A System for Building Semantic Maps
of Indoor Environments
Exploiting the Concept of Building Typology

Matteo Luperto, Alberto Quattrini Li, and Francesco Amigoni

Politecnico di Milano, Milano, Italy
{matteo.luperto,alberto.quattrini,
francesco.amigoni}@polimi.it

Abstract. Semantic mapping of indoor environments refers to the task of building representations of these environments that associate spatial concepts with spatial entities. In particular, semantic labels, like 'rooms' and 'corridors' are associated to portions of an underlying metric map, to allow robots or humans to exploit this additional knowledge. Usually, the classifiers that build semantic maps process data coming from laser range scanners and cameras and do not consider the specific type of the mapped building. However, in architecture it is well known that each building has a specific typology. The concept of *building typology* denotes the set of buildings that have the same function (e.g., being a school building) and that share the same structural features. In this paper, we exploit the concept of building typology to build semantic maps of indoor environments. The proposed system uses only data from laser range scanners and creates a specific classifier for each building typology, showing good classification accuracy.

Keywords: semantic mapping, building typology, line segment maps.

1 Introduction

A *semantic map* of an environment provides human-level knowledge about its structure, for example about the type of the rooms. Automatic building of semantic maps by means of mobile robots has received significant attention in the last years (e.g., [1, 2]). Usually semantic maps are built from metric maps, which represent the physical structure of environments (e.g., the locations of obstacles), and by labeling specific areas (e.g., as 'kitchen' or as 'bathroom'). The knowledge embedded in a semantic map could be exploited in several robotic tasks; for example in a domestic scenario, a semantic map could be used by a robot to find an object or to reason about what actions can be performed in a specific room (e.g., kitchen, bedroom). The labeling is often performed by a classifier that has been previously trained and that is fed with data coming from a metric map and from robot sensors. To the best of our knowledge, no approach attempts to exploit information on the typology of the buildings in classifying places of indoor environments.

In this paper, we contribute in this direction by proposing a semantic mapping system for indoor environments that considers the concept of building typology. The *building*

typology, as studied in architecture [3], denotes a set of buildings that have the same function (e.g., being a school or an office building) and that share the same structural features (e.g., the fact that classrooms are not directly connected to gymnasiums). The system proposed in this paper exploits these common features to improve the classification of places of buildings with specific typologies.

More precisely, in our approach we create and train classifiers specific for each building typology. These classifiers are used for labeling rooms of indoor environments with five semantic labels: 'small room', 'medium room', 'big room', 'corridor', 'hall'. The input data for the classifiers are a number of features that characterize a room and that are extracted from a metric map represented as a set of line segments. Such a metric map could be easily built by mobile robots using laser range scanners.

The main original contribution of this paper is to provide an initial study and evaluation of the intuition of considering building typology for semantic labeling of rooms. In particular, we propose to exploit *a priori* knowledge on a building for performing a sort of "informed" semantic mapping that uses classifiers specific for each building typology. Our experimental results show that these classifiers present a better classification accuracy than a generic classifier that does not refer to any building typology. Moreover, they show performance comparable with that of state-of-the-art classifiers that require much richer input data (for example, coming from cameras) and that use a smaller number of semantic labels.

2   Related Work

Semantic mapping refers to the "process of building a representation of the environment which associates spatial concepts with spatial entities" [4]. A semantic map usually associates semantic labels to places and objects present in an underlying metric map. In this paper, we are interested in labeling portions of a metric map with the corresponding tag, like 'corridor' or 'room'. This problem is known as *place classification* and can be further specialized as place recognition, when the classifier is used in the same environment in which it was trained, and place categorization, when the classifier is used in previously unseen environments [4]. The system presented in this paper performs place categorization.

Approaches to classification of places can be organized according to the sensor modalities they employ, basically laser range scanners and cameras. Since the system we propose in this paper relies on data that are obtained from laser range scanners, in the following we analyze in detail the systems belonging to this class.

