In "Experimental Robotics VIII", B. Siciliano and P. Dario, eds., Springer-Verlaag, 2003, 146-155

Localization for Mobile Robot Teams: A Distributed MLE Approach

Andrew Howard, Maja J Matarić, and Gaurav S Sukhatme

Robotics Research Laboratories, Department of Computer Science,
University of Southern California, Los Angeles, CA 90089-0781.
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. We assume that robots are equipped with sensors that allow them to measure the relative pose and identity of nearby robots, as well as sensors that allow them to measure changes in their own pose. Using a combination of maximum likelihood estimation and distributed numerical optimization, we can, for each robot, estimate the relative range, bearing, and orientation of every other robot in the team. This paper describes the basic formalism, its distributed 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. We also describe a distributed implementation of this method that has the potential to scale to large teams, and to be robust to the failure or destruction of individual robots.

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 four basic assumptions. First, we assume that each robot is equipped with a proprioceptive *motion detector* such that it can measure changes in its own pose. Suitable motion detectors can be constructed using either odometry or inertial measurement units. Second, we assume that each robot is equipped with a *robot* *detector* 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 bar-codes). 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 messages to every other robot in the team. Standard IEEE 802.11b wireless network adapters can be used for this purpose.

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 $X = \{x\}$ in which each element x 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 $M = \{m\}$ in which each element m represents an observation made by a motion detector, and a set of observations $O = \{o\}$ in which each element o represents an observation made by a robot detector. Finally, we use numerical optimization to determine the set of estimates X that is most likely to give rise to the combined set of observations (M, O). Note that this method effectively 'unrolls' the time component, creating a static estimation problem over some bounded time interval. Note also that we do not expect robots to use the set of pose estimates X directly; these estimates are defined with respect to an arbitrary coordinate system whose relationship with the external environment is undefined. Instead, each robot uses these estimates to compute the pose of every other robot *relative to itself*, and uses this egocentric viewpoint to coordinate activity.

The localization method described above can be implemented in an entirely distributed manner. Each robot is given responsibility for maintaining and optimizing its own pose estimates, while broadcast communication is used to ensure consistency between the pose estimates generated by different robots. In effect, the algorithm partitions the set of poses X into N non-intersecting subsets (one for each robot), which are then optimized in parallel. The final result is comparable to that obtained using a single centralized optimization algorithm.

In the remainder of paper, we describe both the basic formalism and its distributed 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 [10,15,4], or localizing a single robot while simultaneously building a map [17,11,18,2,5,1]. Recently, some authors have also considered the related problem of map building with multiple robots [16]. All of these authors make use of statistical or probabilistic techniques of one sort or another; the

In "Experimental Robotics VIII", B. Siciliano and P. Dario, eds., Springer-Verlaag, 2003, 146-155

common tools of choice are Kalman filters, maximum likelihood estimation (MLE), expectation maximization (EM, and Markovian techniques.

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 for this set of relationships.

Among those who have considered the problem of *cooperative* localization are Roumeliotis and Bekey [14] and Fox *et al.* [3]. Roumeliotis and Bekey present an approach 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. They show how this centralized Kalman filter can be broken down into N separate Kalman filters (one for each robot) to allow for distributed processing. In a similar vein, Fox *et al.* describe an approach 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 tractable, it does result in some loss of expressiveness.

Finally, a number of authors [9,13,6] have described approaches in which team members actively coordinate their movements 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.

3 Formalism

We formulate the team localization problem as follows. Let x_t^r denote the *absolute pose estimate* for robot r at time t, and let X denote the set of all such estimates. Let $m_{tt'}^r$ denote an observation made by a motion detector describing the change in pose of robot r between times t and t'. Let M denote the set of all such observations. Let $o_t^{rr'}$ denote an observation made by a robot detector at time t, in which robot r measures the relative pose of robot r'. Let O denote the set of all such observations. These definitions are illustrated in Figure 1. Each estimate x_t^r can be thought of as a node in a graph, and each observation $m_{tt'}^r$ or $o_t^{rr'}$ can be thought of as a link between two nodes. Thus, 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.

