Lifelong Localization and Dynamic Map Estimation in Changing Environments

Gian Diego Tipaldi[1], Daniel Meyer-Delius[1,2],
Maximilian Beinhofer[1], and Wolfram Burgard[1]

[1] Autonomous Systems Lab, University of Freiburg, D-79110 Freiburg, Germany
[2] KUKA Laboratories GmbH, D-86165 Augsburg, Germany

Fig. 1. A robot navigating in a parking lot at noon (top) and at 6 pm (bottom). Note that despite being at the same spot in both cases, the observations will be substantially different due to the changed number of parked cars

Long term operation of mobile robots in changing environments has become a major focus of interest in robotics research in recent years, as this ability is required for robots navigating in the real world. One of the most challenging tasks in this context is that of dealing with the dynamic aspects of the environment. Many existing approaches to robot navigation assume that the world is static and apply models which treat dynamic objects as outliers . For highly dynamic objects like moving people or cars, these methods typically work quite well, but they are less effective for *semi-static* objects. By semi-static objects we mean objects which change their state slowly or seldom, like parked cars, doors, pallets in warehouses or furniture which can be moved. In many realistic scenarios (see Figure 1), in which robots operate for extended periods of time, semi-static objects are ubiquitous and ignoring them can substantially deteriorate the navigation performance.

In this paper, we present a novel approach to lifelong localization in changing environments, which explicitly takes into account the dynamics of the environment. The approach is able to distinguish among objects that presents fast dynamic behaviors, e.g., cars and people, objects that can be moved around and change configuration, e.g., boxes, shelfs, doors, and objects that are static and do not move around, e.g., walls. To represent the environment, we use a *dynamic occupancy grid* , which employs hidden

Table 1. Global localization and pose tracking experiments

Experiment	MCL-GT			RBPF			MCL-S			MCL-TM		
	Failure	Error^2	σ^2	Failure	Error^2	σ^2	Failure	Error^2	σ^2	Failure	Error^2	σ^2
Global	100%	0.11	0.19	52%	0.11	0.18	36%	0.17	0.22	NA	NA	NA
Tracking	100%	0.03	0.01	98%	0.04	0.02	73%	0.10	0.08	93%	0.41	0.25

Markov models on a two-dimensional grid to represent the occupancy and the corresponding transition probabilities for each cell of this grid. We first learn the parameters of this model using a variant of the expectation maximization (EM) algorithm and then use this information to jointly estimate the pose of the robot and the state of the environment during global localization. We employ a Rao-Blackwellized particle filter (RBPF), in which the robot pose is represented by the sampled part of the filter, and the occupancy probability of a cell is represented in the analytical part of the factorization. We further propose a map management method, which uses a local map representation that is able to forget changes in a sound probabilistic way, by considering the mixing times of the associated Markov chain, and to minimize memory requirements.

Compared to previous approaches, our algorithm has several desirable advantages. First, it improves the robustness and accuracy of the pose estimate of the robot. Second, our method is able to provide an up-to-date map of the environment. Finally, our map management method considerably reduces the runtime of the process whilst minimizing the memory requirements. As a result, our approach allows a robot to simultaneously perform the estimation of its pose and the potentially changing state of the environment in an online fashion. To the best of our knowledge, this is the first approach to address this problem in a general and systematic way. Previous attempts either focused on how to filter spurious observations due to dynamic objects, focused only on limited areas or individual elements of the map or specifically addressed the problem of pose tracking.

We tested our proposed approach using a data set collected with a MobileRobots Powerbot equipped with a SICK LMS laser range finder. In the data set, the robot has been steered through a general parking lot, performing a run every full hour from 7am until 6pm during one day. The range data obtained from the twelve runs (data sets d_1 through d_{12}) corresponds to twelve different configurations of the parked cars, including an almost empty parking lot (data set d_1) and a relatively occupied one (data set d_8).

In order to assess the performances of the localization approach, we compared it to standard Monte Carlo localization both in a global localization and pose tracking setting. Furthermore, we compared our approach with the temporary maps of Meyer-Delius *et al.* but only in the pose tracking case, since it employs standard MCL before the filter reaches convergence.

The results of the experiments are shown in Table 1. The results demonstrate that our model significantly outperforms standard Monte-Carlo localization on static maps. This makes our method more suitable for long-term operation of mobile robots in changing environments.