

Supervised Learning of Topological Maps using Semantic Information Extracted from Range Data

Óscar Martínez Mozos Wolfram Burgard

University of Freiburg, Department of Computer Science, D-79110 Freiburg, Germany.

Email: {omartine|burgard}@informatik.uni-freiburg.de

Abstract— This paper presents an approach to create topological maps from geometric maps obtained with a mobile robot in an indoor-environment using range data. Our approach utilizes AdaBoost, a supervised learning algorithm, to classify each point of the geometric map into semantic classes. We then apply a segmentation procedure based on probabilistic relaxation labeling on the resulting classifications to eliminate errors. The topological graph is then extracted from the individual different regions and their connections. In this way, we obtain a topological map in the form of a graph, in which each node indicates a region in the environment with its corresponding semantic class (e.g., corridor, or room) and the edges indicate the connections between them. Experimental results obtained with data from different real-world environments demonstrate the effectiveness of our approach.

I. INTRODUCTION

Topological maps have been quite popular in the robotics community because they are believed to be cognitively more adequate, since they can be stored more compactly than geometric maps and can be also communicated more easily to users of a mobile robot. In the past, many researchers have considered the problem of building topological maps of the environment from the data gathered with a mobile robot. The question of how to augment such maps by semantic information, however, is virtually unexplored. Whenever robots are designed to interact with their users, semantic information about places can be important. A robot that possesses this knowledge can be instructed, for example, to go to the kitchen.

In this paper, we consider the problem of learning topological maps with semantic information from geometric maps that were obtained with a mobile robot in an indoor-environment using range data. Our approach is based on the assumption that indoor environments, like the one depicted in the left image of Figure 1, can be typically decomposed into areas with different functionalities such as rooms, corridors and doorways, and that these areas build the vertices of a topological graph. The connections of the vertices are then given by the neighborhood of the regions in the occupancy map. For example, a doorway is typically connected to two rooms, two corridors, or to a room and a corridor.

Throughout this paper we assume that the robot is given a map of the environment in the form of an occupancy grid [12]. Our approach then determines for each unoccupied cell of such a grid its semantic class. This is achieved by simulating a range scan of the robot given it is located in that particular cell, and then classifying this scan into one of the semantic classes.




Fig. 1. The left image shows a geometric map of a typical indoor environment with rooms, doorways, and a corridor depicted in colors/grey levels. The middle images show two simulated range scans in the geometric map. The right image depicts the corresponding semantic-topological map.

Examples for typical simulated range scans obtained in an office environment are shown in the middle images of Figure 1. The classification is then done using a sequence of classifiers learned with the AdaBoost algorithm [15]. These classifiers are built in a supervised fashion from simple geometric features that are extracted from range scans simulated in a previously labeled map of a similar environment. To remove noise and clutter from the resulting classifications, we apply an approach denoted as probabilistic relaxation labeling [13]. From the resulting labeling we construct a graph whose nodes correspond to the regions of identically labeled poses and whose edges represent the connections between them. Additionally, each node contains geometrical information about the region it represents, like the area, the centroid and the orientation. A typical topological map obtained with our approach is shown in the right image of Figure 1. Experimental results shown in this paper illustrate that our method can determine the semantic-topological map of an environment with a recognition rate of more than 98%. We also present results that illustrate that our approach can even construct a topological map of an environment from which no training data was available. Furthermore, we compare the extended set of simple features used in our AdaBoost algorithm with the one previously applied by Martínez Mozos *et al.* [10]. Experimental results illustrate that our features provide better classification results.

This paper is organized as follows. After discussing related work in the following section, Section III describes the semantic AdaBoost classifier. Then, Section IV introduces the new set of simple features used in the system described here. Furthermore, Section V describes the probabilistic relaxation approach. In Section VI we then explain the method used to

extract semantic regions and to create the final topological map. Finally, in Section VII we present experimental results obtained using our approach.

