2010 IEEE International Conference on Robotics and Automation
Anchorage Convention District
May 3-8, 2010, Anchorage, Alaska, USA

Improving Indoor Navigation of Autonomous Robots by an Explicit Representation of Doors

Matthias Nieuwenhuisen, Jörg Stückler, and Sven Behnke

Abstract— **In the last decades, tremendous progress has been made in the field of autonomous indoor navigation for mobile robots. However, these approaches assume the structural part of the environment to be completely static. In practice, movable parts of scenes, e.g. doors, frequently violate this assumption which leads to poor performance. Also, mobile manipulation capabilities can only be utilized, if the robot knows about the movability of objects.**

In this paper, we address an important part of these problems by the explicit representation of doors as door leaves and joints. We propose to augment standard approaches to navigation like 2D occupancy grid mapping and Monte-Carlo-Localization. Our algorithm detects doors during mapping and represents their movability adequately in the map. During localization, the state of doors is estimated from measurements while it is simultaneously used to improve localization robustness and accuracy. In experimental results we demonstrate superior performance of our method compared to a state-of-the-art approach to localization.

I. Introduction

In recent years, robust and efficient approaches to autonomous robot navigation in indoor environments have been developed. Research in this field has produced a vast set of algorithms and tools to address the problems of simultaneous localization and mapping (SLAM) and motion planning. A state-of-the-art approach to indoor navigation is to use precise laser range finders (LRFs) to acquire a map, to represent the map in a 2D occupancy probability grid, and to apply probabilistic filters like particle filters for robust state estimation.

However, the majority of these approaches assume the environment to be completely static. Thus, movable parts of the environment, e.g. doors, violate this core assumption and may lead to poor performance or even failure. To cope with dynamic objects, e.g. people, previous work investigated to detect which measurements are caused by dynamic objects and to neglect them in further processing stages.

Although this procedure would also increase robustness with movable objects, the explicit representation of the movable parts of the environment instead could further increase navigation performance: The robot can localize itself through measurements to consistently estimated movable objects. Also, if the robot has mobile manipulation capabilities, the robot may utilize information about the movability of objects to deliberatively achieve its goal, e.g. by opening doors.

This work has been supported partially by grant BE 2556/2-3 of German Research Foundation (DFG). The authors are with Computer Science Institute VI, Autonomous Intelligent Systems, University of Bonn, Roemerstrasse 164, 53117 Bonn, Germany {nieuwenh,stueckler}@ais.uni-bonn.de, behnke@cs.uni-bonn.de

Fig. 1: *Doors are typical examples for movable parts that have significant impact on the environment structure.*

Doors are typical examples for movable parts in indoor environments. Especially in corridors, the opening of a door has significant impact on the structure of the environment (s. Fig.1). Thus, we propose to extend a standard approach to indoor navigation by an explicit representation of doors.

We represent the static parts of the environment in a 2D occupancy grid map. In addition, doors are modeled in parametrized form as linear segments hinged on a vertical axis. We also attribute an opening angle range to each door which is inferred from observations. A 2D LRF measures bearing and distance to reflecting surfaces. From these measurements, we build the map through standard occupancy grid mapping. We detect doors where linear segments in the scan mismatch with the map. Measurements of dynamic objects that do not correspond to doors are not used for mapping. For localization, we present an extension of the Monte-Carlo-Localization (MCL) algorithm. While the robot localizes itself with respect to the static parts of the map and doors, it concurrently estimates the door states from the measurements.

In experiments, we demonstrate that our approach yields superior results compared to a state-of-the-art localization scheme that assumes a static map.

The remainder of this paper is structured as follows: After a brief discussion of related work, we describe our door detection and mapping algorithm in Sec. III. We detail how to localize with respect to this map and how to simultaneously estimate door states in Sec. IV. In Sec. V, we present experimental results. We conclude the paper with a discussion of

978-1-4244-5040-4/10/$26.00 ©2010 IEEE

our approach and future work.

II. Related work

In the context of navigation and environment perception, movable objects have attracted some attention in the past.

To improve localization performance in the presence of movable objects, Stachniss and Burgard [13] proposed an approach that uses local grid maps of typical configurations of the environment. Schulz and Burgard [12] estimate the state of dynamic objects, especially doors, in a dual approach to robot localization and object state estimation. They represent the belief over robot pose and object states as distinct particle sets. Instead, we estimate both quantities in a joint distribution to consider their correlations. We also describe how to detect doors and how to augment maps automatically. An approach to estimate binary door states (i.e. open or closed) and to use them to improve localization has been presented by Avots et al. [3]. They assume that a map containing doors is a priori known.

