A Constructive Machine Learning Approach for Robot Exploration and Search

Matteo Luperto and Francesco Amigoni[1]

***Abstract*— *Semantic knowledge*, more specifically *semantic maps* that associate semantic concepts (labels like 'room' and 'corridor') to spatial entities, has been employed to improve the performance of (multi-)robot planning tasks, such as search and exploration. However, although current semantic mapping approaches are very effective in labeling the parts of environments already visited by the robots, they are usually unable to predict the labels and, more generally, the structure of *unvisited* parts of environments. In this contribution, following a Constructive Machine Learning (CML) approach, we discuss the use of a generative method that is able to model and predict the topological structure and the labels of rooms for an indoor, previously unknown (or partially observed) environment. While this approach is not always able to find a perfect prediction of the structure of a given unknown environment, it seems nevertheless able to capture some fundamental structural properties. We explicitly note that the purpose of this paper is not to show any definitive results (although we provide a detailed example), but advocate the potential of using high-level semantic knowledge to predict the structure of unknown parts of indoor buildings in order to improve exploration and search.**

I. INTRODUCTION

In the last years, it has been shown how the use of *semantic knowledge* of environments could improve the performance for several (multi-)robot planning tasks, such as exploration and search [1]. This semantic knowledge takes usually the form of a *semantic map* of the environment, that associates semantic labels to spatial portions of environments [2]–[4]. For example, some areas can be labeled as 'rooms', while other areas could be labeled as 'corridors', performing *place classification* [4]. For instance, [1] proposes a method for speeding up multi-robot exploration using semantic knowledge. Most of the methods for place classification follow an approach that starts from the data perceived by the sensors mounted on-board mobile robots (e.g., laser range scanners and cameras), extracts some features from these data, and classifies the area from which the data have been acquired using (supervised) machine learning techniques. This approach has been shown to be very effective in labeling parts of environments already visited by the robots, but usually does not address the problem on inferring new knowledge on the labels and, more generally, the structure of *unvisited* parts of environments (except from some remarkable examples, as [5]), which are considered as completely unknown.

More recently, techniques for improving performance in exploration and search that use semantic knowledge not only relative to the visited portions of the environment, but also on the unvisited ones have been proposed. In [6], it is shown how search methods could be improved using a probabilistic model of the search environment able to perform place classification and make local predictions on the locally connected unexplored spaces as well. In [7], it is shown how a correct prediction of the labeling of the unexplored parts of an environment could improve the exploration performance of a team of robots.

[1]Matteo Luperto and Francesco Amigoni are with the Artificial Intelligence and Robotics Laboratory, Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133 Milano, Italy {matteo.luperto,francesco.amigoni}@polimi.it

II. OUR CML APPROACH

While (approximate) knowledge on the unexplored space is useful for many online planning applications, it is still unclear how such knowledge could be obtained. An attempt is that of [5], [6] that show a method that could perform *local* predictions, i.e., probabilistic predictions of the labels of the unvisited rooms *directly connected* to the places already visited (and semantically labeled) by a robot.

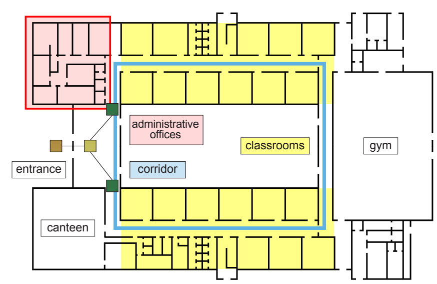

(a) Floor plan of a real school.

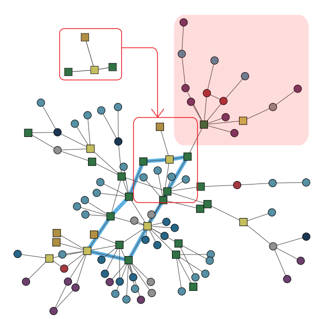

(b) Predicted topological structure.

Fig. 1: Example of a prediction of the structure of a real school with our method presented in [8].

In this paper, we propose a general framework for obtaining semantic knowledge on the *global* structure of a previously unknown (or partially visited) indoor building, as

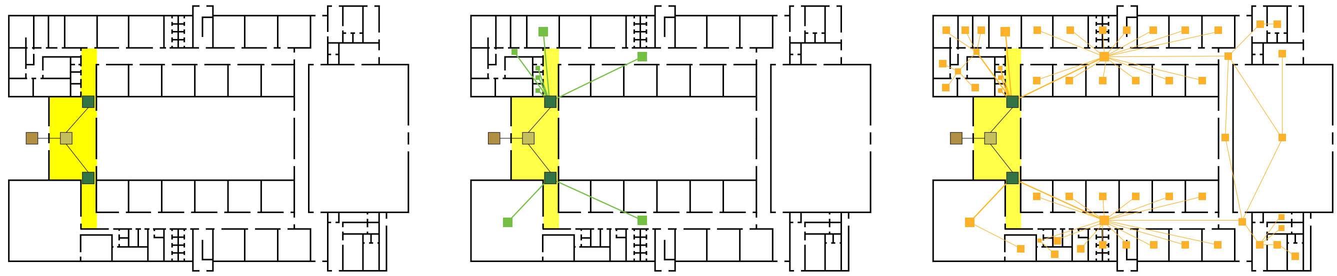

(a) Explored/known part, in yellow.

(b) Only local prediction, in green.

