonomous mobile robots to efficiently operate in industrial, office, and domestic applications. When robots build such maps, some areas of interest could be inaccessible, for instance, due to closed doors. As a consequence, these areas are not represented in the maps, possibly causing limitations in robot localization and navigation. In this paper, we provide a method that completes 2D grid maps by adding the predicted layout of the rooms behind closed doors. The main idea of our approach is to exploit the underlying geometrical structure of indoor environments to estimate the shape of unobserved rooms. Results show that our method is accurate in completing maps also when large portions of environments cannot be accessed by the robot during map building. We experimentally validate the quality of the completed maps by using them to perform planning tasks.

Email addresses: matteo.luperto@unimi.it (Matteo Luperto), federico.amadelli@mail.polimi.it (Federico Amadelli), moreno.diberardino@mail.polimi.it (Moreno Di Berardino), francesco.amigoni@polimi.it (Francesco Amigoni)

1. Introduction

In recent years, ground mobile robots have been successfully employed in several indoor applications in industrial, office, and domestic environments [1]. When a robot is deployed in a new setting, it often autonomously builds a *map* representing the environment in which it operates. Then, the robot exploits the map to efficiently localize, navigate, and plan paths and tasks in the environment. Sometimes, the robot building the map and the robot using the map are different. 2D metric maps, like *grid maps*, are widely employed since they can be built from data coming from pervasive and relatively cheap sensors like 2D laser range scanners by using consolidated 2D SLAM methods [2]. Moreover, such maps are rather robust to events like day/night light changes, the presence of people, and objects moving around [1].

Ideally, a map should represent the entire operational environment of the robot. However, during the process of map building, it could happen that some areas of interest for the robot's activity are inaccessible, due to temporary conditions that are beyond the robot's control, like a blocked path or a closed door. As a consequence, these areas are not represented on the map, and this can limit the autonomy and operations of the robot exploiting the map. For example, if the robot is unaware of the presence of some rooms behind closed doors, it has no means to plan in advance the actions to be performed when the doors are opened (e.g., in order to visit them).

In this paper, we provide an initial contribution towards solving the above problem, by presenting a method that completes robot 2D grid maps with the predicted layouts (i.e., the geometrical shapes) of unobserved rooms behind closed doors, which we call *closed rooms*. The main idea of our approach is to exploit the underlying geometrical structure of indoor environments that can be detected from the walls to provide knowledge about parts of the environment that are not directly observable at mapping time.

This estimated knowledge, although approximated [3], could provide meaningful insights to the robot about the structure of the environment and could

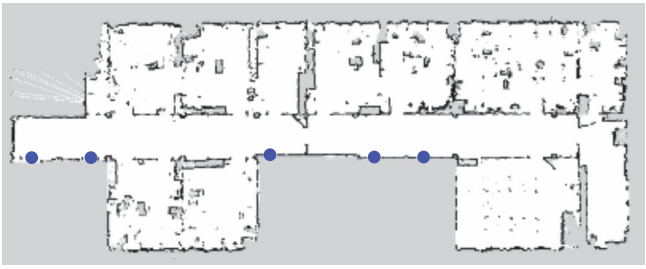

(a) Map built when some doors are closed (blue dots).

(b) Completed map with the predicted layout of closed rooms.

Figure 1: An example run of our approach for predicting layouts of rooms that are behind doors closed at mapping time, where we simulated 5 closed doors in a map from [5].

be exploited in tasks such as exploration [4], localization, task planning, and reasoning. In particular, the availability of a reliable estimate about the shape of closed rooms can improve robot performance in tasks that involve *offline* planning on a map; this way, the robot can plan visits to closed rooms, in the same way it does with the portion of the environment it actually observed. In the following, we show how our completed maps are enough reliable to be used to plan paths for full *coverage* of the environments also when several rooms are closed when the map has been built.

We assume a robot that is able to build a 2D grid map of an indoor environment in which some doors are closed and that is able to detect the positions of such closed doors. Since detecting doors (e.g., from vision) is not the purpose of this paper, we assume that the robot employs a method like [6, 7]. Given a grid map (Fig. 1a) and the positions in the map of the closed doors, our method identifies the main structural features of the environment by detecting the walls. The directions of the walls are associated with representative lines that are used

to partition the map into a number of polygonal faces. Although we assume that most of the walls can be approximated by straight lines, which is the case for the vast majority of indoor environments, we do not enforce any Manhattan structure, but we use the walls' main directions, directly retrieved from the map, for making predictions. The predicted layout of a closed room is the set of faces that maximize an evaluation function that accounts for the consistency with the known portion of the environment. Finally, the predicted layouts of closed rooms are inpainted within the grid map (Fig. 1b). An interesting feature of our method is that it can jointly predict the layouts of multiple adjacent closed rooms (e.g., when all rooms along the same side of a corridor are closed). Experimental evaluation is performed by considering both large-scale simulated environments and real-world grid maps from publicly available datasets. Results show that our method successfully predicts the layout of closed rooms and accurately completes grid maps even when large portions of the environments are not accessible at mapping time. Moreover, we show that the completed map can be efficiently used to plan offline a path for covering the environment. The planning is performed by solving a Traveling Salesperson Problem (TSP) on a graph derived using a Voronoi segmentation of the completed map of the environment, after the inpaiting of the predicted layout of closed rooms.

This paper extends our previous work [8]. Specifically, we add the application of our approach to planning a coverage path (Section 5). Moreover, we develop a variant of our method that uses the robust structural features computation of [9] to identify representative lines, showing that the performance of our prediction method are rather independent of how the structural features are detected (Section 4.2).

2. Related Work

In this section, we survey some techniques developed across different fields to retrieve structural knowledge in indoor environments and to predict their unobserved parts.

Within the field of mobile robotics, a popular approach towards structure identification in indoor environments is room segmentation, where rooms are identified by dividing a metric map into semantically meaningful parts. A survey that compares 2D room segmentation methods is reported in [10]. Authors of [11] present a room segmentation method that uses structural line features similar to the representative lines we use in this paper. Another popular problem is the structure identification in indoor environments from 3D point clouds [12, 13, 14, 15, 16]. We take inspiration from the structure identification steps of [16], but we adapt them to noisy 2D grid maps. While the above methods identify structures within the map representing the observed part of the environment, in this paper we are interested in predicting the structure of unobserved rooms.

The method of [17], starting from a 2D metric map of an indoor environment, reconstructs a geometrical representation of (partially) observed rooms using Markov Logic Networks and data-driven Markov Chain Monte Carlo (MCMC) sampling. The shape of a room, approximated by a polygon, is obtained using a set of logic rules identifying the desired properties (e.g., perpendicular walls, box model). In [18], we propose a method to complete partially observed rooms. In this paper, we assume to have no knowledge about the closed rooms whose layout we predict.

The idea of obtaining knowledge on unobserved parts of environments has recently been addressed using heterogeneous approaches. Some of them predict unknown features of an environment by exploiting knowledge coming from other maps previously acquired in the same or in other environments. For instance, [19] predicts loop closures in a metric map. The parts of the metric map that are yet unknown are completed by superimposing matching maps from a database of previously observed environments. The method of [20] uses a library of map structures to predict the unknown parts of a map incrementally explored by a team of robots. Another recent method is that of [21], where a variational autoencoder (VAE) is employed to predict unobserved regions of an environment starting from a partial map. However, [21] considers buildings that are very similar to each other (see [22] for a discussion of a similar prob-

lem in a different, but related, setting), thus providing an approach that seems difficult to generalize to other different environments. Several methods predict the presence of specific elements in the unobserved parts of environments using neural networks trained on similar environments. For instance, [23] trains a convolutional neural network (CNN) on a set of images representing building floor plans and uses it to predict the locations of emergency exits. In [24], U-nets, a type of CNNs, are used to expand egocentric RGB-D observations to infer the occupancy state beyond the visible regions. In doing so, the robot can anticipate the next sensorial readings. The method of [25] uses Conditional Neural Process for predicting the local map of the unobserved parts of an environment to improve online trajectories planning.

