

## Localization for Mobile Robot Teams Using Maximum Likelihood Estimation

Andrew Howard, Maja J Matarić and Gaurav S Sukhatme

*Robotics Research Laboratory, Computer Science Department, University of Southern California  
ahoward@usc.edu, mataric@usc.edu, gaurav@usc.edu*

### Abstract

*This paper describes a method for localizing the members of a mobile robot team, using only the robots themselves as landmarks. That is, we describe a method whereby each robot can determine the relative range, bearing and orientation of every other robot in the team, without the use of GPS, external landmarks, or instrumentation of the environment. Our method assumes that each robot is able to measure the relative pose of nearby robots, together with changes in its own pose; using a combination of maximum likelihood estimation (MLE) and numerical optimization, we can subsequently infer the relative pose of every robot in the team. This paper describes the basic formalism, its practical implementation, and presents experimental results obtained using a team of four mobile robots.*

### 1 Introduction

This paper describes a method for localizing the members of a mobile robot team, using only the robots themselves as landmarks. That is, we describe a method whereby each robot can determine the relative range, bearing and orientation of every other robot in the team, without the use of GPS, external landmarks, or instrumentation of the environment. Our approach is motivated by the need to localize robots in hostile and sometimes dynamic environments. Consider, for example, a search-and-rescue scenario in which a team of robots must deploy into a damaged structure, search for survivors, and guide rescuers to those survivors. In such environments, localization information cannot be obtained using GPS or landmark-based techniques: GPS is generally unavailable or unreliable due to signal obstructions or multi-path effects, while landmark-based techniques require prior models of the environment that are either unavailable, incomplete or inaccurate. In contrast, by using the robot themselves as landmarks, the method described in this paper can generate good localization information in almost any environment, including those that are undergoing dynamic structural changes. Our only requirement is that the robots are able to maintain at least intermittent line-of-sight contact with one-another.

We make three basic assumptions. First, we assume that each robot is equipped with a proprioceptive *motion sensor* such that it can measure changes in its own pose. Suitable motion sensors can be constructed using either odometry or inertial measurement units. Second, we assume that each robot is equipped with a *robot sensor* such that it can measure the relative pose and identity of nearby robots. Suitable sensors can be constructed using either vision (in combination with color-coded markers) or scanning laser range-finders (in combination with retro-reflective tags). We further assume that the identity of robots is always determined correctly, which eliminates what would otherwise be a combinatorial labeling problem. Finally, we assume that each robot is equipped with some form of transceiver that can be used to broadcast information back to a central location, where the localization is performed. Standard 802.11b wireless network adapters can be used for this purpose. We note in passing that while the implementation described in this paper is entirely centralized, distributed implementations are also possible; see [8].

Given these assumptions, the team localization problem can be solved using a combination of maximum likelihood estimation and numerical optimization. The basic method is as follows. First, we construct a set of estimates  $H = \{h\}$  in which each element  $h$  represents a pose estimate for a particular robot at a particular time. These pose estimates are defined with respect to some *arbitrary* global coordinate system. Second, we construct a set of observations  $O = \{o\}$  in which each element  $o$  represents an observation made by either a motion sensor (in which case  $o$  is the measured change in pose of a single robot) or a robot sensor (in which case  $o$  is the measured pose of one robot, relative to another). Finally, we use numerical optimization to determine the set of estimates  $H$  that is most likely to give rise to the set of observations  $O$ .

Note that, in general, we do not expect robots to use the set of pose estimates  $H$  directly; these estimates are defined with respect to an arbitrary coordinate system whose relationship with the external world is undefined. Instead, each robot uses these estimates to compute the relative pose of the other robots, and uses this ego-centric


viewpoint to coordinate activity

In the remainder of this paper, we describe the basic formalism, its practical implementation, and present results from a controlled experiment conducted with a team of four mobile robots.