In "Experimental Robotics VIII", B. Siciliano and P. Dario, eds., Springer-Verlaag, 2003, 146-155

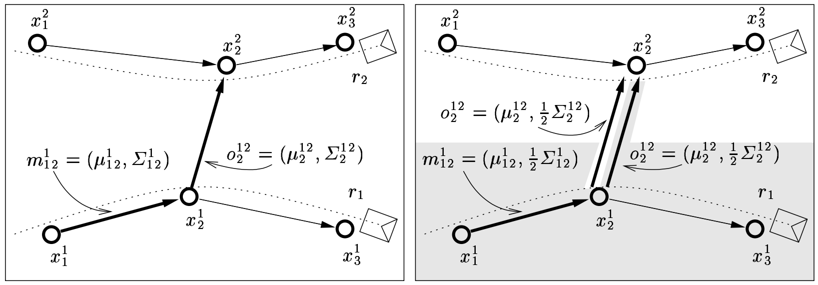

(a) **(b)**

Fig. 1. (a) 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 $r = 2$). **(b)** Graph decomposition for the distributed implementation. Note that each robot carries a duplicate of the robot observation o_2^{12}

Our aim is to determine the set of pose estimates X that maximizes the probability of obtaining the set of observations (M, O); i.e., we seek to maximize the conditional probability $P(M, O \mid X)$. If we assume that observations are statistically independent, we can write this probability as:

$$P(M, O \mid X) = \prod_{m_{tt'}^r \in M} P(m_{tt'}^r \mid x_t^r, x_{t'}^r) \prod_{o_t^{rr'} \in O} P(o_t^{rr'} \mid x_t^r, x_t^{r'}) \tag{1}$$

where $P(m_{tt'}^r \mid x_t^r, x_{t'}^r)$ is the probability of obtaining the individual motion observation $m_{tt'}^r$, given that the estimated initial pose for robot r is x_t^r, and the estimated final pose for the same robot is $x_{t'}^r$. Note that we have made the additional (but not unreasonable) assumption that this probability is independent of other pose estimates. In a similar vein, $P(o_t^{rr'} \mid x_t^r, x_t^{r'})$ specifies the probability of obtaining the individual robot observation $o_t^{rr'}$, given that the estimate pose for the robot r making the observation is x_t^r, and the estimated pose for the robot r' being observed is $x_t^{r'}$. Taking the logarithm of 1, we can write:

$$U(M, O \mid X) = \sum_{m_{tt'}^r \in M} U(m_{tt'}^r \mid x_t^r, x_{t'}^r) + \sum_{o_t^{rr'} \in O} U(o_t^{rr'} \mid x_t^r, x_t^{r'}) \tag{2}$$

where

$$\begin{aligned} U(M, O \mid X) &= -\log P(M, O \mid X) \\ U(m_{tt'}^r \mid x_t^r, x_{t'}^r) &= -\log P(m_{tt'}^r \mid x_t^r, x_{t'}^r) \\ U(o_t^{rr'} \mid x_t^r, x_t^{r'}) &= -\log P(o_t^{rr'} \mid x_t^r, x_t^{r'}) \end{aligned} \tag{3}$$

This latter form is somewhat more convenient for numerical optimization. Our aim now is to find the set of estimates X that *minimizes* $U(M, O \mid X)$, for which we must determine the form of the conditional log-probabilities $U(m^r_{tt'} \mid x^r_t, x^r_{t'})$ and $U(o^{rr'}_t \mid x^r_t, x^{r'}_t)$.