The system in [1] extracts a topological map of the environment, which is a graph where every node represents a room and every edge represents a passage connecting two rooms. This method is based on partitioning the environment into a set of open spaces connected by narrow passages. The partitioning is done using a fuzzy grid map of the emptiness of the environment, from which the rooms (assumed to be rectangular) are extracted using a fuzzy mathematical morphological technique.

In [5], a topological map of the environment, where edges have semantic values ('corridor', 'room', 'door'), is used by a behavior-based robot. However, the construction of the map is not illustrated in detail.

The system of [6] classifies a single laser range scan as belonging to a 'room', to a 'corridor', or to a 'hallway'. The classifier is based on AdaBoost and considers both features that are extracted from raw sensor data (e.g., average beam length) and features that are calculated after approximating raw data with a polygon (e.g., area of the approximating polygon). Experimental results show that the system is able to correctly classify scans taken in environments different of those used for learning the classifier. A similar approach, which uses also features extracted from cameras, is presented in [7].

Authors in [8] describe a method for integrating the robot map with a human representation of the environment. This goal is reached using a semantic map based on data obtained from a laser range scanner and using the concept of region and locations, statistically derived from sensor data. Every room is characterized by features, such as its area, its major and minor axes, and its excentricity. These features are used to affix a semantic label ('small', 'medium size', or 'large') to rooms.

In [9], a topological map of the environment is extracted using Voronoi Random Fields (VRF). To build this map, a Voronoi graph is extracted from an occupancy grid map. Each point of the Voronoi graph is a node of a conditional random field, and the resulting VRF estimates the label ('room', 'doorway', 'hallway', 'junction', 'other') of each node using features from both the grid map and the Voronoi graph.

The system proposed in [10] extracts semantic information from 2D and 3D maps of outdoor environments, using Hidden Markov Models (HMM) and Support Vector Machines (SVM) classifiers, and labels terrains as 'traversable' or 'non-traversable' and parts of urban environments as 'street' or as 'sidewalk'.

Basically, all these approaches, which use only data coming from laser range scanners, have been shown to have good performance on the same data set used for training and with a small number of labels (often two or three).

Semantic labeling of places has been also addressed using visual data obtained from cameras. Some works from computer vision make exclusive use of images (e.g., [11]). In robotics, a combination of camera and laser range scanner is usually employed to improve performance of place classification. For example, the system in [2] uses SVMs to combine multiple visual cues and laser range data. Another example is [12], which integrates a line segment-based metric map (built using data from laser range scanners) with higher-level maps that contain information about places and objects. These maps are built on the basis of the metric map, of the perceived images, and of an ontology representing knowledge about indoor environments (e.g., the fact that fridges are usually located in kitchens).

The system we propose in this paper exploits a metric map composed of a set of line segments, which can be easily built from points acquired by laser range scanners. We do not use any visual information. However, we exploit the *a priori* knowledge about the building typology of the indoor environment being mapped. To the best of our knowledge, no other approach attempts to distinguish between different building typologies. Assuming that this background information is available, we show that it can improve place classification.

3   The Proposed System

3.1   Building Typology

A building is an artifact created for a specific function, its purpose, because it is built by people for people who inhabit it. This simple observation is often neglected when designing mobile robots that operate in indoor environments. Indeed, these environments are usually treated as natural, fixed, immutable entities, rather than as a cultural product, the result and summary of centuries of social evolution, as the modern buildings are. As a consequence, robotics designers consider the environments in which the robots move as structured, without fully exploiting the implications of their structure.

The function of a building imposes its structure, its floor plan, and the structure of its rooms. Each building, having a precise function, shares some structural features with all other buildings with the same purpose. A *building typology* is a set of buildings that have the same function [3]. A building typology can be associated to a model that represents the structural features shared by buildings belonging to the typology.

Main features of this model are well known (as cultural facts) to humans, who use them for localization and orientation, as largely studied in architecture [13]. For example, office buildings are structured in the following way: at the ground floor, there are the rooms for public relations and social spaces, and possibly a canteen and conference rooms; at the upper floors, there are office rooms for the back-end activities.

