

Active Monte Carlo Localization in Outdoor Terrains using Multi-Level Surface Maps

Rainer Kümmerle¹, Patrick Pfaff¹, Rudolph Triebel², and Wolfram Burgard¹

¹Department of Computer Science, University of Freiburg, Germany,
{kuemmer1,pfaff,burgard}@informatik.uni-freiburg.de

²Autonomous Systems Lab, ETH Zurich, Switzerland, rudolph.triebel@mavt.ethz.ch

Abstract. In this paper we consider the problem of active mobile robot localization with range sensors in outdoor environments. In contrast to passive approaches our approach actively selects the orientation of the laser range finder to improve the localization results. It applies a particle filter to estimate the full six-dimensional state of the robot. To represent the environment we utilize multi-level surface maps which allow the robot to represent vertical structures and multiple levels. To efficiently calculate the optimal orientation for the range scanner, we apply a clustering operation on the particles and only evaluate potential orientations based on these clusters. Experimental results obtained with a mobile robot in an outdoor environment indicate that the active control of the range sensor leads to more efficient localization results.

1 Introduction

The problem of mobile robot localization with range sensors in outdoor environments arises whenever GPS signals are missing due to occlusions caused by buildings, bridges, or trees. In such situations, a mobile robot typically has to estimate its position in the environment using its exteroceptive sensors and a map of the environment. However, when a robot attempts to perceive its environment to localize itself, the choice of the direction of the perception can substantially influence the accuracy of the position estimate. An example situation is shown in Figure 1. In the left image, the range sensor of the robot is oriented parallel to the floor plane as in most robot scenarios with a 2D sensor setup. This has the effect, that the vertical object shown in the image can not be sensed by the robot. However, this vertical object might be crucial for localization, because it might allow the robot to reduce its uncertainty. In contrast, the right image shows a robot with a slightly different sensor orientation so that the vertical object can be perceived. Accordingly, the robot can achieve a more accurate position estimate by actively orienting its sensor. This is why the technique is called *active localization*. In this paper, we consider the problem of active localization in outdoor environments by matching laser range measurements to a given map of the environment. In a former approach [1], we already applied multi-level surface (MLS) maps [2] to model the environment for passive localization with a fixed mounted laser range finder. The MLS maps can be regarded as an extension of the classical elevation maps [3,4,5,6] as they additionally represent intervals corresponding to vertical objects in the environment.




Fig. 1. Robot with the standard orientation of the range sensor where the robot is unable to sense the vertical object (left). In contrast to this the autonomously adapted sensor orientation allows the robot to sense the vertical object (right).

A further advantage of MLS maps is that they can represent multiple levels. This is important when mobile robots are deployed, e.g., in environments with bridges or underpasses.

The paper is organized as follows. After discussing related work in the next section, we briefly describe the general Monte Carlo localization technique in Section 3. The details of our active Monte Carlo localization are presented in Section 4. Finally, in Section 5, we present experimental results illustrating the advantages of applying active localization in outdoor environments.

2 Related Work

In general, the problem of active localization can be described as generating robot actions that particularly aim at improving its position estimate. In the past, this problem has been addressed by several authors. For example, Kaelbling et al. [7] and Koenig and Simmons [8] used a partially observable Markov decision process to model actions in the environment. The action selector chooses the action that minimizes the expected entropy after the next control action or maximizes the expected total reward, respectively. Jensfelt and Kristensen [9] applied multi-hypothesis localization to topological maps for active global localization. Davison and Kita [10] described a vision-based localization in which the robot actively senses the features based on visibility and information gain. Recently, Porta et al. [11] proposed an entropy-based criterion for action selection within a localization algorithm using a stereo vision system, which allows the robot to recover its location in the initial stages or within a failure recovery procedure more efficiently. To use fine-grained grid maps and laser range finders, Fox et al. [12] proposed an approach for active localization based on Markov localization. Whereas their approach is able to increase the efficiency of the localization by minimizing the expected entropy, Markov localization has high computational demands.