If we assume that the uncertainty in motion observations is normally distributed in *some* coordinate system, we can describe each motion observation using a tuple of the form $m^r_{tt'} = (\mu^r_{tt'}, \Sigma^r_{tt'})$ where $\mu^r_{tt'}$ is a relative pose measurement and $\Sigma^r_{tt'}$ is a covariance matrix representing the uncertainty in this measurement. The conditional log-probability for such observations is given by:

$$U(m^r_{tt'} \mid x^r_t, x^r_{t'}) = \frac{1}{2}(\mu^r_{tt'} - y^r_{tt'})^T \Sigma^r_{tt'} (\mu^r_{tt'} - y^r_{tt'}) \qquad (4)$$

where $y^r_{tt'}$ is a *relative pose estimate* describing the estimated change in pose of robot r between times t and t'. The relative pose estimate is derived from the absolute pose estimates x^r_t and $x^r_{t'}$ via some coordinate transform Γ_m:

$$y^r_{tt'} = \Gamma_m(x^r_t, x^r_{t'}) \qquad (5)$$

The specific form of Γ_m depends on the coordinate system chosen to represent the absolute pose estimates X and the motion observations M. For the standard *planar* localization problem, each pose estimate x^r_t is 3-vector describing the robot's position and orientation, and each motion observation $m^r_{tt'}$ is a 3-vector describing the range, bearing, and orientation of the robot at time t', relative to its pose at time t. Hence the coordinate transform Γ_m maps from polar to Cartesian coordinates. See [8] for a detailed treatment of the planar localization problem.

Robot observations are handled in a similar manner; each observation is described using a tuple of the form $o^{rr'}_t = (\mu^{rr'}_t, \Sigma^{rr'}_t)$ where $\mu^{rr'}_t$ is the relative pose of robot r', as measured by robot r at time t; $\Sigma^{rr'}_t$ is the covariance matrix representing the uncertainty in this measurement. The conditional log-probability is given by:

$$U(o^{rr'}_t \mid x^r_t, x^{r'}_t) = \frac{1}{2}(\mu^{rr'}_t - y^{rr'}_t)^T \Sigma^{rr'}_t (\mu^{rr'}_t - y^{rr'}_t) \qquad (6)$$

where $y^{rr'}_t$ is the estimated pose of robot r' relative to robot r at time t. The relative pose estimate is derived from the absolute pose estimates x^r_t and $x^{r'}_t$ via some coordinate transform Γ_o:

$$y^{rr'}_t = \Gamma_o(x^r_t, x^{r'}_t) \qquad (7)$$

As with the motion observations, the specific form of Γ_o depends on the coordinate systems being used. In the planar localization problem, $o^{rr'}_t$ is a 3-vector describing the range, bearing and orientation of robot r' relative to robot r, and hence Γ_o maps from polar to Cartesian coordinates.

Given appropriate definitions for Γ_m and Γ_o, one can determine the optimal set of pose estimates X using a standard numerical optimization algorithm. The selection of an appropriate algorithm is driven largely by the form of these coordinate transforms, which are, in general, non-linear by differentiable. This rules out fast linear algorithms, but allows gradient-based algorithms such as steepest descent. In

In "Experimental Robotics VIII", B. Siciliano and P. Dario, eds., Springer-Verlaag, 2003, 146-155

practice, we use a conjugate gradient algorithm [12] for optimization; this algorithm is somewhat more complex than a steepest descent algorithm, but has the advantage of being much faster. See [7] for details.

4 Distributed Implementation

The distributed implementation of this approach relies heavily on communication between the robots. Consider once again the graph-based visualization shown in Figure 1(a). We can decompose this graph into a set of subgraphs – one for each robot – as shown in Figure 1(b). Each robot maintains a set nodes representing its own pose at various points in time, and a set of links representing its motion observations. Each robot also maintains a set of links representing *robot* observations in which it was either the observer or the observed; these links connect the otherwise separate sub-graphs. The robot making the observation is responsible for transmitting this fact to the observed, which then generates its own copy of the link. For the sake of mathematical consistency, each copy of the duplicated observation is weighted by a factor of 0.5.

Using this decomposed representation, the optimization task can also be decomposed: each robot is responsible for optimizing $U(M, O \mid X)$ with respect to its own pose estimates, while treating the remaining pose estimates as fixed. To maintain consistency between subgraphs, robots periodically broadcast their updated pose estimates. This is, in effect, a simple parallel optimization algorithm, with each robot attempting to optimize $U(M, O \mid X)$ for a certain subset of X. We assert (without proof) that this distributed algorithm produces results that are comparable to those obtained using a centralized optimization algorithm.

