Semantic Labeling of Places using Information Extracted from Laser and Vision Sensor Data

Óscar Martínez Mozos*, Axel Rottmann*, Rudolph Triebel*, Patric Jensfelt† and Wolfram Burgard*
*University of Freiburg, Department of Computer Science, Freiburg, Germany
†Royal Institute of Technology, Center for Autonomous Systems, Stockholm, Sweden
Email: {omartine|rottmann|triebel|burgard}@informatik.uni-freiburg.de, patric@nada.kth.se

***Abstract*— Indoor environments can typically be divided into places with different functionalities like corridors, kitchens, offices, or seminar rooms. The ability to learn such semantic categories from sensor data enables a mobile robot to extend the representation of the environment facilitating the interaction with humans. As an example, natural language terms like "corridor" or "room" can be used to communicate the position of the robot in a map in a more intuitive way. In this work, we first propose an approach based on supervised learning to classify the pose of a mobile robot into semantic classes. Our method uses AdaBoost to boost simple features extracted from range data and vision into a strong classifier. We present two main applications of this approach. Firstly, we show how our approach can be utilized by a moving robot for an online classification of the poses traversed along its path using a hidden Markov model. Secondly, we introduce an approach to learn topological maps from geometric maps by applying our semantic classification procedure in combination with a probabilistic relaxation procedure. We finally show how to apply associative Markov networks (AMNs) together with AdaBoost for classifying complete geometric maps. Experimental results obtained in simulation and with real robots demonstrate the effectiveness of our approach in various indoor environments.**

I. Introduction

In the past, many researchers have considered the problem of building accurate 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 improve the human-robot communication. From the point of view of humans, terms like "corridor" or "room" give a more intuitive idea of the position of the robot than using, for example, the 2D coordinates in a map.

In this work, we address the problem of classifying places of the environment of a mobile robot using range finder and vision data, as well as building topological maps based on that knowledge. Indoor environments, like the one depicted in Figure 1, can typically be divided into areas with different functionalities such as laboratories, office rooms, corridors, or kitchens. Whereas some of these places have special geometric structures and can therefore be distinguished merely based on laser range data, other places can only be identified according to the objects found there like, for example, monitors in a laboratory. To detect such objects, we use vision data acquired by a camera system.

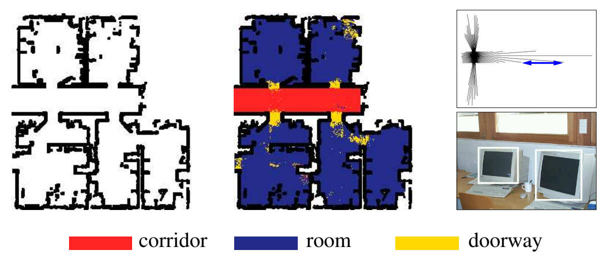

Fig. 1. The left image shows a map of a typical indoor environment. The middle image depicts the classification into three semantic classes as colors/grey levels. For this purpose the robot was positioned in each free pose of the original map and the corresponding laser observations were simulated and classified. The right images show typical laser and image observations together with some extracted features, namely the average distance between two consecutive beams in the laser and the number of monitors detected in the image.

The key idea is to classify the pose of the robot based on the current laser and vision observations. Examples for typical observations obtained in an office environment are shown in the right images of Figure 1. The classification is then done applying a sequence of classifiers learned with the AdaBoost algorithm [18]. These classifiers are built in a supervised fashion from simple geometric features that are extracted from the current laser scan and from objects extracted from the current images as shown in the right images of Figure 1. As an example, the left image in Figure 1 shows a typical indoor environment and the middle image depicts the classification obtained using our method.

We furthermore present two main applications of this approach. Firstly, we show how to classify the different poses of the robot during a trajectory and improve the final classification using a hidden Markov model. Secondly, we introduce an approach to learn topological maps from geometric maps by applying our semantic classification in combination with a probabilistic relaxation procedure. In this last case we compare the results when using an associative Markov networks (AMNs) with those obtained with AdaBoost.

The rest of this work is organized as follows. Section II presents related work. In Section III, we describe the sequential AdaBoost classifier. In Section IV, we present the application of a hidden Markov model to the online place classification with a moving robot. Section V contains our approach for topological map building. In Section VI we present some results when using a range finder with a restricted field of view. Finally, Section VII presents experimental results obtained using our methods.