In contrast to the former approaches we focus on reducing the computational demands of the active localization. The goal of this paper is to develop an active localization method which is able to deal with large outdoor environments.

3 Monte Carlo Localization

To estimate the pose $\mathbf{x} = (x, y, z, \varphi, \vartheta, \psi)$ of the robot in its environment, we consider probabilistic localization, which follows the recursive Bayesian filtering scheme. The

key idea of this approach is to maintain a probability density $p(\mathbf{x}_t | \mathbf{z}_{1:t}, \mathbf{u}_{0:t-1})$ of the robot's location \mathbf{x}_t at time t given all observations $\mathbf{z}_{1:t}$ up to time t and all control inputs $\mathbf{u}_{0:t-1}$ up to time $t - 1$. This probability is updated as follows:

$$bel(\mathbf{x}_t) = p(\mathbf{x}_t | \mathbf{z}_{1:t}, \mathbf{u}_{0:t-1}) = \alpha \cdot p(\mathbf{z}_t | \mathbf{x}_t) \cdot \int p(\mathbf{x}_t | \mathbf{u}_{t-1}, \mathbf{x}_{t-1}) \cdot p(\mathbf{x}_{t-1}) d\mathbf{x}_{t-1}. \quad (1)$$

Here, α is a normalization constant ensuring that $p(\mathbf{x}_t | \mathbf{z}_{1:t}, \mathbf{u}_{0:t-1})$ sums up to one over all \mathbf{x}_t . The terms to be described in Eqn. (1) are the *prediction model* $p(\mathbf{x}_t | \mathbf{u}_{t-1}, \mathbf{x}_{t-1})$ and the *sensor model* $p(\mathbf{z}_t | \mathbf{x}_t)$. For the implementation of the described filtering scheme, we use a sample-based approach which is commonly known as *Monte Carlo localization* (MCL) [13]. Monte Carlo localization is a variant of particle filtering [14] where each particle $\mathbf{x}^{[i]}$ corresponds to a possible robot pose and has an assigned weight $w^{[i]}$. The *belief update* from Eqn. (1) is performed by the following two alternating steps:

1. In the **prediction step**, we draw for each particle with weight $w^{[i]}$ a new particle according to $w^{[i]}$ and to the prediction model $p(\mathbf{x}_t | \mathbf{u}_{t-1}, \mathbf{x}_{t-1})$.
2. In the **correction step**, a new observation \mathbf{z}_t is integrated. This is done by assigning a new weight $w^{[i]}$ to each particle according to the sensor model $p(\mathbf{z}_t | \mathbf{x}_t)$.

The details of the particle filter implementation in combination with the MLS maps can be found in our previous work [1].

4 Active Monte Carlo Localization

The purpose of our active localization approach is to find the orientation of the laser range finder which reduces the uncertainty of the current pose estimate as much as possible. To achieve this, we apply the greedy approach of Fox et al. [12]. We assume that at a given time step t the robot is able to execute a discrete set of actions \mathcal{A} . The benefit of a sensing action $a \in \mathcal{A}$ can be determined by considering the uncertainty of the posterior $p(\mathbf{x}_{t+1} | a, \mathbf{z}_{t+1})$. The uncertainty of the pose estimate is represented by the entropy

$$h(\mathbf{x}_t) = - \int_{\mathbf{x}_t} bel(\mathbf{x}_t) \log bel(\mathbf{x}_t) d\mathbf{x}_t. \quad (2)$$

