Collaborative Probabilistic Constraint-Based Landmark Localization

Ashley W. Stroupe[1], Tucker Balch[2]

[1]*Carnegie Mellon University, Pittsburgh, Pennsylvania,* ashley@ri.cmu.edu
[2]*Georgia Institute of Technology, Atlanta, Georgia,* tucker@cc.gatech.edu

Abstract

We present an efficient probabilistic method for localization using landmarks that supports individual robot and multi-robot collaborative localization. The approach, based on the Kalman-Bucy filter, reduces computation by treating different types of landmark measurements (for example, range and bearing) separately. Our algorithm has been extended to perform two types of collaborative localization for robot teams. Results illustrating the utility of the approach in simulation and on a real robot are presented.

1. Introduction

Current approaches to robot localization, such as Kalman-Bucy filter and Monte Carlo Localization (MCL), have demonstrated success on a number of mobile robot systems [8]. However, in some high-performance, dynamic environments the necessary computational resources may not be available to support computationally demanding methods. Environments with few landmarks or minimal sensor information regarding the landmarks may not provide good results (see section 2). On the other hand, simple geometric approaches (such as triangulation) are computationally very fast, but do not usually provide quantitative estimates of positional uncertainty. Such estimates are often critical for enabling a robot to realize it is "lost."

We present an approach to robot localization, Collaborative Probabilistic Constraint-Based Landmark Localization (CPCBL) that overcomes these limitations. CPCBL is similar to the Kalman-Bucy filter, but has reduced computational requirements. We reduce computation though several means. Separating different measurement types (range and bearing) on landmarks simplifies computation by considering one aspect at a time. Additionally, Gaussian distributions are used to model sensor and pose uncertainty. This representation enables the approach to also evaluate pose estimate uncertainty with minimal overhead. Lastly, single updates are made per cycle: heading or position.

In addition to gains in efficiency, separate treatment of different types information from landmarks leads to flexibility. Different landmarks may provide range and bearing of very different quality or one without the other (such as distant landmarks too far for range but visible for bearing). Separating these computations allows focus on high quality or available information.

In addition, two types of collaborative localization are investigated: active (integrating the observations of own position by teammates) and passive (using teammates as additional landmarks). Active collaborative localization enables robots with lower quality pose sensors to benefit from the observations of robots with more expensive, high quality position sensors. One robot, for example, might be equipped with differential GPS, while its teammates use only odometry for localization. Collaboration between the teammates would enable the entire team to benefit from the GPS sensor.

2. Related Work

The most commonly used approaches to robot localization are triangulation, Kalman-Bucy filters and MCL [8]. Triangulation provides fast, geometric updates with no estimate of pose uncertainty [7][12]. Probabilistic approaches use triangulation and geometry to determine a likely pose, to which uncertainty is added. Kalman-Bucy filters use Gaussian distributions to represent pose parameters, allowing fairly efficient representation and movement and sensor updates [9][10][13][14]. MCL allows arbitrary distributions to be represented by samples with high sample density in regions of high likelihood. Bayesian updates are made on each sample, leading to complex updates [3][11].

Implementations of Kalman-Bucy filters and MCL typically use all available information to perform sensor updates and may not perform well with few landmarks or poor information on landmarks. For example, our group implemented MCL on a robot team for robot soccer. We found that it did not work well using the soccer goal corners, which could not provide reliable range measurements, only bearing. This motivates consideration of alternative localization strategies.

Some approaches, geared towards outdoor environments with landmarks at large distances, use only landmark bearing information [2]. These approaches, however, do not make use of range information, even if it is

available, and are typically computationally expensive, again driving a need for an alternative approach.

Triangulation, Kalman-Bucy filters, and MCL have been applied in cooperative localization. The most common approach uses temporarily stationary teammates as additional landmarks and performs triangulation or Kalman-Bucy updates [6][12][13]. Triangulation and Gaussian distributions can be efficiently exchanged for collaboration. Active and passive collaboration can be included with Kalman-Bucy filters; Roumeliotis' group incorporates the dependence between subsequent observations into covariance, adding to mathematical complexity and computational resource requirements [13]. MCL can also use active collaboration [5]. This requires transmission of relatively large amounts of information.