II. Related Work

In the past, several authors considered the problem of adding semantic information to places. Buschka and Saffiotti [5] describe a virtual sensor to identify rooms from range data. Koenig and Simmons [9] apply a pre-programmed routine to detect doorways. Finally, Althaus and Christensen [1] use sonar data to detect corridors and doorways. Learning algorithms have additionally been used to identify objects in the environment. For example, Anguelov *et al.* [2], [3] apply the EM algorithm to cluster different types of objects from sequences of range data and to learn the state of doors. Limketkai *et al.* [12] use relational Markov networks to detect objects like doorways based on laser range data. Finally, Torralba and colleagues [23] use hidden Markov models for learning places from image data.

Compared to these approaches, our algorithm is able to combine arbitrary features extracted from different sensors to form a sequence of binary strong classifiers to label places. Our approach is also supervised, which has the advantage that the resulting labels correspond to user-defined classes.

On the other hand, different algorithms for creating topological maps have been proposed. Kuipers and Byun [11] extract distinctive points in the map defined as local maxima of a distinctiveness measure. Kortenkamp and Weymouth [10] fuse vision and ultrasound information to determine topologically relevant places. Shatkey and Kaelbling [19] apply a HMM learning approach to learn topological maps. Thrun [22] uses the Voronoi diagram to find critical points, which minimize the clearance locally. Choset [7] encodes metric and topological information in a generalized Voronoi graph to solve the SLAM problem. Additionally, Beeson *et al.* [4] used an extension of the Voronoi graph for detecting topological places. Zivkovic *et al.* [26] use visual landmarks and geometric constraints to create a higher level conceptual map. Finally, Tapus and Siegwart [20] used fingerprints to create topological maps.

In contrast to these previous approaches, the technique described in this paper applies a supervised learning method to identify complete regions in the map like corridors, rooms or doorways that have a direct relation with a human understanding of the environment. This knowledge about semantic labels of places is used then to build topological maps with a mobile robot. We also apply associative Markov networks (AMNs) together with AdaBoost to label each point in a geometric map.

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 [18]. The input to the algorithm is a set of labeled training examples $(x_n, y_n), n = 1, \ldots, N$, 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 laser and vision observations. In several iterations the algorithm repeatedly selects a weak classifier using a weight distribution over the training examples. The final strong classifier is a weighted majority vote of the best weak classifiers.

Throughout this work, we use the approach presented by Viola and Jones [25] in which the weak classifiers depend on single-valued features $f_j \in \Re$. For a more detail description see [17].

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.* [14] and create a sequential multi-class classifier using $K-1$ binary classifiers, where K is the number of classes we want to recognize. The classification output of the decision list is then represented by a histogram z. Each bin of z stores the probability that the classified example belongs to the k-th class. The order of the classifiers in the decision list can be selected according to different methods as described in [13] and [14].

A. Features from Laser and Vision Data

In this section, we describe the features used to create the weak classifiers in the AdaBoost algorithm. Our robot is equipped with a 360 degree field of view laser sensor and a camera. Each laser observation consists of 360 beams. Each vision observation consists of eight images which form a panoramic view. Figure 1 shows a typical laser range reading as well as one of the images from the panoramic view taken in an office environment. Accordingly, each training example for the AdaBoost algorithm consist of one laser observation, one vision observation, and its classification.

Our method for place classification is based on single-valued features extracted from laser and vision data. All features are invariant with respect to rotation to make the classification of a pose dependent only on the position of the robot and not on its orientation. Most of our laser features are standard geometrical features used for shape analysis as the one shown in Figure 1. In the case of vision, the selection of the features is motivated by the fact that typical objects appear with different probabilities at different places. For example, the probability of detecting a computer monitor is larger in an office than in a kitchen. For each type of object, a vision feature is defined as a function that takes as argument a panoramic vision observation and returns the number of detected objects of this type in it. This number represents the single-valued feature f_j as explained in Section III. As an example, Figure 1 shows one image of a panoramic view and its detected monitors. A more detailed list of laser and image features is contained in our previous work [14].

IV. Probabilistic Classification of Trajectories

The approach described so far is able to classify single observations only but does not take into account past classifications when determining the type of place the robot is currently at. However, whenever a mobile robot moves through an environment, the semantic labels of nearby places are typically identical. Furthermore, certain transitions between classes are unlikely. For example, if the robot is currently in a kitchen then it is rather unlikely that the robot ends up in an office given it moved a short distance only. In many environments, to get from the kitchen to the office, the robot has to move through a doorway first.