The ideal action would allow the robot to find out its position with a high certainty. In other words, the posterior would become a single peaked distribution with a very low entropy. Therefore the information gain $g_t(a)$ of an action a to change the orientation of the laser range finder is defined by:

$$g_t(a) = h(\mathbf{x}_t) - h(\mathbf{x}_{t+1} | a, \mathbf{z}_{t+1}), \quad (3)$$

where $h(\mathbf{x}_{t+1} | a, \mathbf{z}_{t+1})$ defines the entropy after the integration of a laser measurement according to the action a . In general we do not know which range measurement the robot will obtain after changing the sensor orientation according to the action a . Therefore, we instead consider the *expected entropy* by integrating over all possible measurements \mathbf{z}_{t+1} :

$$E[g_t(a)] = h(\mathbf{x}_t) - E[h(\mathbf{x}_{t+1} | a, \mathbf{z}_{t+1})], \quad (4)$$

where $E[h(\mathbf{x}_{t+1} | a, \mathbf{z}_{t+1})]$ defines the expected entropy after the integration of a laser measurement obtained by executing action a . If we now take into account that changing the orientation of the range sensor does not change the location of the robot, then according to the reasoning by Fox et al. [12] the expected entropy is calculated as follows:

$$E[h(\mathbf{x}_{t+1} | a, \mathbf{z}_{t+1})] = - \int_{\hat{\mathbf{z}}} \int_{\mathbf{x}_t} p(\hat{\mathbf{z}} | \mathbf{x}_t) bel(\mathbf{x}_t) \log \frac{p(\hat{\mathbf{z}} | \mathbf{x}_t) bel(\mathbf{x}_t)}{p(\hat{\mathbf{z}} | a)} d\mathbf{x}_t d\hat{\mathbf{z}} \quad (5)$$

Now the action \hat{a} can be selected out of the action set \mathcal{A} which maximizes the information gain as follows:

$$\hat{a} = \operatorname{argmax}_{a \in \mathcal{A}} E[g_t(a)] = \operatorname{argmax}_{a \in \mathcal{A}} h(\mathbf{x}_t) - E[h(\mathbf{x}_t | a, \mathbf{z}_{t+1})]. \quad (6)$$

The calculation of the expected entropy $E[h(\mathbf{x}_t | a, \mathbf{z}_{t+1})]$ can be achieved by performing ray casting operations in the given MLS map. The ray casting operation approximates a possible range measurement of the robot. So we do not have to consider all possible range measurements that our sensor may generate. Using ray casting operations seems to be a good approximation for a laser range finder, as we figured out in several experiments. The result of a ray casting operation depends on the position \mathbf{x} and the action a . Furthermore this approximation allows us to compute the probability $p(\mathbf{z}_{t+1} | \mathbf{x}_t)$ which is required to calculate the expected entropy (5).

Performing a ray casting operation for each particle of our particle set would result in high computational demands. To reduce the required computation time, we only simulate the range beams on a subset calculated by a clustering operation. This is motivated by the fact that typically the particles are located in a small number of areas of high probability. Each cluster of particles represents such an area. To cluster the particle set into subsets, we apply a technique known as QT-Clustering [15]. This technique allows us to specify the maximal extent of a cluster beforehand and thus ensures that the centroid of each cluster represents the whole cluster well. In our current implementation the maximal diameter of a cluster is set to 1 m. The clustering algorithm yields a cluster set $\mathcal{K} = \bigcup_{j=1}^J (\mathbf{m}_j, \mathcal{I}_j)$, where \mathbf{m}_j is the center of mass and \mathcal{I}_j is an index set of the particles contained in cluster j . Each particle is a member of exactly one clustered subset. The range measurement generated for a subset is weighted according to the sum of the normalized weights of the particles contained in the subset. Following this approximation and also considering the underlying particle filter implementation, the expected entropy (5) is calculated as follows:

$$E[h(\mathbf{x}_{t+1} | a, \mathbf{z}_{t+1})] = - \sum_{j=1}^{|\mathcal{K}|} \sum_{i=1}^N w(\mathcal{I}_j) p(\mathbf{z}_{\mathbf{m}_j}^a | \mathbf{x}^{[i]}) \cdot w^{[i]} \log \frac{p(\mathbf{z}_{\mathbf{m}_j}^a | \mathbf{x}^{[i]}) \cdot w^{[i]}}{p(\mathbf{z}_{\mathbf{m}_j}^a | a)}, \quad (7)$$

where $w(\mathcal{I}_j)$ refers to the sum of the normalized weights of the particles contained in cluster j and $\mathbf{z}_{\mathbf{m}_j}^a$ stands for a ray casting operation whose simulated laser beams originate from the center of mass \mathbf{m}_j with a sensor orientation described by the sensing action a .