Our approach is based on the simpler probabilistic framework of Kalman-Bucy filters, with two differences. Separating range and bearing information reduces the complexity of computing distribution means and allows focus on higher quality landmark measurements while ignoring missing measurement types. Separating active and passive collaborative localization allows leveraging higher quality sensors and reduces potential for pose update interdependence.

3. Constraint-Based Landmark Localization

Constraint-Based Landmark Localization (CBL) uses movement and sensor models to update pose estimates according to differential motion and sensor data. The update process is illustrated in Figure 1 with a position update example. In Figure 1, pose estimates are indicated in light gray and with subscript 0 and actual pose is indicated in dark gray without subscript. Small solid circles indicate point landmarks.

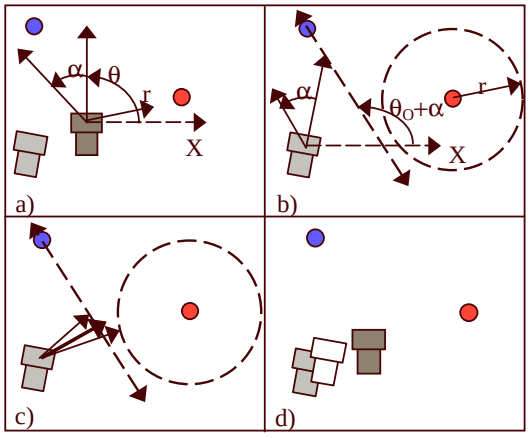


Figure 1. *Position update (pose dark gray, estimate light gray): a) sensing, b) constraints, c) closest constraint points and resultant, d) update result (white).*

A complete pose update proceeds as follows:

1) Differential motion is used to update pose.

2) Individual measurements (range and/or bearing) are made on landmarks (Figure 1a)

3) Measurements and map geometry are used to compute constraints on the possible position or heading of the robot (Figure 1b).

4) Sensor-based estimates consistent with constraints and closest to the robot's current pose estimate are determined and combined (Figure 1c).

5) Sensor-based position or heading estimates are combined with the robot's current pose estimate to produce an update (Figure 1d).

Updates of heading and position are alternated to improve efficiency. Localization updates are permitted after a specified minimum time passes since the last localization of the same type; this improves efficiency and reduces the possibility of dependent updates.

Different types of measurements on landmarks lead to different constraints. Range to a point landmark places the robot on a circle around the landmark with a radius equivalent to the range measurement (like the right landmark in Figure 1). Bearing to a point landmark (assuming heading is correct) places the robot on a line, the slope of which is defined by the bearing and the robot's heading (left landmark in Figure 1). Bearing to a point landmark also defines a single heading, assuming position is correct. A perpendicular distance to a surface defines a curve parallel to the surface at the measured distance (such as a line parallel to a wall).

For each types of constraint, the point on the constraint closest to the robot's current pose estimate is determined using simple geometric equations. Details of computation can be found in previous work [15].

The basic implementation of CBL is non-probabilistic and uses a weighted average of sensor-based pose and current pose estimate to update. Non-probabilistic updates using weighted averages produce estimates equivalent to those that would be obtained if the observations were represented using circular Gaussian distributions; constant weighting is equivalent to constant relative magnitude of covariance values.

The effectiveness of CBL for individual robot localization has been demonstrated in laboratory trials and in a robot soccer tournament. These results are provided in section 6. However, the non-probabilistic implementation of CBL provides no estimate of uncertainty and requires hand tuning of weighting parameters. This leads us to consider how to appropriately integrate uncertainty into our approach.