To incorporate such spatial dependencies between the individual classes, we apply a hidden Markov model (HMM) and maintain a posterior $Bel(l_t)$ about the type of the place l_t the robot is currently at

$$Bel(l_t) \quad = \quad \alpha P(z_t \mid l_t) \sum_{l_{t-1}} P(l_t \mid l_{t-1}, u_{t-1}) Bel(l_{t-1}) \qquad (1)$$

In this equation, α is a normalizing constant ensuring that the left-hand side sums up to one over all l_t. To implement this HMM, three components need to be known. First, we need to specify the observation model $P(z_t \mid l_t)$ which is the likelihood that the classification output is z_t given the actual class is l_t. Second, we need to specify the transition model $P(l_t \mid l_{t-1}, u_{t-1})$ which defines the probability that the robot moves from class l_{t-1} to class l_t by executing action u_{t-1}. Finally, we need to specify how the belief $Bel(l_0)$ is initialized.

In our current system, we choose a uniform distribution to initialize $Bel(l_0)$. The quantity $P(z_t|l_t)$ has been obtained by a statistics about the classification output of the AdaBoost algorithm given that the robot was at a place corresponding to l_t. To realize the transition model $P(l_t|l_{t-1}, u_{t-1})$ we only consider the two actions $u_{t-1} \in \{MOVE, STAY\}$. The transition probabilities were estimated by running 1000 simulation experiments. A more complete description is given in [17].

V. Topological Map Building

A second application of our semantic place classification is the extraction of topological maps from geometric maps. Throughout this section we assume that the robot is given a map of the environment in the form of an occupancy grid [15]. 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 labeling this scan into one of the semantic classes. To remove noise and clutter from the resulting classifications, we apply an approach denoted as probabilistic relaxation labeling [16]. This method takes into account the labels of the neighborhood when changing (or maintaining) the label of a given cell. 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 we apply a heuristic region correction to the topological map to increase the classification rate. A typical topological map obtained with our approach is shown in the Figure 7. For more detail see [14].

A. Semantic Classification of Maps using Associative Markov Networks

The improvement on the labeling of free cells given by our AdaBoost approach can also be seen as a collective classification problem [6]. In this approach, the labeling of each free cell in the map is also influenced by the labeling of other cells in the vicinity. One popular method for the task of collective classification are relational Markov networks (RMNs) [21]. In addition to the labels of neighboring points, RMNs also consider the relations between different objects. E.g., we can model the fact that two classes A and B are more strongly related to each other than, say, classes A and C. This modeling is done on the abstract class level by introducing clique templates [6]. Applying these clique templates to a given data set yields an ordinary Markov network (MN). In this MN, the result is a higher weighting of neighboring points with labels A and B than of points labeled A and C. Additionally, each node in the network is associated a set of features.

The whole process of labeling is composed of two steps. First, a supervised learning process is used to learn the parameters of the RMN used as a training set. Second, a new network is classified using these parameters. This last step is also called inference. In this work, we will use a special type of RMNs known as associative Markov networks (AMNs). Efficient algorithms are available for learning and inference in AMNs (for more detail see [24]).

In our case we create an AMN in which each node represents a cell in the geometric map. Each node is given a semantic label corresponding to the place in the map (corridor, doorway or room). We also create a 8-neighborhood for each cell. Furthermore, a set of features is calculated for each cell. These features correspond to the geometric ones extracted from a simulated laser beam as explained in Section III-A. To reduce the number of features during the training and inference steps, we select a subset of them. This selection is done using the AdaBoost algorithm [13].

VI. Laser Observations with Restricted Field of View

In this section we present some practical issues when classifying a trajectory using range data with a restricted field of view. Specifically, we explain how to extract features when using a laser range finder which only covers 180^o in front of the robot. This is one of the most common configurations when using mobile robots. As an example, if a robot is looking at the end of a corridor, then it is not able to see the rest of the corridor, as is the case with an additional rear laser. This situation is shown in Figure 2. When classifying a trajectory we propose to maintain a local map around the robot as shown in the right image of Figure 2. This local map can be updated during the movements of the robot and then used to simulate the rear laser beams. In Section VII we show some results when learning and classifying a place using this method.