II. RELATED WORK

In the past, different algorithms for creating topological maps have been proposed. Kuipers and Byun [9] extract distinctive points in the map, which are defined as the local maximum of some measure of distinctiveness. Kortenkamp and Weymouth [7] fuse the information obtained with vision and ultrasound sensors to determine topologically relevant places. Shatkey and Kaelbling [16] apply a HMM learning approach to learn topological maps in which the nodes represent points in the plane. Thrun [18] uses the Voronoi diagram to find critical points, which minimize the clearance locally. Additionally, Kuipers and Beeson [8] apply different learning algorithms to calculate topological maps of environments. These former approaches only identify points in the map that have special properties but they do not include any means for extracting the types of places or even regions. Our approach, in contrast, is able to identify complete regions in the map like corridors, rooms or doorways, which have a direct relation with a human understanding of the environment.

Additionally, several authors considered the problem of identifying certain types of places. For example, Buschka and Saffiotti [2] describe a virtual sensor that is able to identify rooms from range data. Also Koenig and Simmons [6] use a pre-programmed routine to detect doorways from range data. Althaus and Christensen [1] use sonar data to detect corridors and doorways. With respect to place classification, our approach is most similar to the algorithm proposed by Martínez Mozos *et al.* [10] in that we also apply the AdaBoost algorithm. However, we use a probabilistic variant of the classifier and additionally we apply a probabilistic relaxation labeling to incorporate similarity constraints between neighboring points and to eliminate false classifications.

III. SEMANTIC CLASSIFICATION OF POSES USING ADABOOST

Boosting is a general method for creating an accurate strong classifier by combining a set of weak classifiers. The requirement to each weak classifier is that its accuracy is better than a random guessing. In this work we will use the boosting algorithm AdaBoost in its generalized form presented by Schapire and Singer [15]. The input to the algorithm is a set of labeled training examples $(x_n, y_n), n = 1, \dots, N$, where each x_n is an example and each $y_n \in \{+1, -1\}$ is a value indicating whether x_n is positive or negative respectively. In our case, the training examples are composed by simulated laser observations as described in Section IV. In several iterations $t = 1, \dots, T$, the algorithm repeatedly selects a weak classifier $h_t(x)$ using a distribution D over the training examples. The selected weak classifier is expected to have a small classification error in the training data. The idea of the algorithm is to modify the distribution D increasing the weights of the most difficult training examples in each round. The final strong classifier is a weighted majority vote of the

T best weak classifiers. Throughout this work we will use the approach by Viola and Jones [19] in which each weak classifier h_j depends on a single-valued feature $f_j \in \mathbb{R}$

$$h_j(x) = \begin{cases} +1 & \text{if } p_j f_j(x) < p_j \theta_j \\ -1 & \text{otherwise,} \end{cases} \quad (1)$$

where θ_j is a threshold and p_j is either -1 or $+1$ and thus represents the direction of the inequality. The parameters θ_j and p_j are determined during the training process of the algorithm.

The generalized Adaboost is only able to predict the label $y \in \{+1, -1\}$ of an example as positive or negative. To additionally estimate the probability of a particular label, we use the method suggested by Friedman *et al.* [3]. It uses the output of AdaBoost to determine a confidence value $C \in [0, 1]$ for a positive classification of an example $C = P(y = +1 | x)$.

The so far described method is able to distinguish between two classes of examples, namely positives and negatives. In practical applications, however, we want to distinguish between more than two classes. To create a multi-class classifier we used the approach applied by Martínez Mozos *et al.* [10], which creates a sequential multi-class classifier using $K - 1$ binary classifiers, where K is the number of classes we want to recognize. Additionally, we use the method by Stachniss *et al.* [17], in which the classification output of the decision list is represented by a histogram z . Each bin of z stores the probability that the classified example belongs to the k -th class according to the sequence

$$z^{[k]} = C_k \prod_{j=1}^{k-1} (1 - C_j), \quad (2)$$

where $C_k = P_k(y = +1 | x)$ and $C_K = 1$.

IV. LASER DATA FEATURES

As explained in the introduction, in order to classify each free cell in the occupancy grid map we simulate a range scan in its position ray-casting in the map. The simulated scans we use correspond to a robot equipped with a 360 degree field of view laser sensor. Each simulated laser observation then consists of 360 beams. Each training example for the AdaBoost algorithm consists of one laser observation and its classification. Throughout this paper we assume that the classification of the training examples is given in advance. The single-valued features used in the Adaboost algorithm are geometrical features used for shape analysis. The features must be rotationally invariant to make the classification of a pose dependent only on the (x,y) -position of the robot and not of its orientation. We define a feature f as a function that takes as argument one observation $o \in O$ and returns a real value: $f(o) \in \mathbb{R}$. Compared to the features employed by Martínez Mozos *et al.* [10], we use an extended feature-set additionally containing the features shown in Table I. In our practical experiments we found out that these additional features greatly improve the robustness of the resulting classifier.

