

# Position Estimation for Mobile Robots in Dynamic Environments

Dieter Fox, Wolfram Burgard, Sebastian Thrun<sup>‡</sup>, Armin B. Cremers

Computer Science Department III  
University of Bonn  
Bonn, Germany

<sup>‡</sup>School of Computer Science  
Carnegie Mellon University  
Pittsburgh, PA

## Abstract

For mobile robots to be successful, they have to navigate safely in populated and dynamic environments. While recent research has led to a variety of localization methods that can track robots well in *static* environments, we still lack methods that can robustly localize mobile robots in dynamic environments, in which people block the robot's sensors for extensive periods of time or the position of furniture may change. This paper proposes extensions to Markov localization algorithms enabling them to localize mobile robots even in densely populated environments. Two different filters for determining the "believability" of sensor readings are employed. These filters are designed to detect sensor readings that are corrupted by humans or unexpected changes in the environment. The technique was recently implemented and applied as part of an installation, in which a mobile robot gave interactive tours to visitors of the "Deutsches Museum Bonn." Extensive empirical tests involving datasets recorded during peak traffic hours in the museum demonstrate that this approach is able to accurately estimate the robot's position in more than 98% of the cases even in such highly dynamic environments.

## Introduction

To operate autonomously, mobile robots must know where they are. *Mobile robot localization*, that is the process of determining and tracking the position (location) of a mobile robot relative to its environment, has received considerable attention over the past few years. Accurate localization is a key prerequisite for successful navigation in large-scale environments, particularly when global models are used, such as maps, drawings, topological descriptions, and CAD models (Kortenkamp, Bonasso, & Murphy 1998). As demonstrated by a recent survey of localization methods (Borenstein, Everett, & Feng 1996), the number of existing approaches is diverse. Mobile robot localization techniques can be categorized at least along the following two dimensions: local vs. global approaches, and approaches for static vs. dynamic environments:

**1. Local vs. global localization** Local approaches to localization are designed to compensate odometric error based on sensor data. They usually require that the initial lo-

cation of the robot is known, and are only capable of *tracking* the location of a robot. The majority of existing localization approaches falls into this category. Global approaches are more general. They can localize a robot globally, that is, they can determine its location without knowledge of the initial location and thus can handle the "kidnaped robot problem" (Engelson 1994). Recently, several researchers proposed a new localization paradigm, called *Markov localization*, which enables robots to localize themselves under global uncertainty. Global approaches have two important advantages over local ones: First, the initial location of the robot does not have to be specified and, second, they provide an additional level of robustness, due to their ability to recover from localization failures.

**2. Static vs. dynamic environments** A second dimension along which localization methods can be grouped is concerned with the nature of the environment which they can master. The majority of approaches can only cope with *static* environments, that is, environments where, according to the robot's sensors, the only aspect that may change over time is the robot's own location. However, these approaches are typically brittle in environments where the dynamics are perceived through the robot's sensors. The approach of (King & Weiman 1990) uses cameras pointed towards the ceiling and thus cannot perceive most of the changes that occur in typical office environments. Unfortunately, such an approach is only applicable if the ceiling contains enough structure for accurate position estimation. Thus, the development of methods that can localize a robot in dynamic environments is still an important goal of research on mobile robot navigation.

This paper proposes a localization algorithm that can localize robots in the most difficult of all situations, namely localization under global uncertainty and in highly dynamic environments. The approach is based on Markov localization. Like Markov localization, it localizes robots *probabilistically*, that is, it maintains multiple hypotheses as to where the robot might be, weighted by a numerical probability factor. As a consequence, our approach inherits from Markov localization the ability to localize a robot under global uncertainty (see (Burgard *et al.* 1996)). Markov localization is based on the assumption that the position of the robot is the only state in the world. Unfortunately, this assumption

is violated if not all aspects of the environment are covered by the world model, which is the case for most dynamic environments. Thus, although Markov localization has been found to be robust to occasional dynamical effects (such as people walking by or doors being closed), it typically fails to localize a robot in densely crowded environments. Unlike Markov localization, however, our approach actively filters sensor data to eliminate the damaging effect of sensor data corrupted by external (unmodeled) dynamics. In this paper, we propose and compare two such filters, one that filters sensor data based on entropy change, and one that incorporates additional knowledge concerning the nature of possible corruptions.

