Merging Gaussian Distributions for Object Localization in Multi-Robot Systems

Ashley W. Stroupe, Martin C. Martin, and Tucker Balch
Robotics Institute, Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
{ashley, mcm, tcb}@ri.cmu.edu

Abstract: We present a method for representing, communicating, and fusing distributed, noisy, and uncertain observations of an object by multiple robots. The approach relies on re-parameterization the two-dimensional Gaussian distribution to correspond more naturally to a robots' observation space. The approach enables two or more observers to achieve greater effective sensor coverage of the environment and improved accuracy in object position estimation. We demonstrate empirically that, using our approach, more observers achieve more accurate object position estimates. The method is tested in three application areas: object location, object tracking, and ball position estimation for robot soccer. We provide quantitative evaluation of the technique on mobile robots.

1. Introduction

Typically, individual robots can only observe part of their environment at any moment in time. In dynamic environments, information previously collected about currently unobservable parts of the environment grows stale and becomes inaccurate. Sharing information among robots increases the effective instantaneous visibility of the environment, allowing for more accurate modeling and more appropriate response. If processed effectively, information collected from multiple points of view can provide reduced uncertainty, improved accuracy, and increased tolerance to single point failures in estimating the location of observed objects.

In order to meet the time demands of a highly dynamic environment (e.g. robotic soccer), both the information transmitted between robots and the computational demands to combine observations must be minimal. Our approach makes use of a few easily obtainable parameters describing an observation and simple computations to meet these needs. We use two-dimensional statistical representations of target location observations generated by individual robots. Each robot independently combines the multiple observations provided it in order to produce improved estimates of target locations. The local computation allows for robust use of all available information without failure due to communication loss.

2. Background and Related Work

One common method of object position estimation, including robot location, uses Kalman filters to track objects [17, 21, 22, 23] or robots [13, 15, 19] through time, updating estimates based on previously known positions. Data fusion of multiple sensors (such as odometry, sonar, cameras, and laser range scanners) on a single robot, using Kalman filters, has been used for robot localization [14, 15, 18, 19]. The technique is especially useful in dynamic applications such as robotic soccer. The CS Freiburg RoboCup team of Germany, for example, estimates position using odometry and by finding the field borders in laser range scans [10]. These two estimates are fused using a Kalman filter to localize the robots. Recent developments in localization in unknown environments by simultaneously mapping and localizing rather than relying on previously generated maps [3, 8]. The specific issues of multi-robot localization in a mapped environment are also investigated [9].

The ability to rapidly share distributed observations is critical in distributed dynamic tasks like robotic soccer. Most robot soccer team approaches use vision and/or sonar to localize and vision to locate objects in the environment. Some share information for planning and dynamic role assignment (ART [16]). Others fill-in blank areas in the world model with shared data (CS Freiburg [10, 11], RMIT [4], 5dpo [7]). Other distributed sensing approaches include merging independent grid cell occupancy probabilities measured by multiple robots (possibly distributed in time) [5, 6], and curve fitting of models and observations by multiple robots [12].

The task we address is distinct from the others described above. We focus on fusing multiple simultaneous observations of the same object from distributed vantage points (as opposed to observations from the same vantage point over multiple instants in time). Our objective is to provide more accurate instantaneous estimations of the location of dynamic objects that are simultaneously visible by multiple robots without relying on historical data. Additionally, most probabilistic methods rely on decomposing the space into discrete cells [5, 6, 9, 14, etc]. Our approach does not require discretization, working in the continuous spatial domain.

3. Fusing Gaussian Distributions

3.1. Overview

We represent a single observation of an object as a two-dimensional Gaussian distribution (Figure 1). The center, or mean, of the distribution is the estimated location of the object and the standard deviations along the major and minor axes of the distribution correspond to estimates of the uncertainty (or noise) in the observation along each axis. The distribution corresponds to the conditional probability that the object is in that location, given the observation.

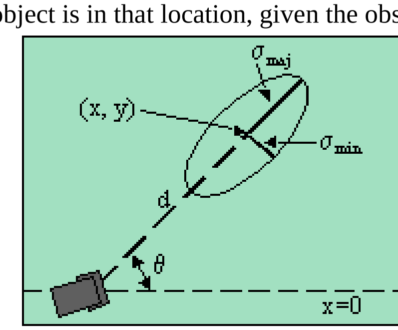

Figure 1. Distribution parameter definitions: mean (x, y), angle of major axis (θ), major and minor axis standard deviations (σ_{maj}, σ_{min}), distance to mean (d).

