

Robust Onboard Localization in Changing Environments Exploiting Text Spotting

Nicky Zimmerman Louis Wiesmann Tiziano Guadagnino Thomas Läbe Jens Behley Cyril Stachniss

Abstract—Robust localization in a given map is a crucial component of most autonomous robots. In this paper, we address the problem of localizing in an indoor environment that changes and where prominent structures have no correspondence in the map built at a different point in time. To overcome the discrepancy between the map and the observed environment caused by such changes, we exploit human-readable localization cues to assist localization. These cues are readily available in most facilities and can be detected using RGB camera images by utilizing text spotting. We integrate these cues into a Monte Carlo localization framework using a particle filter that operates on 2D LiDAR scans and camera data. By this, we provide a robust localization solution for environments with structural changes and dynamics by humans walking. We evaluate our localization framework on multiple challenging indoor scenarios in an office environment. The experiments suggest that our approach is robust to structural changes and can run on an onboard computer. We release an open source implementation of our approach¹, which uses off-the-shelf text spotting, written in C++ with a ROS wrapper.

I. INTRODUCTION

Localization in a given map is a fundamental capability required by most autonomous robots operating in indoor environments, such as office or hospitals. These environments are often populated by people, also undergoing “quasi-static” changes such as closing of doors, objects temporarily standing at some place, or moved furniture that is not reflected in the given map that was recorded at a different point in time. Such changes, which we refer to as “quasi-static” in contrast to dynamic ones such as moving people, result in sensor observations that substantially differ from the map and can lead to localization failure, as illustrated in Fig. 1, where closed doors in a corridor remove localization cues that can lead to ambiguities.

To overcome such localization challenges, readily available sources of information can be exploited to aid pose estimation. One example is using WiFi signal strength [16] from existing access points to aid the localization. Another example is using textual information that is part of the building infrastructure. Textual cues are often used by humans to navigate in the environment and are therefore available in most buildings designed for humans. With the recent advances in deep learning-based text recognition [29],

All authors are with the University of Bonn, Germany. Cyril Stachniss is also with the Lamarr Institute for Machine Learning and Artificial Intelligence, Germany.

This work has partially been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy, EXC-2070 – 390732324 – PhenoRob and by the European Union’s Horizon 2020 research and innovation programme under grant agreement No 101017008 (Harmony). ¹ <https://github.com/PRBonn/tmcl>

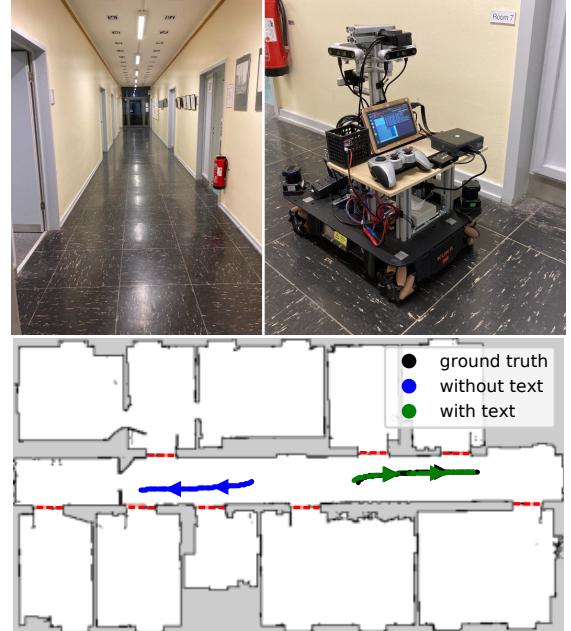


Fig. 1: Top Left: The corridor in which the experiment took place in. Top right: The Kuka YouBot platform that was used for data collection, equipped with 2D LiDAR scanners and cameras that cover the complete 360° field-of-view we utilize for text spotting. Bottom: The results of of localization in a corridor with closed doors (indicated by red lines), which are not reflected in the map, with and without textual cues.

we can reliably and efficiently decode textual content from images and utilize these hints in our localization approach. Surprisingly, there exist only a few approaches [9][28] in the robotics community to exploit text spotting or optical character recognition (OCR) for robot localization. The main contribution of this paper is a localization framework that integrates text spotting into a particle filter to improve localization. To this end, we build maps indicating the likelihood of detecting room numbers across the environment. The locations with high likelihood for successful detection are then used to inject particles when a known sign is detected. The textual cues allow us to globally localize with a small number of particles enabling online performance on mobile robots with limited computational resources. In our experiments, we show that our approach is able to (i) localize in quasi-static environments, (ii) localize in an environment with low dynamics, (iii) localize in different maps types – a featureless floor plan-like map, and LiDAR-based, feature rich map. Furthermore, our approach runs online on an onboard computer.

II. RELATED WORK

Localization of mobile platforms is a well researched area in robotics [8][32][37]. Probabilistic methods that estimate the robot's state have proven to be exceptionally robust, and include the extended Kalman filter (EKF) [19], Markov localization by Fox et al. [13] and particle filters often referred to as Monte Carlo localization (MCL) by Dellaert et al. [11]. These seminal works focused on localization using range sensors such as 2D LiDARs and sonars, as well as cameras. For cameras, the global localization task is framed under the visual place recognition framework, for which multiple algorithms have been proposed [3][10].