In an experimental study, our extended Markov localization is compared to the original Markov localization without filtering. These experiments are conducted using data gathered during a six-day deployment of our mobile robot RHINO in the “Deutsches Museum Bonn” shown in Figure 1(a) (see also (Burgard *et al.* 1998)). Our comparisons show that in such situations, conventional Markov localization fails to track the robot’s location. Our filter techniques, in contrast, successfully accommodate the environment’s dynamics. Additionally, our approach can reliably recover from localization errors.




Fig. 1. (a) RHINO surrounded by visitors and (b) a highly corrupted sensor scan.

The remainder of this paper is organized as follows. After introducing Markov localization in the following section, we will describe our extension, followed by experimental comparisons and a discussion section.

## Markov localization

This section briefly outlines the basic Markov localization algorithm upon which our approach is based. The key idea of Markov localization which has recently been applied with great success at various sites (Nourbakhsh, Powers, & Birchfield 1995; Simmons & Koenig 1995; Kaelbling, Cassandra, & Kurien 1996; Burgard *et al.* 1996) is to compute a probability distribution over all possible locations in the environment. Let  $l = \langle x, y, \theta \rangle$  denote a location in the state space of the robot, where  $x$  and  $y$  are the robot’s coordinates in a world-centered Cartesian reference frame, and  $\theta$  is the robot’s orientation. The distribution  $Bel(l)$  over all locations  $l$  expresses the robot’s subjective belief for being at position  $l$ . Initially,  $Bel(l)$  reflects the initial state of knowl-

edge: if the robot knows its initial position,  $Bel(l)$  is centered on the correct location; if the robot does not know its initial location,  $Bel(l)$  is uniformly distributed to reflect the global uncertainty of the robot. As the robot operates,  $Bel(l)$  is incrementally refined.

Markov localization applies two different probabilistic models to update  $Bel(l)$ , an action model to incorporate movements of the robot into  $Bel(l)$  and a perception model to update the belief upon sensory input.

**Robot motion** is modeled by a conditional probability  $p(l | l', a)$  specifying the probability that a measured movement action  $a$ , when executed at  $l'$ , carries the robot to  $l$ .  $Bel(l)$  is then updated according to the following general formula coming from the domain of Markov chains (Chung 1960):

$$Bel(l) \leftarrow \sum_{l'} P(l | l', a) \cdot Bel(l') \quad (1)$$

The term  $p(l | l', a)$  represents a model of the robot’s kinematics. In our implementation we assume the errors of the odometry to be normally distributed.

**Sensor readings** are integrated according to the well-known Bayesian update formula. Let  $s$  denote a sensor reading and  $p(s | l)$  the likelihood of perceiving  $s$  given that the robot is at position  $l$ , then  $Bel(l)$  is updated according to the following rule:

$$Bel(l) \leftarrow \alpha p(s | l) Bel(l) \quad (2)$$

Here  $\alpha$  is a normalizer ensuring that  $Bel(l)$  sums up to 1 over all  $l$ .

Strictly speaking, both update steps are only applicable if the problem is *Markovian*, that is, if past sensor readings are conditionally independent of future readings given the location of the robot. The Markov assumption thus assumes that the world is static. While in practice, the approach has been applied even in environments that contained people and hence violate the Markov assumption, the experiments reported here indicate that it does not scale to densely populated environments.

In this paper we use a fine-grained grid-based representation of the state space, just like the approach described in (Burgard *et al.* 1996). In all our experiments, the resolution of robot orientation was  $2^\circ$ , and the spatial resolution was 15cm. Different optimization techniques described in (Fox, Burgard, & Thrun to appear) allow the robot to efficiently update such large state spaces in real-time, without restricting the power of the approach in any noticeable way. The primary advantage of the high resolution are the resulting accuracy of position estimates and the ability to incorporate raw data of proximity sensors, which were required in the specific application domain described below.

## Localization in Highly Dynamic Environments

The standard Markov localization approach has been found to be robust in static environments. However, as argued in

the introduction to this paper (and demonstrated in the results section), it is prone to fail in densely populated environments which violate the underlying Markov assumption. In the museum, where the robot is naturally accompanied by crowds of people, this assumption is clearly violated. To illustrate this point, Figure 1(b) shows a typical example situation where RHINO has been projected into the map at its correct position. The lines indicate the current proximity measurements and the different shading of the measurements indicates the two classes they belong to: the black values correspond to static obstacles that are part of the map, whereas others are caused by humans and thus violate the Markov assumption (max-range measurements are not shown).