Some other approaches, like the one we present in this paper, do not rely on external knowledge but only on the content of the partial map. In this sense, a method that shares some similarities with our approach is that of [26], which reconstructs the neighborhood of a frontier between known and unknown portions of a map by identifying similar structures in the known map. If a match is found, the matching portion of the known map is superimposed to the frontier, thus providing an estimate of the structure of its neighborhood.

Some methods predict the existence and the semantic labels of unobserved rooms, but they do not predict their geometrical shapes nor update the metric map with the prediction, as we do in this paper. Examples are the systems proposed in [27], that uses semantic knowledge in the form of chain graphs to predict the existence of a room (and its label) in the unexplored space, and in [28], where the prediction of the existence of a new room is made by using sum-product networks. The approach of [29] predicts the topology and the labels of unobserved rooms by matching the observed part of the environment (represented as a labeled graph) to a database of environments. Also [30] predicts the presence of new rooms in partially observed environments by reasoning on graphs using graph kernels.

To the best of our knowledge, no method addresses our specific problem, that of predicting the geometrical shape of rooms behind closed doors.

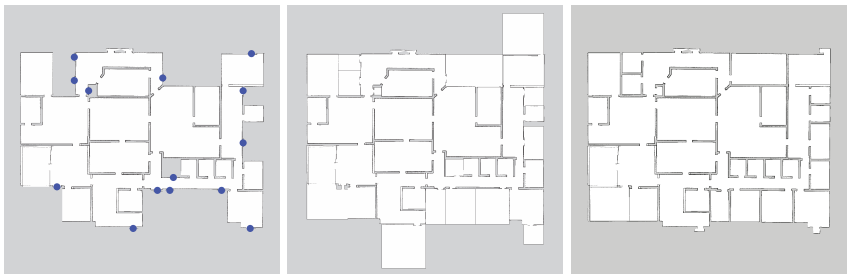

(a) Partial map with closed doors.

(b) Inpainted map with predicted closed rooms.

(c) Full map.

Figure 2: A partial map of a large-scale indoor environment with 14 closed doors (2a), the map completed by our method (2b), and the actual full map (2c).

3. Our Method

Our method predicts the layout of rooms behind closed doors in a purely geometrical way, without learning models from other environments. It starts from a 2D *grid map* M of an indoor environment obtained by a robot through a SLAM mapping process using data acquired by a laser range scanner. This map is composed of identical square cells that are labeled as known or unknown according to the fact that they have been perceived or not by the robot. Known cells are further labeled as either free or obstacle, according to the occupancy of the corresponding area. The grid map is assumed partial in the sense that some rooms could not be accessed by the robot during the mapping process due to some *closed doors* and, as a consequence, are not included in the map. We propose a method that predicts the possible layout of these *closed rooms*, which is then used to complete the grid map.

Our method assumes that the robot can detect the position of closed doors inside the environment (blue dots in Fig. 2a), for example by using existing computer vision methods like [6, 7]. Consequently, the initial input of our method is a grid map M and a set of closed door locations $d \in D$. Although doors are represented as line segments in 2D maps, our method considers their middle points. Hence, each $d = (x, y)$ represents the mid-point coordinates (coordinate system of M) of the line segment corresponding to a closed door in

the map. We assume that each closed room has exactly one door.

Our method predicts the possible layouts of closed rooms by leveraging the fact that, due to the structured nature of indoor environments, their geometrical shapes have some common features with other rooms and walls in the metric map. The method is based on a sequence of steps that are detailed in the next sections using the map of Fig. 2 as a reference.

3.1. Structural features

The first step relies on the method of [18] to extract the structure of indoor environments by identifying the direction of walls in the metric map M and to partition M using those directions. The method starts by extracting a set of line segments from M by using the Canny edge detection [31] and the probabilistic Hough line transform [32]. Line segments are clustered together in two phases. First, the mean shift algorithm [33] clusters together line segments with similar angular coefficients. Then, for each angular cluster, all line segments that are also collinear (along the same line) are clustered together by performing spatial clustering. Full details are omitted for brevity, please refer to [18].

Each spatial cluster is then associated to a *representative line*, in red in Fig. 3, which indicates the direction of collinear, but possibly spatially separated, walls. The result is the detection of a (hopefully small) number of representative lines that describe the direction of all the walls within the environment. Four additional *boundary lines* are added at the extremity of M at a fixed distance from the bounding box of the map and with the same angular coefficients of the two largest angular clusters of line segments. We do not assume Manhattan environments, as the directions of the representative lines are directly recovered from the map. However, in many real-world indoor environments, most walls are perpendicular (e.g., see [34]) and, consequently, the representative lines used for map segmentation are often perpendicular. For example, the Intel Lab map from [5], shown in Fig. 4, contains curved walls, but representative lines corresponding to most walls are perpendicular. The resulting approximation is adequate for accurately predicting the geometrical shape of closed rooms, as

shown in Section 4.

The intersections of all the representative lines partition the map into *faces*. A face $f \in \mathcal{F}$ is a polygon having as *edges* the line segments obtained by the intersections of the representative lines (Fig. 3). The faces with an edge belonging to a boundary line are called *border faces*.

Finally, we separate the faces that are inside the part of the environment observed by the robot from those that belong to the unobserved parts of the environment, as only the latter ones will be considered when predicting the layout of a closed room. Specifically, we keep only faces $f \in \mathcal{F}_u \subseteq \mathcal{F}$ where at least a 30% of their area is unknown.

3.2. Closed room locations

The prediction of the geometrical shape of closed rooms starts from the faces that are immediately behind closed doors. We associate each closed door $d = (x, y)$ to its closest edge e_d (on a representative line), thus determining the door orientation as collinear to e_d. As edges are the boundaries between faces, we consider the two faces that share an edge e_d: one of the two faces is inside the known part of M and the other one belongs to the unknown part of M. The latter face is the one that is behind the closed door d and is inserted in the set of initial faces $I \subseteq \mathcal{F}_u$, which are used as seeds to estimate the layouts of the rooms behind the closed doors. Initial faces of closed rooms for the map of Fig. 2 are shown in Fig. 3.

A particular case arises when there are two closed doors, d_1 and d_2, that are associated with the same edge e_d along a representative line ℓ (and, consequently, that have the same initial face). As we assume that there is only one door for each closed room, we artificially add a representative line perpendicular to ℓ and passing at equal distance from the two doors' positions d_1 and d_2. In this way, e_d is split into e_{d_1} and e_{d_2}. This allows us to address situations, as in the corridors of Fig. 6, with multiple closed rooms adjacent to each other.

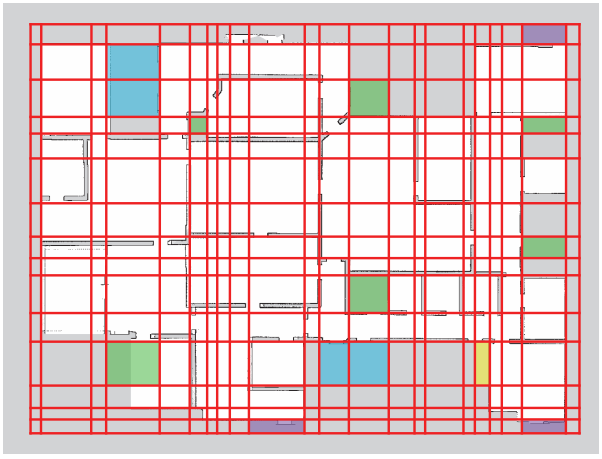

Figure 3: Representative lines (red) and faces obtained from segmenting the map of Fig. 2. Initial faces of closed rooms are shown with different colors: green for independent rooms and light blue for dependent rooms, while initial faces of closed rooms on the border (border faces) are in purple. A particular case is the initial face in yellow, which represents a room that is initially independent but, expanding, becomes dependent by touching the predicted layout of a nearby closed room (in light blue at its left).