For mobile manipulation purposes, Petrovskaya and Ng [11] developed a method to precisely localize a robot relative to doors. They model doors as polygons and, similar to our approach, apply a Rao-Blackwellized particle filter to localize the robot and estimate the state of the door. In our approach, we do not require to precisely model the position and shape of doors in advance. Also, as we model to measure the distance to doors, our observation model is very similar to the well known end-point model of laser range beams. Thus, it can be easily and consistently integrated into approaches that apply the end-point model.

To perceive movable objects it is desirable to be able to detect them automatically during or after the mapping process. Biswas et al. [4] detect objects that moved between scans of the environment. They describe the shape of detected objects as local grid maps and use these to recognize objects in future scans. The approach of Anguelov et al. [2] combines laser range sensing with vision to learn models of doors in corridor environments. The map is represented as line segments with shape, color, and movability attributes, which are extracted from data in a batch process through expectation-maximization. Our approach detects doors online during the map building process and provides a representation directly suitable for localization purposes.

Perception of articulated objects requires accurate kinematic models. Several approaches to learn kinematic models and topologies of articulated objects using e.g., vision [15] or 3D range data [1], have been presented in the literature. As our approach focusses on state estimation of door leaves hinged on a rotational joint, we use an appropriate predefined model to explain the observed rigid body transforms.

Several approaches to SLAM have been designed to operate in dynamic environments. Hähnel et al. [8] apply expectation-maximization to classify if measurements correspond to static or dynamic objects. Only measurements of the static environment are finally used for SLAM. Wang and Thorpe [17] acquire grid maps through scan-matching based SLAM. They detect dynamic objects where mismatches between the current scan and the map occur, and remove them from the map of static objects. Compared to these approaches, we do not aim to detect measurements of dynamic objects to solely exclude them from the map. Instead, we explicitly represent a subset of such objects, namely doors, in the map as movable features of the environment and utilize them for localization. We also reject measurements of dynamic objects that do not correspond to doors during the mapping process.

III. Door Detection and Mapping

We assume that the environment consists of static parts, movable doors, and other dynamic objects. Movable objects typically appear static in a sequence of scans, but they may rapidly change their state during or between observations. Especially doors change their state frequently between environment snapshots. Such changes can modify the environment structure as measured by laser range finders significantly. We detect these changes during the mapping process and recognize doors as linear features that mismatch between scan and map. As a byproduct we exclude measurements of other dynamic objects from the mapping process.

Occupancy grid maps are an efficient mean to represent the static structure of an environment. They do not impose strong assumptions on the environment structure in contrast to feature-based maps. However, the independent treatment of the grid cells makes the representation of dynamic objects that span over multiple grid cells difficult. Thus, we propose to augment occupancy grid mapping by detecting door features in the map. We remove them from the grid map and replace them with a parametrized door feature.

In occupancy grid mapping, the environment is discretized at a fixed resolution. Each cell contains a probabilistic estimate of its binary occupancy state. In addition, we assign further attributes to each cell, similar to the notation in [9]: A *transient* cell has once been estimated to be unoccupied with a high likelihood, and thus may only cover transient objects. This is indicated by the occupancy probability falling below some threshold. Complementarily, a cell is denoted *seen-occupied*, if its occupancy probability ever exceeds a second threshold. Small scan registration errors can cause spurious measurements in cells close to static features in the environment. Thus, only cells within a neighborhood of cells with low occupancy probability are marked transient.

We use these attributes to classify new measurements as either static or dynamic. Measurements that fall into transient cells are collected in the set of dynamic measurements (s. Fig. 2a). The remaining static measurements are further processed in occupancy grid mapping. It is possible that doors and other dynamic objects appear static since their first observation, but change their state during the map building process. For this reason, we detect when previously seen-occupied cells attain low occupancy probability. The mean points of these cells are added to a second set of dynamic points.