In this paper, we propose a system that attempts to formalize and build the model of a building typology, in order to provide it to a mobile robot. We consider three building typologies: *House* (residential buildings and houses), *School* (school buildings), and *Office* (office buildings and open spaces). These three building typologies have been chosen because they are particularly significant in the real world and because they are usually considered for experimental activities of autonomous mobile robots. Note that, when comparing two school buildings, they are, at an initial sight, different from each other. However, they share a common model that, for example, represents the fact that rooms are typically connected in certain ways. Recognizing this common model helps humans in acting properly in schools they never entered before.

Finally, it is worth noting that the concept of building typology is debated in architecture and is considered part of an analytical approach to architecture that has produced a huge amount of data on structures of buildings [14].

3.2   Construction of the Classifiers

Our system semantically labels portions of a metric map using different classifiers that are trained on data belonging to different building typologies. In this way, by knowing the typology of the building where a robot operates, the correct classifier can be selected. Note that having different classifiers for different building typologies allows to have explicit models of the specific building typologies, which could be used for different applications. In this section, we illustrate how we construct and train the classifiers for place classification.

We represent indoor environments by a line segment-based metric map, built using data from laser range scanners, as shown in the next section. From this metric map, our

system derives a topological map, represented as a graph. A node represents a room. Doorways are associated to edges that connect two nodes (rooms). A description of the method for the extraction of the topological map from the metric map, including doorway identification, can be found in the next section. Each room is associated to a *semantic label*. We consider the following set of semantic labels: *'small room'*, *'medium room'*, *'big room'*, *'corridor'*, *'hall'*. They have been chosen to be simple enough to be immediately understandable by humans, but sufficiently descriptive of types of environments operated by robots. The number of labels we use is higher than those typically used in state-of-the-art semantic mapping approaches using only laser range scanners (e.g., [6] uses 'corridor', 'hall', and 'room' and [8] uses 'small', 'medium size', and 'large' room).

To associate a semantic label, each room is described by a vector of features, used by classifiers. The features are chosen to capture (some of) the characteristics of the model underlying a specific building typology. The features can be divided in two groups. The first group of features captures the shape of a room and consists of the area a of the room and the axes ratio $rt = M/m$ of the major axis M and minor axis m. The details on how the latter ratio is calculated are given in the next section. Note that these features are similar to those used in [8] and constitute a subset of those used in [6]. The second group of features represents the structure of the building and the connections of the room with the rest of the environment (in particular, with adjacent rooms), as this is one of the distinguishing characteristics of a building typology. The features of this second group include the number d of doorways present in a room r and the labels of the rooms directly connected to r. More precisely, for each semantic label s (s could be S, M, B, C, H for 'small room', 'medium room', 'big room', 'corridor' and 'hall', respectively), we consider the number of doorways l_s that connect r to a room with semantic label s. If one of the adjacent rooms has not yet a semantic label, a temporary label is associated to that room with a naïve classifier that uses only the features $\langle a, rt, d \rangle$ and that has been manually trained. A room r is hence described by the feature vector $F_r = \langle a, rt, d, l_S, l_M, l_B, l_C, l_H \rangle$.

According to these features, rooms are semantically labeled by using classifiers specifically trained for each building typology. We consider four well-known classification algorithms: rule induction (RI), multi-layer perceptron (MLP), decision tree (DT), and k-nearest neighbor (K-NN) [15], in order to compare their performance. These algorithms have been chosen because they are general and not tailored for the task of classifying places. In this way, we do not introduce any bias in the evaluation of our approach.

All the classifiers are trained using a supervised learning approach starting from four data sets, that have been manually built. Each entry in these data sets consists of a vector of features F_r (describing the structure of a room r) and of a corresponding semantic label. Three data sets refer to specific building typologies: D_H refers to House, D_S refers to School, D_O refers to Office. The idea is that the classifiers trained using one of these data sets, called *typology classifiers*, embed an implicit model of the corresponding building typology. The fourth data set D_{H+S+O} is composed of all the elements in D_H, D_S, and D_O and does not refer to any specific building typology, as commonly assumed in literature. The classifier trained with D_{H+S+O} is called *standard classifier*.