The proximity of people usually increases the robot’s belief of being close to modeled obstacles, which has the effect that the robot frequently loses track of its position when relying on *all* sensor measurements. Approaches for concurrent estimation of the state of the world and of the position of the robot as proposed in (Gutmann & Schlegel 1996; Lu & Milius 1997; Thrun, Fox, & Burgard to appear), unfortunately, require too many computational resources to be applied on-line or even in real-time. Our approach to solve this problem is to develop filters which select those readings of a complete scan which with high likelihood are not due to static obstacles in the map thus making the system more robust against such kind of noise.

In the following two sections we will describe two different filters aiming at detecting corrupted readings and thus allowing the robot to keep track of its location even in considerably difficult situations, in which more than 50% of all readings are misleading. The first filter is a general method for filtering sensor data in dynamic environments. It selects only those readings that increase the robot’s certainty, which is measured by the entropy of the belief  $Bel(l)$ . The second filter is especially designed for proximity sensors, as it attempts to filter such readings which with high probability are shorter than expected according to the model of the environment and the current belief state  $Bel(l)$  of the robot.

### Entropy filter

The first filter used in our implementation is called *entropy filter*. The *entropy*  $H(l)$  of a distribution over  $l$  is defined by

$$H(l) = - \sum_l Bel(l) \log Bel(l). \quad (3)$$

Entropy is a measure of uncertainty: The larger the entropy, the higher the robot’s uncertainty as to where it is. The *entropy filter* measures the relative change of entropy upon incorporating a sensor reading into the belief  $Bel(l)$ . More specifically, let  $s$  denote the measurement of a sensor (in our case a single range measurement). The change of the *entropy* of  $Bel(l)$  given  $s$  is defined as:

$$\Delta H(l | s) := H(l) - H(l | s). \quad (4)$$

While a positive change of entropy indicates that after incorporating  $s$ , the robot is less certain about its position, a negative change indicates an increase in certainty.

RHINO’s *entropy filter* uses only such sensor measurements  $s$  for which  $\Delta H(l | s) \geq 0$ . Thus, the entropy filter makes robot perception highly selective, in that it considers only sensor readings confirming the robot’s current belief.

### Novelty filter

While the entropy filter makes no assumptions about the nature of the sensor data and the kind of disturbances to expect in dynamic environments, the second filter is especially designed for proximity sensors and detects additional obstacles in the environment. This filter is called *novelty filter*, since it selects sensor readings based on the degree of their “novelty.” To be more specific, a measurement  $s$  is assumed to be “novel” to the robot, if it is reflected by an obstacle not represented in the map. Obviously, for proximity measurements such a case can only be detected if the measurement is *shorter* than expected.

The novelty filter removes those proximity measurements which with probability higher than  $\theta$  (this threshold is set to 0.99 in all experiments) are shorter than expected. Suppose  $d_1, \dots, d_n$  is a discrete set of possible distances measured by a proximity sensor. Let  $p(d_j | l)$  denote the probability of measuring distance  $d_j$  if the sensor detects the next obstacle in the map. This distribution describes the *expected* measurement, and a distribution for laser-range finder given the distance  $o_l$  to the next obstacle is shown by the dashed line in Figure 2 (see (Fox, Burgard, & Thrun to appear) for further discussion of the proximity sensor models we use). Now we can derive  $p_n(d_i | l)$ , namely the probability of  $d_i$  being shorter than expected, by the following equation (c.f. 2):

$$p_n(d_i | l) = 1.0 - \sum_{j < i} p(d_j | l). \quad (5)$$

In order to deal with situations in which the position of the robot is not known with absolute certainty, we average over all possible locations of the robot to obtain the probability that  $d_i$  is shorter than expected:

$$p_n(d_i) = \sum_l p_n(d_i | l) Bel(l). \quad (6)$$

The selection scheme of the novelty filter is to exclude all measurements  $d$  with  $p_n(d) > \theta$ .

The difference between both filters can be characterized as follows: the entropy filter always tries to confirm the current belief of the robot (whether this is right or wrong) while the novelty filter also forces the incorporation of very unlikely sensor measurements (especially too long readings).