(c) Global prediction, in orange.

Fig. 2: Example of a prediction of the structure of a real school, under the assumption of making perfect predictions.

similarly done in [9]. For this purpose, we propose to use a generative method, that we present in detail in [8], that is able to model the topological structures and the labeling schemas of buildings and to predict the structure and the schema for an hypothetical new building. While in [8] we abstractly describe the generative method in all of its steps, in this paper we want to discuss one of its possible applications. Our generative model follows the general pattern of Constructive Machine Learning (CML), where the ultimate goal of learning is not to find good models of the data but instead to find one or more particular instances of the domain which are likely to exhibit desired properties, while selectively sampling from an infinite or exponentially large instance space. A (floor of a) building is represented as an undirected graph, whose nodes represent labeled rooms and edges represent direct physical connections between two rooms. Our approach segments graphs representing some buildings for finding significant subgraphs, which are then clustered according to their similarity. Finally, a graph representing a new building or an unseen part of a building is generated by sampling subgraphs from clusters and connecting them. The generated building shares the same structure and semantic labeling of the original buildings. This generative method, similarly to [9], can be thought as a move from a room-level perspective, modeling the semantic relations between perceived features and rooms, to a building-level perspective, modeling the connections between rooms. Our approach distinctively focuses on *building types*, namely on specific classes of buildings that share the same functions and, consequently, the same structures, like school or office buildings [10], [11].

This generative method could be used to predict a tentative structure and a semantic labeling schema of the unknown parts of an indoor building that could be exploited in order to speed up search and exploration tasks. We now discuss an example of the possible improvements in planning tasks that could be achieved using our method as source of knowledge to predict a possible configuration of the unvisited environment. Note that, although the predicted configuration might not correspond to that of any real building, it could exhibit relevant features of the real buildings, thus providing some *global semantic knowledge* of the unknown space.

III. EXAMPLE OF APPLICATION

In this section we do not present any definitive experimental result, but we provide an example of the possible use of our framework. In Fig. 1, we assume to have a robot that is exploring an unknown school building and that has already discovered the structure and the room's labels of the entrance, which is the 4-edge graph reported in Fig. 1a (with the same notation of [8]). Given this initial knowledge, and using our method, the robot gets a prediction of the global structure the school, thus obtaining new semantic knowledge (Fig. 1b, where the known part is enclosed in a red box on top left). This generated semantic knowledge is represented as a graph, where every node is a room and an edge is placed wherever two rooms are connected. The color of each node indicates the label of the room. The relevant features of a real school building are present in the predicted semantic structure of Fig. 1b, and are highlighted in both images: a loop of corridors constituting the skeleton of the building (light blue) and the administrative section (red overlay). Of course, our method is not able to perfectly predict the structure of a partially visited real environment, but both figures show roughly the same number of rooms and the same distribution of labels. Network properties of the sampled building (such as graph centrality, degree distribution, and length of the shortest path between two nodes) are consistent with the corresponding values of the real building.

Fig. 2 shows the amount of knowledge available for planning, when we consider the semantic knowledge obtained from the explored parts of the building (Fig. 2a), a local prediction obtained with a method like [5] (Fig. 2b), and the global prediction obtained with our approach (Fig. 2c). For simplicity, for both local and global prediction, we suppose to have a perfect prediction on the unknown part of the environment. From the initial knowledge of visited rooms, a method able to perform local prediction, such the one used in [5], could infer the presence of the rooms and of the corresponding labels as in Fig. 2b, in green. A method like ours, performing a global prediction, could infer the presence of the rooms and of the corresponding labels for all the environment, as shown in Fig. 2c, in orange.

References

[1] C. Stachniss, O. Mozos, and W. Burgard, "Speeding-up multi-robot exploration by considering semantic place information," in *Proc. ICRA*, 2006, pp. 1692–1697.

[2] S. Hemachandra, M. Walter, S. Tellex, and S. Teller, "Learning spatial-semantic representations from natural language descriptions and scene classifications," in *Proc. ICRA*, 2014, pp. 2623–2630.

[3] O. Mozos, C. Stachniss, and W. Burgard, "Supervised learning of places from range data using AdaBoost," in *Proc. ICRA*, 2005, pp. 1730–1735.

[4] A. Pronobis, O. Mozos, B. Caputo, and P. Jensfelt, "Multi-modal semantic place classification," *Int J Robot Res*, vol. 29, no. 2-3, pp. 298–320, 2010.

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

[6] A. Aydemir, A. Pronobis, M. Gobelbecker, and P. Jensfelt, "Active visual object search in unknown environments using uncertain semantics," *IEEE T Robotic Autom*, vol. 29, no. 4, pp. 986–1002, 2013.

[7] A. Quattrini Li, R. Cipolleschi, M. Giusto, and F. Amigoni, "A semantically-informed multirobot system for exploration of relevant areas in search and rescue settings," *Auton Robot*, 2015, to appear.

[8] M. Luperto and F. Amigoni, "A generative spectral model for semantic mapping of buildings," in *Proc. IROS*, 2015.

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

[10] M. Luperto, A. Quattrini Li, and F. Amigoni, "A system for building semantic maps of indoor environments exploiting the concept of building typology," in *Proc. RoboCup*, 2013.

[11] M. Luperto and F. Amigoni, "Exploiting structural properties of buildings towards general semantic mapping systems," in *Proc. IAS-13*, 2014.