The data sets have been built as follows. We selected the floor plans of dozens of buildings belonging to the typologies House, School, and Office from eleven monographic books used for the design and analysis of buildings in architecture (for example, [16] is one of the books used for the School building typology). In these books, rooms are labeled according to their function and these labels are assigned by the architects who designed the buildings. The labeled floor plans have been digitally represented using a custom CAD-like software program that allows, for each floor plan, to identify rooms and to semantically label them, as shown, for example, in Fig. 1.

Fig. 1. An example of a floor plan of a house made with AutoCAD (left) and with our CAD-like software (right). Colors indicate semantic labels (green is corridor, yellow is small room, red is big room, brown is medium room).

All the above steps have been manually performed. Then, each labeled floor plan is fed into another software program that automatically extracts the features F_r for each room r. These features F_r and the corresponding label for r are used to create entries to populate the data sets D_H, D_S, and D_O. In the following table, some entries of data set D_H with features and semantic labels are shown.

Label	a	d	rt	l_S	l_M	l_B	l_C	l_H
S	3.0	1	1.0	0	0	0	1	0
B	44.0	2	1.4	1	1	0	0	0
C	13.0	8	3.0	0	1	5	2	0
M	15.0	1	1.2	1	0	0	0	0
C	9.0	6	2.0	0	0	4	1	1

These three building typology-specific data sets are composed of approximately 2,000 rooms (entries) each. Hence, D_{H+S+O} contains around 6,000 rooms.

In summary, we have 16 different classifiers (12 typology classifiers and 4 standard classifiers) obtained from all the possible combinations of algorithms (RI, MLP, DT, and K-NN) and of data sets (D_H, D_S, D_O, and D_{H+S+O}) used for training. These classifiers, given a vector of features as input, return a semantic label.

3.3   Use of the Classifiers

The classifiers are used to associate a semantic label to rooms detected by a robot moving in the environment. In this section, we illustrate how the vector of features for a room to be classified is constructed from sensor data.

We assume to have a mobile robot equipped with two laser range scanners that are mounted back-to-back so that they cover an area of $360°$ around the robot. The point-based scans acquired by the sensors are approximated with line segments using the

split-and-merge algorithm [17]. The set of line segments representing a scan is then merged in the metric map, starting from an initial guess provided by odometry, by using a simple scan matching approach similar to that of [18]. The free space already explored can be represented by a polygon (possibly with holes) whose edges are either line segments representing obstacles or line segments representing frontiers, namely the boundaries between the known and the unknown portions of the environment. We use line segments instead of more common occupancy grids to have a more compact representation of the environment.

The metric map built as illustrated above is then used to extract the vectors of features. Initially, we eliminate short line segments (in our experiments, shorter than 30 cm) to remove from the map line segments representing parts of furniture and small obstacles and to keep only the line segments representing the structure of the rooms. Fig. 2 shows the filtered metric map of an environment, where objects and furniture were present, as built by our mapping system. The map preserves the main walls and, thus, the structure of the environment.
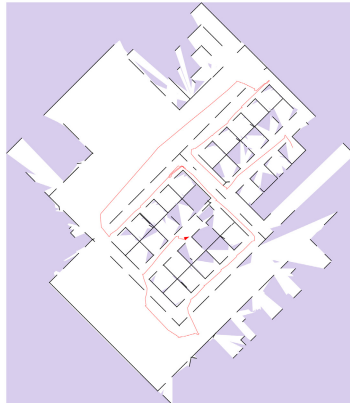
Fig. 2. Line segment map built by our system (white area: known free space, gray area: unknown space, black line segments: obstacles, red line: path of the robot)