TABLE I
NEW LASER FEATURES

<ol style="list-style-type: none"> 1) Average and standard deviation of the fraction between the length of two consecutive beams. 2) Average and standard deviation of the fraction between the length of two consecutive beams divided by the maximum beam length. 3) Circularity. Let \mathbf{P} be the perimeter of the area covered by the beams and \mathbf{A} be the area covered by the beams. The circularity is defined as \mathbf{P}^2/\mathbf{A}. 4) Average and standard deviation of the distance from the centroid of \mathbf{A} to the shape boundary of \mathbf{A}, divided by the maximum distance to the shape boundary. 5) Number of gaps. Two consecutive beams form a gap if the fraction between the first and the second is smaller than a threshold. 6) Kurtosis. The kurtosis is defined as $\frac{\sum_{i=1}^N (\text{length}(\text{beam}_i) - \bar{l})^4}{N \cdot \sigma^4} - 3$ <p>where \bar{l} is the average beam length and σ the corresponding standard deviation.</p>

V. PROBABILISTIC RELAXATION LABELING

One of the key problems that need to be solved in order to learn accurate topological maps, in which the nodes correspond to the individual rooms in the environment, is to eliminate classification errors. In this section, we apply the probabilistic relaxation labeling, which has been introduced by Rosenfeld *et al.* [13], to smooth the AdaBoost classifications based on neighborhood relations.

Probabilistic relaxation labeling is defined as follows. Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a graph consisting of nodes $\mathcal{V} = \{v_1, \dots, v_N\}$ and edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$. Let furthermore $\mathcal{L} = \{l_1, \dots, l_L\}$ be a set of labels. We assume that every node v_i stores a probability distribution about its label which is represented by a histogram P_i . Each bin $p_i(l)$ of that histogram stores the probability that the node v_i has the label l . Thus, $\sum_{l=1}^L p_i(l) = 1$.

For each node v_i , $\mathcal{N}(v_i) \subset \mathcal{V}$ denotes its neighborhood which consists of the nodes $v_j \neq v_i$ that are connected to v_i . Each neighborhood relation is represented by two values. Whereas the first one describes the compatibility between the labels of two nodes, the second one represents the influence between the two nodes. The term $\mathcal{R} = \{r_{ij}(l, l') \mid v_j \in \mathcal{N}(v_i)\}$ defines the compatibility coefficients between the label l of node v_i and the label l' of v_j . And $\mathcal{C} = \{c_{ij} \mid v_j \in \mathcal{N}(v_i)\}$ is the set of weights indicating the influence of node v_j on node v_i .

Given an initial estimation for the probability distribution over labels $p_i^{(0)}(l)$ for the node v_i , the probabilistic relaxation method iteratively computes estimates $p_i^{(r)}(l)$, $r = 1, 2, \dots$, based on the initial probabilities $p_i^{(0)}(l)$, the compatibility coefficients \mathcal{R} , and the weights \mathcal{C} in the form

$$p_i^{(r+1)}(l) = \frac{p_i^{(r)}(l) \left[1 + q_i^{(r)}(l) \right]}{\sum_{l'=1}^L p_i^{(r)}(l') \left[1 + q_i^{(r)}(l') \right]}, \quad (3)$$

where

$$q_i^{(r)}(l) = \sum_{j=1}^M c_{ij} \left[\sum_{l'=1}^L r_{ij}(l, l') p_j^{(r)}(l') \right]. \quad (4)$$

Note that the compatibility coefficients $r_{ij}(l, l') \in [-1, 1]$ do not need to be symmetric. A value $r_{ij}(l, l')$ close to -1 indicates that label l' is unlikely at node v_j when label l occurs at node v_i , whereas values close to 1 indicate the opposite. A value of exactly -1 indicates that the relation is not possible and a value of exactly 1 means that the relation always occurs.