Localization in feature-rich maps, often constructed by range sensors, is well-established [24]. However, there are advantages for using sparse maps such as floor plans for localization. Floor plans are often available for buildings and do not require prior mapping with LiDARs or other range sensors. Their sparsity also means they do not need to be updated as frequently as detailed maps that include possibly moving objects, such as furniture. Their downside is their lack of details, which can render global localization challenging when faced with multiple identical rooms. Another issue is a possible discrepancy between the plans and the construction [6]. Boniardi et al. [7] localize in floor plans with a camera by inferring the room layout and match it against the floor plans. Li et al. [20] introduce a new state variable, scale, to address the scale difference between floor plans and the actual structures.

A problem arises when dynamic objects are detected in the scans and observations cannot be correctly matched to a given map. Sun et al. [31] propose to detect those dynamic objects as outliers using a distance filter. Thrun et al. [32] also incorporate the appearance of unexpected objects in the sensors model. Another aspect of scene dynamics is changes that are longer-lasting and not as fast to appear and disappear like moving objects. These long-lasting changes can be closing and opening of passages, transferring large packages from one place to another and shifting of large furniture. Since those changes are more constant, standard filtering technique will fail to remove them. Stachniss and Burgard [30] specifically address the case of closing and opening doors, by trying to detect areas of the map that can have different configurations and learn the possible environmental states in these corresponding areas. Another approach by Krajnik et al. [18] try to capture periodic changes by representing every cell in the occupancy map as a periodic function. Another challenge are seasonal changes in an environment, which were addressed by Vysotska et al. [36] and Milford et al. [22]. In our approach, we do not assume to have prior knowledge on changes that may occur in the map, nor do we require long sequences of images to match against.

To tackle more general semi-permanent changes, Valencia et al. [34] suggest using multiple static maps, each corresponding to a different time scale. Biber et al. [5] also propose to update a short-term map online. In the work of

Tipaldi et al. [33], a Hidden Markov model [2] is assigned to every grid cell, creating dynamics occupancy grids that can be updated. These methods require continuous update of the map, while we handle changes without altering the map.

To assist global localization, additional modalities were considered. Ito et al. [16] use WiFi signal strength to estimate the initial pose, based on signal strength maps that were previously constructed. Joho et al. [17] suggest a sensor model for RFID that combines the likelihood of detecting a tag at a given pose and the likelihood of receiving a specific signal strength. We take inspiration from these papers for building our text likelihoods/priors but apply it for a different modality.

Considerable amount of information is helping humans navigate, from publicly available maps to direction signs. Vysotska et al. [35] use publicly available maps, like Open Street Map, to localize with LiDAR. However, exploitation of text for localization is not commonly explored. It was suggested by Radwan et al. [28] but considers outdoor environment and usage of Google Maps, while our approach tackles indoor environments. Another implementation of text spotting in a MCL framework is presented by Cui et al. [9], who rely on text detection as its only sensor model. This differs from our work, which uses a 2D LiDAR-based sensor model and only leveraged text to improve global localization. The advantage of our method is that we are able to localize even in the absence of textual cues. Furthermore, in the work of Cui et al., text spotting is trained specifically for spotting parking space numbers, while we use a generic, off-the-shelf text spotting that performs well on a variety of textual cues [29].

III. OUR APPROACH

Our goal is to globally localize in an indoor environment that can undergo significant structural changes using 2D LiDAR scanners, cameras and wheel odometry. In sum, we achieve this by building upon the Monte Carlo localization framework. To aid with global localization and recover from localization failures, we use a text spotting approach inferred from camera images to detect room numbers of an human oriented environment. To integrate the textual cues, we create text likelihood maps, which indicates the likelihood of detection of each room number as a function of the robot position. We inject particles corresponding to the locations suggested by the text likelihood.

A. Monte Carlo Localization

Monte Carlo localization [11] is a probabilistic method for estimating a robot's state x_t given a map m and sensor readings z_t at time t . As we operate in an indoor environment, the robot's state x_t is given by the 2D coordinates (x, y) and the orientation $\theta \in [0, 2\pi]$. In our case an observation z is composed of K beams z_k and the map m is represented by an occupancy grid map [25].

We use a particle filter to represent the belief about the robot's state $p(x_t \mid z_{1:t}, m)$, where each particle $s_t^{(i)} = (x_t^{(i)}, w_t^{(i)})$ is represented by a state $x_t^{(i)}$ and a

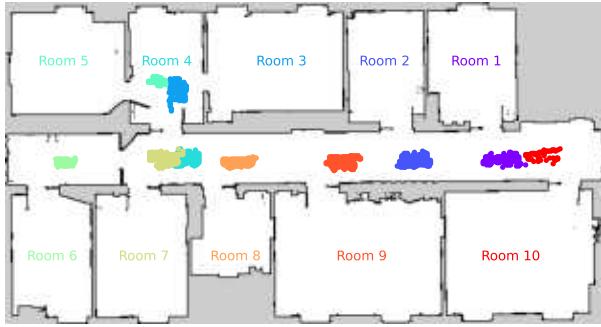


Fig. 2: The text likelihood maps, based on the collected data, indicate the locations in which detection of each room number is likely. The likelihood maps are used for particle injection when a detection of a known text cues occurs.

weight $w_t^{(i)}$. When odometry is available, successive states are sampled from a proposal distribution represented by holonomic motion model with odometry noise $\sigma_{\text{odom}} \in \mathbb{R}^3$. For each observation, each particle is weighted according to the likelihood of the observation given its state, i.e., $w_t^{(i)} = p(z_t | x_t^{(i)}, m)$.