## 2 Related Work

Localization is an extremely well studied area in mobile robotics. The vast majority of this research has concentrated on two problems: localizing a single robot using an a priori map of the environment [10, 14, 4], or localizing a single robot while simultaneously building a map [16, 11, 17, 2, 5, 1]. Recently, some authors have also considered the related problem of map building with multiple robots [15]. All of these authors make use of statistical or probabilistic techniques; the common tools of choice are Kalman filters, maximum likelihood estimation, expectation maximization, and Markovian techniques (using grid or sample-based representations for probability distributions). The team localization problem described in this paper bears many similarities to the simultaneous localization and map building problem, and is amenable to similar mathematical treatments. In this context, the work of Lu and Milios [11] should be noted. These authors describe a method for constructing globally consistent maps by enforcing pairwise geometric relationships between individual range scans; relationships are derived either from odometry, or from the comparison of range scan pairs. MLE is used to determine the set of pose estimates that best accounts this set of relationships. Our mathematical formalism is very similar to that described by these authors, even though it is directed towards a somewhat different objective, i.e., the localization of mobile robot teams, rather than the construction of globally consistent maps.

Among those who have considered the specific problem of team localization are [13] and [3]. Roumeliotis and Bekey present an approach to multi-robot localization in which sensor data from a heterogeneous collection of robots are combined through a single Kalman filter to estimate the pose of each robot in the team. It should be noted, however, that this method still relies entirely on external landmarks; no attempt is made to sense other robots or to use this information to constrain the pose estimates. In contrast, Fox *et al.* describe an approach to multi-robot localization in which each robot maintains a probability distribution describing its own pose (based on odometry and environment sensing), but is able to refine this distribution through the observation of other robots. This approach extends earlier work on single-robot probabilistic localization techniques [4]. The authors avoid the curse of dimensionality (for  $N$  robots, one must maintain a  $3N$  dimensional distribution) by factoring the distribution into  $N$  separate components (one for each robot). While this step makes the algorithm



**Figure 1:** An illustration of the basic formalism. The figure shows two robots,  $r_1$  and  $r_2$ , traveling from left to right and observing each other exactly once. The robots' activity is encoded in the graph, with nodes representing pose estimates and arcs representing observations. Two observations are highlighted: a motion observation for robot  $r_1$  (between times  $t_1$  and  $t_2$ ) and a robot observation at time  $t_2$  (between robots  $r_1$  and  $r_2$ ).

tractable, it also results in some loss of expressiveness.

Finally, a number of authors [9, 12, 6] have considered the problem of team localization from a somewhat different perspective. These authors describe *cooperative* approaches to localization, in which team members actively coordinate their activities in order to reduce cumulative odometric errors. While our approach does not require such explicit cooperation on the part of robots, the accuracy of localization can certainly be improved by the adoption of such strategies; we will return to this topic briefly in the final sections of the paper.

## 3 Formalism

We formulate the team localization problem as follows. Let  $h$  denote the pose estimate for a particular robot at a particular time, and let  $H = \{h\}$  be the set of all such estimates. Similarly, let  $o$  denote an observation made by some sensor, and let  $O = \{o\}$  be the set of all such observations. Our aim is to determine the set of estimates  $H$  that maximizes the probability of obtaining the set of observations  $O$ ; i.e., we seek to maximize the conditional probability  $P(O | H)$ . If we assume that observations are statistically independent, we can write this probability as:

$$P(O | H) = \prod_{o \in O} P(o | H) \quad (1)$$

where  $P(o | H)$  is the probability of obtaining the individual measurement  $o$ , given the estimates  $H$ . Taking the log of both sides, we can rewrite this equation as:

$$U(O | H) = \sum_{o \in O} U(o | H) \quad (2)$$

where  $U(O | H) = -\log P(O | H)$  and  $U(o | H) = -\log P(o | H)$ . This latter form is somewhat more efficient for numerical optimization. Our aim is now to find the set of estimates  $H$  that *minimizes*  $U(O | H)$ .