## Experimental Results

Localization using the entropy filter was a central component of the tour-guide robot in the Deutsches Museum Bonn. Accurate position estimation was a crucial component, as many




Fig. 2. Probability  $p_n(d_i | l)$  that the sensing  $s_i$  is shorter than the expected sensing.




Fig. 3. Path of the robot in the second dataset.




Fig. 4. Percentage of noisy sensor measurements.

of the obstacles were “invisible” to the robot’s sensors (such as glass cages, metal bars, staircases, and the alike). Only through accurate localization could collisions with those obstacles be avoided (Fox *et al.* 1998). Using Markov localization with entropy filters, our approach led only to a single software-related collision, which involved an “invisible” obstacle and which was caused by a localization error that was slightly larger than a 30cm safety margin. During six days of operation, RHINO traversed approximately 18.5 km at an average speed of approximately 36.6 cm/sec. In this application, the entropy filter was used and the novelty filter was developed post fact, based on an analysis of the collision reported above, in an attempt to prevent similar effects in future installations.

The evidence from the museum is purely anecdotal. We also investigated the merit of the approaches proposed here more systematically, and under even more extreme conditions. In particular, we were interested in the localization results (1) when the environment is densely populated and (2) when the robot suffers from extreme dead-reckoning errors.

## Datasets

Two datasets were used in our comparison, which both were recorded in the museum, and which mainly differed by the amount of disturbances.

1. The first dataset was collected during 2.0 hours of robot motion, during which the robot traversed as much as 1,000 meters. This data was collected when the museum was closed, and the robot guided only remote internet-visitors through the museum. The robot’s top speed was limited to 50cm/sec. Thus, this dataset was “ideal” in that the environment was only extremely sparsely populated, and the robot moved slowly.

2. Figure 3 shows the second dataset, which represents 1,540 meters of robot motion through dense crowds over a period of 4.8 hours. This dataset was collected during peak traffic hours on the most crowded day during the entire exhibition. When collecting this data, the robot was frequently faced with situations as illustrated in Figure 1(a) and (b). The top speed in this dataset was 80cm/sec.

Both datasets consist of logs of odometry and laser-range finder scans collected while the robot moved through the mu-

seum. Using the time stamps in the logs, all tests have been conducted in real-time simulation on a SUN-Ultra-Sparc 1 (177-MHz). The first dataset contained more than 32,000, and the second dataset more than 73,000 laser scans. The reader may notice that only the obstacles shown in black in Figure 3 were actually used for localization; the others were either invisible, or could not be detected reliably.

Figure 4 shows the estimated percentage of corrupted sensor readings over time for both datasets. The dashed line corresponds to the first dataset while the solid line illustrates the corruption of the second (longer) dataset. In the second dataset, more than half of all measurements were corrupted for extended durations of time. These numbers are estimates only; they were obtained by analyzing each laser reading as to whether it could be “explained” by the obstacles represented in the map.

To evaluate the different localization methods, we generated two *reference paths* through nine independent runs for each filter on the datasets (with small random disturbances) to determine the location of *all* sensor scans. Visual inspection made us believe that the resulting reference locations were indeed correct and accurate enough. In order to estimate the *accuracy* of the methods on the second dataset, we selected 118 representative reference positions, for which we manually determined the robot’s location as closely as possible through careful comparison of sensor scans, the robot’s path, and the environment.

## Localization in densely populated environments

In our first series of experiments, we were interested in comparing the localization performance of all three approaches — plain Markov localization, localization with entropy filters, and localization with novelty filters — under normal working conditions.

Table 1 summarizes the results obtained for the different approaches. The first row provides the percentage of *failures* (including 95% confidence intervals) for the different filters on the first dataset. Position estimates were considered as a “failure,” if the estimated location deviated from the reference path by more than 45cm. All three approaches worked nicely for tracking the robot’s position in the empty museum (first dataset), exhibiting only negligible errors in localiza-