Then, the filtered metric map is partitioned in rooms in the following way. Walls are identified by finding out collinear line segments. Two line segments are considered collinear if their angular coefficients are equal, up to a threshold value (10 degrees). Doorways are identified as passages in a wall, which connect rooms. In particular, a doorway is found by identifying the gaps between the collinear line segments of a wall or between the endpoints of a wall and another wall. If the gaps have the size of a doorway (e.g., more than 30 cm and less than 150 cm), they are recognized as such. Note that, although more complex situations could be handled, we consider static environments, in which doors can be either open or closed (as many semantic mapping systems do, with few notable exceptions [19]). Portions of area delimited by walls and doorways are identified as rooms using a Monte Carlo method. A random sample of points

is thrown within the space. For each point p, the set of line segments $\{L_p\}$ visible (in straight line) from p is calculated. Line segments $\{L_p\}$ are tentatively considered as delimiting the room in which p is placed. The points are checked for straight line visibility between them. Mutually visible points p, p' belong to the same room, and $\{L_p\}$ and $\{L_{p'}\}$ are merged accordingly (see Fig. 3).
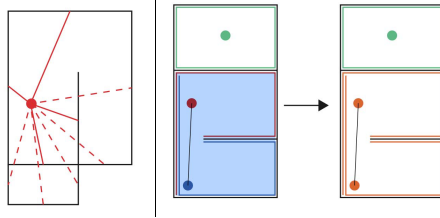
Fig. 3. On the left, line segments directly visible from the point (red lines) and not directly visible from the point (dotted lines). On the right, two mutually visible points in the same room.

A graph representing the topological map is then created by considering each room as a node and connecting nodes through edges that represent doorways. Finally, for each identified room r, the corresponding vector of features F_r is calculated. In particular, the area a is calculated as the average of the area of the inner and outer quadrilaterals that best approximate walls of the room. The major axis M is calculated as the longest wall of the room. The minor axis m is calculated as the longest wall of the room perpendicular to M. The rest of the features, representing the connection graph, are directly derived from the graph. Semantically unlabeled rooms in the graph are temporarily labeled using a naïve rule classifier, as described in the previous section.

Each feature vector F_r becomes an input for the classifiers, which associate a semantic label to the corresponding room r; the result of this last step is a topological map with semantically labeled nodes, namely a semantic map.

4   Experimental Activity

4.1   Experimental Setting

The implementation of our semantic mapping system is divided in two parts. The first one is responsible for the extraction of the features from the metric map. This part has been developed in C++ and tested in environments simulated in USARSim [20] using a Pioneer P3AT robot equipped with two laser range scanners, as described before. Results are satisfactory, but are not shown here because this paper focuses on the classifiers. The second part refers to the training and evaluation of the classifiers. The classifiers have been trained using algorithms implemented in OpenCV [21] and Rapid-Miner [22]. In particular, the training phase and the testing phase of the classifiers are carried out on the same data set, using 10-fold cross-validation. The results of the experimental evaluation of the classifiers are presented in the following. We measured the
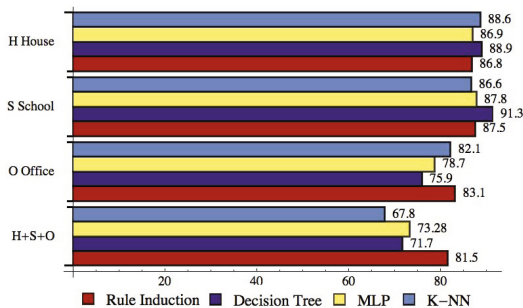
accuracy of classification comparing typology classifiers with standard classifiers. The accuracy of classification is defined, according to [2, 6, 8], as the percentage of rooms that have been properly classified with the right semantic label.

Our data sets D_H, D_S, D_O, and D_{H+S+O} are composed of simulated data. This allowed us to use data sets with thousands of rooms, an amount of data larger than those typically used in state-of-the-art approaches, usually limited to few buildings. Available real data sets (like those in Radish [23]) do not cover all the building typologies we consider (e.g., houses are largely underrepresented). With our data sets, we preliminarily assess the usefulness of the concept of building typology for semantic mapping by evaluating it under ideal conditions, without measurement and mapping errors.