Note that this distributed implementation requires the transmission of two kinds of information: robot observations must be sent from the observer to the observed, and pose estimates must be broadcast to the team as a whole. Hence the bandwidth requirements are relatively modest (on the order of a few hundred bytes per robot per second), and the total required bandwidth scales linearly with team size.

5 Experiments

This section presents the results of a controlled experiment aimed at determining the accuracy of the distributed team localization algorithm. 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 2(a); these totem-poles can be detected from a wide range of angles using the SICK lasers. 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° ambiguity. This arrangement does *not* allow individual robots to be identified. Given the ambiguity in both orientation and identity, it was necessary, for this experiment,

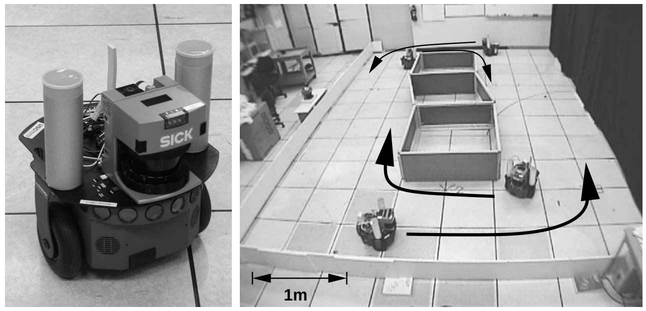

(a) **(b)**

Fig. 2. (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

to manually label the data. Ground-truth information was provided by an external laser-based tracking system.

The team was placed into the environment shown in Figure 2(c) with each robot executing 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 one 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 relative pose estimates $\{y_t^{rr'}\}$ with their corresponding true values (as determined by the ground-truth system). That is, at each time t, we determine how accurately, on average, each robot has estimated the relative pose of all the other robots. Hence our accuracy is measured by an error term of the form:

$$(\epsilon_t)^2 = \frac{1}{n(n-1)} \sum_r \sum_{r'} (y_t^{rr'} - z_t^{rr'})^T (y_t^{rr'} - z_t^{rr'}) \tag{8}$$

where $z_t^{rr'}$ is the true pose of robot r' relative to robot r at time t. The error is averaged over all pairs of robots. Note that we make no attempt to compare the *absolute* pose estimates x_t^r against some 'true' value; the arbitrary nature of the global coordinate system renders such comparison meaningless.

The qualitative results for this experiment are summarized in Figure 3, 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 snap-shot is largely incoherent; at this time, the relative pose of the robots is completely unknown, and their absolute poses have been chosen randomly. In the interval $0 < t < 80$ sec, the robots commence wall

In "Experimental Robotics VIII", B. Siciliano and P. Dario, eds., Springer-Verlaag, 2003, 146-155

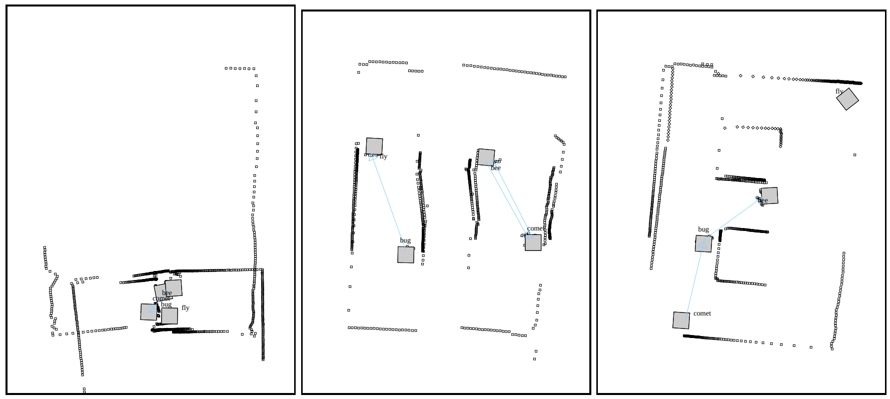

(a) t = 0 sec **(b)** t = 82 sec **(c)** t = 277 sec

Fig. 3. Experimental snap-shots. Each sub-figure shows the estimated pose of the robots at a particular time, with the corresponding laser scan data overlaid