Probabilistic relaxation provides a framework for smoothing but does not specify how the compatibility coefficients are computed. In this work, we apply the coefficients as defined by Yamamoto [20]

$$r_{ij}(l, l') = \begin{cases} \frac{1}{1-p_i(l)} \left(1 - \frac{p_i(l)}{p_{ij}(l|l')} \right) & \text{if } p_i(l) < p_{ij}(l|l') \\ \frac{p_{ij}(l|l')}{p_i(l)} - 1 & \text{otherwise,} \end{cases} \quad (5)$$

where $p_{ij}(l|l')$ is the conditional probability that node v_i has label l given that node $v_j \in \mathcal{N}(v_i)$ has label l' . Each of the values $p_i(l)$ and $p_{ij}(l|l')$ are pre-calculated only once and remain the same during the iterations of the relaxation process. Thus, the coefficients \mathcal{R} remain the same as well.

So far we described the general method for relaxation labeling. It remains to describe how we apply this method for spatial smoothing of the classifications obtained by our AdaBoost classifier. To learn a topological map, we assume a given two-dimensional occupancy grid map in which each cell $m_{(x,y)}$ stores the probability that it is occupied. We furthermore consider the eight-connected graph induced by such a grid. Let $v_i = v_{(x,y)}$ be a node corresponding to a cell $m_{(x,y)}$ from the map. Then we define a neighborhood $\mathcal{N}_8(v_{(x,y)})$ using the 8-connected cells to $v_{(x,y)}$ as described in [4].

For the initial probabilities $p_{(x,y)}^{(0)}(l)$, we use the output z of the classifier as described in Section III. Our set of labels is $\mathcal{L} = \{\text{corridor}, \text{room}, \text{doorway}, \text{wall}\}$. For each node $v_{(x,y)}$ in the free space of the occupancy grid map, we calculate the expected laser scan by ray-casting in the map. We then classify the observation and obtain a probability distribution z over all the possible places according to Equation (2). The classification output z for each pose (x, y) is used to initialize the probability distribution $P_{(x,y)}^{(0)}$ of node $v_{(x,y)}$. For the nodes lying in the free space, the probability $p_{(x,y)}^{(0)}(\text{wall})$ of being a wall is initialized with 0. Accordingly, the nodes corresponding to occupied cells in the map are initialized with $p_{(x,y)}^{(0)}(\text{wall}) = 1$.

Each of the weights $c_{ij} \in \mathcal{C}$ is initialized with the value $\frac{1}{8}$, indicating that all the eight neighbors v_j of node v_i are equally important. The compatibility coefficients are calculated using Equation (5). The values $p_i(l)$ and $p_{ij}(l|l')$ are obtained from statistics in the given (occupancy grid) map corresponding to the training data as will be described in Section VII.

VI. REGION EXTRACTION AND TOPOLOGICAL MAPPING

We define a region λ_l on a adjacency graph \mathcal{A} as a set of 8-connected nodes with the same label l . For each label

$l \in \{\text{corridor, room, doorway}\}$, regions are extracted from the adjacency graph using the algorithm by Rosenfeld and Pfaltz [14]. Each region λ_l is assigned a different identifier. The connections between regions are extracted using a similar algorithm [4]. Finally, a topological graph $\mathcal{T} = (\mathcal{V}_{\mathcal{T}}, \mathcal{E}_{\mathcal{T}})$ is constructed in which each node $v_i \in \mathcal{V}_{\mathcal{T}}$ represents a region and each edge $e_s \in \mathcal{E}_{\mathcal{T}}$ represents a connection. Additionally we add to each node v_i information about the properties of the region λ_l which represents: area, centroid, and major and minor axis of the ellipse approximation of λ_l . The major and minor axis are vectors which represent the elongation of the region and its orientation. The topological graph, together with the region properties form the final topological map. We finally apply a heuristic region correction to the topological map to increase the classification rate:

- 1) We mark each region corresponding to a room or a corridor whose size does not exceed a given threshold of 1m^2 compared to the training set as classification error and assign the label of one of its connected regions.
- 2) We mark each region labeled as doorway whose size does not exceed a given threshold of 0.1m^2 square meters or that is connected to only one region as false classification and assign the label of one of its connected regions.