4.2   Experimental Results

Fig. 4 shows the accuracy of classification of the 12 typology classifiers compared with the classification accuracy of the 4 standard classifiers and Table 1 shows the confusion matrices of each building type. The average (over algorithms) accuracy of the typology classifiers is 87.8% (4.9 is the standard deviation) for building typology House, 88.3% (4.8) for building typology School, and 81.6% (4.6) for building typology Office. The lower performance on the Office building typology with respect to classifiers trained and tested on the other two building typologies may be due to a further division that can be made internally in this category, namely between small offices, located in mixed office/residential buildings, and large open spaces. By contrast, standard classifiers have a lower average accuracy of 73.6% (4.7). These results seem to suggest the effectiveness of using building typology in semantic mapping. The performance of the standard classifiers is rather good because they exploit the several rooms labeled as S and C, whose features (large number of doors and small area) are similar in all the building typologies. Analyzing the algorithms, it can be seen that, for all building typologies, their results are similar, with the RI algorithm performing slightly better than the others over all the typologies.

A statistical analysis of some of the features of the rooms in the data sets provides an explanation and a motivation for the above results. In this regard, mean and standard

Fig. 4. Accuracy of classification (average and standard deviation) of the typology and standard classifiers

Table 1. Confusion matrices of the typology and the standard classifiers (averaged over the algorithms), where all numbers are in percentage, rows are true labels, columns are predicted labels (C 'corridor'; H 'hall'; S 'small room'; M 'medium room'; B 'big room'), TP are true positives, FP false positives, TN true negatives, and FN false negatives

H House

Label	C	H	S	M	B	TP	FP
C	92.3	1.2	1.7	3.5	1.3	92.3	7.7
H	3.0	88.0	0.0	0.7	8.2	88.0	12.0
S	3.8	0.3	91.4	4.5	0.0	91.4	8.5
M	3.7	0.4	8.1	83.4	4.3	83.4	16.6
B	1.0	7.1	0.0	5.4	86.5	86.5	13.5
TN	75.7	98.5	96.3	96.0	97.5		
FN	24.2	1.5	3.7	3.9	2.5		

S School

Label	C	H	S	M	B	TP	FP
C	92.0	3.3	4.2	0.3	0.1	92.0	8.0
H	7.4	80.0	1.3	5.2	6.1	80.0	20.0
S	3.1	0.3	94.0	2.6	0.0	94.0	6.0
M	0.6	3.5	6.5	85.5	3.9	85.5	14.5
B	0.0	21.7	2.9	18.8	56.5	56.5	43.5
TN	97.4	98.2	42.4	97.4	98.8		
FN	2.6	1.8	57.6	2.6	1.2		

O Office

Label	C	H	S	M	B	TP	FP
C	91.9	3.6	0.2	2.6	1.7	91.9	8.1
H	7.8	80.9	0.0	6.5	4.8	80.9	19.1
S	2.0	0.0	92.3	5.4	0.3	92.3	7.7
M	2.6	1.2	5.6	70.9	19.6	70.9	29.1
B	3.6	5.9	0.1	25.3	65.1	65.1	34.8
TN	79.3	98.1	96.8	92.4	90.1		
FN	20.7	1.9	3.2	7.6	9.9		

H+S+O

Label	C	H	S	M	B	TP	FP
C	85.7	6.9	2.1	2.8	2.4	85.7	14.2
H	7.9	75.8	1.0	7.1	8.0	75.8	24.1
S	4.0	0.2	75.4	17.6	2.7	75.4	24.6
M	2.3	3.0	11.9	63.3	19.6	63.3	36.7
B	4.6	9.3	3.6	15.8	66.7	66.7	33.3
TN	96.4	96.9	64.7	87.9	91.1		
FN	3.6	3.1	35.3	12.1	8.9		

Table 2. Characteristics of the rooms in the data sets, where % is the percentage of semantic labels (C 'corridor'; H 'hall'; S 'small room'; M 'medium room'; B 'big room') present in the data set and μ and σ are the mean and the standard deviation of the corresponding feature, respectively

H House

Label	%	Area a		Doorways d		rt=M/m	
		μ	σ	μ	σ	μ	σ
C	23.9	14.0	11.4	4.0	1.8	3.4	3.6
H	8.0	61.5	25.1	3.3	1.5	1.8	0.8
S	28.0	6.0	2.6	1.0	0.2	1.6	0.6
M	26.3	17.4	4.8	1.3	0.4	2.0	1.0
B	13.7	33.7	7.8	1.8	1.0	1.5	0.4