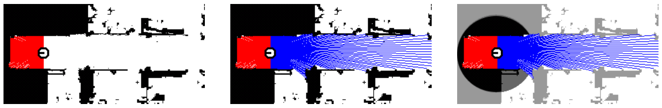

Fig. 2. The left image shows a robot at the end of a corridor with only a front laser (red). In the middle image the robot has an additional rear laser (blue). The right image depicts an example local map (shaded area).

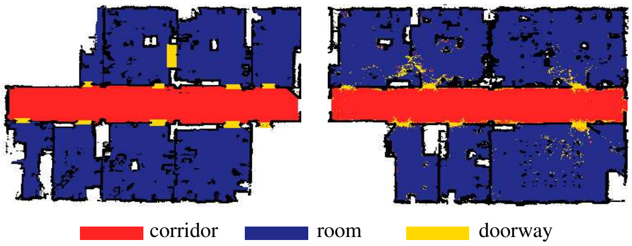

Fig. 3. The left image depicts the training data. The right image shows the test set with a classification rate of 97.3%. The training and test data were obtained by simulating laser range scans in the map.

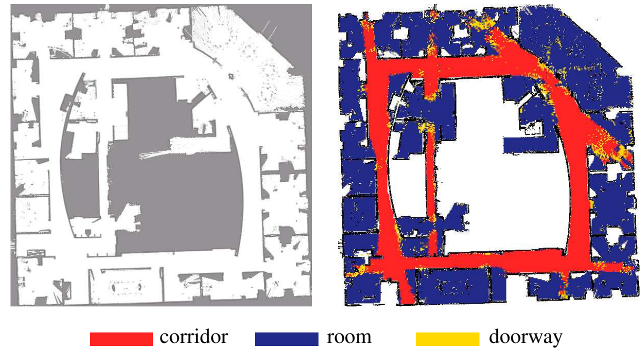

Fig. 4. The left map depicts the occupancy grid map of the Intel Research Lab and the right image depicts the classification results obtained by applying the classifier learned from the environment depicted in Figure 1 to this environment. The fact that 83.0% of all places could be correctly classified illustrates that the resulting classifiers can be applied to so far unknown environments.

VII. Experiments

The approaches described above have been implemented and tested on real robots as well as in simulation. The robots used to carry out the experiments were an ActivMedia Pioneer 2-DX8 equipped with two SICK lasers, an iRobot B21r robot equipped with a camera system and an ActivMedia PowerBot equipped only with a front laser.

The goal of the experiments is to demonstrate that our simple features can be boosted to a robust classifier of places. Additionally, we analyze whether the resulting classifier can be used to classify places in environments for which no training data was available. Furthermore, we demonstrate the advantages of utilizing the vision information to distinguish between different rooms like, e.g., kitchens, offices, or seminar rooms. Additionally, we illustrate the advantages of the HMM filtering for classifying places with a moving mobile robot. We also present results applying our method for building semantic topological maps. Finally, we show experiments using a robot with only a front laser.

A. Results with the Sequential Classifier using Laser Data

The first experiment was performed using simulated data from our office environment in building 79 at the University of Freiburg. The task was to distinguish between three different types of places, namely rooms, doorways, and a corridor based on laser range data only. In this experiment, we applied the sequential classifier without any filtering. For the sake of clarity, we separated the test from the training data by dividing the overall environment into two areas. Whereas the left part of the map contains the training examples, the right part includes only test data (Figure 3). The optimal decision list for this classification problem, in which the robot had to distinguish between three classes, is room-doorway. This decision list correctly classifies 97.3% of all test examples (right image of Figure 3). Additionally, we performed an experiment using a map of the entrance hall at the University of Freiburg which contained four different classes, namely rooms, corridors, doorways, and hallways. The optimal decision list is corridor-hallway-doorway with a success rate of 89.5%. The worst configurations of the decision list are those in which the doorway classifier is in the first place. This is probably due to the fact, that doorways are hard to detect because typically most parts of a range scan obtained in a doorway cover the adjacent room and the corridor. The high error in the first element of the decision list then leads to a high overall classification error.

B. Transferring the Classifiers to New Environments

The second experiment is designed to analyze whether a classifier learned in a particular environment can be used to successfully classify the places of a new environment. To carry out this experiment, we trained our sequential classifier in the left map of Figure 1, which corresponds to the building 52 at the University of Freiburg. The resulting classifier was then evaluated on scans simulated given the map of the Intel Research Lab in Seattle depicted in Figure 4. Although the classification rate decreased to 83.0%, the result indicates that our algorithm yields good generalizations which can also be applied to correctly label places of so far unknown environments. Note that a success rate of 83.0% is quite high for this environment, since even humans typically cannot consistently classify the different places.