4. Adding Probabilistic Pose Representation

To provide uncertainty estimates and reduce the need for hand tuning parameters, a probabilistic approach is added to CBL. The probabilistic distribution of pose is represented by Gaussian probability density functions. Gaussian distributions provide a flexible representation of uncertainty because they can be oriented arbitrarily with respect to the global axes and can be non-circular

to indicate a bias along a particular axis. The probabilistic implementation requires experimental estimation of odometry and sensor noise models and a geometric mapping of vision parameters to Gaussian distribution parameters in the robot's frame.

Updates proceed similarly to basic CBL, with additions to account for the representation of uncertainty:

1) Differential motion updates the pose as before. Uncertainty estimates include uncertainty in the motion model and the previous uncertainty.

2) Individual measurements are made on landmarks. Uncertainty in measurements is estimated using sensor models.

3) Constraints are computed as before.

4) Constraint points closest to robot are computed as before and form the Gaussian distribution mean. The major and minor axes of the distribution are set according to empirical models of sensor noise.

5) Sensor-based estimates and motion update results are combined with Gaussian multiplication.

Odometry updates add differential odometry to the current position estimate. Differential odometry mean is computed directly from desired heading and distance. Uncertainty is computed by compounding uncertainty using a differential odometry model and the current estimate. Compound uncertainty is computed as JCJ^T, where C is the joint covariance matrix of the current estimate and the differential motion (each rotated to align with global axes) and J is the Jacobian of the transformation between differential and global frames. After aligning with the global frame, JCJ^T reduces to a sum of uncertainties along each axis.

Uncertainty on landmark measurements, in terms of percent of range and angle, is converted into a two-dimensional Gaussian that is rotated to align with the global frame. The robot's vision model and associated Gaussian distribution parameters are shown in Figure 2, and the conversion from vision to Gaussian is shown in Equations 1 and 2. Typically there is more uncertainty in sensed range to an observed object than in azimuth. Thus, the Gaussian distribution is an elongated ellipsoid with major axis along a line from the robot to the object.

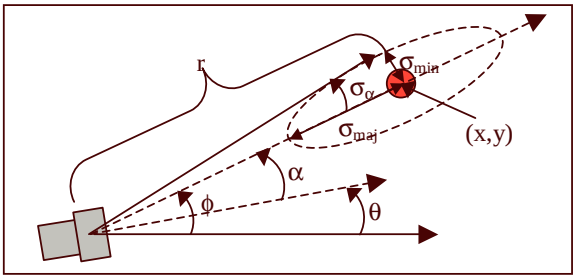


Figure 2. *Geometric definition of vision parameters ($r, \alpha, \sigma_r, \sigma_\alpha$) and distribution parameters ($x, y, \sigma_{maj}, \sigma_{min}, \phi$).*

$$x = r\cos(\alpha) \quad (1a)$$

$$y = r\sin(\alpha) \quad (1b)$$

$$\sigma_{maj} = \sigma_r \quad (2a)$$

$$\sigma_{min} = r\sin(\sigma_\alpha) \quad (2b)$$

In CBL, measurement uncertainty leads to uncertainty in constraint location, and thus to uncertainty around the sensor-based pose estimate. A Gaussian distribution, centered at the point computed by the CBL algorithm, approximates this uncertainty.

Error in range adds "thickness" to the circle constraint. We set the width of the distribution to one standard deviation of the range measurement. Uncertainty along the direction tangent to the circle can be selected arbitrarily; the constraint itself merely places the robot somewhere on the circle (Figure 3a).

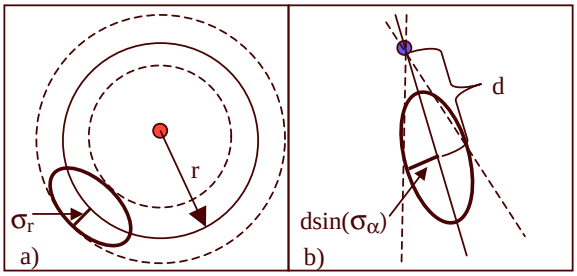


Figure 3. *a) Range constraint estimate uncertainty; Gaussian width is the uncertainty in range. b) Bearing constraint estimate uncertainty; Gaussian width is uncertainty in position at given distance from landmark.*

Error in landmark bearing leads to error perpendicular to the line, adding "thickness" to the constraint. Like the computation of σ_{maj}, the standard deviation of error perpendicular to the line is the possible offset caused by an error in angle of one standard deviation at the distance the robot is from the landmark (Figure 3b).