In both sets we detect lines by first clustering the points

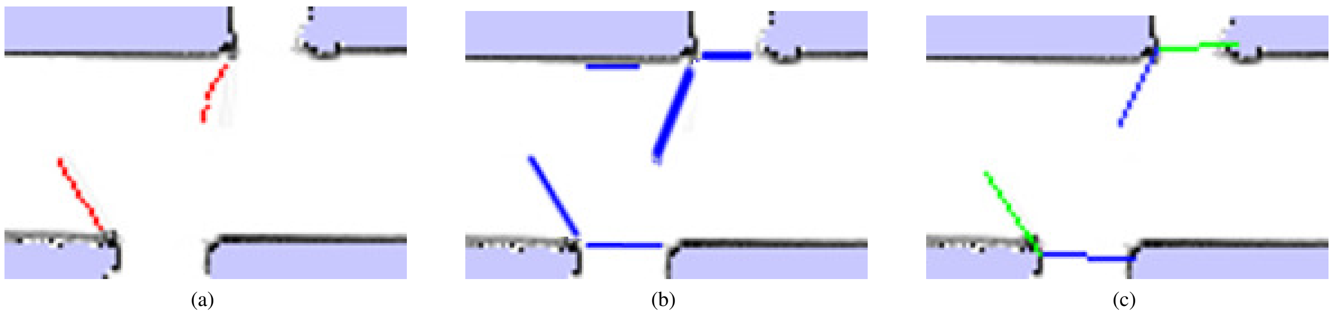

Fig. 2: (a) Scan points (red) falling into transient cells are classified as measurements of dynamic objects. (b) In each scan, lines (blue) are extracted from dynamic measurements. (c) Doors are extracted where lines (blue/green) in multiple scans rotate about a common joint.

by proximity:

$$C_i = \{p \mid \exists q \in C_i : \|p - q\| < \epsilon\}. \tag{1}$$

Each cluster C_i potentially corresponds to an object. Therefore, it is necessary that the distance of the closest measurements of different objects in a scan is greater than ϵ.

In each cluster, we determine all possible lines with a minimal length δ,

$$L_i = \{(p, q) \in C_i \times C_i \mid \|p - q\| > \delta\}. \tag{2}$$

If the maximal angle between any two lines extracted from cluster points of the same time instance exceeds a threshold, we consider the cluster as non-line object. Otherwise, we determine the mean line segment. Its length is given by the longest segment in L_i. The line object is added to a persistent set of potential door leaves, if its length is within a specific range $[\lambda_{min}, \lambda_{max}]$. Fig. 2b shows lines detected in the dynamic measurements of a scan.

In each iteration of the map building process, we extract door objects from the set of potential door leaves. As we model doors that rotate about a fixed hinge joint, we determine line objects with close end points. We require that a door is perceived in at least two distinguishable opening angles, i.e. the angle between two candidate line objects has to be significantly large. By this we can avoid false classifications of other linear appearing dynamic objects, as long as they do not rotate about one of their endpoints. Additionally, we exclude objects which rotate more than $180°$ about the hinge axis to improve the robustness of our approach in office environments. Though, this constraint has to be relaxed in environments with e.g. swinging doors.

We represent an extracted door as the position of the door's hinge joint, the length of the door leaf, a reference opening angle in global coordinates, and an observed opening angle range.

IV. Simultaneous Localization and Door State Estimation

In standard approaches to mobile robot localization like Monte-Carlo-Localization with occupancy grid maps, the map is assumed static. Thus, only the robot pose is estimated during localization. However, typical indoor environments contain movable parts that violate the assumption of a static environment.

To improve localization performance, we propose an extension to MCL that localizes the robot with respect to our map representation proposed in Sec. III. For this purpose, localization and door state estimation have to be performed simultaneously. We formulate this problem as the estimation of the joint probability distribution $p(x_{1:t}, d_t | z_{1:t}, u_{1:t}, m)$, where $x_{1:t}$ denotes the trajectory of the robot until time step t, d_t is a vector of door opening angles, and m is the static map. In each time step, the robot acquires measurements z_t in the form of laser range scans. Its motion actions are summarized in the control inputs u_t.

Similar to the FastSLAM approach [10] to simultaneous localization and mapping, we factor this distribution as

$$\begin{aligned} &p(x_{1:t}, d_t | z_{1:t}, u_{1:t}, m) = \\ &\qquad p(x_{1:t} | z_{1:t}, u_{1:t}, m) \prod_k p(d_{k,t} | x_{1:t}, z_{1:t}, u_{1:t}, m) \end{aligned} \tag{3}$$

into a trajectory estimation problem and the estimation of individual door states conditioned on the trajectory. For this factorization we make the assumption that door states are stochastically independent of each other given the robot trajectory and the observations.

We solve this estimation problem with a Rao-Blackwellized particle filter: We apply particle filtering to the trajectory estimation problem. In addition to a pose sample, each particle maintains individual normal distributed door state estimates. For the estimation of the door opening angles, we use linear Kalman filters [18].

In each time step, the particle filter algorithm proceeds as follows:

1) Sampling: The sample pose $x_t^{[i]}$ of each particle i is propagated according to the robot motion model $p(x_t | x_{t-1}, u_t)$. Afterwards, the trajectory estimate of the new particle set is distributed according to the proposal distribution $p(x_{1:t}^{[i]} | z_{1:t-1}, u_{1:t}, m)$. The door states are also updated with a simple state transition model $d_{k,t} = d_{k,t-1} + \epsilon_t$ with $\epsilon_t \sim \mathcal{N}(0, \sigma_{d,t}^2)$. We assume here that the state of doors may change unpredictably over time and that the robot does not actively influence this state.