C. Classification of Trajectories using HMM Filtering

The third experiment was performed using real laser and vision data obtained in an office environment, which contains six different types of places, namely offices, doorways, a laboratory, a kitchen, a seminar room, and a corridor. The true classification of the different places in this environments is shown in Figure 5. The classification performance of the classifier along a sample trajectory taken by a real robot is shown in the left image of Figure 6. The classification rate in this experiment is 82.8%. If we additionally apply the HMM

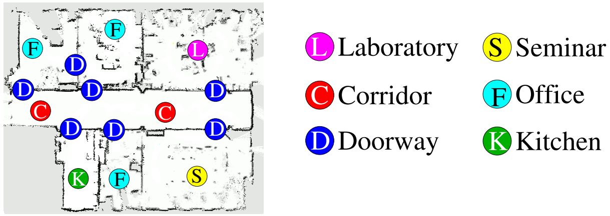

Fig. 5. Ground truth labeling of the individual areas in the environment.

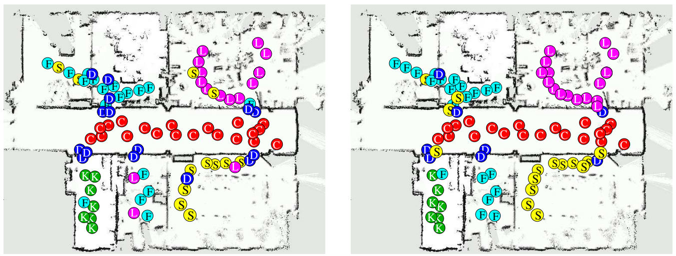

Fig. 6. The left image depicts a typical classification result for a test set obtained using only the output of the sequence of classifiers. The right image shows the resulting classification in case a HMM is additionally applied to filter the output of the sequential classifier.

for temporal filtering, the classification rate increases up to 87.9% as shown in the right image of Figure 6.

A further experiment was carried out using test data obtained in a different part of the same building. We applied the same classifier as in the previous experiment. Whereas the sequential classifier yields a classification rate of 86.0%, the combination with the HMM generated the correct answer in 94.7% of all cases. A two-sample t-test applied to the classification results obtained along the trajectories for both experiments showed that the improvements introduced by the HMM are significant on the $\alpha = 0.05$ level. Furthermore, we classified the same data based solely on the laser features and ignoring the vision information. In this case, only 67.7% could be classified correctly without the HMM. The application of the HMM increases the classification performance to 71.7%. These three experiments illustrate that the HMM significantly improves the overall rate of correctly classified places. Moreover, the third experiment shows that only the laser information is not sufficient to distinguish robustly between places with similar structure (see “office” and “kitchen” in Figure 6).

D. Building Topological Maps

The next experiment is designed to analyze our approach to build topological maps. It was carried out in the office environment depicted in the motivating example shown in Figure 1. The length of the complete corridor in this environment is approx. 20 m. After applying the sequential AdaBoost classifier (see middle image in Figure 1), we applied the probabilistic relaxation method together with the heuristics explained in Section V. The resulting topological map is shown in Figure 7. The final result gives a classification rate of 98.0% for all data points. The doorway between the two rightmost rooms under the corridor is correctly detected. Therefore,

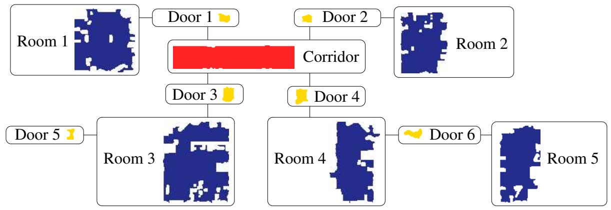

Fig. 7. Final topological map with of building 52 at Freiburg University.

the rooms are labeled as two different regions in the final topological map.

E. Learning Topological Maps of Unknown Environments

This experiment is designed to analyze whether our approach can be used to create a topological map of a new unseen environment. To carry out the experiment we trained a sequential AdaBoost classifier using the training examples of the maps shown in Figure 3 and Figure 1 with different scales. The resulting classifier was then evaluated on scans simulated in the map denoted as “SDR site B” in Radish [8]. This map represents an empty building in Virginia, USA. The corridor is approx. 26 meters long. The whole process for obtaining the topological map is depicted in Figure 8. The Adaboost classifier gives a first classification of 92.4%. As can be seen in Figure 8(d), rooms number 11 and 30 are actually 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 difficult to detect and the majority of them dissappear after the relaxation process because they are very sparse. In the final topological map 96.9% of the data points are correctly classified.