Constraints are computed relative only to a known landmark position. Therefore, the only error to be considered is the sensing error.

Combining current pose with sensor-based updates is based on Kalman-Bucy filter updates. Gaussian distributions, computed simply using constraints and error models, are multiplied to perform updates [14].

$$C' = C_1 - C_1\left[C_1 + C_2\right]^{-1} C_1 \equiv \begin{bmatrix} A & B \\ B & D \end{bmatrix} \quad (3)$$

$$\hat{X}' = \hat{X}_1 + C_1\left[C_1 + C_2\right]^{-1}\left(\hat{X}_2 - \hat{X}_1\right) \quad (4)$$

To improve speed, Gaussian distributions may be circular (rather than as in Figures 2 and 3). This reduces mean and covariance to weighted averages. While this reduces accuracy, it still estimates quality of the pose.

A limitation of this approach is that using only landmarks for localization leaves those robots with poor sensors more vulnerable to becoming lost and it does not take advantage of teammate capabilities.

5. Supporting Collaborative Localization

Each robot on a team has a means to estimate pose. Typically, each robot also has a set of sensors that may support collaborative localization. This collaboration can improve teammates' ability to estimate their poses. Collaborative localization could be especially helpful for a heterogeneous team, if some robots have better sensors than others. To explore this possibility, we consider two forms of collaboration: active and passive.

A robot observing a teammate that is known (via communication) to be at a location with some level of uncertainty can use that teammate as an additional landmark. This is called passive collaborative localization. If a teammate observes a robot and tells the robot where it sees it, the robot can incorporate this estimate as an additional measurement on pose. This is called active collaborative localization. Both types of collaboration can be used in probabilistic or non-probabilistic implementations.

In the probabilistic implementation, uncertainty estimates must be added to the distribution means. In each case, an observed teammate's uncertainty is computed and aligned with the global frame as with other sensor measurements. When a teammate is used passively as a landmark, the uncertainty on the sensor measurement is compounded with the uncertainty on the landmark's position before the constraint is computed. When a teammate's position is observed for active collaborative localization, sensor measurement uncertainty is compounded with the observing robot's pose uncertainty prior to broadcasting the pose estimate.

6. Experimental Results

Several experiments were performed on real robots and in simulation to test implementations of our approach.

6.1 Robot Experiments

The constraint-based approach for an individual robot was implemented on a Minnow robot[1], a research platform used for RoboCup middle-size soccer. The Minnow (Figure 4) is a differential drive robot (with a trailer) controlled using *TeamBots* architecture. The Minnow uses a USB camera for vision and *CMVision* for image processing. Experimental camera calibration provided vision model standard deviations of approximately 3% error in range and 2° error in angle.

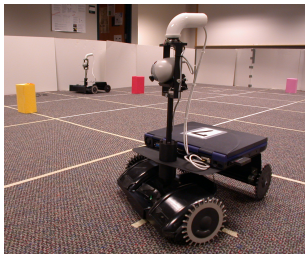


Figure 4. *Minnow robots and colored block landmarks.*

To accommodate soccer dynamics and high minimum cycle time on Minnow, the non-probabilistic implementation was employed. Sensor estimates were weighted with respect to each other using estimated quality: inverse distance. The result was a static weighted average of sensor-based and current estimates.

In addition to qualitative improvement seen on the soccer field, quantitative results were evaluated in the lab. The robot was directed to drive a circuit for five loops (Figure 5). At waypoints, on-board odometry, localization, and ground truth were recorded. Landmark measurement availability varied for each experiment to discover the value of separating measurement types:

A) Four landmarks with range for position updates and bearing for heading updates,

B) Four landmarks with range and bearing for position updates and four landmarks with bearing for heading updates,

C) Four landmarks with range and bearing for position and bearing for heading and four landmarks with bearing for position and heading.