2) Importance: The particles are weighted with the mismatch between target and proposal distribution:

$$w_t^{[i]} = \frac{p(x_{1:t}^{[i]} | z_{1:t}, u_{1:t}, m)}{p(x_{1:t}^{[i]} | z_{1:t-1}, u_{1:t}, m)} \tag{4}$$

Applying Bayes rule and standard Markov assumptions we arrive at

$$\begin{aligned} w_t^{[i]} &= \eta \ p(z_t | x_{1:t}^{[i]}, z_{1:t-1}, m) \\ &= \eta \ \prod_j p(z_{j,t} | x_{1:t}^{[i]}, z_{1:t-1}, m) \end{aligned} \tag{5}$$

where we assume stochastic independence between individual beams z_j in a scan.

For beams that measure the static part of the environment the observation likelihood is given by $p(z_{j,t}|x_t^{[i]}, m)$. When a beam measures a door d_k, the door state estimate has to be incorporated into the observation likelihood by marginalization:

$$\begin{aligned} & p(z_{j,t} | x_{1:t}^{[i]}, z_{1:t-1}, m) \\ &\quad = \int p(z_{j,t}, d_k | x_{1:t}^{[i]}, z_{1:t-1}) dd_k \\ &\quad = \int p(z_{j,t} | x_t^{[i]}, d_k) \ p(d_k | x_{1:t-1}^{[i]}, z_{1:t-1}) dd_k \end{aligned} \tag{6}$$

We model observations in the occupancy grid map with the endpoint model. It assumes that the measurement of the distance of a beam's endpoint to the closest occupied cell is normal distributed. Analogously, our observation model of doors measures the distance of the beam's endpoint to the door leaf. The convolution of the observation likelihood with the estimated door state probability in eq. (6) can not be computed in closed form. Instead we approximate it through first order taylor expansion and error propagation. As both models measure the same quantity, we can use maximum likelihood data association to determine the correspondence of a beam to either a door or to the static map.

Note that we represent the doors' hinge joints both as occupied cells in the grid map and as endpoints of door-leaves. By our association method we prevent double integration of information.

3) Door State Update: The door state estimates are maintained in each particle individually and they are conditioned on the pose sample of the particle. Thus, we find line segments in the scan z_t and associate them for each particle to corresponding door features in the map from the particle's pose. We employ a linear Kalman filter to update the door state estimates with the observed door angles.

Although the door states have to be updated for each particle individually, the update can be implemented very efficiently. We detect lines only once in a scan with the Douglas-Peucker algorithm [5] (s. Fig. 3) which efficiently extracts a polygon on the scan points. For each particle, this polygon is transformed to the sample pose. Then, for each door, we find lines that have an endpoint close to the hinge positions of the door, have similar length, and are oriented within the door's opening angle range.

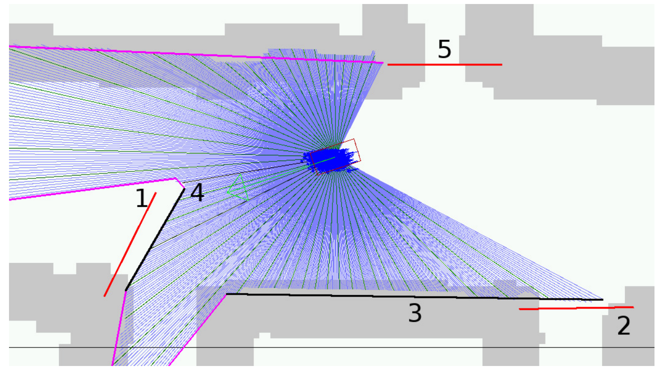

Fig. 3: *Doors (red) are observed as line segments (black) in a scan (blue). To detect lines we extract a polygon (purple) that approximates the shape of the current laser scan. As each particle maintains its own door state estimate, the lines are transformed to the particle poses. For each door, corresponding line observations are determined. In this example, segment 4 is an observation of door 1, as it ends close to the door's hinge and has a similar length. Segment 3 meets none of the two requirement. As the line segment corresponding to door 2 is merged with adjacent walls, our algorithm still uses line 3 as observation of door 2. No segment corresponds to door 5.*

Doors that are parallel to adjacent walls may not be distinguishable from these walls. Thus, if no line segment could be found for a door, a potentially longer segment is searched with low point-to-line-distance to the door hinge position. This segment can still be used to measure the angle of the door.

4) Resampling: To concentrate the particles on the relevant parts of the state space, a new particle set is drawn from the current set. The probability of a particle to be contained in the new set is proportional to its weight. We apply low-variance-sampling [16] to reduce the chance of particle depletion.