F. Learning Topological Maps using Associative Markov Networks (AMNs)

In this experiment, we classify the map of the building 79 at the University of Freiburg applying the learning and inference process for AMNs as explained in Section V-A. We divide the map in two parts and use one of them for training (see left image in Figure 3) and the second one for testing. In this experiment we reduce the resolution of the maps to 20cm. The reason is that the original resolution of 5cm generates a huge network which exceeds the memory resources of our computers during the training step of the corresponding AMN. The left image of Figure 9 shows the results of the classification using AMNs. The classification rate using AMNs was 98.8%. We compare this method with the classification obtained using our sequential AdaBoost together with the probabilistic relaxation procedure. The right image of Figure 9 depicts the classification results. In this case only 92.1% of the cells were correctly classified. As we can see, one consequence of changing the resolution to 20cm, is that the classification rate decreases (see right image of Figure 3). We think this is due to the worse quality of the simulated beams in such a granulated map. On the other hand, AMNs

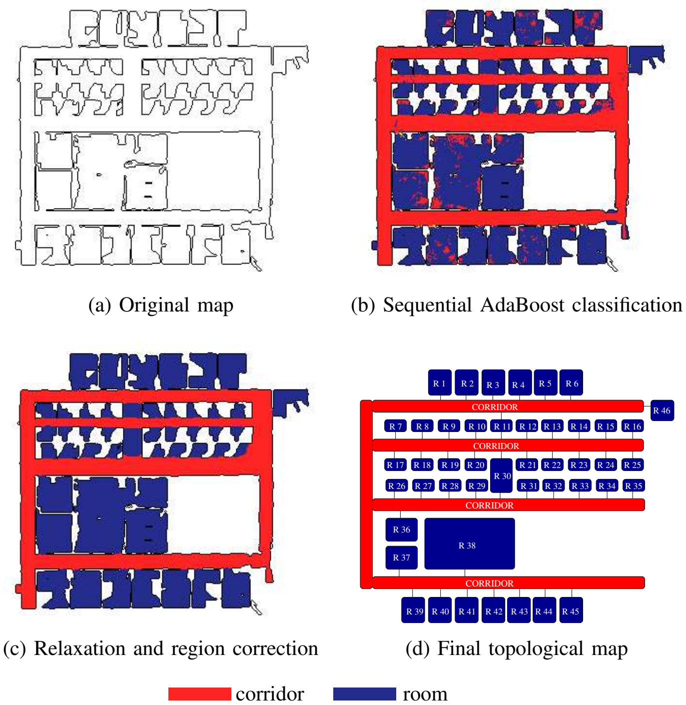

(a) Original map

(b) Sequential AdaBoost classification

(c) Relaxation and region correction

(d) Final topological map

Fig. 8. This figure shows (a) the original map of the building, (b) the results of applying the sequential AdaBoost classifier with a classification rate of 93%, (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.

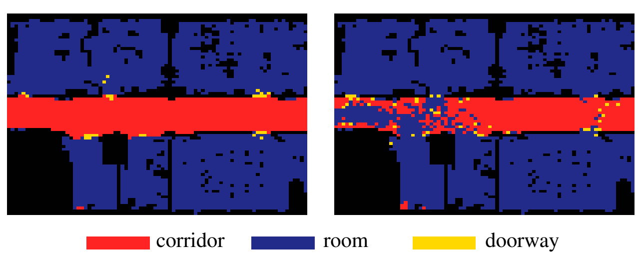

Fig. 9. The left image depicts a classification of 98.8% of the building 79 at University of Freiburg using AMNs. The right image shows the classification of the same building using the sequential AdaBoost classifier together with the probabilistic labeling method. In this case the classification rate was 92.1%. The training and test data were obtained by simulating laser range scans in the left map of Figure 3.

seems to be more robust to changes in resolution and give better classifications results.