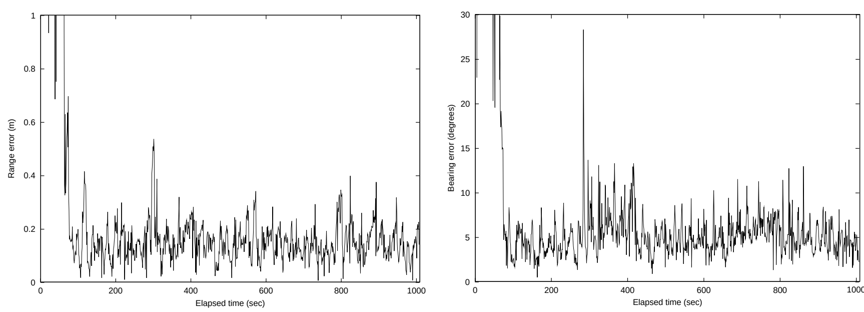

Fig. 4. Experimental results. The plots show the range and bearing components of the relative pose error ϵ_t as a function of time

following. By time $t = 80$ sec, both of the robots following the outer wall have obserovered 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 observed each other directly. At time $t = 265$ sec, the environment is modified, with the inner wall being restructured to form two separate islands. The two robots following the inner wall now follow different paths. The localization is un-affected by this change, as shown in the snap-shot at time $t = 272$ sec. A key feature of the method described in this paper is that it is largely indifferent to such structural changes in the environment.

The quantitative results for this experiment are summarized in Figure 4, which plots the range and bearing components of the total error ϵ_t as a function of time. At time $t = 0$ sec, the relative pose of the robots is completely unknown, and the errors are correspondingly high. By time $t = 80$ 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 and bearing errors oscillate around 9.7 ± 14 cm and $3.2 \pm 2.3°$, respectively. The magnitude of these errors can be attributed to two factors. First, there is some uncertainty in the relative pose measurements made by the laser-range-finder/retro-reflector combination. The uncertainty is difficult to characterize precisely, but is of the order of ± 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 100 ms; in this time, the robot may travel 2 cm and/or rotate through $3°$, which will significantly affect the results.

We ascribe the *variation* seen in the error plots to three 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. Second, we expect that errors will fall during those periods when robots are observing one another. This fall, however, may be proceeded 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 100 ms) to incoporate new data and relax to a new set of pose estimates. 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 two of the robots when the environment was changed for the first time. As a result of this collision, the robots had to be manually re-positioned, leading to gross errors in both robots' odometry. Nevertheless, as the plot indicates, the system was able to quickly recover.

6 Conclusion and Further Work

The experiment described in the previous section demonstrates several key capabilities of the team localization method described in this paper: this method does not require external landmarks, nor does it require that any of the robots remain stationary; it is robust to changes in the environment and to poor motion sensing; and robots can infer the pose of robots they have never seen. The accuracy of the localization is more than adequate to facilitate many forms of coorperative behavior.

There remain many aspects of both the general method and of our distributed implementation that require further experimental analysis. With regards to the method, we have not yet analyzed the impact of local minima (which necessarily plague any non-trivial numerical optimization problem). With regards to the distributed implementation, we are yet to measure how the algorithm scales with team size, although we suspect that both computation and bandwidth requirements scale linearly.

Acknowledgments

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

In "Experimental Robotics VIII", B. Siciliano and P. Dario, eds., Springer-Verlaag, 2003, 146-155

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 Sourthern California, 2001.
8. A. Howard, M. J. Matarić, and G. S. Sukhatme. Team localization: A maximum likelihood approach. Technical Report IRIS-01-415, Institute for Robotics and Intelligent Systems Technical Report, University of Sourthern California, 2002.
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. W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, 2nd edition, 1999.
13. 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.
14. 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.
15. 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.
16. S. Thrun. A probabilistic online mapping algorithm for teams of mobile robots. *International Journal of Robotics Research*, 20(5):335–363, 2001.
17. 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.
18. 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.