V. Experimental Results

We evaluate the applicability and the performance of our approach in localization and mapping experiments. We compare our method with an implementation of Monte-Carlo-Localization with occupancy grid maps. For data acquisition we used our robot platform Dynamaid [14] and the simulation environment Stage of the Player/Stage project [6]. Dynamaid is equipped with a SICK S300 laser range finder. The experiments with Dynamaid have been conducted in an office environment at the University of Bonn.

A. Door Detection and Mapping

Our approach augments occupancy grid mapping in known poses with the detection and modeling of doors. To obtain a trajectory for mapping we apply the FastSLAM 2.0 implementation GMapping [7] which is robust to dynamic changes in the environment to some degree. To collect test data, the robot moved through the corridor in Fig. 1 and Fig. 4 back and forth. The door states were changed manually in between the two passages. Our method correctly discovers all doors that appear in at least two distinguishable opening angles. As our approach depends on the classification of cells as transient or seen-occupied, these regions have to be observed sufficiently often. By this it is possible that door detections are missed in the initial mapping stages, when they only move one time during mapping.

B. Localization Accuracy

Accurate localization requires that distance to features is distinctively measurable along perpendicular directions. In a corridor, walls and doors are the main features of the environment. For standard approaches to localization, movable objects must be removed from the map of static objects. Thus, they can mainly use walls as features for localization.

In most buildings, walls are either parallel or perpendicular to each other. If the robot moves along a corridor, it measures wall segments mainly aligned with the corridor direction. Thus, when only walls can be used for localization, its accuracy along the corridor direction is lower than in the orthogonal direction. If doors are present, our approach can measure distance at doors along the corridor direction, as their orientation is not restricted to the main orientations of the building. In our experiments, we denote the corridor direction as x- and its orthogonal direction as y-direction.

Another source of inaccuracy are errors in the model. If the current states of doors are not represented in the static map, the pose estimate is distorted towards poses that yield erroneously a higher scan likelihood.

In our experiments, the simulated robot moved along the corridor. We evaluated the performance of standard MCL with doors manually removed from the map and our approach for different opening states of doors. The parameter sets used in both approaches were the same, with the exception of parameters related to door state estimation. To obtain comparable results we didn't alter the door configuration in the corridor during test runs. See Fig. 4 for an example configuration of our corridor.

Fig. 5 shows the average localization error over 10 runs, when the doors are closed. Both approaches can measure the end of the corridor. However, while our approach has a constantly low localization error of about 0.05m, the error in standard MCL fluctuates between 0.05m and 0.22m. This is due to the fact that in standard MCL, the pose estimate is biased towards positions, in which large parts of the scan match with the walls in the static map. In our approach, the doors are correctly estimated as closed, such that range can be consistently measured on doors.

If doors are open, they may occlude parts of the static environment. On the other hand, parts of adjacent rooms become visible through open doorways. From Fig. 6 it can be seen that our approach outperforms standard MCL under such conditions. Fig. 7 shows similar results when the doors are partially opened.

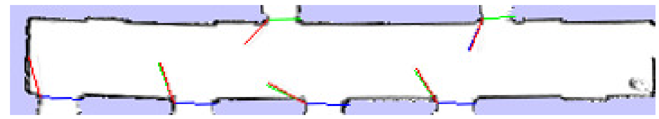

(a)

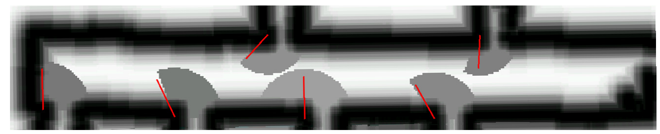

(b)

Fig. 4: *Top: Grid map of the environment used in the experiments with minimum and maximum door angles. Bottom: Likelihood field of the static map with estimated door states.*

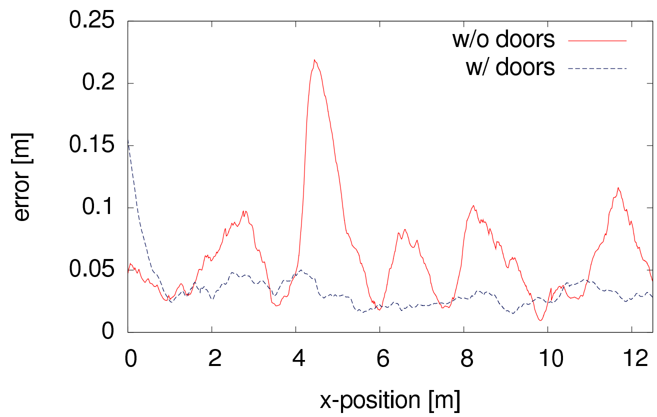

Fig. 5: *Localization error along a corridor with closed doors in simulation: While both standard MCL (red/solid) and our approach (blue/dashed) can measure the end of the corridor, the error in standard MCL fluctuates. As movable objects are missing in the standard MCL map, the pose estimate is biased towards poses with higher scan likelihood.*

When the doors are opened, the nearly opposing doors on the right side of our corridor compel the robot to drive two tight curves. This entails a fast increase in the pose error caused by the robot's odometry. Hence, the localization accuracy temporarily drops until the pose belief converges to the correct pose again.