VII. EXPERIMENTS

The approach described above has been implemented and tested using occupancy maps obtained from real environments. The laser range data used for the training and classification were simulated using the Carnegie Mellon Robot Navigation Toolkit (CARMEN) [11]. The goal of the experiments is to demonstrate that we can construct a semantic-topological map of typical indoor environments using only laser range data. Additionally, we analyze whether our method can be used to create a topological map of an environment for which no training data were available. Furthermore, we analyze the improvement of the sequential AdaBoost classifier using our new set of features.

A. Indoor Environments

The first experiment was performed using data obtained in the office environment of building 79 at the University of Freiburg. This environment contains rooms, doorways and a corridor, which has a length of approximately 22m. For the sake of clarity we give the result of the obtained classification by separating the environment into two parts. The left half of the environment contains the poses used as training examples (see Figure 2(a)), and the right half of the environment was used for test classification and for the topological map creation. We used the sequential classifier corridor-room which correctly classifies 97.27% of the test examples. The classification is depicted as colors/grey levels in Figure 2(b). After the sequential classification, the probabilistic relaxation method explained in Section V is applied for 50 iterations. This method generates more compact regions and eliminates noise. The result is illustrated in the Figure 2(c). Finally,

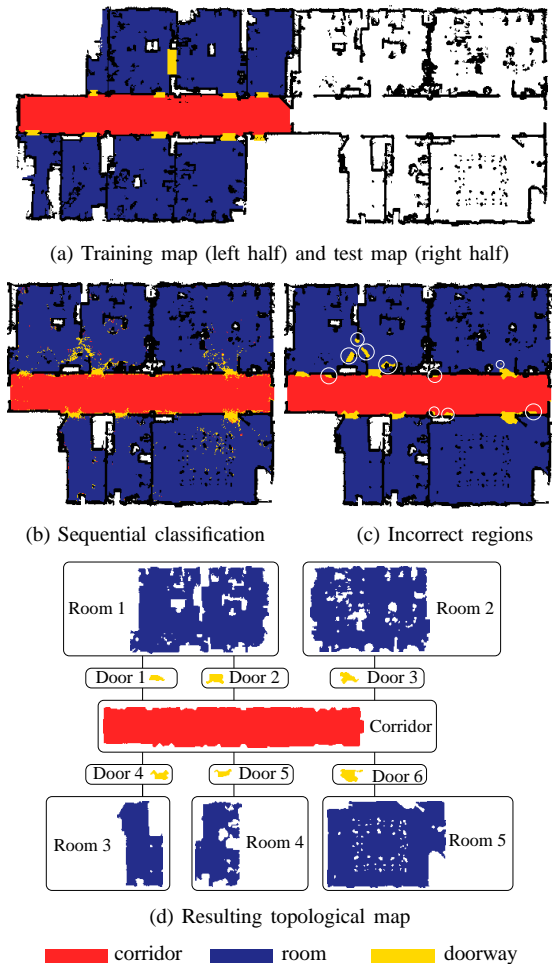




Fig. 2. This figure shows (a) the training and test map of the building 79 at the University of Freiburg, (b) the result of applying the sequential AdaBoost with a classification rate of 97.27%, (c) the result of applying relaxation and the detection of incorrect labeled regions (marked with circles), and (d) the final topological map with the corresponding regions.

the topological map is created using the connections between regions. As can be seen in Figure 2(c), some regions detected as doorways (marked with circles) do not correspond to real doorways. After applying the steps described in Section VI on the corresponding topological map, these false doorways are eliminated. Furthermore, the two left rooms situated above the corridor are detected as only one region. That is due to the fact that the doorway in between was not completely detected. Thus, the two rooms remain connected and are classified as only one region. The final topological map, depicted in Figure 2(d), has a final classification rate of 98.95% of the data points.


In a second experiment we created a topological map of the right part of the office environment of building 52 at the University of Freiburg (see Figure 3(a)). The length of the corridor in this environment is approximately 20m. After applying the sequential AdaBoost classifier room-corridor, the classification of the test set was 97%. Like in the previous experiment, we applied the relaxation process for 50 iterations as well as the operations for region correction. The final result



(a) Training map (left half) and test map (right half)



(b) Sequential classification (c) Incorrect regions



(d) Final topological map

corridor room doorway

Fig. 3. This figure shows (a) the training and test map of the building 52 at the University of Freiburg; (b) the result of applying the sequential AdaBoost with a classification rate of 97%, (c) the result of applying relaxation and the detection of incorrect labeled regions (marked with circles), and (d) the final topological map with the corresponding regions.

gives a classification rate of 98.66% of the data points. The different steps of the process are illustrated as colors/grey levels in Figure 3. As opposed to the previous experiment, the doorway between the two right-most rooms under the corridor is correctly detected (Figure 3(c)). Therefore, the rooms are labeled as two different regions in the final topological map.