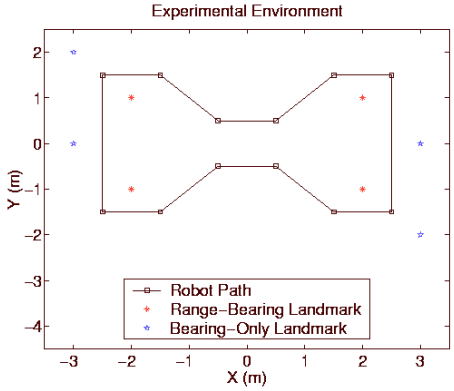


Figure 5. *Experimental environment: robot path with waypoints and landmarks (range-bearing landmarks used in A, B, and C; bearing landmarks used in C).*

Table 1 shows total position error and heading magnitude error by run. Odometry is represented by the single run that produced the least odometry error.

Table 1. *Pose Estimate Error Results for Single Robot*

	Odometry	Localized A	Localized B	Localized C
Mean Position Error	0.5592m	0.0994m	0.0719m	0.0627m
$\sigma_{\Delta x}$	0.4378m	0.0822m	0.0578m	0.0562m
$\sigma_{\Delta y}$	0.5624m	0.0920m	0.0620m	0.0482m
Maximum Position Error	1.7275m	0.3459m	0.2326m	0.1652m
Mean Heading Error	29.6°	5.1°	2.3°	1.7°
Maximum Heading Error	42°	11°	5°	4°

The robot remained localized, improving performance as more landmark information was provided. This demonstrates not only the minimal localization's ability to keep the robot localized, but also demonstrates the flexibility in measurement use. For instance, the robot can take advantage of only range or bearing information for particular landmarks, depending on availability.

6.2 Simulation Experiments

Simulation experiments were conducted in *TeamBots* using simulated Minnow robots. Two robots are present: a navigating robot and a stationary observing robot. Four fixed landmarks were used and the test robot followed a looping path (the same as that for the real robot experiments) for five loops.

Odometry and vision error models on the navigating robot were based on empirical data but were artificially inflated to test abilities of the approach. The observing robot represents a single sensor with very high accuracy that may provide information to teammates.

Five localization configurations were explored (along with odometry), each using probabilistic representation and all four landmarks (except as specified):

A) Odometry: odometry only, no localization,
B) Non-Probabilistic: CBL with weighted average,
C) Landmark: landmark only localization,
D) Passive: landmarks and passive collaboration on teammate,
E) Active: landmarks and active collaboration by teammate,
F) Active and Passive: landmarks and both active and passive collaboration with a teammate,

Here, active collaboration is used for position only, as robots cannot visually detect a teammate's heading. A baseline experiment (B) used weighted average updates with no uncertainty, as on the real robot. The same vision and odometry simulation parameters were used, and weighting matched the real robot experiments.

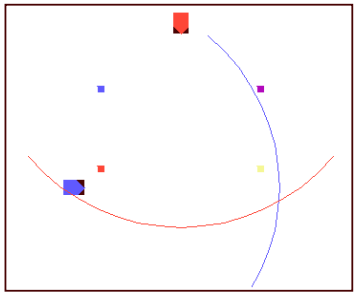


Figure 6. *TeamBots simulator showing the navigating robot (bottom) and the observing robot (top) with four landmarks. Visual fields of view are shown as arcs.*

Two robots were present in each experiment. One, the navigating robot, drove the looping path described above. The second, the observing robot, remained stationary at a location from which it could observe the entire path of the navigating robot. The environment is shown in Figure 6, with navigating robot at bottom left; and observing robot at top center. Arcs show each robot's field of view. Squares symbolize landmarks.