C. Global Localization

In the preceding section we have evaluated the localization accuracy of an initially localized robot. Another important aspect of a localization approach is the ability to perform global localization if the robot's initial pose is unknown. To successfully estimate the correct pose it is essential to sample the target distribution sufficiently dense to obtain hypotheses that approximate the correct pose. Also, to keep the distribution of particle poses consistent, a correct map is necessary.

We evaluate the robustness of our approach during global localization and compare it with MCL. Higher success rates at specific numbers of particles indicate higher robustness. The number of particles is adapted in relation to the uncer-

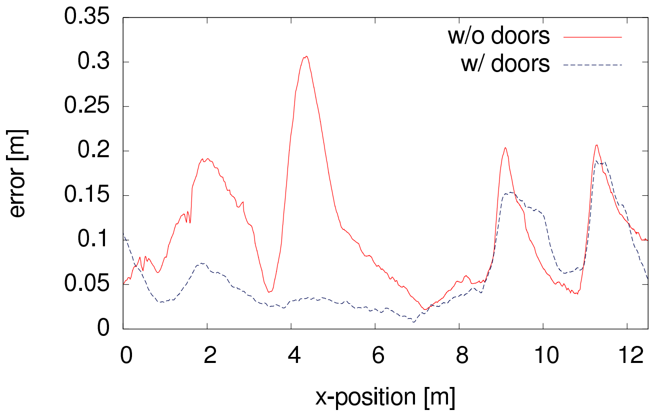

Fig. 6: *Localization error along a corridor with open doors in simulation: Open doors occlude parts of the static map. On the other hand, adjacent rooms become visible through open doorways. Our approach (blue/dashed) outperforms standard MCL (red/solid) in this environment.*

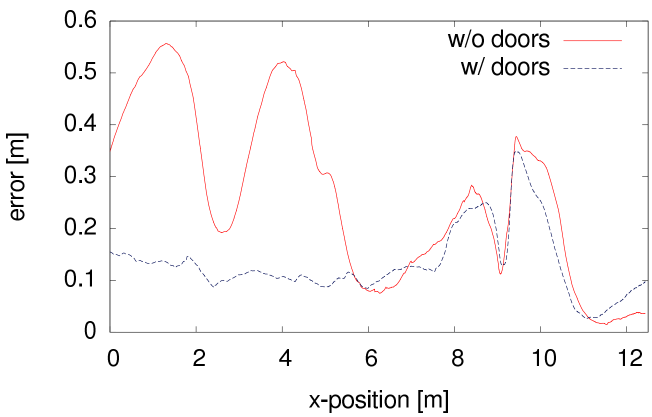

Fig. 7: *Localization error along a corridor with partially opened doors in simulation: Our approach (blue/dashed) is more accurate than standard MCL (red/solid) under these conditions.*

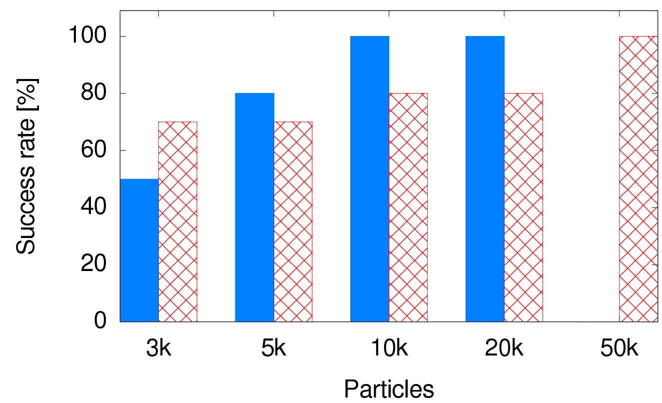

Fig. 8: *Percentage of successful global localization attempts with closed doors in 10 simulated runs. Our approach (blue/solid) performs mostly better than MCL (red/striped).*