As observation model $p(z_t | x_t, m)$, we use a beam-end model [32]. The product of likelihood model assumes scan points are independent of each other. With the high angular resolution of our LiDAR this assumption does not hold. To address the overconfidence problem of the product of likelihood model, we decided to use the product of experts model [23], where the weight of each particle is computed as the geometric mean of all scan points

$$p(z_t | x_t, m) = \prod_{k=0}^K p(z_t^k | x_t, m)^{\frac{1}{K}}, \quad (1)$$

where

$$p(z_t^k | x_t, m) = \frac{1}{\sqrt{2\pi}\sigma_{\text{obs}}} \exp\left(-\frac{EDT(\hat{z}_t^k)^2}{2\sigma^2}\right). \quad (2)$$

In Eq. (2), \hat{z}_t^k is the end point of the beam in the map m , and EDT is the Euclidean distance transform [12] that indicates the distance to an occupied cell in the occupancy map. We truncated the EDT at a predefined maximal range, r_{\max} .

For resampling, we chose low-variance resampling [32] with an efficient sample size criteria [1][4] of $N/2$, where N is the number of particles. Furthermore, our implementation of MCL [11] is asynchronous – we sample from the motion model every time we get an odometry input, and we compute the weights whenever an observation is available and the robot traveled a predefined minimum distance (d_{xy}, d_θ).

B. Text Spotting

Text spotting can traditionally be split into text detection, i.e., localizing a bounding box that includes text, and text recognition, i.e., decoding the image patches extracted from the bounding boxes, to text. Text recognition is essentially a classification problem, therefore only the characters that are introduced during training can be inferred. The last decade’s progress in object detection and text recognition allows us to use deep learning models for text spotting.

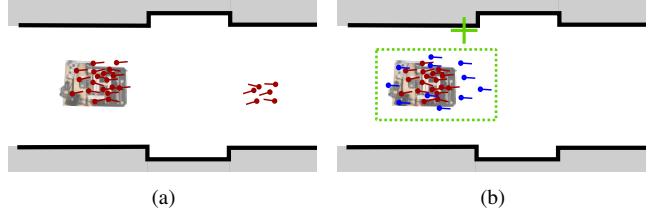


Fig. 3: Particle injection with text spotting. (a) Before detection, we have a situation with multi-modal distribution of particles (shown in red) as the corridor with closed doors is a symmetric situation that cannot be resolved just using the LiDAR scans. (b) With the first text detection (indicated by the green cross), we can inject new particles inside the bounding box extracted from the text map. We replace low weighted particles by new particles (shown in blue) that are uniformly distributed inside the corresponding bounding box of the text detection (shown by a dashed green line).

For the text spotting, we used the differentiable binarization text detector proposed by Liao et al. [21]. The backbone is a ResNet18 [15] neural network, which is powerful but also efficient enough to allow for fast inference.

The text recognition model is based on the work of Shi et al. [29], who proposed the CRNN architecture, that combines convolution, recurrent and transcription layers. This model can handle text of arbitrary length, is end-to-end trainable without requiring fine-tuning and is relatively small while maintaining accuracy. We use four cameras, with a coverage of 360° , to spot text.

C. Text Likelihood Maps

To incorporate text spotting into the MCL framework, we build a likelihood function of where the robot might detect a specific room number by collecting data that included image streams and the robot’s pose. We apply the text spotting pipeline on the recorded images, assuming that the textual cues we are interested in follow a specific pattern (“Room X”) but it can generally be used for any textual content of interest.

We compute 2D histograms for each room number, of locations where successful detections were made. The sampled locations give a sparse description of the text spotting likelihood, which we refer to as text likelihood maps (Fig. 2). As we are interested in a dense representation for the likelihood, we chose a simple strategy – for each text tag, we compute an axis-aligned bounding box around all sampled locations where the detection rate is above threshold τ for this textual cue. We approximate the likelihood of text detection with an uniform distribution within the bounding box.

D. Integration of Textual Cues

When a room number is detected, we store the room number and from which camera it was observed. Upon first detection, we inject particles into the corresponding area of the map (Fig. 3). If the last detection was made from the same camera and of the same room number, we do not inject particles. The number of particles injected is defined by the injection ratio, ρ , the number of injected particles divided by the total number of particles.

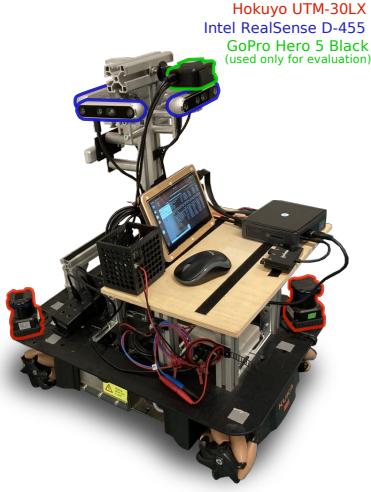


Fig. 4: The data collection platform, an omnidirectional Kuka YouBot, with 2D LiDAR scanners (marked by a red outline) and with 4 cameras (marked by a blue outline) providing 360° coverage. The up-ward facing camera (marked by a green outline) is only used for generating the ground truth via AprilTag detections.

In the injection process, for a particle filter with N particles, we first remove ρN particles with the lowest weights, and then inject an equal number of particles uniformly into the bounding box corresponding to that room number. The orientation o_i of the injected particle $s_t^{(i)}$ depends on which camera spotted the text. We assumed that the camera detecting the text facing the room number at perpendicular angle. Thus, we inject particles with corresponding orientation and add Gaussian noise, $\sigma_{\text{inject}} = 0.05$. The injection ratio ρ was chosen to be 0.5. A very high injection ratio could lead to localization failure if a wrong room number is detected. A low injection ratio has limited impact on the pose estimation. The injection of particles is done asynchronously, whenever a textual cue is available, and new particles are initialized with weight $w^{(i)} = \frac{1}{N}$.