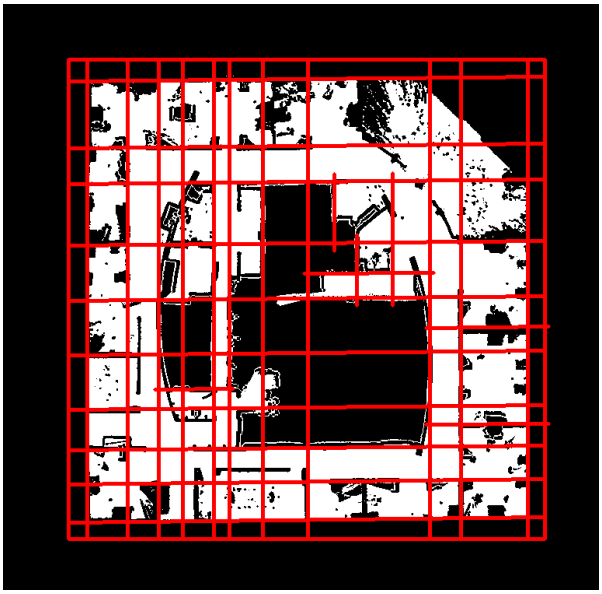

Figure 4: Representative lines for the Intel Lab map from [5].

3.3. Closed rooms expansion

The predicted layout of a closed room r in the environment is composed of one or more faces $f \in \mathcal{F}_u$ and is obtained by selecting the most likely set of faces from $\mathcal{F}_u$, adjacent to its initial face i_r, according to the surrounding environment.

We define as $d(f, f')$ the topological distance between faces f and f'. For instance, if two faces have one common edge, their distance is 1. The process of identifying the predicted layouts of closed rooms is performed greedily by jointly iteratively expanding them by considering an increasingly larger set of faces. More precisely:

(1) We initialize $k = 1$.

(2) For each closed room r, we select a set of candidate faces from $\mathcal{F}_u$ as:

$$F_r^k = \{f : f \in \mathcal{F}_u \mid \exists f' \in \hat{F}_r^{k-1} \text{ and } d(f, f') = 1\},$$

where $\hat{F}_r^0 = \{i_r\}$ (with $i_r \in I$) is the initial face behind door d of room r. Calling $\mathcal{P}(\cdot)$ the power set and $\Phi(\cdot)$ an evaluation function (described in Section 3.4), we select the best layout for room r at step k as the set of faces F that, together with $\hat{F}_r^{k-1}$ (thus expanding the layout at step $k-1$), maximize $\Phi(\cdot)$:

$$\hat{F}_r^k = \underset{F \in \mathcal{P}(F_r^k)}{\arg\max}\, \Phi(\hat{F}_r^{k-1} \cup F).$$

We remove the faces in $\hat{F}_r^k$ from $\mathcal{F}_u$ (so that a face belongs to the predicted layout of at most one closed room) and we consider the next closed room r. For a given k, we consider closed rooms ordered from the smallest to the largest F_r^k (ties are broken randomly). However, we empirically observed that room ordering has a small impact on the final result.

(3) We increase $k \leftarrow k+1$ and we repeat from (2), until no faces are left in $\mathcal{F}_u$ or a threshold for k is reached. For each room, the set of faces $\hat{F}_r^* = \hat{F}_r^k$ selected in the last step is considered as the predicted layout.

At a generic step k, the predicted layout $\hat{F}_r^k$ of a closed room r is thus updated from the predicted layout $\hat{F}_r^{k-1}$ at the previous step. The idea is that we jointly expand the predicted layouts of all closed rooms until a good estimate is found for each one of them. This is motivated by the fact that closed rooms can belong to two categories: *independent* closed rooms, whose predicted layout is not adjacent to the predicted layout of any other closed room; and *dependent* closed rooms, which have at least a face of their predicted layout that is adjacent to a face of the predicted layout of another closed room. Examples of initial faces of independent (dependent) rooms are shown in green (light blue) in Fig. 3. Note that, with the increase of k, some independent closed rooms may become dependent; this happens when the predicted layouts $\hat{F}_r^k$ and $\hat{F}_{r'}^k$ of two rooms r and r' are expanded in opposing direction, eventually sharing adjacent faces. An example is the room in yellow of Fig. 3. The reason for separating these two closed room types is that, while the predicted layouts of independent rooms should be consistent only with M, the predicted layouts of dependent rooms should be jointly estimated with that of the nearby closed rooms.

3.4. Candidate layout evaluation

A possible predicted layout of a room r, represented as a set of faces $\hat{F}_r^{k-1} \cup F$, is scored using an evaluation function $\Phi(\hat{F}_r^{k-1} \cup F)$. The function embeds competing objectives, like to maximize the area of the room and to maximize the coherence of the room structure wrt that of nearby rooms. Because of that, the objective function is a weighted sum of some different components, that are now described. In what follows, with a slight abuse of notation, we use F to denote the layout $\hat{F}_r^{k-1} \cup F$.

The first component, $area(F)$, is the room layout area.

The second component is the convex hull ratio *CHR*, that prefers regular room layouts:

$$CHR(F) = CH(F)/area(F),$$

where *CH* is the area of the room's predicted convex hull.

The third and fourth components are designed to maximize the similarity between the predicted layout of the room and the rest of the map. In particular, the third component minimizes the edges of a predicted layout that touch the unknown parts of the map. More precisely, the ratio of the free edges *FER* is defined as the ratio between the sum of the length of the edges of F that are also edges of a face $f \in \mathcal{F}_u$ (where $\mathcal{F}_u$ is the set of unobserved faces remaining after the application of the algorithm in the previous section) and the sum of the length of all edges along the external contour of F. The fourth component penalizes the predicted layouts that are not regular. More precisely, free faces penalty *FFP* is defined as the number of faces $f \in F$ that have at least *two* edges in common with faces $f \in \mathcal{F}_u$.

The fifth component is the room proportion P, the ratio between the two main dimensions of the room's bounding box, which is intended to favor regular predicted layouts.

We define two different evaluation functions, Φ_{ind} and Φ_{dep}, for independent and dependent rooms, respectively:

$$\Phi_{ind}(F) = \omega_1 \cdot \sqrt{area(F)} - \omega_2 \cdot CHR(F) \\ - \omega_3 \cdot FER(F) + \omega_4 \cdot FFP(F),$$

$$\Phi_{dep} = \omega_1 \cdot \sqrt{area(F)} - \omega_2 \cdot CHR(F) - \omega_5 \cdot FER(F) \\ - \max(\omega_4 \cdot (FFP(F) - 1), 0) - \omega_6 \cdot (P(F) \cdot \min(FFP(F), 1)).$$

In the case of multiple adjacent closed rooms, the last term in Φ_{dep} tends to prevent that the expansion is stopped before all these rooms have a similar shape.

We do not enforce a square or rectangular shape for the predicted layouts (as shown, for example, in Fig. 6), but our evaluation function aims at predicting accurate room shapes according to the observed map. However, since real-world indoor environments are inherently structured and most walls are perpendicular also in non-Manhattan environments (Section 3.1 and Fig. 4), good predictions are usually rectangular.

The next two steps address special cases.