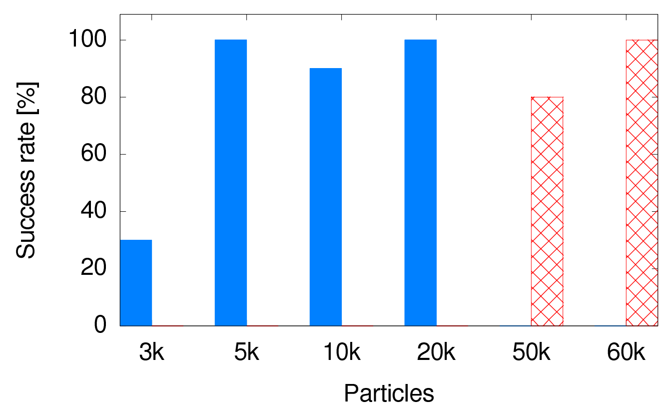

Fig. 9: *Percentage of successful global localization attempts with partially open doors in 10 simulated runs. Our approach (blue/solid) outperforms MCL (red/striped) clearly.*

tainty in the robot's pose. In our experiments we assume the robot to be localized if the localization error and the number of particles falls below a threshold. Each initial particle set size is evaluated in 10 global localization runs.

In a first simulation experiment we compare the number of particles necessary to perform global localization in a corridor environment with closed doors. Although global localization with MCL is possible in 80% of the tests with 10,000 particles, our approach succeeded in all cases with this number of particles (s. Fig. 8).

In experiments with 50,000 particles the best particle was correctly localized with our approach. However, due to performance issues, the update rate was low and the amount of particles could not be reduced sufficiently. With this large amount of particles, MCL achieves 100% success rate.

In a second simulation experiment, the doors are partially open. This changes the corridor environment significantly. While MCL only succeeds with large particle set sizes over 50,000 particles, our approach demonstrates robust global localization with only 5,000 particles.

To demonstrate the applicability of our system in real world scenarios we used our approach on our mobile robot platform Dynamaid. We evaluated the ability to localize globally with initial particle sets of 1,000 to 10,000 particles. For MCL, we also tested 20,000 particles. Each configuration was evaluated in eight test runs. We assumed the global localization to be finished if the pose estimate converged to a unimodal distribution peaked at the robot location.

Fig. 10 shows the number of successful localization attempts. The results obtained during the simulation runs are confirmed by the experiments with the real robot. Our approach succeeded to globally localize the robot in five out of eight runs with only 1,000 particles. MCL achieves a comparable success rate with much more particles. While our method succeeded in every test with 10,000 particles, 20,000 particles seem not to suffice for robust global localization with MCL.

VI. Discussion

In this paper, we presented an approach to mobile robot localization and mapping that utilizes doors as movable features of the environment. Our method enhances occupancy

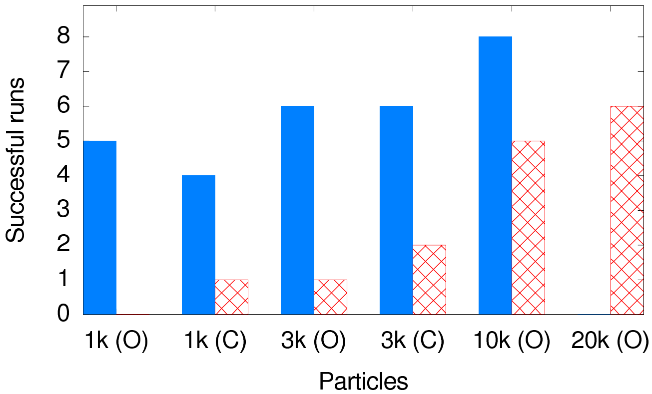

Fig. 10: *Percentage of successful global localization attempts on a real robot. Our approach (blue/solid) outperforms MCL (red/striped) and was successful in all test runs with 10,000 particles. We evaluated the approaches with partially open doors (O) and closed doors (C).*

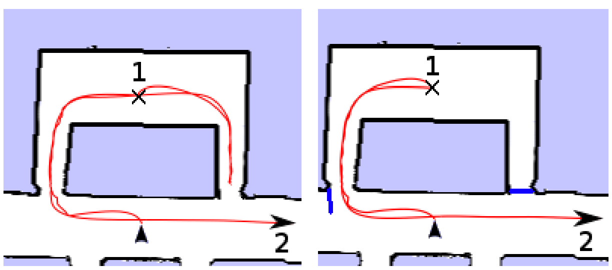

Fig. 11: *Trajectories of a simulated mobile robot navigating to a goal 2 via a subgoal 1 in a room. The room is accessible through two doors of which only the left is open. Left: without considering door states. Right: with our approach incorporating door states (blue).*

grid maps with parametrized models of doors. We propose a particle filter algorithm that enables a robot to localize with respect to our augmented map and to concurrently estimate the state of doors.