We make the following definitions. Let each estimate  $h$  be denoted by a tuple of the form  $h = (\hat{q}; r, t)$  where  $\hat{q}$  is the *absolute pose estimate* for robot  $r$  at time  $t$ . Note that it is the value of  $\hat{q}$  that we are trying to estimate;  $r$  and  $t$  are constants used for book-keeping purposes only. Let each observation  $o$  be denoted by a tuple of the form  $o = (\mu, \Sigma; r_a, t_a; r_b, t_b)$  where  $\mu$  is the measured pose of robot  $r_b$  at time  $t_b$ , relative to robot  $r_a$  at time  $t_a$ ; henceforth, we will refer to  $\mu$  as a *relative pose measurement*. The  $\Sigma$  term is a covariance matrix representing the uncertainty in this measurement. Recall that each robot is assumed to be equipped with both motion and robot sensors. Each measurements from the motion sensors can be encoded using an observation of the form  $o = (\mu, \Sigma; r_a, t_a, r_a, t_b)$  where  $\mu$  is the measured change in pose for robot  $r_a$  between times  $t_a$  and  $t_b$ . Similarly, each measurement from the robot sensors can be encoded using an observation of the form  $o = (\mu, \Sigma; r_a, t_a, r_b, t_a)$  where  $\mu$  is the measured pose of robot  $r_b$  relative to robot  $r_a$ , for a measurement taken at time  $t_a$ .

One can visualize these definitions using a directed graph, as shown in Figure 1. We associate each estimate  $h$  with a node in the graph, and each observation  $o$  with an arc. Each node may have both outgoing arcs, corresponding to observations in which the node is the *observer*, and incoming arcs, corresponding to observations in which the node was the *observee*. Motion observations join nodes representing the same robot at two different points in time, while robot observations join nodes representing two different robots at the same point in time, as indicated in the figure.

If we assume that the measurement uncertainties are normally distributed, the conditional log-probability  $U(o | H)$  must be given by the quadratic expression:

$$U(o | H) = \frac{1}{2}(\mu - \hat{\mu})^T \Sigma (\mu - \hat{\mu}) \quad (3)$$

where  $\mu$  is the relative pose measurement defined above, and  $\hat{\mu}$  is the corresponding *relative pose estimate*; i.e.  $\hat{\mu}$  is the *estimated* pose of robot  $r_b$  at time  $t_b$ , relative to robot  $r_a$  at time  $t_a$ . The relative pose estimate  $\hat{\mu}$  is derived from a pair of absolute pose estimates  $\hat{q}_a$  and  $\hat{q}_b$  via some coordinate transformation  $\Gamma$ :

$$\hat{\mu} = \Gamma(\hat{q}_a, \hat{q}_b) \quad (4)$$

where  $\hat{q}_a$  and  $\hat{q}_b$  describe the absolute pose estimates for robot  $r_a$  at time  $t_a$ , and robot  $r_b$  at time  $t_b$ , respectively. The specific form of  $\Gamma$  depends on the dimensionality of the localization problem (e.g., 2D versus 3D) and on the particular representation chosen for both absolute and relative poses (e.g., Cartesian versus polar coordinates, or cylindrical versus spherical coordinates).

Given Equations 2 and 3, together with an appropriate definition for  $\Gamma$ , one can determine the set of poses  $\hat{q}$

that minimizes  $U(O | H)$  using standard numerical optimization techniques. The selection of an appropriate algorithm is driven largely by the form of  $\Gamma$ , which is generally non-linear but differentiable. This rules out fast linear techniques, but does permit gradient-based techniques such as steepest descent or conjugate gradient algorithms. In practice, we have found both these algorithms to be highly effective, with the conjugate gradient algorithm having the advantage of being significantly faster (albeit at the expense of greater complexity).

The formalism described above is quite general, and can be applied to localization problems in two, three, or more dimensions. The specific problem of localization in a plane (in which robots have two degrees of translational freedom and one of rotation) can be solved using a straight-forward application of this general formalism; see [7] for details.