B. Application to a New and Unknown Indoor Environment

This experiment is designed to analyze whether our approach can be used to create a topological map of a new environment from which no training data were available. To carry out the experiment we trained a sequential AdaBoost classifier using the training examples of the maps shown in Figure 2 and Figure 3 with different scales. In this way, we obtained a classifier with a better generalization. The resulting classifier was then evaluated on scans simulated in the map denoted as “SDR site B” in Radish [5]. This map represents an empty building in Virginia, USA. The corridor is approx. 26 meters long. The whole process for obtaining the

TABLE II

CLASSIFICATION RESULTS OF OUR FEATURE SET COMPARED TO THE ORIGINAL FEATURE SET PROPOSED BY MARTÍNEZ MOZOS *et al.* [10].

Classifier Sequence	Our feature set [%]	Original feature set [%]
room-door	96.94	93.94
room-corridor	97.26	93.31
corridor-room	97.27	93.16
corridor-door	87.73	80.10
door-corridor	87.21	80.10
door-room	86.60	80.49

topological map is depicted in Figure 4. We use the sequence corridor-doorway which gives a first classification of 92.36%. As can be seen in Figure 4(d), rooms number 11 and 30 are originally part of the corridor, and thus falsely classified. Moreover, the corridor is detected as only one region, although humans potentially would prefer to separate it into six different corridors: four horizontal and two vertical ones. Doorways are very difficult to detect by the sequential classifier. The majority of poses detected as doorways disappear after the relaxation process because they are very sparse. The main reason for the problem of doorway detection is that the maps have different sizes and resolutions, and the features are not scale invariant. In the final topological map, 96.94% of the data points are correctly classified.

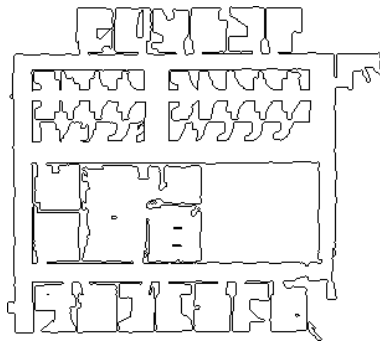
We also analyzed the results obtained without applying the relaxation process. This had several effects. First, omitting the relaxation procedure reduces the classification rate. Furthermore, the finally obtained regions typically are more sparse and do not represent the original ones as well as with relaxation. Finally, omitting the relaxation procedure increases the number of errors in the resulting topological map. For example, the map for the building in Virginia contained four incorrect nodes without relaxation, whereas there were only two incorrect nodes when we used the probabilistic relaxation.

C. Comparison with previous approaches


In this final experiment, we compare our feature set described in Section IV with the one proposed by Martínez Mozos *et al.* [10]. For this purpose, we trained a sequential AdaBoost classifier for each of the feature sets using the training set shown in Figure 2(a). The different sequential classifiers were then applied to the test set depicted in Figure 2(b). The obtained classification results are shown in Table II. As can be seen, our extended feature-set provides better results in all of the experiments.

VIII. CONCLUSIONS


In this paper, we presented a novel approach to create topological maps from indoor environments. Our approach applies AdaBoost to learn a strong classifier for categorizing places into semantic classes such as rooms, doorways, and corridors. We then apply a probabilistic relaxation process on the resulting classifications to reduce classification errors. Finally, we extract regions and their connections. The advantage of this approach is that the nodes of the resulting graph correspond to the individual semantic regions.




(a) Original map



(b) Sequential Adaboost classification



(c) Relaxation and region correction



(d) Final topological map

Fig. 4. This figure shows (a) the original map of the building, (b) the results of applying the sequential AdaBoost classifier with a classification rate of 92.36%, (c) the resulting classification after the relaxation and region correction, and (d) the final topological map with semantic information. The regions are omitted in each node. The rooms are numbered left to right and top to bottom with respect to the map in (a). For the sake of clarity, the corridor-node is drawn maintaining part of its region structure.