Provided two observations are independent and drawn from normal distributions, the observations can be merged into an improved estimate by multiplying the distributions. To meet cycle time requirements of a highly reactive system, an efficient method of multiplying distributions is necessary. We use a two-dimensional statistical approach based on Bayes' Rule and Kalman filters, first introduced by Duffin [1]. In this approach, multi-dimensional Gaussian distributions can be combined using simple matrix operations. Since multiplying Gaussian distributions results in a Gaussian distribution, the operation is symmetric, associative, and can combine any number of distributions in any order.

Our approach, illustrated in Figure 2, is to collect observations of multiple robots, and then merge the corresponding Gaussian distributions to yield a better estimate of the location and uncertainty of the observed object.

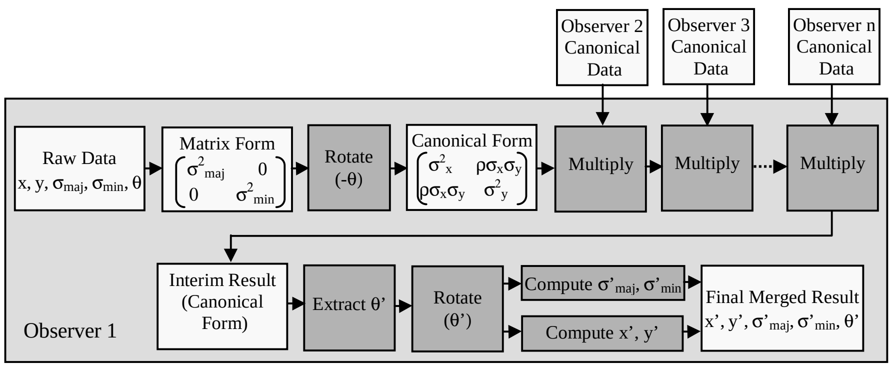

Figure 2. Block diagram of the multi-distribution merging process. Multiplication is conducted using the mathematical formulation described above. Each subsequent distribution is merged with the previous result, and the final parameters are extracted.

The canonical form of the two-dimensional Gaussian distribution depends on standard deviations, σ_x and σ_y, a covariance matrix, C, and the mean, as shown [20]:

$$p(X) = \frac{1}{2\pi\sqrt{|C|}} \exp\left(-\frac{1}{2}\left(X - \overline{X}\right)^T C^{-1}\left(X - \overline{X}\right)\right), \text{ where } C = \begin{bmatrix} \sigma_x^2 & \rho\sigma_x\sigma_y \\ \rho\sigma_x\sigma_y & \sigma_y^2 \end{bmatrix} \quad (1)$$

The parameterization of the Gaussian distribution in this representation does not correspond the parameters of our observations (Figure 1). We address the problem through a transformation of parameters from observation form to canonical form. In this form, the distributions can be merged using matrix operations. After all observations are merged, we extract the mean and standard deviations from the merged result (these correspond to estimated location and uncertainty of the object).

3.2. Mathematical Details

We wish to determine the mean, standard deviations, and angle of the combined distribution to estimate object position and characterize the quality of the estimate. We can compute these parameters from sensor readings and models of sensor error (deviations). Thus, we require a method of determining combined parameters from those of individual distributions. The formulation we use is adopted from Smith and Cheeseman [20] and derivations are provided in more detail in a technical report [21]. Since the mean, standard deviations, and orientation of the major axis are independent of scaling, they can be extracted from the resulting merged covariance matrices without considering absolute probability values.

The covariance matrix, C, of an observation relative to coordinates aligned with the major and minor distribution axes is initially determined from the major and minor axis standard deviations in the local coordinate frame (designated L).

$$C_L = \begin{bmatrix} \sigma_{maj}^2 & 0 \\ 0 & \sigma_{min}^2 \end{bmatrix} \tag{2}$$

Since observations may be oriented arbitrarily with respect to the global coordinate frame (angle θ relative to global x-axis), they must be transformed to this frame. Rotation of X in equation 1 by leads to the following relationship.

$$C^{-1} = R(-\theta)^T C_L^{-1} R(-\theta) \quad \Rightarrow \quad C = R(-\theta)^T C_L \; R(-\theta) \tag{3}$$

Once the observation is in canonical form, we combine individual covariance matrices into a covariance matrix representing the combined distribution.

$$C' = C_1 - C_1 \left[C_1 + C_2 \right]^{-1} C_1 \tag{4}$$

The mean of the resulting merged distribution, X, is computed from the individual distribution means and covariance matrices.

$$\hat{X}' = \hat{X}_1 + C_1 \left[C_1 + C_2 \right]^{-1} \left(\hat{X}_2 - \hat{X}_1 \right) \tag{5}$$

The principal axis angle is obtained from the merged covariance matrix entries:

$$\theta' = \frac{1}{2} \tan^{-1} \left(\frac{2B}{A - D} \right) \tag{6}$$

A, B, and D are top left, top right/lower left, and lower right entries, respectively.