IV. EXPERIMENTAL EVALUATION

The main focus of this work is an efficient, robust localization algorithm that leverage text information to better handle significant changes in the environment. We present our experiments to show the capabilities of our method and to support our claims, that we can (i) localize in quasi-static environments, (ii) localize in an environment with low dynamics, (iii) localize in different maps types – a featureless floor plan-like map, and LiDAR-based, feature rich map. These capabilities can be run online on our robot.

A. Experimental Setup

To benchmark the performance of our approach, we recorded a dataset in an indoor office environment. To this end, we equipped a Kuka YouBot platform with 2 Hokuyo UTM-30LX LiDAR sensors, 4 sideways-looking Intel RealSense RGB-D (D455), and an upward-looking GoPro Hero5 Black that is used only for evaluation purposes,

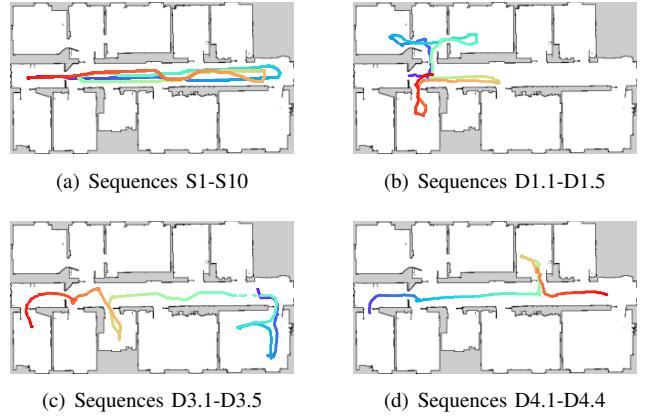


Fig. 5: Visualization of the different sequences used for evaluating our approach. Sequences S1-S10 correspond to the scenario where all doors are closed. Sequences D1-D4 were recorded with all the doors open, and with moderate amount of humans moving around. The color of the trajectory correspond to the time, where purple is the beginning and red corresponds to the end of the sequence.

TABLE I: Algorithm parameters

σ_{odom}	σ_{obs}	r_{max}	τ	ρ	d_{xy}	d_{θ}
(0.02 m, 0.02 m, 0.02 m)	2.0	15.0 m	0	0.5	0.05 m	0.05 rad

as shown in Fig. 4. We recorded the data including wheel odometry for different scenarios.

We recorded different scenarios. A long recording was made in the corridor with all doors closed, and sequences S1-S10 are randomly sampled from that data. Similarly, the sequences starting with D are sampled from recordings D1-D4, where doors were open but contain fast-moving dynamics. We include a plot of the trajectory of the scenarios in Fig. 5.

To determine the ground truth pose of the robot, we use precisely localized AprilTags [26], densely placed on the ceiling of every room and corridor. The AprilTags were detected using an up-facing camera that is used solely for this purpose. The AprilTags allow to continuously and accurately localize the robot with the dedicated sensor even under dynamic changes.

We explore two map representations, a floor plan-like map and a 2D LiDAR-based occupancy grid produced by GMapping [14], both illustrated in Fig. 6. The sparse, floor plan-like map was extracted from a high resolution terrestrial FARO laser scan, by slicing the dense point cloud at a fixed height. For all experiments, we use a map resolution of 0.05 m/cell and the parameters specified in Tab. I.

As baseline, we compare against AMCL [27], which is a publicly available and highly-used ROS package for MCL-based localization, and our implementation of MCL that does not rely on textual cues. Additionally, we implemented two sensor models for integrating textual information into the MCL framework, referred to as SM1 and SM2. SM1 assigns all particles within the bounding box a high weight, $w_t^k = 1.0$, and a low weight, $w_t^k = 0.1$, to particles elsewhere. SM2 converts the bounding box into a likelihood map



(a) Map constructed by horizontally slicing a 3D point cloud captured with a FARO Focus that was aligned to the FARO scan.
(b) Occupancy grid map from GMapping [14] that was aligned to the FARO scan, based on the recordings from the corridor scenario.

Fig. 6: Different maps used in the experiments: (a) floor plan-like maps and (b) LiDAR-based maps. (c) map built using GMmapping, based on the recordings from the corridor scenario, which significantly deviates from the maps provided for localization.

TABLE II: ATE after convergence, for each sequence for the corridor scenario, using the **sparse map** with 300 particles. Angular error in radians / translational error in meters.

Method	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10
AMCL	-/-	-/-	-/-	-/-	0.01/0.110	-/-	-/-	-/-	2.426/10.468	-/-
MCL	2.241/9.592	-/-	0.287/0.594	0.045/0.581	-/-	1.795/7.803	0.625/2.036	2.465/11.550	1.859/9.262	1.011/1.640
SM1	0.405/2.329	-/-	0.01/0.537	0.045/0.563	-/-	1.795/7.803	0.625/2.022	2.465/11.550	1.855/9.393	1.011/1.501
SM2	1.116/1.795	0.777/2.597	1.227/3.214	0.706/2.611	0.861/1.704	1.600/3.782	0.118/5.644	0.321/0.563	1.388/2.276	1.148/1.466
MCL+Text	0.063/0.250	0.01/0.245	0.063/ 0.246	0.063/ 0.266	0.01/0.246	0.179/0.369	0.045/0.221	0.077/0.343	0.493/0.332	0.01/0.184

TABLE III: Errors averaged over the trajectory, after convergence, for the mostly static environment scenario, using the **sparse map** with 300 particles. Angular error in radians / translational error in meters.