3.5. Joint rooms layout prediction

At the end of the expansion (Section 3.3) it could be the case that two adjacent dependent rooms have different shapes (e.g., this happens when one of the two rooms is by chance initially expanded in the direction of the other room's initial face, thus limiting the expansion of the second room). To adjust such situations, we give to adjacent dependent closed rooms the possibility to swap one or more faces between them. Given two sets of faces $\hat{F}^*_r$ and $\hat{F}^*_{r'}$, representing the predicted layouts of dependent rooms r and r', we compute E as the set of faces in $\hat{F}^*_r$ (or in $\hat{F}^*_{r'}$) that have an edge in common with a face in $\hat{F}^*_{r'}$ ($\hat{F}^*_r$) and that, consequently, could be exchanged between the two rooms. We jointly evaluate all the possible combinations of face assignments $\mathcal{P}(E)$ (in one assignment, some faces of E are assigned to r, the other ones to r') by evaluating the corresponding rooms' predicted layouts $\bar{F}^*_r$ $\bar{F}^*_{r'}$ using the following function:

$$\Phi_{joint}(\bar{F}^*_r, \bar{F}^*_{r'}) = \omega_7 \cdot \sqrt{\min(area(\bar{F}^*_r), area(\bar{F}^*_{r'}))/\max(area(\bar{F}^*_r), area(\bar{F}^*_{r'}))} - \omega_8 \cdot (FFP(\bar{F}^*_r) + FFP(\bar{F}^*_{r'}))$$

We eventually select the face assignment that maximizes Φ_{joint} and we swap the corresponding faces between r and r'. If the adjacent rooms are more than two, they are considered in pairs.

3.6. Closed rooms on the borders

The layout of a closed room r may extend outside the bounding box of the current map. In that situation, we cannot directly use faces and representative lines to predict the layout of r, as M does not provide any knowledge on one of the dimensions of the room. This happens when the initial face i_r of a room r is one of the border faces (e.g., that in purple in Fig. 3). To provide a layout also in this case with limited information, we roughly predict the shape of the

# env	1	1	1	2	2	4	9
$\max \lvert D \rvert$	7	9	10	11	12	13	15

Table 1: Number of simulated environments and corresponding $\max |D|$ number of closed doors.

room as a square (of the same size as the edge e_d of the initial face i_r, see Section 3.2). If there are multiple dependent rooms in this condition, we adjust (by averaging) their outwards dimension to the same value.

3.7. Inpainting predicted layouts into the grid map

In this last step, we inpaint the predicted layouts of closed rooms into the map M. This is done by creating open passages corresponding to the positions of the doors D in the map (door width is a customizable parameter, which we set to 80 cm) and by changing the value of cells in M from unknown to free or obstacle, according to the fact that they correspond to the inner area of a predicted layout or to one of its external edges. As a result, a completed (predicted) grid map M_{pred} is eventually available to the robot. The predicted map for the partial map of Fig. 2a is shown in Fig. 2b. Note that the largest difference wrt the actual map of Fig. 2c is in the rough predictions of closed rooms on the borders.

4. Experimental Evaluation

In this section we present the experimental activities performed to evaluate the proposed method to predict the layouts of closed rooms in indoor environments. We present both quantitative results obtained in simulation and qualitative results obtained by applying our method to real-world maps from public datasets. We also investigate the robustness of our prediction approach by evaluating the impact of using a way different from that of Section 3.1 to extract structural features from the metric map M.

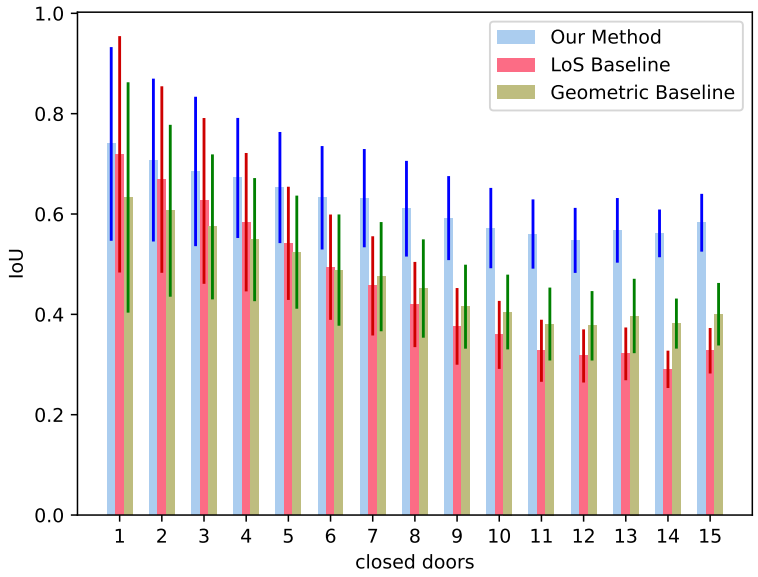

Figure 5: Average and standard deviation of the IoU of the predicted layout of closed rooms, wrt the number of closed doors in the environment.

4.1. Predicted map evaluation

We start presenting results obtained in 20 simulated indoor environments (office and school environments) in which we consider up to 15 closed doors. Maps are obtained by running the ROS implementation[1] of the GMapping algorithm [35] on data collected by a robot equipped with a laser range scanner during the autonomous exploration of the buildings simulated in Stage[2]. The environments have different sizes and, accordingly, different maximum numbers of closed doors max $|D|$ that a robot can find (see Table 1). We limit the number of possible closed doors to 15 even for the larger environments in order to have a balanced evaluation of the method performance.

For each environment, we repeat 15 times the following procedure: we build

[1] `http://wiki.ros.org/gmapping`
[2] `http://wiki.ros.org/stage`

$N = \max |D|$ different maps by incrementally closing $1, 2, \ldots, N$ doors (if a door is closed in a map where i doors are closed, it is closed also for all maps in which $i + 1, \ldots, N$ doors are closed). Closed doors are selected randomly. For each map obtained in this way, we run our method in order to predict the shape of the closed rooms. Overall, we evaluated $3,885$ maps (for a total of $24,045$ predicted room shapes). In each run, our method receives in input a grid map M and a set of closed doors D. We empirically set values of weights $[\omega_i]$ to [0.06, 10, 7, 10, 2.5, 2, 1, 2] and the maximum number of expansion steps k of Section 3.3 to 9. Experiments are performed on a commercial laptop and each run requires less than 2 minutes for all maps.

Given the predicted layout of a closed room $\hat{F}_r^*$ and that of its ground truth counterpart F_r^* (obtained from the floor plan of the simulated environment), we compute their *Intersection over Union* (IoU) as:

$$\text{IoU}(\hat{F}_r^*, F_r^*) = \frac{\hat{F}_r^* \cap F_r^*}{\hat{F}_r^* \cup F_r^*}.$$

An high IoU indicates that the geometric prediction $\hat{F}_r^*$ accurately resembles F_r^* (IoU is commonly used for this type of evaluation, as in [14]).

Since, as discussed in Section 2, we are not aware of any other method that predicts the layout of closed rooms, we compare our method against two baseline methods. The first one is called *line of sight* (*LoS*) *baseline* and predicts the layout of a closed room as the free area that could be observed in line of sight from the corresponding closed door d. This method is based on the assumption that all the unobserved area behind a closed door is part of the closed room. The predicted layout of the room is spatially limited by the bounding box of the map. The second method is called *geometric baseline* and adds to the predicted layout of a room all the faces $f \in \mathcal{F}_u$ that are in line-of-sight from the door d, until boundary faces are met.

Fig. 5 shows the performance of the proposed method against the two baselines. Our method obtains stable and accurate predictions of closed rooms' layouts even when a large number of doors are closed across the environments. On the other side, baseline approaches perform well when few rooms are closed

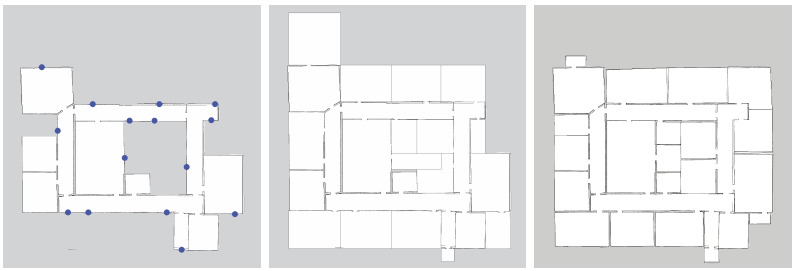

(a) Partial map. (b) Predicted map. (c) Full map.