For these experiments, the observing robot had ten times less sensor error than the navigating robot. This difference was chosen to help illustrate the potential use of our approach in heterogeneous teams. The navigating robot's vision had standard deviations 5% of range and 5.0 degrees, while the observing robot had standard deviations of 0.5% of range and 0.5 degrees. To reflect real time constraints, probabilistic computation was reduced by using circular distributions. Each type of localization was evaluated over five complete runs (5 path loops). Results for position are shown in Figure 7.

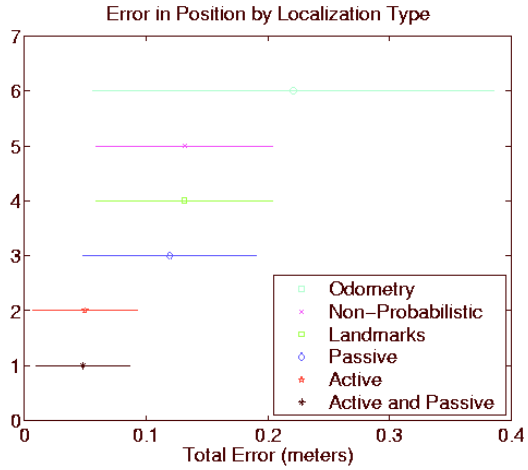


Figure 7. *Position error for each experiment type. Mean is a point, bars are one standard deviation. Note clear improvement at initiation of active collaboration.*

An additional landmark (the observing robot) only slightly improves performance; error range and mean do not change significantly as the landmark teammate is added (with or without active collaboration). With high levels of odometry and vision noise, robots only using their own sensors occasionally bumped landmarks (necessitating a restart). However, use of the single external higher-quality sensor tremendously improves performance; mean error in position significantly drops with the addition of active collaboration. This effect is seen most clearly in position since active collaboration only directly updates position. Improvements are still seen in heading, however, as heading updates depend on position and the position estimates are much improved. Lastly, our simplified treatment of uncertainties is validated: errors remained within 3-σ of actual values.

Non-probabilistic pose estimates are of a quality similar to that without collaboration and with only passive collaboration. Therefore, it can be useful when very fast updates are needed. However, no estimate of result quality is provided for fault detection and recovery.

7. Discussion

We present a computationally low-cost method for landmark localization designed for individual localization and multi-robot collaborative localization. The approach is targeted for dynamic environments with demanding time constraints and environments with limited available landmark information. This approach can be implemented with weighted average updates or with simple probabilistic representation and updates. Collaborative localization can also be incorporated. We introduce, and demonstrate two types of collaborative localization in this framework: passive and active.

This simplified approach to probabilistic localization can produce good results, keeping robots well localized and providing reasonable estimates of pose quality. Simplifying the computation of sensor distribution means reduces complexity for computation and allows using different types of information from different types of landmarks. This distinction allows the algorithm to use only aspects that can be measured accurately, or to weight weaker aspects of the measurement less.

This simplified approach allows for both passive and active collaboration in localization. While the most commonly used approach to cooperative localization uses teammates as landmarks, the separation of passive from active demonstrates that greater improvement can be gained by directing collaboration from robots with better sensors (or current estimates) to those with worse sensors (or current estimates). This improvement can be seen even when observing teammates cannot fully observe robot pose, as in these experiments. Limiting collaboration to one-directional (high quality to low quality) need not significantly diminish performance, but reduces bandwidth requirements and eliminates inter-dependency of position estimates.

The CBL approach is adaptable to different the needs of different systems and environments. Complexity can be reduced by replacing oval, arbitrarily oriented Gaussian distributions with circular distributions; even further reduction can be achieved by using non-probabilistic representations with weighted average updates.

A remaining limitation of this approach is that uncertainty is modeled as Gaussian and is, therefore, approximate. Uncertainties only provide estimates of the pose quality. This is true of all Kalman-Bucy filter approaches. This may be adequate to estimate the results of different behaviors or to indicate a lost condition, particularly in highly reactive systems. A second limitation is that landmarks must be predefined; an extension to allow mapping would be possible, based on Kalman-Bucy filter mapping approaches [4].