Eqn. (7) specifies how to compute the future expected entropy of a sensing action a based on our clustered subsets. Plugging this into Eqn. (6) we are able at any point in time to select the best action \hat{a} to be executed next.




Fig. 2. MLS map and mobile robot used for the localization experiments. The area represented by this map spans approximately 195 by 146 meters. The blue / dark grey line shows the localized robot poses. The yellow / light grey line shows the pure odometry. The traversed trajectory has a length of 284 meters. The top right part depicts the robot Herbert used for the experiments.

5 Experiments

Our approach has been implemented and tested in real world experiments. The robot used for the experiments is a Pioneer II AT system equipped with a SICK LMS laser range scanner and an AMTEC wrist unit, which is used as a pan/tilt device for the laser (see Figure 2). The experiments are designed to investigate if the active localization approach facilitates mobile robot localization and whether it improves localization performance.

The first set of experiments is designed to evaluate the performance of the active localization approach during a position tracking task. To obtain the data, we steered along a 284 meter long loop in our campus environment. Figure 2 depicts a top view of the MLS Map. The blue / dark grey line shows the localized robot poses. The yellow / light grey line shows the pure odometry. The left image of Figure 3 depicts the average localization error for a tracking experiment with 1,000 particles. As can be seen from the figure, our active approach achieves the same performance as the passive approach as long as the laser range finder is not tilted more than 3° downwards. This is due to the fact that at higher tilt angles the robot perceives less vertical objects and therefore misses the map features that are important for a correct position estimate. The right image depicts the frequency of the tilt angles chosen by our active approach. As can be seen from the histogram our active approach prefers upwards to downwards orientations of the range sensors. Thus the active approach enables the robot to sense the important vertical features of the environment and avoids to obtain range measurements from the ground.

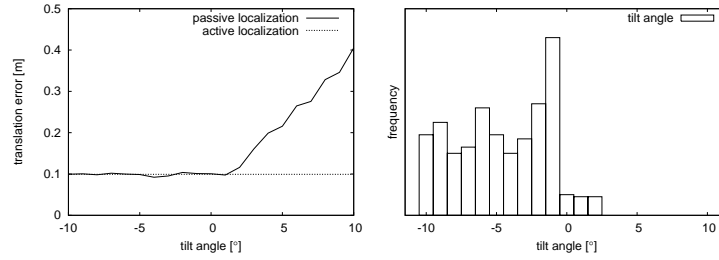


Fig. 3. Localization error of the passive approach using different tilt angles and of the active approach, respectively (left). The right image depicts the frequency of the tilt angles chosen by the active approach.

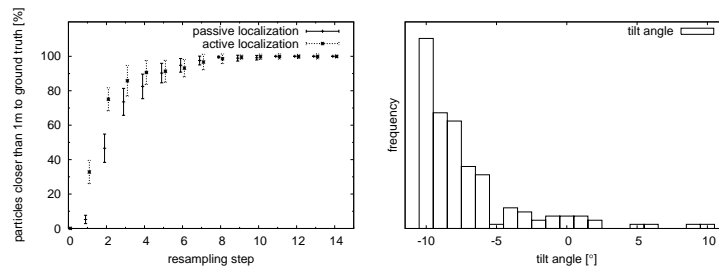


Fig. 4. The left image depicts the convergence of the particles to the true position of the robot with 500,000 particles. The right image depicts the frequency of the tilt angles chosen by the active approach during the global localization task.

Additionally we also carried out experiments, in which we evaluated the convergence of the active localization approach during global localization in a map that spans approximately 195 by 146 meters. Figure 4 depicts the convergence of the particles to the true position of the robot with 500,000 particles. Whereas the x-axis corresponds to the resampling step, the y-axis shows the number of particles in percent that are closer than 1 m to the true position. The figure shows the evolution of these numbers for the passive and active MCL. A t-test revealed that it is significantly better to apply the active approach than the passive one for the global localization task.