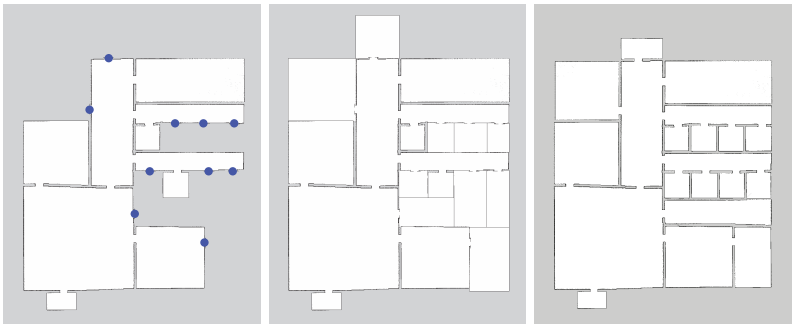

(d) Partial map. (e) Predicted map. (f) Full map.

Figure 6: Two examples in which our method predicts the layouts of 15 and 10 closed rooms.

(because they basically flood-fill gaps in the maps), but have a dramatic drop in performance as the number of closed doors increases.

Fig. 6 shows that our method can complete metric maps also when large parts of the buildings are not explored (15 closed doors). For instance, it provides a rather accurate prediction of all the closed rooms connected to the upper corridor in Fig. 6b. Fig. 6e shows a similar result where, despite the presence of multiple closed doors connected to the same corridor, our method provides a sound estimate of the environment map.

In Fig. 7 we show two examples where large portions of the environments are not mapped: only the corridors and few rooms are observed by the robot, while most rooms are closed. Specifically, in these environments 14 (respectively, 8) doors are closed, and only 6 (7) rooms are mapped in M. As it can be seen from Figs. 7b and 7e, our method can reliably predict the entirety of the map even in these challenging settings. The predicted maps are similar to the actual maps

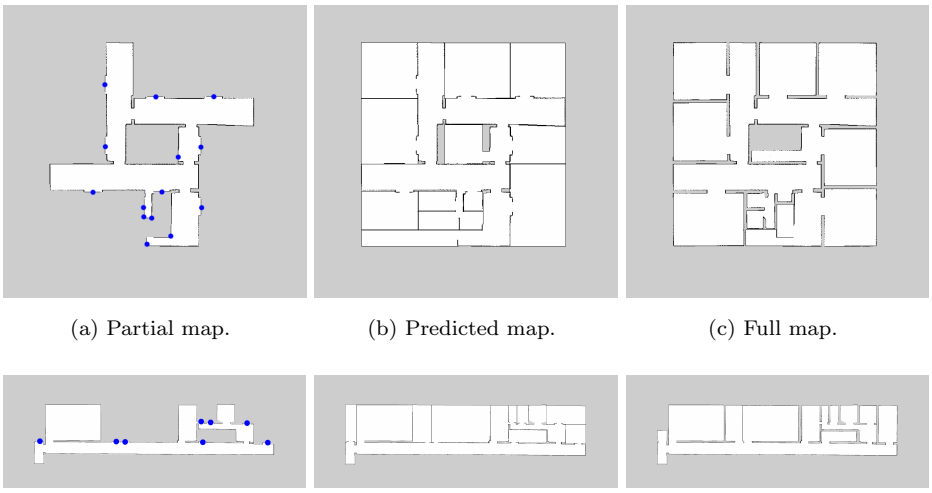

(a) Partial map. (b) Predicted map. (c) Full map.

(d) Partial map. (e) Predicted map. (f) Full map.

Figure 7: Two examples in which our method predicts the layouts of closed rooms starting from little knowledge, mainly about the structure of the environments.

of the environments (Figs. 7c and 7f) and thus the predicted maps represent a reliable source of knowledge by the robot.

Finally, Figs. 1 and 8 show how our method can complete real-world partial maps (obtained from publicly available datasets [5, 36]) with multiple closed doors. For these results, we manually remove some rooms from the original maps and we predict their possible layouts using our method. (Note that, although some maps in [5] have multiple closed doors, their locations are not provided.) Despite large missing portions of the map, our method provides a valid estimate of layouts of closed rooms even in the presence of clutter and inaccuracies. Further results on both simulated and real-world maps are available in a video[3].

4.2. Results with a different identification of structural features

In this section, we show that our approach to predict the layout of closed rooms is rather independent of the method used to extract structural features

[3] `https://amigoni.faculty.polimi.it/research/ECMR2021-completing-maps-closed-rooms.html`

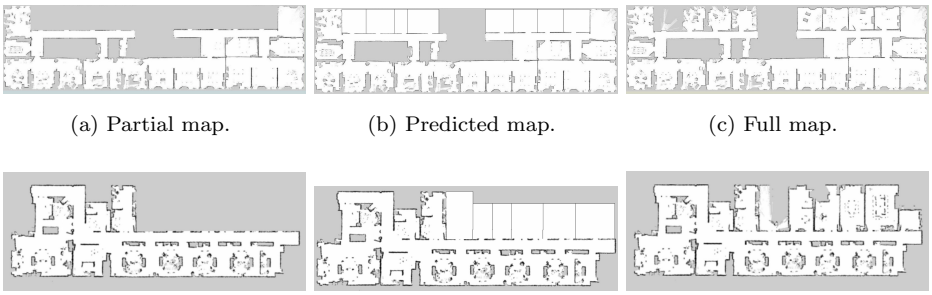

(a) Partial map. (b) Predicted map. (c) Full map.

(d) Partial map. (e) Predicted map. (f) Full map.

Figure 8: Application of our method to publicly available real-world maps from [5, 36].

from M. All the above results are obtained with structural features identified as described in Section 3.1.

We introduce a variant of our proposed approach that employs the feature extraction method of [9], called *ROSE*. *ROSE* identifies the *dominant directions* of M, namely the main directions of straight lines using a frequency-based technique. In the variant, we initially filter the metric map M keeping only the line segments along the dominant directions identified by *ROSE* and we then align the representative lines to the dominant directions. This variant provides robustness when dealing with cluttered and noisy maps (as *ROSE* is explicitly designed to extract robust structural features in such settings).

To compare the two variants, we evaluate the IoU using the same experimental setting of Section 4.1. The results are shown in Fig. 9 and show no noticeable statistical difference in performance.

This shows that our approach to predict the layout of closed rooms could be applied in different settings with different low-level processing of metric maps.

5. An Application to Coverage

In this section, we show how the availability of the estimated layout of closed rooms, inpainted in the grid map M_{pred} returned by our approach, can be used by a robot that plans a sequence of actions to be performed in an environment.

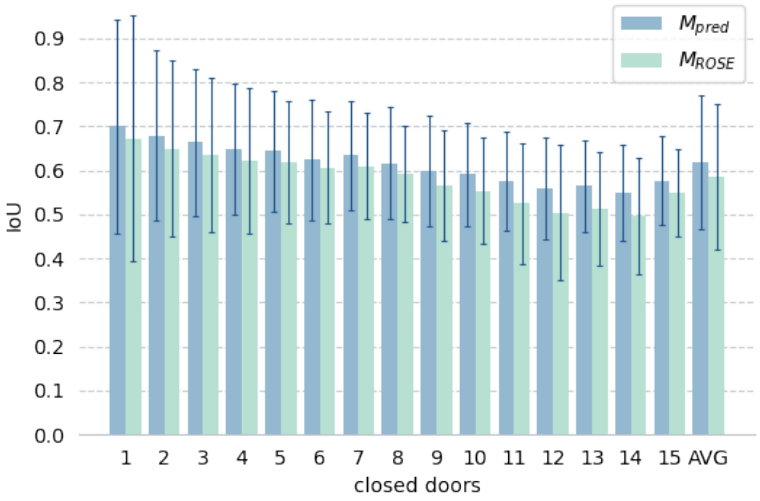

Figure 9: Average and standard deviation of the IoU of the predicted layout of closed rooms for two variants of our approach (M_{pred} extracts structural features as per Section 3.1 and M_{ROSE} extracts structural features as per [9]).