### 3.1 Practical Implementation

Since the dimensionality of optimization problem that must be solved scales linearly with the size of  $H$ , and the computational cost of each step in this optimization process scales linearly with the size of  $O$ , it is necessary, in practice, to bound both the number of estimates in  $H$  and the number of observations in  $O$ . We use three basic methods for constraining the size of these sets: we remove any estimates or observations that have exceeded a certain age, we discard similar observations, and we limit the rate at which pose estimates are generated. The first two methods are both simple and well-defined; information that is very old or highly repetitive can often be discarded with minimal impact on localization accuracy. The third of these methods, however, is somewhat more complicated, and involves some extensions to the formalism described in the previous section. Rather than attempting to estimate the pose of each robot at every point in time, we instead estimate the pose of each robot at only a few discrete points in time, and use information from the motion sensors to ‘fill the gaps’ between these estimates. In effect, we assume that the motion sensors produce relatively good pose estimates that only require occasional corrections.

Let  $\hat{p}$  be the *interpolated pose estimate* for robot  $r$  at time  $t$ ; this estimate is given by:

$$\hat{p} = \Gamma^{-1}(\hat{q}, m) \quad (5)$$

where  $\hat{q}$  is the most recent absolute pose estimate for robot  $r$  in  $H$ , and  $m$  is the *measured change in pose* that has occurred since that estimate was generated;  $\Gamma^{-1}$  is a coordinate transformation that maps from relative to absolute coordinates. This definition is illustrated in Figure 2. With this extension, most of the observations made by the robots will *not* occur at the times represented by the pose estimates in  $H$ . We must therefore extend our observation model by modifying definition of  $\hat{\mu}$  given in



**Figure 2:** An illustration of the extended formalism. The figure shows two robots,  $r_1$  and  $r_2$ , traveling from left to right and observing each other exactly once. The robots’ activity is encoded in the graph, with nodes representing pose estimates and arcs representing observations. Also shown are the interpolated pose estimates  $\hat{p}_1$  and  $\hat{p}_2$  for each of the robots at time  $t$ .

Equation 4. Specifically, we must replace the absolute pose estimates  $(\hat{q}_a, \hat{p}_b)$  with interpolated pose estimates; i.e.:

$$\hat{\mu} = \Gamma(\hat{p}_a, \hat{p}_b) \quad (6)$$

where  $\hat{p}_a$  and  $\hat{p}_b$  are the interpolated pose estimates for robot  $r_a$  at time  $t_a$  and robot  $r_b$  and time  $t_b$ , respectively.

This extended formalism has the attractive feature of allowing us approximate the information provided by the motion sensors to an arbitrary degree of fidelity (rather than simply discarding the information). Thus we are free to trade-off dimensionality (and hence optimization speed) against localization accuracy.

## 4 Validation Experiment

We have conducted a controlled experiment aimed at determining the accuracy of the team localization algorithm described in this paper. The experiment was conducted using a team of four Pioneer 2DX mobile robots equipped with SICK LMS200 scanning laser range-finders. Each robot was also equipped with a pair of retro-reflective ‘totem-poles’ as shown in Figure 3(a). These totem-poles can be detected from a wide range of angles using the SICK lasers (which can be programmed to return intensity information in addition to range measurements). This arrangement allows each robot to detect the presence of other robots and to determine both their range (to within a few centimeters) and bearing (to within a few degrees). Orientation can also be determined to within a few degrees, but is subject to a  $180^\circ$  ambiguity. This arrangement does *not* allow individual robots to be identified. Given the ambiguity in both orientation and identity, it was necessary to manually label the data for this experiment.