Method	D1.1	D1.2	D1.3	D1.4	D2.1	D3.1	D3.2	D3.3	D4.1	D4.2
AMCL	2.022/6.031	0.063/ 0.135	-/-	-/-	0.010/0.095	-/-	-/-	-/-	-/-	-/-
MCL	-/-	0.413/0.690	1.253/2.708	0.893/1.961	-/-	1.439/3.298	2.284/4.743	2.090/3.945	-/-	0.703/1.021
SM1	-/-	1.970/3.637	1.255/2.715	0.893/1.961	-/-	1.537/4.274	2.284/4.743	2.090/3.945	-/-	1.007/6.035
SM2	1.315/3.942	2.341/5.634	1.346/4.275	1.358/2.319	-/-	-/-	1.628/3.336	1.357/2.836	1.524/5.708	1.505/5.609
MCL+Text	0.045/0.158	0.045/0.175	0.077/0.182	0.010/0.152	0.045/0.279	0.010/0.133	0.333/0.697	0.010/0.141	0.045/0.161	0.063/0.197

TABLE IV: Convergence time in seconds, for the corridor scenario, using the **sparse map**, with 300 particles. In parentheses, the length of the sequences in seconds.

Method	S1 (234.4)	S2 (229.6)	S3 (220.8)	S4 (212.1)	S5 (203.2)	S6 (187.7)	S7 (176.3)	S8 (152.1)	S9 (145.4)	S10 (123.0)
AMCL	-	-	-	-	54.5	-	-	-	19.6	-
MCL	99.3	-	18.6	0.0	-	75.9	148.8	121.5	11.5	3.2
SM1	15.7	221.9	35.2	0.0	-	75.9	148.8	121.5	11.5	3.2
SM2	0.0	1.5	4.7	0.0	126.5	56.5	135.4	5.6	11.5	3.3
MCL+Text	0.0	1.7	12.6	0.0	10.9	11.2	50.2	5.7	11.5	2.4

and the weight for each particle is proportional to a Gaussian applied on its distance from the bounding box, similar to Eq. (2). All experiments were executed with 300 particles unless mentioned otherwise, and N particles are initialized uniformly across the map.

We consider two metrics, time to convergence and absolute trajectory error (ATE) after convergence. We define convergence as the point where the prediction is within a distance of 0.5 m from the ground truth pose. If convergence did not occur within the first 95% of the sequence, then we consider it a failure, which is marked as -/-.

B. Localization under Changes using a Sparse Map

The first experiment evaluates the performance of our approach and supports the claim that we can localize in changing environment using floor plan-like maps. It is conducted on sequences recorded in a long corridor with all doors closed, while in the map these doors are all open, and it supports our claim of robust localization in face of

quasi-static changes. We consider 10 sequences (S1-S10), each sequence starting at a different location along the corridor. We evaluate the time to convergence and ATE for this challenging scenario on the 10 sequences. As can be seen in Tab. II and Tab. IV, our text-enriched method converges quickly, and outperformed the baselines in all sequences. When the map no longer reflect the environment, it is expected that classic MCL implementations would perform poorly. For text spotting sensor model to affect the pose estimation, a particle must be in the close vicinity of a specific text likelihood bounding box. With relative low number of particles, such as 300, it is unlikely to have enough particles in such a small area. Therefore, the sensor model methods have limited contribution to global localization compared to particle injection. The MCL+Text method also shows exceptional robustness when reducing the number of particles in the filter, as can be seen in Fig. 7. The ATE for our approach is slightly larger for 10,000 particles, due to the formation of multi-modal hypotheses caused by the symmetry of the corridor.

C. Localization under Few Dynamics using a Sparse Map

The second experiment is presented to support the claim that our approach is able to localize in a floor plan-like map (not built using the robot's sensors) when the environment is mostly static. Recordings D1-D4 are taken across the lab, through different office rooms, with a small number of people moving around. In all sequences, all doors are open, and the environment is similar to the map.

TABLE V: Convergence time in seconds, for the mostly static environment scenario, using the **sparse map**, with 300 particles. In parentheses, the length of the sequences in seconds.

Method	D1.1 (171.4)	D1.2 (162.4)	D1.3 (144.8)	D1.4 (130.5)	D2.1 (78.2)	D3.1 (177.7)	D3.2 (160.6)	D3.3 (147.7)	D4.1 (120.0)	D4.2 (100.4)
AMCL	112.4	8.7	-	-	10.6	-	-	-	-	-
MCL	167.6	9.6	80.9	53.6	-	69.2	55.2	17.7	-	7.4
SM1	167.6	136.7	80.9	53.6	-	70.6	55.2	17.7	-	1.8
SM2	2.1	106.4	0.2	18.4	-	-	55.2	22.9	20.4	0.0
MCL+Text	2.2	2.0	0.7	16.3	4.8	64.6	55.9	14.9	36.4	0.0

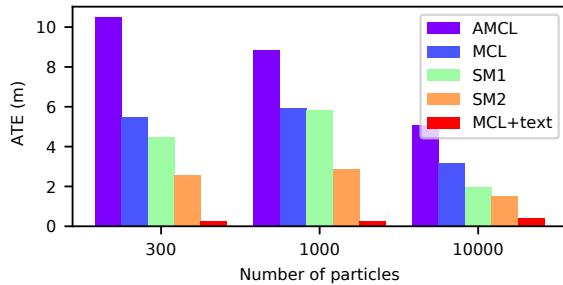


Fig. 7: ATE (xy) averaged over sequences S1-S10 as a function of the number of particles used in the particle filter, for the different methods method. The error for MCL+Text is similar across large range of particle set sizes, exhibiting the robustness of our approach.