More precisely, the robot can use the reliable estimate provided by M_{pred} to plan actions also in the (unobserved) closed rooms that are not represented in the initial metric map M. We focus on a specific task, the *coverage* of a set of locations in the environment [37], where the robot has to plan the shortest path that visits all the locations. This algorithmic task is underlying several real-world applications, like monitoring, cleaning, and patrolling. We naturally model the problem as a Traveling Salesperson Problem (TSP) on an undirected topological graph whose vertices are the locations to visit and edges denote the direct connections between (mutually visible) locations with the corresponding distance. For example, such a graph and the TSP solution on it are used as prior knowledge in [38] for efficiently exploring an unknown environment. In the following, we detail how the topological graph is built and how the TSP solution on the graph is calculated. Then, experimental results are presented.

5.1. Topological graph computation

The topological graph depends on the specific application. For example, its vertices could be close to each other in the case of cleaning and farther apart in the case of patrolling. For our purposes, we build an application-independent topological graph from the *Voronoi graph* [10]. The Voronoi graph of a metric map features vertices and edges that have the maximal distance from at least two obstacles in the map. The Voronoi graph provides a meaningful representation of the structure of the environment and its vertices are (almost) equally distributed over the rooms of the environment (Fig. 10(a)). Thus, visiting all the vertices of such a graph amounts to cover the entire environment. Now we detail the procedure we follow to build the topological graph.

We start from a grid (metric) map. The map is then processed with the method of [9] to filter non-structural features (e.g., clutter) that could introduce unnecessary vertices in the graph. The clean map is then feed to the *skeletonize* function of the Python module Scikit-image[4]. The output is the skeleton of the map, i.e., a 1-pixel wide representation of the map, that is then processed using the NetworkX[5] module in order to obtain the Voronoi graph. Then, like in [38], we keep, as vertices of the topological graph, the vertices in the Voronoi graph that are leaves (i.e., that have degree 1) or those that are junctions of three or more edges (i.e., that have degree at least 3). At this point, we ensure that each edge that connects two vertices does not cross any obstacles in M, so that two connected vertices are visible from each other. Note that, if the initial metric map has closed doors, we also add to the topological graph vertices at the locations of closed doors D (in blue in Fig. 10).

To evaluate the effect of different degrees of knowledge in planning a coverage path, we compute the topological graph starting from different types of metric maps:

- the original map with closed doors M, in which closed rooms are missing;

[4] `https://scikit-image.org/`
[5] `https://networkx.org/`

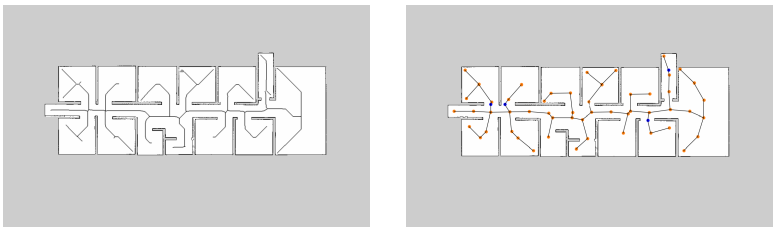

(a) Voronoi graph. (b) Topological graph.

Figure 10: Example of Voronoi graph and of the final topological graph.

- the predicted map M_{pred}, obtained starting from M by applying the method proposed in this paper;
- the ground truth map M_{GT}, the actual map of the environment in which all the doors are open and there are no closed rooms.

Examples of topological graphs built from these three types of metric maps are shown in Fig. 11.

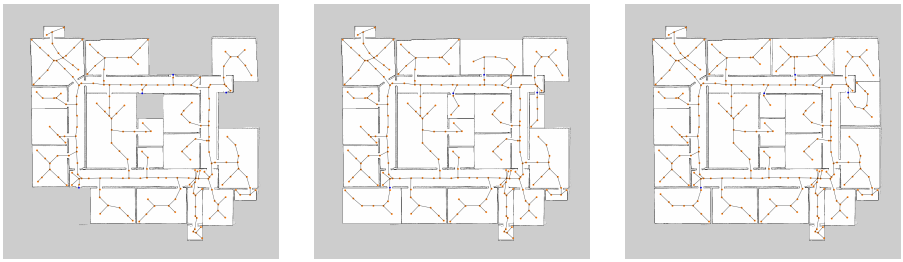

(a) Partial map M. (b) Predicted map M_{pred}. (c) Ground truth map M_{GT}.

Figure 11: Topological graphs computed on three types of maps. The vertices corresponding to the closed doors are in blue.

5.2. TSP solution

The locations of the vertices in the topological graph are given by their positions within the metric map. The graph building process ensures that two vertices are connected only if they are visible from each other. As a consequence, we set the weight w_{ij} of an edge e_{ij} of the topological graph, connecting the

vertices v_i and v_j with coordinates (x_i, y_i) and (x_j, y_j) respectively, equal to the Euclidean distance between the two vertices:

$$w_{ij} = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}.$$

To find a solution to the TSP we rely on the Christofides algorithm [39] using tabu search [40, 41] to escape local minima and possibly reach better solutions. The algorithm is run for 60 seconds and the best solution found during this time is retained as the TSP solution. The implementation is based on the OR-Tools[6] library.

We compute the TSP solution for the topological graphs obtained from the three maps introduced before, M, M_{pred}, and M_{GT}. For M, we add to the topological graph a vertex at the location of each closed door, as explained above. Consequently, we obtain three coverage costs of a given environment, TSP_{closed}, TSP_{pred}, and TSP_{GT}. More precisely,

- TSP_{GT} is the actual cost of the TSP solution as computed from the full map of the environment;
- TSP_{closed} is an understimate of the cost of the TSP solution, as it is obtained from a map with closed rooms;
- TSP_{pred} is an approximation of the cost of the TSP solution informed by our proposed method.

As M covers significantly less area than M_{GT}, TSP_{closed} is a too rough underestimate of the actual cost TSP_{GT}. Consequently, we add a fourth estimate, $TSP_{baseline} = TSP_{closed} + C$, where:

$$C = \frac{TSP_{closed}}{|R|} \cdot |D|$$

represents an estimate of the cost of visiting closed rooms. $|R|$ is the number of observed rooms and $|D|$ is the number of closed doors in M.

[6] `https://developers.google.com/optimization`

Note that we cannot use here the two baseline methods of Section 4, as they predict the layouts of single closed rooms independently and do not guarantee that the predicted layouts of different closed rooms do not overlap. As a consequence, their predicted maps could be geometrically inconsistent and cannot be used for computing a coverage path.

5.3. Results

We consider 15 of the 20 simulated indoor environments used in Section 4.1. Similarly to Section 4.1, for each environment, we perform 10 runs where we close up to 15 doors. For each run, we set a different random seed and assure that the first closed door is different with respect to the other runs. Overall, we consider $1,957$ different maps.

We compare TSP_{closed}, $TSP_{baseline}$, and TSP_{pred} against TSP_{GT}, computing the error in estimating the actual cost of the TSP. For example, for TSP_{pred}:

$$error\%_{pred} = \frac{|TSP_{GT} - TSP_{pred}|}{TSP_{GT}} \cdot 100,$$

and similarly for $error\%_{closed}$ and $error\%_{baseline}$.

Overall, averaging over runs and environments, TSP_{pred} turns out to be a good estimate of the real cost TSP_{GT}, with $error\%_{pred} = 6.21$ ($\sigma = 9.06$), performing better than TSP_{closed} and $TSP_{baseline}$: $error\%_{baseline} = 21.4$ ($\sigma = 23.38$) and $error\%_{closed} = 21.37$ ($\sigma = 16.18$).