Filter	None	Entropy	Novelty
Tracking Ability			
failures <sub>I</sub> [%]	1.6 ±0.4	0.9 ±0.4	0.0 ±0.0
failures <sub>II</sub> [%]	26.8 ±2.4	1.1 ±0.3	1.2 ±0.7
Accuracy in Dataset II			
$\bar{d}$ [cm]	188.9 ±26.9	9.2 ±0.5	11.4 ±3.4
Recovery			
$\bar{t}_{\text{rec}_I}$ [sec]	237 ±27	1779 ±548	188 ±30
$\bar{t}_{\text{rec}_{II}}$ [sec]	269 ±60	1310 ±904	235 ±46

Table 1: Experimental results.

tion. The results obtained for the second, more challenging dataset, however, were quite different. While plain Markov localization failed in 27% of all cases, both filter techniques showed a failure rate well below 2% (see second row). The third row, labeled  $\bar{d}$ , gives the average Euclidean error between the estimated position and the 118 reference positions. Here as well, the gap between conventional Markov localization and our approaches is large. The reader may notice that the accuracy of the filter techniques is higher than the grid resolution of 15cm.

To shed light onto the question as to why Markov localization performs so poorly when compared to the algorithms proposed in this paper, we analyzed the sensor readings that each method considered during localization. Figure 5 shows, for a small fraction of the data, the endpoints of the measurements that were incorporated into the robot’s belief. The figures illustrate that both approaches proposed here manage to focus their attention on the “right” sensor measurements, whereas conventional Markov localization incorporates massive amounts of corrupted (misleading) measurements. Moreover, both filters show similar behavior. As also can be seen in Figure 5, both filter-based approaches produce more accurate results for this example. These results demonstrate that our approach scales much better to populated and dynamic environments than Markov localization.

### Recovery from extreme localization failures

One of the key advantage of the original Markov localization technique lies in its ability to *recover* from extreme localization failures. Re-localization after a failure is often more difficult than global localization from scratch, since the robot has to (1) detect that its current belief is wrong and (2) globally re-localize itself afterwards. Since the filter-based approaches incorporate sensor data selectively, it is not clear that they still maintain the ability to recover from global localization failures.

Our experiments under normal operation conditions did not lead to such failures for the two methods proposed in this paper; thus, we manually introduced such failures into the data to test the robustness of these methods in the extreme. More specifically, in our experiments we “tele-ported” the robot at random points in time to other locations. Techni-

cally, this was done by changing the robot’s orientation by  $180 \pm 90$  degree and shifting it by  $0 \pm 100$  cm, without letting the robot know. These perturbations were introduced randomly, with a probability of 0.005 per meter of robot motion. Obviously, such incidents make the robot lose its position. Each method was tested on 23 differently corrupted datasets. This resulted in an overall of 133 position failures. For each of these failures we measured the time until the methods re-localized the robot correctly. Re-localization was assumed to have succeeded if the distance between the estimated position and the reference position was smaller than 45cm for more than 10 seconds.

The two bottom rows in Table 1 summarize the results for the two datasets.  $\bar{t}_{\text{rec}}$  represents the average time in seconds needed to recover from a situation when the position was lost. Both conventional Markov localization and the extension using novelty filters are relatively efficient in recovering from extreme positioning errors, whereas the entropy filter-based approach is an order of magnitude less efficient. The results illustrate that despite the fact that sensor readings are processed selectively, the novelty filter-based approach recovers as efficiently from extreme localization errors as the conventional Markov approach. These findings are specifically interesting in the light of the fact that in the Deutsches Museum Bonn the entropy-based filter was used, which, according to these results, would have led to poor recovery from extreme failures.

In summary, these experiments suggest that only the localization algorithm with the novelty filter is able to localize robots in densely crowded environments, while retaining the ability to efficiently recover from extreme localization errors.

## Conclusions

This paper proposed an approach for global robot localization that has been demonstrated to reliably localize mobile robots even in extremely challenging dynamic environments. These environments are characterized by the presence of various dynamic effects, such as crowds of people that frequently block the robot’s sensors. Our approach is based on Markov localization, a popular method for mobile robot localization, which provides the ability to recover from arbitrary failures in localization. It extends Markov localization by an approach that filters sensor data, so that the damaging effect of corrupted data is reduced. Two specific filters were proposed and evaluated, one which considers conditional entropy for selecting sensor readings, and one which takes into account additional knowledge about the effects of possible environment dynamics.