G. Laser Observations with Restricted Field of View

In this experiments we show the results of applying our classification methods when the laser range scan has a restricted field of view. No image data was used. We first steered a PowerBot robot equipped with only a front laser along the 6th floor of the CAS building at KTH (right to left). The trajectory is shown in the top image of Figure 10. The data recorded in this floor was used to train the AdaBoost classifier. We then classified a trajectory on the 7th floor in the same building. We started the trajectory in an opposite direction (left to right). The resulting classification rate of 84.4% is depicted in the middle image of Figure 10. We repeated the experiment simulating the rear laser using a local map. The classification decreases slightly to 81.6%. Most of the errors appear in poses where the robot still sees a doorway due to the rear beams. This is not the case when using only a front laser, because the robot only sees a doorway when facing it.

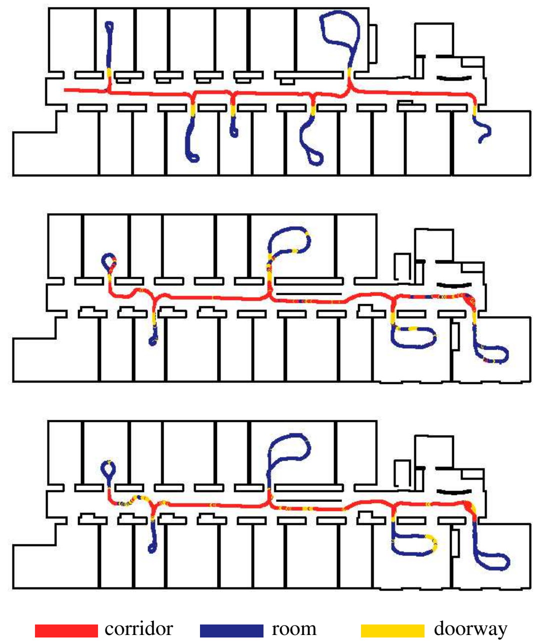

Fig. 10. The top image shows the training trajectory on the 6th floor of the CAS building at KTH. The middle image depicts the labeling of the trajectory of the 7th floor using only a front laser with a classification rate of 84.4%. Finally, the bottom image shows the same labelled trajectory using a complete laser field of view together with a local map. In this case the classification rate decreases slightly to 81.6%.

To verify that the doorways can be the reason of the lack of improvement using local maps, we repeat both experiments, but in this case using only two classes, namely room and corridor. The results are shown in Figure 11. The top image depicts the labeling using only a front laser with a classification rate of 87.3%. The bottom image shows the result of simulating the rear beams using a local map. The classification rate in this case increases to 95.8%.

VIII. Conclusion

In this paper, we presented a novel approach to classify different places in the environment of a mobile robot into semantic classes, like rooms, hallways, corridors, offices, kitchens, or doorways. Our algorithm uses simple geometric features extracted from a single laser range scan and information extracted from camera data and applies the AdaBoost algorithm to form a binary strong classifier. To distinguish between more than two classes, we use a sequence of strong binary classifiers arranged in a decision list.

We presented two applications of our approach. Firstly, we perform an online classification of the positions along the trajectories of a mobile robot by filtering the classification

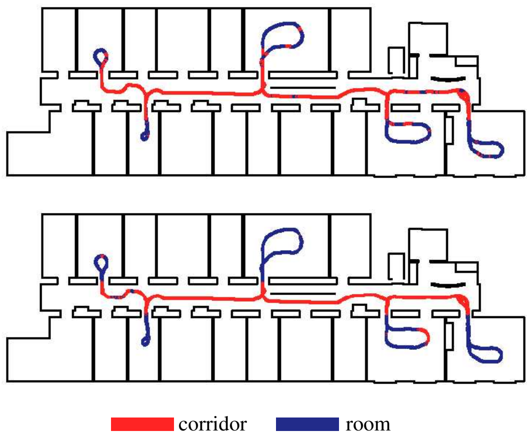

Fig. 11. In this experiments only two classes where used, namely room and corridor. The top image depicts the classification of the trajectory of the 7th floor using only a front laser with a classification rate of 87.3%. The bottom image shows the same trajectory using a complete laser field of view together with a local map. In this case the classification rate increases to 95.8%.

output using a hidden Markov model. Secondly, we present a new approach to create topological graphs from occupancy grids by applying a probabilistic relaxation labeling to take into account dependencies between neighboring places to improve the classifications.

Experiments carried out using real robots as well as in simulation illustrate that our technique is well-suited to reliably label places in different environments. It allows us to robustly separate different semantic regions and in this way it is able to learn topologies of indoor environments. Further experiments illustrate that a learned classifier can even be applied to so far unknown environments.