Fig. 12 shows the details of the comparison between TSP_{closed}, $TSP_{baseline}$, and TSP_{pred} with different numbers of closed doors. Our approach consistently returns an estimate that is close to the actual value of TSP_{GT} and that is stable, namely rather independent of the number of closed doors (i.e., the estimate is accurate in maps that are almost complete and in maps with several closed rooms). The same does not hold for TSP_{closed} and $TSP_{baseline}$, whose errors follow a similar trend and increase with the number of closed doors. In particular, $error\%_{baseline}$ is slightly better when few doors are closed and $error\%_{closed}$ is sometimes better when more than 10 doors are closed.

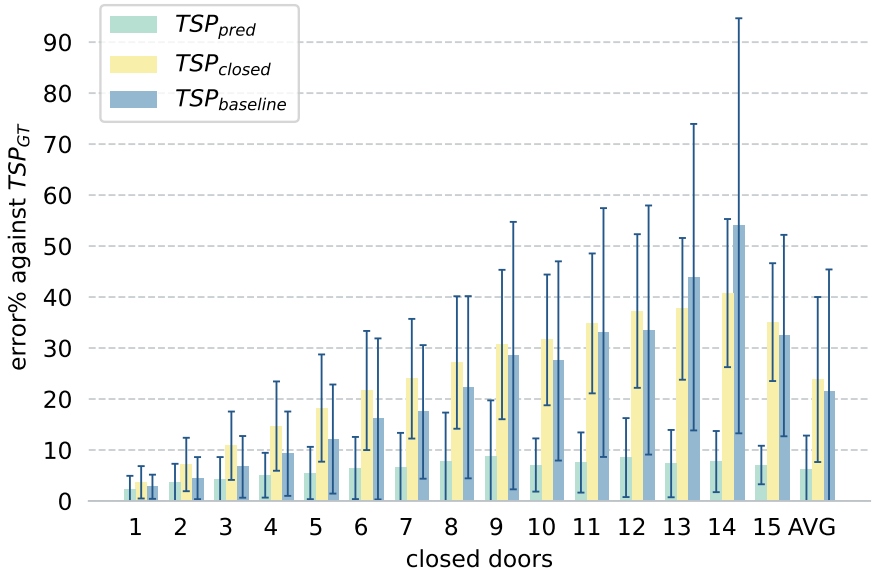

Figure 12: Errors in the length of coverage paths (averaged over the runs and the environments) obtained with different types of maps.

Fig. 13 shows two examples in order to get some insights relative to the above results. The good approximation of TSP_{pred} is achieved because the topological graph build from M_{pred} is similar to the actual topological graph built from M_{GT}. In the first example, shown in Figs. 13a-13c, we have $error\%_{pred} = 0.33$, $error\%_{baseline} = 50.48$, and $error\%_{closed} = 49.21$. In the second example, shown in Figs. 13d-13f, we have $error\%_{pred} = 15.87$, $error\%_{baseline} = 40.54$, and $error\%_{closed} = 35.36$. Note that these two examples are the same of Fig. 7 of Section 4.1, in which most of the environment is not directly observed by the robot when building M. Little more than the main structure of corridors is observed and available to the robot. Nevertheless, our method effectively completes the map and the prediction is used to achieve a reliable estimate of the length of a coverage path.

Finally, Fig. 14 shows that using the *ROSE* variant of Section 4.2 does not impact on the performance, further supporting that our approach for prediction

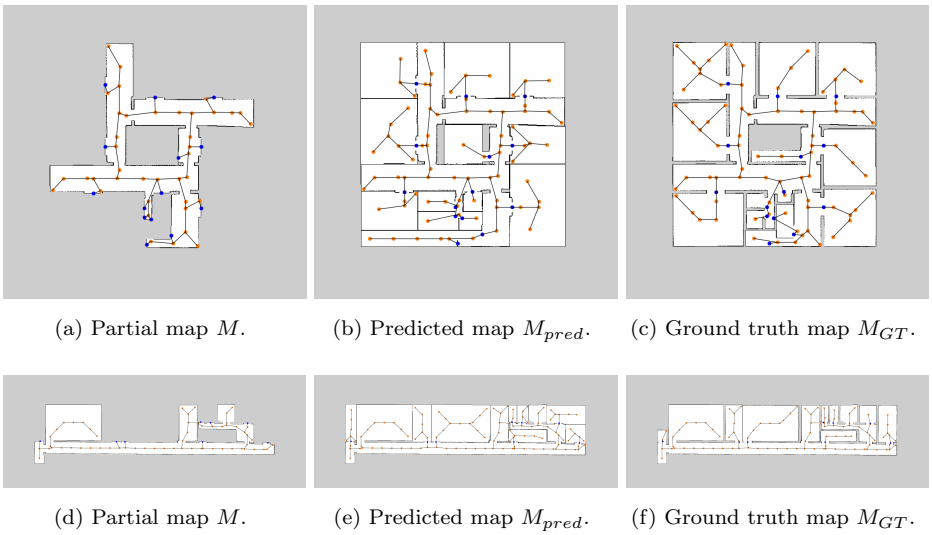

(a) Partial map M. (b) Predicted map M_{pred}. (c) Ground truth map M_{GT}.

(d) Partial map M. (e) Predicted map M_{pred}. (f) Ground truth map M_{GT}.

Figure 13: Topological graphs used to plan coverage paths on the maps of Fig. 7.

is rather independent of the way in which structural features are extracted from the metric map.

6. Conclusions

In this paper, we presented a method for predicting the geometrical shape of closed rooms in indoor environments. The proposed method starts from a grid map in which the positions of closed doors that the robot could not enter are known and exploits the structural regularities of buildings to estimate the layouts of rooms behind such doors. The grid map is then completed by inpainting the layouts of closed rooms. Experiments show the effectiveness of our method, also compared against baseline methods, for large environments with up to 15 closed doors. The method is also shown to depend little on the way in which the structural features are identified. Moreover, we shown that the availability of a reliable estimate of the layout of closed rooms can enable accurate offline planning tasks like that of planning a full coverage path starting from a partially known environment.

In future work we will lift the assumption that a closed room has only one

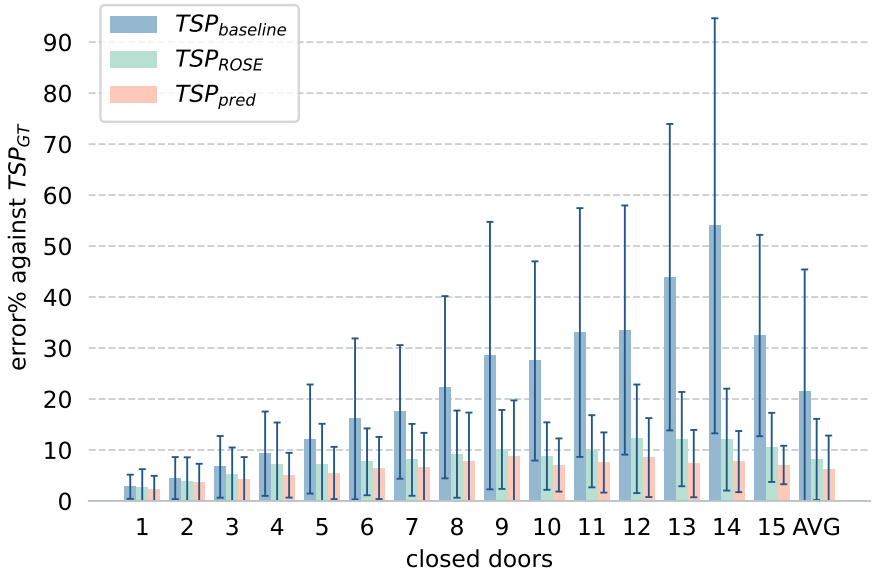

Figure 14: Errors in the length of coverage paths (averaged over runs and environments) obtained with two ways for extracting structural features from metric maps.

door and we will integrate the proposed method in a deployed robot system, using a vision-based system for identifying closed doors and developing a way to update the predicted map as new knowledge is available. We will also investigate the possible uses of our method for search and exploration and for enhancing the understanding of working spaces for collaborative and service robots.