Our approach was essential for operating a robot successfully in a crowded museum. Experimental comparisons using data collected there demonstrated that the technique proposed in this paper is superior to state-of-the-art localization methods. These results also demonstrated that by processing sensor readings selectively, one of the proposed approaches




Fig. 5. Estimated and real paths of the robot along with endpoints of integrated sensor measurements using (a) no filter, (b) entropy filter, and (c) novelty filter.

still retains the ability to recover from global failures in localization. Additional tests in our office environment have shown that our technique can deal with situations in which only an outline of the environment is used as a world model (c.f. 6(a)). In this case sensor readings reflected by the furniture (c.f. 6(b)) are successfully filtered out. We believe that these results are essential for operating mobile robots in highly dynamic and densely populated environments.




Fig. 6. (a) Outline used for localization and (b) environment including furniture.

How specific are these results to the problem of mobile robot localization? We believe that the first filter proposed here, the entropy filter, is applicable to a much wider variety of state estimation (and learning) problems in dynamic environments. Loosely speaking, this filter makes robot perception highly selective, in that only sensor readings are considered that confirm the robot's current belief. This filter rests only on two assumptions: First, that the variable to be estimated is represented probabilistically, and second, that sensor readings can be sorted into two bins, one which only contains corrupted readings, and one that contains authentic (non-corrupted) measurements. A promising application of this filter is the compensation of hardware failures of sensors, a problem addressed in (Murphy & Hershberger 1996). Future work will aim at analyzing quantitatively, to what extent this filter can make robots more robust to sensor failures.

## References

- Borenstein, J.; Everett, B.; and Feng, L. 1996. *Navigating Mobile Robots: Systems and Techniques*. A. K. Peters, Ltd.
- Burgard, W.; Fox, D.; Hennig, D.; and Schmidt, T. 1996. Estimating the absolute position of a mobile robot using position probability grids. In *Proc. of the 13th National Conference on Artificial Intelligence*.
- Burgard, W.; Cremers, A.; Fox, D.; Hähnel, D.; Lakemeyer, G.; Schulz, D.; Steiner, W.; and Thrun, S. 1998. The interactive museum tour-guide robot. In *Proc. of the 15th National Conference on Artificial Intelligence*.
- Chung, K. 1960. *Markov chains with stationary transition probabilities*. Berlin: Springer Publisher.
- Engelson, S. 1994. *Passive Map Learning and Visual Place Recognition*. Ph.D. Dissertation, Department of Computer Science, Yale University.
- Fox, D.; Burgard, W.; Thrun, S.; and Cremers, A. B. 1998. A hybrid collision avoidance method for mobile robots. In *Proc. of the IEEE International Conference on Robotics and Automation*. to appear.
- Fox, D.; Burgard, W.; and Thrun, S. Active markov localization for mobile robots. *Robotics and Autonomous Systems*. accepted for publication.
- Gutmann, J.-S., and Schlegel, C. 1996. Amos: Comparison of scan matching approaches for self-localization in indoor environments. In *Proc. of the 1st Euromicro Workshop on Advanced Mobile Robots*. IEEE Computer Society Press.
- Kaelbling, L.; Cassandra, A.; and Kurien, J. 1996. Acting under uncertainty: Discrete bayesian models for mobile-robot navigation. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- King, S., and Weiman, C. 1990. Helpmate autonomous mobile robot navigation system. In *Proc. of the SPIE Conference on Mobile Robots*, 190–198. Volume 2352.
- Kortenkamp, D.; Bonasso, R.; and Murphy, R., eds. 1998. *Artificial intelligence and mobile robots*. Cambridge, MA: MIT Press.
- Lu, F., and Milios, E. 1997. Globally consistent range scan alignment for environment mapping. *Autonomous Robots* 4:333–349.
- Murphy, R. R., and Hershberger, D. 1996. Classifying and recovering from sensing failures in autonomous mobile robots. In *Proc. of the 13th National Conference on Artificial Intelligence*.
- Nourbakhsh, I.; Powers, R.; and Birchfield, S. 1995. DERVISH an office-navigating robot. *AI Magazine* 16(2).
- Simmons, R., and Koenig, S. 1995. Probabilistic robot navigation in partially observable environments. In *Proc. of the International Joint Conference on Artificial Intelligence*.
- Thrun, S.; Fox, D.; and Burgard, W. to appear. A probabilistic approach for concurrent map acquisition and localization for mobile robots. *Machine Learning / Autonomous Robots*.