S School

Label	%	Area a		Doorways d		rt=M/m	
		μ	σ	μ	σ	μ	σ
C	15.4	36.1	42.2	5.7	3.8	7.4	7.6
H	5.6	205.1	172.9	6.9	4.5	2.0	1.3
S	55.6	10.0	9.7	1.1	0.4	1.6	0.9
M	21.3	54.8	14.8	1.6	0.8	1.3	0.4
B	2.0	102.9	22.3	2.1	0.7	1.7	0.7

O Office

Label	%	Area a		Doorways d		rt=M/m	
		μ	σ	μ	σ	μ	σ
C	20.6	30.5	25.7	6.7	4.8	8.1	7.3
H	6.6	106.9	84.5	6.5	3.6	2.4	1.6
S	23.2	4.5	2.3	1.0	0.1	1.3	0.7
M	34.5	18.5	14.4	1.2	0.3	1.9	0.8
B	14.8	31.3	12.2	1.4	0.9	1.9	1.0

H+S+O

Label	%	Area a		Doorways d		rt=M/m	
		μ	σ	μ	σ	μ	σ
C	20.3	26.4	27.9	5.7	4.1	6.5	6.8
H	6.8	110.9	109.2	3.7	3.7	2.2	1.3
S	32.3	7.1	6.9	1.0	0.4	1.5	0.8
M	29.3	24.5	18.9	1.3	0.6	1.8	0.8
B	11.5	35.0	18.0	1.6	0.9	1.8	0.9

deviation of some of these features have been calculated. Table 2 shows this analysis. Some of the characteristics of the rooms of the same type are rather constant within a building typology. For example, for houses, all the rooms (except corridors and halls) have a standard deviation of the number of doorways not larger than one. The same happens for other building typologies. Rooms with the same label belonging to buildings of different typologies are consistent. For example, a medium room in a house has an area a, on average, of 17.4 square meters, while a medium room in a school building has an average area a of 54.8 square meters. Moreover, a corridor in a house has an average axes ratio ($rt = M/m$) of 3.4, while a corridor in a school or in an office has a ratio of 7.4 and 8.1, respectively. The distinctive characteristics of each building typology are effectively exploited by the typology classifiers.

To further assess the use of building typology, we classified rooms belonging to school buildings (D_S) using House typology classifiers (M_H) and vice versa (D_H/M_S).

As shown in the following table, the classification accuracy is low. This implies that the models embedded in the typology classifiers are actually specific to the corresponding building typology. Classic semantic mapping approaches, not using building typologies, are implicitly developed for a unique typology (e.g., university campus). Their use for a different typology (e.g., houses) may result in a deterioration of their performance by 50%, similarly to our results.

Algorithm	D_H/M_S	D_S/M_H
RI	40.0(3.2)	50.7(2.5)
MLP	35.7(1.1)	56.1(0.5)
DT	43.0(1.8)	53.0(2.2)
K-NN	41.2(0.4)	51.0(0.3)
average	40.0(3.25)	52.7(2.75)

Finally, we qualitatively compare the results obtained by the proposed system with those of the related literature. The use of the concept of building typology in semantic mapping seems to provide better performance compared to systems that use the same type of sensor (i.e., a laser range scanner) as the only source of information, like [2, 8]. The classification accuracy of these systems is usually in a range from 40% and 80%, and using a smaller number of semantic labels for place classification. The classification accuracy we achieve is similar to that of systems that use a camera coupled with a laser range scanner (e.g., [2, 7]). This fact can be explained by saying that knowing *a priori* the building typology (namely, by using the right typology classifier) turns out to be as significant as the use of additional sensors like the camera, but without introducing the computational burden arising from their use. In a sense, the use of information on the structure of the building enriches the data obtained by the laser range scanner.