Our approach has been implemented and evaluated on various maps from real-world environments. Experiments demonstrate that our approach is well-suited to create topological maps from indoor environments even without training the classifier for each environment. Furthermore, we presented results demonstrating that the features we applied for place labeling yield better results than those used in previous approaches.

ACKNOWLEDGMENT

This work has been supported by the European Union under contract number FP6-004250-CoSy.

REFERENCES

- [1] P. Althaus and H. Christensen, "Behaviour coordination in structured environments," *Advanced Robotics*, vol. 17, no. 7, pp. 657–674, 2003.
- [2] P. Buschka and A. Saffiotti, "A virtual sensor for room detection," in *Proc. of the IEEE/RJSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2002, pp. 637–642.
- [3] J. Friedman, T. Hastie, and R. Tibshirani, "Additive logistic regression: a statistical view of boosting," *Annals of Statistics*, vol. 28, no. 2, pp. 337–407, 2000.
- [4] R. Gonzalez and P. Wintz, *Digital Image Processing*. Addison-Wesley Publishing Inc., 1987.
- [5] A. Howard and N. Roy, "Radish: The robotics data set repository." [Online]. Available: <http://radish.sourceforge.net/>
- [6] S. Koenig and R. Simmons, "Xavier: A robot navigation architecture based on partially observable markov decision process models," in *Artificial Intelligence Based Mobile Robotics: Case Studies of Successful Robot Systems*, D. Kortenkamp, R. Bonasso, and R. Murphy, Eds. MIT-Press, 1998, pp. 91–122.
- [7] D. Kortenkamp and T. Weymouth, "Topological mapping for mobile robots using a combination of sonar and vision sensing," in *Proc. of the Twelfth National Conference on Artificial Intelligence*, 1994, pp. 979–984.
- [8] B. Kuipers and P. Beeson, "Bootstrap learning for place recognition," in *Proc. of the Nat. Conf. on Artificial Intelligence (AAAI)*, 2002.
- [9] B. Kuipers and Y. Byun, "A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations," *Robotics and Autonomous Systems*, 8 1981.
- [10] O. Martínez Mozos, C. Stachniss, and W. Burgard, "Supervised learning of places from range data using adaboost," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, Barcelona, Spain, April 2005.
- [11] M. Montemerlo, N. Roy, and S. Thrun, "Perspectives on standardization in mobile robot programming," in *Proc. of the IEEE/RJSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2003.
- [12] H. P. Moravec, "Sensor fusion in certainty grids for mobile robots," *AI Magazine*, pp. 61–74, Summer 1988.
- [13] A. Rosenfeld, R. A. Hummel, and S. W. Zucker, "Scene labeling by relaxation operations," *IEEE Trans. Systems. Man. Cybernet.*, vol. 6, no. 6, pp. 420–433, 1976.
- [14] A. Rosenfeld and J. L. Pfaltz, "Sequential operations in digital picture processing," *Journal of the Association for Computing Machinery*, vol. 13, no. 4, pp. 471–494, 1966.
- [15] R. Schapire and Y. Singer, "Improved boosting algorithms using confidence-rated predictions," *Mach. Learn.*, vol. 37, no. 3, pp. 297–336, 1999.
- [16] H. Shatkey and L. Kaelbling, "Learning topological maps with weak local odometric information," in *Proc. of the Int. Conf. on Artificial Intelligence (IJCAI)*, 1997.
- [17] C. Stachniss, O. Martínez Mozos, A. Rottmann, and W. Burgard, "Semantic labeling of places," in *Proc. of the Int. Symposium of Robotics Research (ISRR)*, San Francisco, CA, USA, October 2005.
- [18] S. Thrun, "Learning metric-topological maps for indoor mobile robot navigation," *Artificial Intelligence*, vol. 99, no. 1, pp. 21–71, 1998.
- [19] P. Viola and M. Jones, "Robust real-time object detection," in *Proc. of IEEE Workshop on Statistical and Theories of Computer Vision*, 2001.
- [20] H. Yamamoto, "A method of deriving compatibility coefficients for relaxation operators," *Compt. Graph. Image Processing*, vol. 10, pp. 256–271, 1979.