This experiment considers localization in a feature-sparse map and in the presence of low dynamics. This presents its own challenges even in a mostly unchanging environment. As seen in Tab. III and Tab. V our text-enriched method performs best. Despite having the doors open, these scenarios include movement in a corridor with very high symmetry. Textual cues can contribute to breaking such symmetries. In addition, there are many details such as furniture, that are not part of the sparse map and can affect the accuracy of LiDAR-only localization. While SM2 shows a rather promising convergence time, the impact of the text-based sensor model is milder than particle injection, leading to divergence later on, and a large ATE.

D. Localization using LiDAR-Based Map Built with the Robot's Sensors

The third experiment is presented to support the claim that our approach is able to localize in a LiDAR-based map, when the environment is structurally changing or when there are a few dynamics in the scene. To ensure our algorithm works sufficiently well in LiDAR-based maps, we constructed a GMapping map based on 2D LiDAR scans. While this map is more detailed, the recordings were made across several weeks, resulting in some differences between the map and the environment. It is still difficult to localize globally with only 300 particles in a big scene, therefore our text-guided method enjoys an advantage. Our approach outperformed the baselines also in corridor scenario, as can be seen in Tab. VI. While the sensor model methods manage to converge in a timely manner (Tab. VIII), they are less stable than our injection technique and result in greater ATE. Similarly, for the mostly static scenario, our approach achieves the best ATE overall (Tab. VII), in addition to its fast convergence, displayed in Tab. IX.

E. Runtime

The next set of experiments has been conducted to support our fourth claim that our approach runs fast enough to execute online on the robot in real-time. We, therefore, tested our approach once using a Dell Precision-3640-Tower and once on an Intel NUC10i7FNK, which we have on our YouBot. The Dell PC has 20 CPU cores at 3.70 GHz and 64 GB of RAM. The Intel NUC has 12 CPU cores at 1.10 GHz and 16 GB of RAM.

Text spotting on the NUC runs at an average of 167 ms, and on the desktop 100 ms. Tab. X summarizes the runtime results for our approach. The numbers support our fourth claim, namely that the computations can be executed fast and in an online fashion.

F. Ablation Study

Additionally, we conducted an ablation study to identify the best way of integrating the textual hints into our MCL framework. In addition to our MCL+Text method, we also explored the following strategies for injecting particles:

- 1) Seed locations: Specific hand-picked locations in the map, which correspond to room number locations, and are used to sample particles around them with a predefined covariance.
- 2) Repeat: Using the text likelihood maps, described in Sec. III-C, we compute a bounding box for each room number plate, and inject particles in that area for every room number detection. If we have multiple consecutive detections of a room number from the same camera, we inject particles each time.
- 3) Conservative: Using the text likelihood maps, we compute a bounding box for each room number plate, and inject particles in that area only once, if the filter's pose estimation mean does not lie in the bounding box. If the mean pose of MCL is within the bounding box, the filter is in line with the tag observations, and we do not inject particles. If we have multiple consecutive detections of a room number from the same camera, we inject particles only in the first detection.

As can be seen in Fig. 8, MCL+Text outperforms the other text-guided methods. MCL+Text also converges faster than the other text-guided methods.

V. CONCLUSION

In this paper, we presented a novel approach to localize a robot in environments that deviate significantly from the provided map, as illustrated in Fig. 6, due to changes in the scene. Our method exploits the readily available human-readable textual cues that assist humans in navigation. This allows us to successfully overcome localization failure in the cases where critical changes to the layout differ greatly from the map. We implemented and evaluated our approach on a dataset collected strictly for simulating such structural alterations, and provided comparisons to other existing techniques and supported all claims made in this paper. The experiments suggest that incorporating human-readable localization cues

TABLE VI: Errors averaged over the trajectory, after convergence, for each sequence for the corridor scenario, using the **GMapping map** with 300 particles. Angular error in radians / translational error in meters.

Method	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10
AMCL	-/-	-/-	0.010/0.087	-/-	-/-	-/-	-/-	2.436/9.151	2.420/10.756	-/-
MCL	-/-	2.399/9.822	0.941/6.763	1.352/7.113	0.010/0.242	0.010/0.547	2.176/7.869	-/-	2.600/9.385	1.280/5.901
SM1	0.010/0.294	0.601/3.727	0.941/6.781	1.407/7.486	2.480/11.063	0.010/0.525	2.176/7.869	-/-	2.600/9.385	1.918/7.056
SM2	1.347/4.181	1.344/4.644	2.014/5.334	1.618/4.257	1.002/3.598	1.566/4.100	0.499/1.915	0.476/1.936	2.312/8.422	1.414/4.625
MCL+Text	0.063/0.191	0.063/0.203	0.063/0.216	0.063/0.228	0.063/ 0.171	0.045/ 0.293	0.632/ 1.788	0.010/0.192	0.697/0.920	0.205/0.196

TABLE VII: Errors averaged over the trajectory, after convergence, for mostly static scenarios, using the **GMapping map** with 300 particles. Angular error in radians / translational error in meters.