Acknowledgment

This work has been partially supported by the EU under contract number FP6-004250-CoSy and by the German Research Foundation under contract number SBF/TR8.

References

[1] P. Althaus and H. Christensen, "Behaviour coordination in structured environments," *Advanced Robotics*, vol. 17, no. 7, pp. 657–674, 2003.

[2] D. Anguelov, R. Biswas, D. Koller, B. Limketkai, S. Sanner, and S. Thrun, "Learning hierarchical object maps of non-stationary environments with mobile robots," in *Proc. of the Conf. on Uncertainty in Artificial Intelligence (UAI)*, 2002.

[3] D. Anguelov, D. Koller, P. E., and S. Thrun, "Detecting and modeling doors with mobile robots," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2004.

[4] P. Beeson, N. K. Jong, and B. Kuipers, "Towards autonomous topological place detection using the extended voronoi graph," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2005.

[5] P. Buschka and A. Saffiotti, "A virtual sensor for room detection," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2002, pp. 637–642.

[6] S. Chakrabarti, B. Dom, and P. Indyk, "Enhanced hypertext categorization using hyperlinks," in *SIGMOD '98: Proceedings of the 1998 ACM SIGMOD international conference on Management of data*. New York, NY, USA: ACM Press, 1998, pp. 307–318.

[7] H. Choset, "Topological simultaneous localization and mapping (SLAM): Toward exact localization without explicit localization," *IEEE Transactions on Robotics and Automation*, 2001.

[8] A. Howard and N. Roy, "Radish: The robotics data set repository." [Online]. Available: http://radish.sourceforge.net/

[9] 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.

[10] 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.

[11] 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.

[12] B. Limketkai, L. Liao, and D. Fox, "Relational object maps for mobile robots," in *Proc. of the Int. Conf. on Artificial Intelligence (IJCAI)*, Edinburgh, Scotland, 2005, pp. 1471–1476.

[13] 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, pp. 1742–1747.

[14] O. Martínez Mozos, C. Stachniss, A. Rottmann, and W. Burgard, "Using adaboost for place labeling and topological map building," in *Proc. of the Int. Symposium of Robotics Research (ISRR)*, San Francisco, CA, USA, October 2005.

[15] H. P. Moravec, "Sensor fusion in certainty grids for mobile robots," *AI Magazine*, pp. 61–74, Summer 1988.

[16] 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.

[17] A. Rottmann, O. Martínez Mozos, C. Stachniss, and W. Burgard, "Place classification of indoor environments with mobile robots using boosting," in *Proc. of the National Conference on Artificial Intelligence (AAAI)*, Pittsburgh, PA, USA, 2005, pp. 1306–1311.

[18] R. Schapire and Y. Singer, "Improved boosting algorithms using confidence-rated predictions," *Mach. Learn.*, vol. 37, no. 3, pp. 297–336, 1999.

[19] H. Shatkey and L. Kaelbling, "Learning topological maps with weak local odometric information," in *Proc. of the Int. Conf. on Artificial Intelligence (IJCAI)*, 1997.

[20] A. Tapus and R. Siegwart, "Incremental robot mapping with fingerprints of places," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, August 2005, pp. 2429– 2434.

[21] B. Taskar, P. Abbeel, and D. Koller, "Discriminative probabilistic models for relational data," in *Proc. Eighteenth Conference on Uncertainty in Artificial Intelligence (UAI)*, Edmonton, Canada, 2002.

[22] S. Thrun, "Learning metric-topological maps for indoor mobile robot navigation," *Artificial Intelligence*, vol. 99, no. 1, pp. 21–71, 1998.

[23] A. Torralba, K. Murphy, W. Freeman, and M. Rubin, "Context-based vision system for place and object recognition," in *Proc. of the Int. Conf. on Computer Vision (ICCV)*, 2003.

[24] R. Triebel, P. Pfaff, and W. Burgard, "Multi-level surface maps for outdoor terrain mapping and loop closing," in *"Proc. of the International Conference on Intelligent Robots and Systems (IROS)"*, 2006.

[25] P. Viola and M. Jones, "Robust real-time object detection," in *Proc. of IEEE Workshop on Statistical and Theories of Computer Vision*, 2001.

[26] Z. Zivkovic, B. Bakker, and B. Kröse, "Hierarchical map building using visual landmarks and geometric constraints," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2005.