References

[1] L. Kunze, N. Hawes, T. Duckett, M. Hanheide, T. Krajnik, Artificial Intelligence for Long-Term Robot Autonomy: A Survey, IEEE RA-L 3 (4) (2018) 4023–4030. doi:10.1109/LRA.2018.2860628.

[2] S. Thrun, W. Burgard, D. Fox, Probabilistic Robotics, The MIT Press, 2005.

[3] M. Luperto, M. Antonazzi, F. Amigoni, N. A. Borghese, Robot exploration of indoor environments using incomplete and inaccurate prior knowledge, Robot Auton Syst 133 (2020) 103622.

[4] M. Luperto, L. Fochetta, F. Amigoni, Exploration of indoor environments through predicting the layout of partially observed rooms, in: Proc. AAMAS, 2021, pp. 836–843.

[5] A. Howard, N. Roy, The robotics data set repository (radish) (2003). URL `http://radish.sourceforge.net/`

[6] S. Prieto, A. Adan, A. Vazquez, B. Quintana, Passing through open/closed doors: A solution for 3D scanning robots, Sensors 19 (21) (2019) 4740.

[7] A. Llopart, O. Ravn, N. A. Andersen, Door and cabinet recognition using convolutional neural nets and real-time method fusion for handle detection and grasping, in: Proc. ICCAR, 2017, pp. 144–149.

[8] M. Luperto, F. Amadelli, F. Amigoni, Completing robot maps by predicting the layout of roms behind closed doors, in: Proc. ECMR, 2021.

[9] T. P. Kucner, M. Luperto, S. Lowry, M. Magnusson, A. J. Lilienthal, Robust frequency-based structure extraction, in: Proc. ICRA, 2021, pp. 1715–1721.

[10] R. Bormann, F. Jordan, W. Li, J. Hampp, M. Hägele, Room segmentation: Survey, implementation, and analysis, in: Proc. ICRA, 2016, pp. 1019–1026.

[11] R. Capobianco, G. Gemignani, D. Bloisi, D. Nardi, L. Iocchi, Automatic extraction of structural representations of environments, in: Proc. IAS-13, 2014, pp. 721–733.

[12] I. Armeni, O. Sener, A. Zamir, H. Jiang, I. Brilakis, M. Fischer, S. Savarese, 3D semantic parsing of large-scale indoor spaces, in: Proc. CVPR, 2016, pp. 1534–1543.

[13] S. Ochmann, R. Vock, R. Wessel, R. Klein, Automatic reconstruction of parametric building models from indoor point clouds, Comput Graph 54 (2016) 94–103.

[14] R. Ambruş, S. Claici, A. Wendt, Automatic room segmentation from unstructured 3-D data of indoor environments, IEEE RA-L 2 (2) (2017) 749–756.

[15] S. Oesau, F. Lafarge, P. Alliez, Indoor scene reconstruction using feature sensitive primitive extraction and graph-cut, ISPRS J Photogramm 90 (2014) 68–82.

[16] C. Mura, O. Mattausch, A. Villanueva, E. Gobbetti, R. Pajarola, Automatic room detection and reconstruction in cluttered indoor environments with complex room layouts, Comput Graph 44 (2014) 20–32.

[17] Z. Liu, G. von Wichert, A generalizable knowledge framework for semantic indoor mapping based on Markov logic networks and data driven MCMC, Future Gener Comp Sy 36 (2014) 42–56.

[18] M. Luperto, F. Amigoni, Predicting the layout of partially observed rooms from grid maps, in: Proc. ICRA, 2019, pp. 6898–6904.

[19] D. Perea Ström, I. Bogoslavskyi, C. Stachniss, Robust exploration and homing for autonomous robots, Robot Auton Syst 90 (2017) 125 – 135.

[20] A. Smith, G. Hollinger, Distributed inference-based multi-robot exploration, Auton Robot 42 (8) (2018) 1651–1668.

[21] R. Shrestha, F. Tian, W. Feng, P. Tan, R. Vaughan, Learned map prediction for enhanced mobile robot exploration, in: Proc. ICRA, 2019, pp. 1197–1204.

[22] M. Luperto, F. Amigoni, Exploiting structural properties of buildings towards general semantic mapping systems, in: Proc. IAS-13, 2014, pp. 375–387.

[23] J. Caley, N. Lawrance, G. Hollinger, Deep learning of structured environments for robot search, in: Proc. IROS, 2016, pp. 3987–3992.

[24] S. Ramakrishnan, Z. Al-Halah, K. Grauman, Occupancy anticipation for efficient exploration and navigation, in: Proc. ECCV, 2020, pp. 400–418.

[25] A. Elhafsi, B. Ivanovic, L. Janson, M. Pavone, Map-predictive motion planning in unknown environments, in: Proc. ICRA, 2020, pp. 8552–8558. doi:10.1109/ICRA40945.2020.9197522.

[26] J. Chang, G. Lee, Y. Lu, C. Hu, P-SLAM: Simultaneous localization and mapping with environmental-structure prediction, IEEE T Robot 23 (2) (2007) 281–293.

[27] A. Pronobis, P. Jensfelt, Large-scale semantic mapping and reasoning with heterogeneous modalities, in: Proc. ICRA, 2012, pp. 3515–3522.

[28] K. Zheng, A. Pronobis, R. Rao, Learning graph-structured sum-product networks for probabilistic semantic maps, in: Proc. AAAI, 2018, pp. 4547–4555.

[29] A. Aydemir, P. Jensfelt, J. Folkesson, What can we learn from 38,000 rooms? Reasoning about unexplored space in indoor environments, in: Proc. IROS, 2012, pp. 4675–4682.

[30] M. Luperto, F. Amigoni, Predicting the global structure of indoor environments: A constructive machine learning approach, Auton Robot 43 (4) (2019) 813–835.

[31] J. Canny, A computational approach to edge detection, IEEE T Pattern Anal (6) (1986) 679–698.

[32] N. Kiryati, Y. Eldar, A. M. Bruckstein, A probabilistic Hough transform, Pattern Recogn 24 (4) (1991) 303–316.

[33] D. Comaniciu, P. Meer, Mean shift: A robust approach toward feature space analysis, IEEE T Pattern Anal 24 (5) (2002) 603–619.

[34] T. Kucner, M. Luperto, S. Lowry, M. Magnusson, A. Lilienthal, Robust frequency-based structure extraction, in: Proc. ICRA, 2021.

[35] G. Grisetti, C. Stachniss, W. Burgard, Improved techniques for grid mapping with Rao-Blackwellized particle filters, IEEE T Robot 23 (2007) 34–46.

[36] N. Hawes, C. Burbridge, F. Jovan, L. Kunze, B. Lacerda, L. Mudrova, J. Young, J. Wyatt, D. Hebesberger, T. Kortner, et al., The Strands project: Long-term autonomy in everyday environments, IEEE RAM 24 (3) (2017) 146–156.

[37] H. Choset, Coverage for robotics: A survey of recent results, Ann Math Artif Intell 31 (1-4) (2001) 113–126.

[38] S. Oßwald, M. Bennewitz, W. Burgard, C. Stachniss, Speeding-up robot exploration by exploiting background information, IEEE RA-L 1 (2) (2016) 716–723.

[39] N. Christofides, Worst-case analysis of a new heuristic for the travelling salesman problem, Tech. rep., Carnegie-Mellon University, Management Sciences Research Group (1976).

[40] F. Glover, Tabu search - part I, ORSA Journal on Computing 1 (3) (1989) 190–206.

[41] F. Glover, Tabu search - part II, ORSA Journal on Computing 2 (1) (1990) 4–32.