Method	D1.1	D1.2	D1.3	D1.4	D2.1	D3.1	D3.2	D3.3	D4.1	D4.2
AMCL	0.547/3.490	-/-	-/-	-/-	-/-	2.389/17.461	-/-	-/-	-/-	-/-
MCL	1.714/2.794	1.895/3.982	0.215/1.230	0.495/1.144	0.262/1.864	1.424/3.427	0.885/10.403	0.812/3.825	0.991/1.190	0.632/9.345
SM1	0.425/3.683	1.895/3.982	0.215/1.230	0.495/1.145	0.265/1.828	1.422/3.453	0.045/0.225	1.599/5.664	0.991/1.190	1.335/7.131
SM2	0.778/1.843	1.881/4.875	1.169/2.799	1.083/1.772	0.118/0.347	1.497/5.087	1.593/2.155	1.628/5.366	1.404/6.117	0.704/1.477
MCL+Text	0.010/0.109	0.010/0.109	0.077/0.099	0.010/0.116	0.045/0.161	0.010/0.156	0.010/0.169	0.045/0.165	0.063/0.160	0.045/0.131

TABLE VIII: Convergence time in seconds, for the corridor scenario, using the **GMapping map**, with 300 particles. In parentheses, the length of the sequences in seconds.

Method	S1 (234.4)	S2 (229.6)	S3 (220.8)	S4 (212.1)	S5 (203.2)	S6 (187.7)	S7 (176.3)	S8 (152.1)	S9 (145.4)	S10 (123.0)
AMCL	-	-	113.0	-	-	-	-	22.8	19.7	-
MCL	-	29.5	21.8	0.2	159.1	122.8	7.9	-	8.9	1.1
SM1	18.1	9.7	21.8	0.2	173.2	122.2	7.9	-	8.9	1.1
SM2	6.5	1.7	3.2	0.2	114.1	10.9	7.9	47.8	8.9	1.1
MCL+Text	0.0	1.5	3.6	0.2	8.4	9.7	7.9	31.2	8.9	1.1

TABLE IX: Convergence time in seconds, for the mostly static scenarios, using the **GMapping map**, with 300 particles. In parentheses, the length of the sequences in seconds.

Method	D1.1 (171.4)	D1.2 (162.4)	D1.3 (144.8)	D1.4 (130.5)	D2.1 (78.2)	D3.1 (177.7)	D3.2 (160.6)	D3.3 (147.7)	D4.1 (120.0)	D4.2 (100.4)
AMCL	0.0	-	-	-	-	35.2	-	-	-	-
MCL	111.2	10.1	30.7	50.9	21.5	67.8	133.4	31.0	51.6	64.2
SM1	115.7	10.1	30.7	50.9	21.5	67.8	87.2	39.2	51.6	8.3
SM2	2.2	73.9	30.2	16.2	7.1	72.0	50.3	15.6	26.5	0.5
MCL+Text	2.1	1.8	0.2	16.0	4.7	71.5	50.6	19.3	21.4	0.0

in mobile robot localization systems provides considerable improvement in robustness.

ACKNOWLEDGMENTS

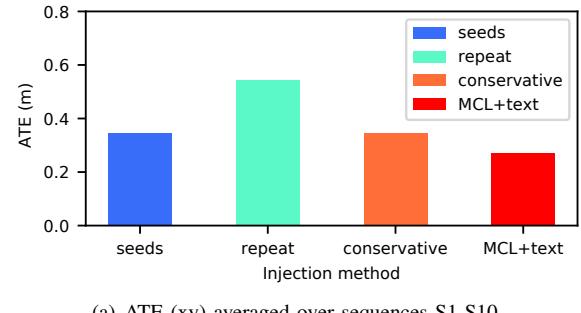
We thank Holger Milz, Michael Plech, and Ralf Becker for their contribution in assembling our mobile platform.

REFERENCES

- [1] M. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing (TSP)*, 50(2):174–188, 2002.
- [2] L.E. Baum and T. Petrie. Statistical inference for probabilistic functions of finite state Markov chains. *The annals of mathematical statistics*, 37(6):1554–1563, 1966.
- [3] M. Bennewitz, C. Stachniss, W. Burgard, and S. Behnke. Metric Localization with Scale-Invariant Visual Features using a Single Perspective Camera. In H. Christensen, editor, *European Robotics Symposium 2006*, volume 22 of *STAR Springer Tracts in Advanced Robotics*, pages 143–157. Springer Verlag, 2006.
- [4] N. Bergman. *Recursive Bayesian Estimation : Navigation and Tracking Applications*. PhD thesis, 1999.
- [5] P. Biber and T. Duckett. Dynamic Maps for Long-Term Operation of Mobile Service Robots. In *Proc. of Robotics: Science and Systems (RSS)*, 2005.
- [6] F. Boniardi, T. Caselitz, R. Kümmerle, and W. Burgard. Robust LiDAR-based localization in architectural floor plans. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2017.
- [7] F. Boniardi, A. Valada, R. Mohan, T. Caselitz, and W. Burgard. Robot localization in floor plans using a room layout edge extraction network. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2019.
- [8] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. Leonard. Past, Present, and Future of Simultaneous Localization And Mapping: Towards the Robust-Perception Age. *IEEE Trans. on Robotics (TRO)*, 32:1309–1332, 2016.
- [9] L. Cui, C. Rong, J. Huang, A. Rosendo, and L. Kneip. Monte-Carlo Localization in Underground Parking Lots Using Parking Slot Numbers. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2021.
- [10] M. Cummins and P. Newman. FAB-MAP: Probabilistic localization

TABLE X: Average inference time in ms for the sensor model on the NUC as a function of the number of particles.