6 Conclusions

In this paper, we presented an approach to active Monte Carlo localization with a mobile robot using MLS maps. Our approach actively selects the orientation of the laser range finder to improve the localization results. To speed up the entire process, we apply a clustering operation on the particles and only evaluate potential orientations based on these clusters. In experiments obtained with a robot in a real outdoor environment we analyzed the active control of the range sensor and demonstrated that it leads to more efficient localization results.

Acknowledgment

This work has partly been supported by the German Research Foundation (DFG) within the Research Training Group 1103 and under contract number SFB/TR-8 and by the EC under contract number FP6-IST-027140-BACS.

References

1. R. Kümmerle, R. Triebel, P. Pfaff, and W. Burgard. Monte carlo localization in outdoor terrains using multi-level surface maps. In *Proc. of the International Conference on Field and Service Robotics (FSR)*, Chamonix, France, 2007.
2. R. Triebel, P. Pfaff, and W. Burgard. Multi level surface maps for outdoor terrain mapping and loop closing. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2006.
3. J. Bares, M. Hebert, T. Kanade, E. Krotkov, T. Mitchell, R. Simmons, and W. R. L. Whittaker. Ambler: An autonomous rover for planetary exploration. *IEEE Computer Society Press*, 22(6):18–22, 1989.
4. M. Hebert, C. Caillas, E. Krotkov, I.S. Kweon, and T. Kanade. Terrain mapping for a roving planetary explorer. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, pages 997–1002, 1989.
5. S. Lacroix, A. Mallet, D. Bonnafous, G. Bauzil, S. Fleury and; M. Herrb, and R. Chatila. Autonomous rover navigation on unknown terrains: Functions and integration. *International Journal of Robotics Research*, 21(10-11):917–942, 2002.
6. C. Parra, R. Murrieta-Cid, M. Devy, and M. Briot. 3-d modelling and robot localization from visual and range data in natural scenes. In *1st International Conference on Computer Vision Systems (ICVS)*, number 1542 in LNCS, pages 450–468, 1999.
7. L. P. Kaelbling, A. R. Cassandra, and J. A. Kurien. Acting under uncertainty: Discrete Bayesian models for mobile-robot navigation. In *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Osaka, Japan, 1996.
8. S. Koenig and R. Simmons. Xavier: A robot navigation architecture based on partially observable markov decision process models. In D. Kortenkamp, R. Bonasso, and R. Murphy, editors, *Artificial Intelligence Based Mobile Robotics: Case Studies of Successful Robot Systems*, pages 91–122. MIT Press, 1998.
9. P. Jensfelt and S. Kristensen. Active global localisation for a mobile robot using multiple hypothesis tracking. *IEEE Transactions on Robotics and Automation*, 17(5):748–760, October 2001.
10. A Davison and N. Kita. 3d simultaneous localisation and map-building using active vision for a robot moving on undulating terrain. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition, Kauai*. IEEE Computer Society Press, December 2001.
11. J. M. Porta, J. J. Verbeek, and B. J. A. Kröse. Active appearance-based robot localization using stereo vision. *Auton. Robots*, 18(1):59–80, 2005.
12. D. Fox, W. Burgard, and S. Thrun. Active markov localization for mobile robots. *Journal of Robotics & Autonomous Systems*, 25:195–207, 1998.
13. F. Dellaert, D. Fox, W. Burgard, and S. Thrun. Monte Carlo localization for mobile robots. In *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 1999.
14. A. Doucet, N. de Freitas, and N. Gordon, editors. *Sequential Monte-Carlo Methods in Practice*. Springer Verlag, 2001.
15. L. J. Heyer, S. Kruglyak, and S. Yooseph. Exploring expression data: Identification and analysis of coexpressed genes. *Genome Res.*, 9(11):1106–1115, 1999.