The team was placed into the environment shown in Figure 3(b) and each robot executed a simple wall following

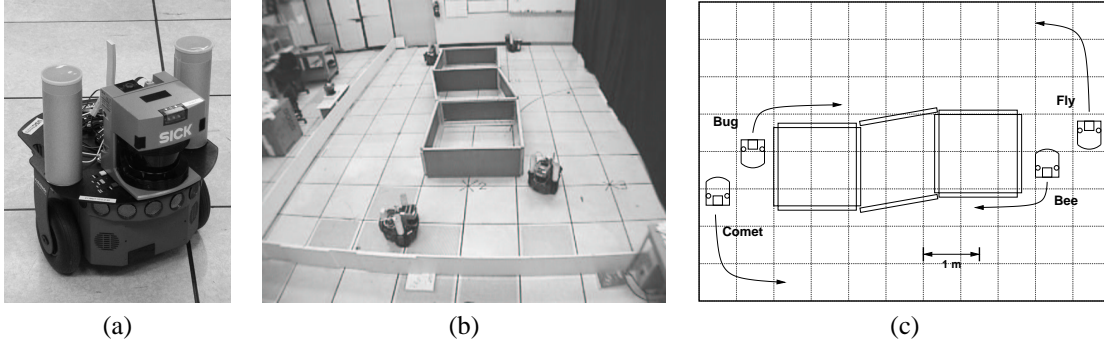
algorithm. Two robots followed the inner wall, and two followed the outer wall. The robots were arranged such that at no time were the two robots on outer wall able to directly sense each other. The structure of the environment was modified a number of times during the course of the experiment. At time  $t = 265$  sec, for example, the inner wall was modified to form two separate ‘islands’, with one robot circumnavigating each. The original structure was later restored, then broken, then restored again.

The accuracy of the algorithm was determined by comparing the robot’s *relative pose estimates* with their corresponding true values (as determined by an external ground-truth system). Thus, we define the *average range error*  $\epsilon_r$  to be:


$$\epsilon_r^2(t) = \frac{1}{N(N-1)} \sum_{r_a} \sum_{r_b} (\hat{\mu}_r - \bar{\mu}_r)^2 \quad (7)$$

where  $\mu_r$  is the estimated range of robot  $r_b$  relative to robot  $r_a$  at time  $t$ , and  $\bar{\mu}_r$  is the true range of robot  $r_b$  relative to robot  $r_a$  at the same time. The summation is over all pairs of robots and the result is normalized by the number of robots  $N$  to generate an average result. One can define similar measures for the bearing error  $\epsilon_\psi$  and orientation error  $\epsilon_\phi$ . Collectively, these error terms measure the average accuracy with which robots are able to determine each other’s relative pose. Note that we make no attempt to compare the *absolute* pose estimates  $\{h\}$  against some ‘true’ value; these estimates are defined with respect to an arbitrary coordinate system which renders such comparison meaningless.

The qualitative results for this experiment are summarized in Figure 4, which contains a series of ‘snap-shots’ of the experiment. Each snap-shot shows the estimated pose of the robots at a particular point in time, overlaid with the corresponding laser scan data. Note that these are snap-shots of *live* data, not cumulative maps of *stored* data. At time  $t = 0$ , the relative pose of the robots is completely unknown, the snap-shot at this time is therefore incoherent; the pose of the robots is largely random, and the laser scans are completely mis-aligned. In the interval  $0 < t < 12$  sec, the robots commence wall following. The robots Fly and Comet follow the outer wall, while Bee and Bug follow the inner wall. By time  $t = 12$  sec, both of the robots following the outer wall have observed both of the robots following the inner wall. As the snap-shot from this time indicates, there is now sufficient information to fully constrain the relative poses of the robots, and to correctly align the laser scan data. It should be noted that the two robots on the outer wall can correctly determine each other’s pose, even though they have never seen each other. At time  $t = 265$  sec, the environment is modified, with the inner wall being re-structured to form two separate islands. The two robots following the inner wall now follow different paths, but the localization is unaffected, as shown in the



**Figure 3:** (a) A Pioneer 2DX equipped with a SICK LMS200 scanning laser range-finder and a pair of retro-reflective totem-poles. (b) The arena for the experiment; the central island is constructed from partitions that can be re-arranged during the course of the experiment. The dimensions of the environment are  $7\text{m} \times 5\text{m}$ . (c) Robot behavior: robots Fly and Comet follow the outer wall, robots Bee and Bug follow the inner wall(s).