	300	500	1 000	10 000
NUC10i7FNK	30	51	106	1027
Dell Precision-3640-Tower	24	40	80	793



(a) ATE (xy) averaged over sequences S1-S10.

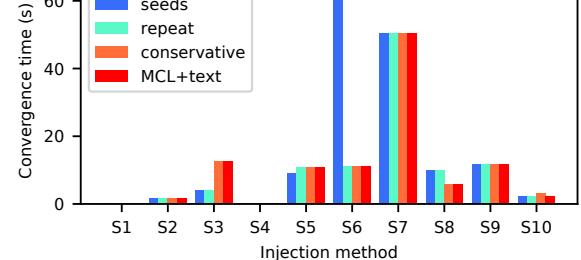


Fig. 8: Results for the ablation study exploring different injection strategies, with the sparse map and 300 particles.

- and mapping in the space of appearance. *Intl. Journal of Robotics Research (IJRR)*, 27(6):647–665, 2008.
- [11] F. Dellaert, D. Fox, W. Burgard, and S. Thrun. Monte Carlo Localization for Mobile Robots. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 1999.
 - [12] P.F. Felzenszwalb and D.P. Huttenlocher. Distance Transforms of Sampled Functions. *Theory of Computing*, 8(1):415–428, 2012.
 - [13] D. Fox, W. Burgard, and S. Thrun. Markov localization for mobile robots in dynamic environments. *Journal of Artificial Intelligence Research (JAIR)*, 11:391–427, 1999.
 - [14] G. Grisetti, C. Stachniss, and W. Burgard. Improving Grid-based SLAM with Rao-Blackwellized Particle Filters by Adaptive Proposals and Selective Resampling. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2005.
 - [15] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.
 - [16] S. Ito, F. Endres, M. Kuderer, G. Tipaldi, C. Stachniss, and W. Burgard. W-RGB-D: Floor-Plan-Based Indoor Global Localization Using a Depth Camera and WiFi. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2014.
 - [17] D. Joho, C. Plagemann, and W. Burgard. Modeling RFID signal strength and tag detection for localization and mapping. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2009.
 - [18] T. Krajinik, J.P. Fentanes, M. Hanheide, and T. Duckett. Persistent localization and life-long mapping in changing environments using the frequency map enhancement. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2016.
 - [19] J. Leonard and H. Durrant-Whyte. Mobile robot localization by tracking geometric beacons. *IEEE Trans. on Robotics and Automation*, 7(3):376–382, 1991.
 - [20] Z. Li, M.H. Ang, and D. Rus. Online Localization with Imprecise Floor Space Maps using Stochastic Gradient Descent. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2020.
 - [21] M. Liao, Z. Wan, C. Yao, K. Chen, and X. Bai. Real-time Scene Text Detection with Differentiable Binarization. *arXiv preprint*, 1911.08947, 2019.
 - [22] M. Milford and G. Wyeth. SeqSLAM: Visual route-based navigation for sunny summer days and stormy winter nights. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2012.
 - [23] R. Miyagusuku, A. Yamashita, and H. Asama. Data Information Fusion From Multiple Access Points for WiFi-Based Self-localization. *IEEE Robotics and Automation Letters (RA-L)*, 4(2):269–276, 2019.
 - [24] H. Moravec and A. Elfes. High resolution maps from wide angle sonar. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 1985.
 - [25] H.P. Moravec. Sensor Fusion in Certainty Grids for Mobile Robots. In *Sensor Devices and Systems for Robotics (SDSR)*, 1989.
 - [26] E. Olson. Apriltag: A robust and flexible visual fiducial system. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2011.
 - [27] P. Pfaff, W. Burgard, and D. Fox. Robust Monte-Carlo Localization Using Adaptive Likelihood Models. In *STAR Springer Tracts in Advanced Robotics*, 2006.
 - [28] N. Radwan, G. Tipaldi, L. Spinello, and W. Burgard. Do You See the Bakery? Leveraging Geo-Referenced Texts for Global Localization in Public Maps. In *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*, 2016.
 - [29] B. Shi, X. Bai, and C. Yao. An End-to-End Trainable Neural Network for Image-based Sequence Recognition and Its Application to Scene Text Recognition. *arXiv preprint*, 1507.05717, 2015.
 - [30] C. Stachniss and W. Burgard. Mobile Robot Mapping and Localization in Non-Static Environments. In *Proc. of the National Conference on Artificial Intelligence (AAAI)*, 2005.
 - [31] D. Sun, F. Geißer, and B. Nebel. Towards effective localization in dynamic environments. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2016.
 - [32] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, 2005.
 - [33] G.D. Tipaldi, D. Meyer-Delius, and W. Burgard. Lifelong localization in changing environments. *Intl. Journal of Robotics Research (IJRR)*, 32(14):1662–1678, 2013.
 - [34] R. Valencia, J. Saarinen, H. Andreasson, J. Vallvé, J. Andrade-Cetto, and A.J. Lilienthal. Localization in highly dynamic environments using dual-timescale NDT-MCL. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2014.
 - [35] O. Vysotska and C. Stachniss. Exploiting Building Information from Publicly Available Maps in Graph-Based SLAM. In *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2016.
 - [36] O. Vysotska and C. Stachniss. Lazy Data Association For Image Sequences Matching Under Substantial Appearance Changes. *IEEE Robotics and Automation Letters (RA-L)*, 1(1):213–220, 2016.
 - [37] F. Zafari, A. Gkelias, and K.K. Leung. A Survey of Indoor Localization Systems and Technologies. *IEEE Communications Surveys Tutorials (CST)*, 21(3):2568–2599, 2019.