The next step to continue this research is to implement collaborative and probabilistic CBL on real robots. This will be done using Sony quadruped robots for RoboCup 2002. To accomplish this, team size must be extended to more than two robots. The greatest challenge in moving to real robots and to larger teams is visually locating teammates and distinguishing among them. Also, comparisons of circular and arbitrary Gaussian distributions will be made in this implementation.

In other future work, heterogeneity in sensor quality among teammates will be leveraged, fulfilling the potential demonstrated here. Collaborative sensor readings that cannot greatly improve localization results can be eliminated, reducing bandwidth requirements and reducing interdependency. The judgment of what can improve personal localization may vary with the current quality of position estimates. In teams larger than two, robots can choose which teammates to use for localization based on quality of observations. Dynamically configuring collaborative assistance as needed is an extension that will be investigated.

Acknowledgements

The authors wish to acknowledge the contributions of Jim Bruce, Scott Lenser, Hans Moravec, Kevin Sikorski, and Manuela Veloso. We also wish to acknowledge the support of the DARPA Mobile Autonomous Robot Software Program and the Northrop Grumman Corporation

8. References

[1] T. Balch, "The Minnow Project," http://www-2.cs.cmu.edu/~coral/minnow.

[2] M. Deans and M. Hebert, "Experimental Comparison of Techniques for Localization and Mapping using a Bearings Only Sensor," *Proc ISER '00*, IEEE, 2000.

[3] F. Dellaert, D. Fox, W. Burgard, and S. Thrun, "Monte Carlo Localization for Mobile Robots," *Proc ICRA 1999*, IEEE, 1999.

[4] G. Dissanayake, H. Durant-Whyte, and T. Bailey, "A Computationally Efficient Solution to the Simultaneous Localisation and Map Building (SLAM) Problem," ***Proc ICRA 2000***, IEEE, 2000.

[5] D. Fox, W. Burgard, H. Kruppa, and S. Thrun. "A Probabilistic Approach to Collaborative Multi-Robot Localization," *Autonomous Robots on Heterogeneous Multi-Robot Systems, Special Issue*, vol. 8, no. 3, 2000.

[6] R. Grabowski, L.E. Navarro-Serment, C.J.J. Pareidis, P.K. Khosla, "Heterogeneous Teams of Modular Robots for Mapping and Exploration," *Autonomous Robots, Heterogeneous Multi-Robot Systems Special Iss*, 2000.

[7] C. Graydon, K. Hanson (eds.), *Mountaineering: The Freedom of the Hills*, 6th edition, 1997.

[8] J.-S. Gutmann, W. Burgard, D. Fox, and K. Konolige, "An Experimental Comparison of Localization Methods," *Proc IROS-98*, IEEE, 1998.

[9] R. Hanek and T. Schmitt, "Vision-Based Localization and Data Fusion in a System of Cooperating Mobile Robots," ***Proc IROS***, IEEE, 2000.

[10] T.D. Larsen, K.L. Hansen, N.A. Andersen, and O. Ravn, "Design of Kalman filters for mobile robots; evaluation of the kinematic and odometric approach," ***Proc 1999 IEEE Int Conf on Control Applications***, vol. 2, 1999.

[11] S. Lenser and M. Veloso, "Sensor Resetting Localization for Poorly Modeled Mobile Robots," *Proc ICRA 2000*, IEEE, 2000.

[12] I.M. Rekleitis, G. Dudek, and E.E. Milios, "Multi-Robot Collaboration for Robust Exploration," *Proc ICRA 2000*, IEEE, 2000.
[13] S.I. Roumeliotis and G.A. Bekey, "Distributed Multi-Robot Localization," *Proc 4th DARS*, 2000.
[14] R.C. Smith and P. Cheeseman, "On the Representation and Estimation of Spatial Uncertainty," *Int J Robotics Research*, vol. 5, 1998, 56-68.
[15] A.W. Stroupe, K. Sikorski, and T. Balch, "Constraint-Based Landmark Localization," Submitted to *Proc of RoboCup Symp*, 2002.