5   Conclusions

In this paper we have presented a semantic mapping system that classifies rooms of indoor environments considering typology of buildings where a robot is operating. More precisely, assuming that a robot is moving in a building with a known typology, the proposed system employs classifiers specific for that typology to semantically label rooms identified from data acquired by laser range scanners. Experimental results show that considering building typology is effective for place classification. In particular, availability of *a priori* information about the building typology of an environment seems to have an impact similar to that of using a camera on improving classification accuracy.

Future work will address the further assessment of the proposed system, by considering other building typologies and larger data sets (also acquired by real robots). This will allow to evaluate effects of perception and mapping noise on the classifiers. Moreover, we will investigate the system performance when *a priori* knowledge is not provided, but the system itself has to infer the building typology according to the best classifier. Finally, possible uses of the system will be investigated. For example, semantic maps built by our system could be used to predict labels of rooms not yet discovered during an exploration task or to validate realism of synthetic environments used in robotic simulations, like those used in the RoboCup Rescue Simulation League.

References

1. Buschka, P., Saffiotti, A.: A virtual sensor for room detection. In: Proc. IROS, pp. 637–642 (2002)
2. Pronobis, A., Mozos, O., Caputo, B., Jensfelt, P.: Multi-modal semantic place classification. Int. J. Robot. Res. 29(2-3), 298–320 (2010)
3. Rossi, A.: The architecture of the city. MIT Press (1984)
4. Pronobis, A.: Semantic mapping with mobile robots. PhD thesis. KTH (2011)
5. Althaus, P., Christensen, H.: Behavior coordination in structured environments. Adv. Robotics 17(7), 657–674 (2003)
6. Mozos, O., Stachniss, C., Burgard, W.: Supervised learning of places from range data using AdaBoost. In: Proc. ICRA, pp. 1730–1735 (2005)
7. Mozos, O., Triebel, R., Jensfelt, P., Rottmann, A., Burgard, W.: Supervised semantic labeling of places using information extracted from sensor data. Robot. Auton. Syst. 55(5), 391–402 (2007)
8. Topp, E., Christensen, H.: Topological modelling for human augmented mapping. In: Proc. IROS, pp. 2257–2263 (2006)
9. Friedman, S., Pasula, H., Fox, D.: Voronoi random fields: Extracting the topological structure of indoor environments via place labeling. In: Proc. IJCAI, pp. 2109–2114 (2007)
10. Wolf, D., Sukhatme, G.: Semantic mapping using mobile robots. IEEE T. Robot. 24(2), 245–258 (2008)
11. Quattoni, A., Torralba, A.: Recognizing indoor scenes. In: Proc. CVPR, pp. 413–420 (2009)
12. Zender, H., Mozos, O., Jensfelt, P., Kruijff, G., Burgard, W.: Conceptual spatial representations for indoor mobile robots. Robot. Auton. Syst. 56(6), 493–502 (2008)
13. Lynch, K.: The image of the city. MIT Press (1960)
14. Neufert, E., Neufert, P.: Architects' data. Wiley-Blackwell (2012)
15. Bishop, C.: Pattern recognition and machine learning. Springer (2006)
16. Perkins, B.: Building type basics for elementary and secondary schools. Wiley (2001)
17. Nguyen, V., Gächter, S., Martinelli, A., Tomatis, N., Siegwart, R.: A comparison of line extraction algorithms using 2d range data for indoor mobile robotics. Auton. Robot. 23(2), 97–111 (2007)
18. Lu, F., Milios, E.: Robot pose estimation in unknown environments by matching 2d range scans. J. Intell. Robot. Syst. 18(3), 249–275 (1997)
19. Anguelov, D., Koller, D., Parker, E., Thrun, S.: Detecting and modeling doors with mobile robots. In: Proc. ICRA (2004)
20. Carpin, S., Lewis, M., Wang, J., Balakirsky, S., Scrapper, C.: USARSim: A robot simulator for research and education. In: Proc. ICRA, pp. 1400–1405 (2007)
21. OpenCV: OpenCV (2012), http://opencv.org/
22. RapidMiner: RapidMiner (2012), http://www.rapidminer.com/
23. Howard, A., Roy, N.: The robotics data set repository (Radish) (2003), http://radish.sourceforge.net/