**Figure 4:** Experimental snap-shots. Each sub-figure shows the estimated pose of the robots at a particular point in time, with the corresponding laser scan data overlaid. Arrows denote the observation of one robot by another. Note that these are snap-shots of live data; they are not cumulative maps of stored data.

snap-shot at time  $t = 272$  sec. The algorithm described in this paper is completely indifferent to such structural changes in the environment.


The quantitative results for this experiment are summarized in Figure 5, which plots the average range, bearing and orientation errors for the team. At time  $t = 0$  sec, the relative pose of the robots is completely unknown, and the errors are consequently high. By time  $t = 20$  sec, however, the robots have gathered sufficient information to fully constrain their relative pose, and the errors have fallen to more modest values. Over the duration of the experiment, the range error oscillates around  $5.5 \pm 5.2$  cm, while the bearing and orientation errors oscillate around  $1.7 \pm 0.7^\circ$  and  $1.9 \pm 0.6^\circ$  respectively.

The magnitude of these errors can be attributed to two key factors. First, there is some uncertainty in the relative pose measurements made by the laser-range-finder/retro-reflector combination. These uncertainties are difficult to characterize precisely, but are of the order of  $\pm 2.5$  cm. Second, and more significantly, there are uncertainties associated with the temporal synchronization of the laser and odometric measurements. Our low-level implementation is such that the time at which events occur can only be measured to the nearest 0.1 s; in this time, the robot

may travel 2 cm and/or rotate through  $3^\circ$ , which will significantly affect the results.

We ascribe the *variation* seen in the error plots to two different factors. First, we expect that the error will rise during those periods in which the robots cannot see each other and localization is reliant on odometry alone. The odometric accuracy of the robots used in this experiment varies from quite good to quite poor: drift rates for orientation vary from  $2.5^\circ/\text{revolution}$  on Fly to  $30^\circ/\text{revolution}$  on Bug. Second, we expect that errors will fall during those periods when robots are observing one another. This fall, however, may be preceded by a ‘spike’ in the error term; this spike is an artifact produced by the optimization algorithm, which may take several cycles (each cycle is 0.1 s) to incorporate the new data and generate self-consistent results.

Finally, we note that there is a major spike in the plot at around  $t = 300$  sec. This spike corresponds to a collision that occurred between robots Bee and Bug following the first structural change in the environment. As a result of this collision, the robots had to be manually repositioned, leading to gross errors in both robot’s odometry. Nevertheless, as the plot indicates, the system was able to quickly recover.



**Figure 5:** Plots showing the relative pose error as a function of time. The three plots show the average range, bearing and orientation errors, respectively.

## 5 Conclusion

The experiment described in the previous section validates several key capabilities of the team localization method described in this paper. The method does not require external landmarks, it does not require any of the robots to remain stationary. It is robust to changes in the environment, and robots can use transitive relationships to infer the pose of robots they have never seen. In addition, the accuracy of the localization – while not outstanding – is certainly good enough to facilitate many forms of cooperative behavior.

Several aspects of both the general method require further experimental analysis. For example, we have not yet analyzed the impact of local minima (which necessarily plague any non-trivial numerical optimization problem) and we have not fully characterized the scaling properties of the algorithm (although we have previously demonstrated this algorithm working in simulation with up to 20 robots [7]).

In closing, we note that the mathematical formalism presented in this paper can be extended in a number of interesting directions. We can, for example, define a covariance matrix that measures the relative uncertainty in the pose estimates between pairs of robots. This matrix can then be used as a *signal* to actively control the behavior of robots. Thus, for example, if two robots need to cooperate, but their relative pose is not well known, they can undertake actions (such as seeking out other robots) that will reduce this uncertainty.