In experiments we demonstrate that our approach yields superior results compared to standard Monte-Carlo-Localization with occupancy grid maps. In the presence of doors, it improves localization accuracy significantly, if the assumption on a static environment is violated. For global localization, the measurement of doors additionally reduces the ambiguity of the environment. With our approach, the robot can relocalize itself within shorter time and more robustly.

Action Planning

Information about movable objects like doors can be useful to improve action planning in the navigation or mobile manipulation context. For instance, to plan efficient paths it is valuable to estimate the state of doors in the environment, if the robot has no manipulation capabilities. Fig. 11 compares the trajectories of a simulated robot following a plan with two subgoals with and without door state knowledge. For mobile manipulation purposes, the robot needs the ability to acquire movability properties of an object. To open doors, the robot requires knowledge about the position of the hinge joint, the length of the door leaf, etc., as given in our representation of doors.

Future Work

In future work, we plan to integrate our proposed method into SLAM and to generalize our approach to arbitrary dynamic objects. More generally, the acquisition of adequate models of dynamic objects, the compact representation of object knowledge, and the efficient state estimation of dynamic objects within the SLAM and mobile manipulation context is an interesting topic for future research.

References

[1] D. Anguelov, D. Koller, H.-C. Pang, P. Srinivasan, and S. Thrun. Recovering articulated object models from 3d range data. In *Proc. of the conf. on Uncertainty in Artificial Intelligence, pp. 18-26*, 2004.

[2] D. Anguelov, D. Koller, E. Parker, and S. Thrun. Detecting and modeling doors with mobile robots. In *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2004.

[3] D. Avots, E. Lim, R. Thibaux, and S. Thrun. A probabilistic technique for simultaneous localization and door state estimation with mobile robots in dynamic environments. In *Proceedings of IROS*, 2002.

[4] R. Biswas, B. Limketkai, S. Sanner, and S. Thrun. Towards object mapping in non-stationary environments with mobile robots. In *Proceedings of IROS*, pages 1014–1019, 2002.

[5] D. Douglas and T. Peucker. Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. *Cartographica: The International Journal for Geographic Information and Geovisualization*, 10(2):112–122, 1973.

[6] B. Gerkey, R. Vaughan, and A. Howard. The player/stage project: Tools for multi-robot and distributed sensor systems. In *Proc. of the 11th int. Conference on Advanced Robotics*, pages 317–323, 2003.

[7] G. Grisetti, C. Stachniss, and W. Burgard. Improved techniques for grid mapping with rao-blackwellized particle filters. *IEEE Transactions on Robotics*, 23(1):34–46, 2007.

[8] D. Hähnel, R. Triebel, W. Burgard, and S. Thrun. Map building with mobile robots in dynamic environments. In *IEEE International Conference on Robotics and Automation, pp. 1557–1563*, 2003.

[9] J. Modayil and B. Kuipers. Bootstrap learning for object discovery. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, volume 1, 2004.

[10] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. FastSLAM: A factored solution to the simultaneous localization and mapping problem. In *Proc. of AAAI National Conference on Artificial Intelligence, pp. 593–598*, 2002.

[11] A. Petrovskaya and A. Ng. Probabilistic mobile manipulation in dynamic environments, with application to opening doors. In *Proc. of the Int. Joint Conference on Artificial Intelligence (IJCAI)*, 2007.

[12] D. Schulz and W. Burgard. Probabilistic state estimation of dynamic objects with a moving mobile robot. *Journal of Robotics and Autonomous Systems*, 34(2-3):107–115, 2001.

[13] C. Stachniss and W. Burgard. Mobile robot mapping and localization in non-static environments. In *Proc. of the National Conference on Artificial Intelligence*, volume 20, page 1324, 2005.

[14] J. Stückler and S. Behnke. Integrating Indoor Mobility, Object Manipulation and Intuitive Interaction for Domestic Service Tasks. In *Proc. of the 9th IEEE-RAS Int. Conf. on Humanoid Robots*, 2009.

[15] J. Sturm, C. Stachniss, V. Pradeep, C. Plagemann, K. Konolige, and W. Burgard. Learning Kinematic Models for Articulated Objects. In *Proc. of the Int. Conf. on Artificial Intelligence (IJCAI)*, 2009.

[16] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, 2005.

[17] C. Wang and C. Thorpe. Simultaneous localization and mapping with detection and tracking of moving objects. In *Proc. IEEE International Conference on Robotics and Automation*, pages 2918–2924, 2002.

[18] G. Welch and G. Bishop. An introduction to the Kalman filter. *University of North Carolina at Chapel Hill, Chapel Hill, NC*, 1995.