## Acknowledgments

This work is supported in part by the DARPA MARS Program grant DABT63-99-1-0015, ONR grant N000140110354, and ONR DURIP grant N00014-00-1-0638.

## References

- [1] M. W. M. G. Dissanayake, P. Newman, S. Clark, H. F. Durrant-Whyte, and M. Csorba. A solution to the simultaneous localization and map building (SLAM) problem. *IEEE Transactions on Robotics and Automation*, 17(3):229–241, 2001.
- [2] T. Duckett, S. Marsland, and J. Shapiro. Learning globally consistent maps by relaxation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 4, pages 3841–3846, San Francisco, U.S.A., 2000.
- [3] D. Fox, W. Burgard, H. Kruppa, and S. Thrun. A probabilistic approach to collaborative multi-robot localization. *Autonomous Robots, Special Issue on Heterogeneous Multi-Robot Systems*, 8(3):325–344, 2000.
- [4] D. Fox, W. Burgard, and S. Thrun. Markov localization for mobile robots in dynamic environments. *Journal of Artificial Intelligence Research*, 11:391–427, 1999.
- [5] M. Golfarelli, D. Maio, and S. Rizzi. Elastic correction of dead reckoning errors in map building. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 2, pages 905–911, Victoria, Canada, 1998.
- [6] A. Howard and L. Kitchen. Cooperative localisation and mapping. In *International Conference on Field and Service Robotics (FSR99)*, pages 92–97, 1999.
- [7] A. Howard, M. J. Matarić, and G. S. Sukhatme. Localization for mobile robot teams: A maximum likelihood approach. Technical Report IRIS-01-407, Institute for Robotics and Intelligent Systems Technical Report, University of Southern California, 2001.
- [8] A. Howard, M. J. Matarić, and G. S. Sukhatme. Localization for mobile robot teams: A distributed MLE approach. In *Proceedings of the 8th International Symposium on Experimental Robotics (ISER'02)*, Sant' Angelo d' Ischia, Italy, July 2002. To appear.
- [9] R. Kurazume and S. Hirose. An experimental study of a cooperative positioning system. *Autonomous Robots*, 8(1):43–52, 2000.
- [10] J. J. Leonard and H. F. Durrant-Whyte. Mobile robot localization by tracking geometric beacons. *IEEE Transactions on Robotics and Automation*, 7(3):376–382, 1991.
- [11] F. Lu and E. Milios. Globally consistent range scan alignment for environment mapping. *Autonomous Robots*, 4:333–349, 1997.
- [12] I. M. Rekleitis, G. Dudek, and E. Milios. Multi-robot exploration of an unknown environment: efficiently reducing the odometry error. In *Proc. of the International Joint Conference on Artificial Intelligence (IJCAI)*, volume 2, pages 1340–1345, 1997.
- [13] S. I. Roumeliotis and G. A. Bekey. Collective localization: a distributed kalman filter approach. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 3, pages 2958–2965, San Francisco, U.S.A., 2000.
- [14] R. Simmons and S. Koenig. Probabilistic navigation in partially observable environments. In *Proceedings of International Joint Conference on Artificial Intelligence*, volume 2, pages 1080–1087, 1995.
- [15] S. Thrun. A probabilistic online mapping algorithm for teams of mobile robots. *International Journal of Robotics Research*, 20(5):335–363, 2001.
- [16] S. Thrun, D. Fox, and W. Burgard. A probabilistic approach to concurrent mapping and localisation for mobile robots. *Machine Learning*, 31(5):29–55, 1998. Joint issue with *Autonomous Robots*.
- [17] B. Yamauchi, A. Shultz, and W. Adams. Mobile robot exploration and map-building with continuous localization. In *Proceedings of the 1998 IEEE/RSJ International Conference on Robotics and Automation*, volume 4, pages 3175–3720, San Francisco, U.S.A., 1998.