Lastly, the resulting major and minor axis standard deviations are extracted by rotating the covariance matrix to align with those axes and reversing Equation 2.

$$C' = R(\theta')^T C' R(\theta') \tag{7}$$

3.3. Simulated Example

Two robots observe a target object (Figure 3). Each observation produces a Gaussian distribution of possible locations for the object; typically, each distribution provides greater accuracy along a different direction than the other distributions.

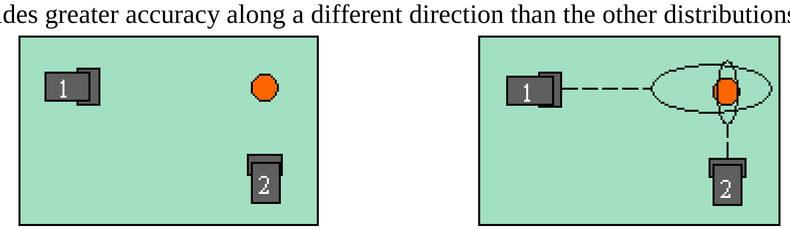

Figure 3. Left: Two (distributed) robots see a target. Right: Observations generate Gaussians; uncertainty (1-σ ovals shown) increases with distance.

For this example, the robots are positioned with relative headings 90 degrees apart and looking directly at the target. The target is located at (10,10). The two simulated robot observations were drawn from a random normal distribution centered at the object's true position. The major and minor axis standard deviations of these distributions were (5,3) for robot 1 and (3,1) for robot 2. Robot 1 reports a mean of (12.34, 9.02) and robot 2 reports a mean of (9.90, 11.69). In Figure 4, the distribution resulting from the single measurements by robot 1 and robot 2 are shown at left. The resulting merged distribution is shown in at right. The narrowing of the distribution indicates that implied uncertainty (standard deviations) is reduced, and the mean is more accurate relative to the actual target position. The merged mean is (9.97, 9.57), with major and minor axis standard deviations (0.89, 0.49).

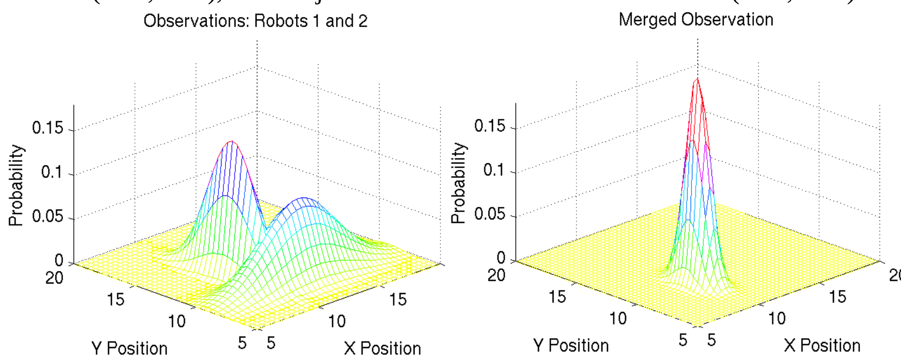

Figure 4. Left: The individual distributions to be merged. Right: Resulting merged distribution, with reduced error and higher accuracy in the mean.

4. Validation On Robots

4.1. Hardware Platform

The hardware platform used is a modified Cye robot, an inexpensive and commercially available platform (Figure 5). This platform consists of a drive section and a trailer. The drive section uses differential drive with two wheels and is equipped with a front bump sensor; the trailer is passive. On board the Cye is a motor controller processor. High level commands and image processing algorithms are implemented in C and Java on a Pentium 266 (running linux) using *TeamBots* [2]. A wide field-of-view NTSC video camera provides sensory input. *CMVision* performs color analysis and color blob detection and merging [2]. Robots communicate with wireless ethernet.

Camera calibration was conducted at two levels. First, *Flatfish* [14] determined the parameters describing the aberrations of the lens. These parameters enable mapping from pixel location to points in three-space.

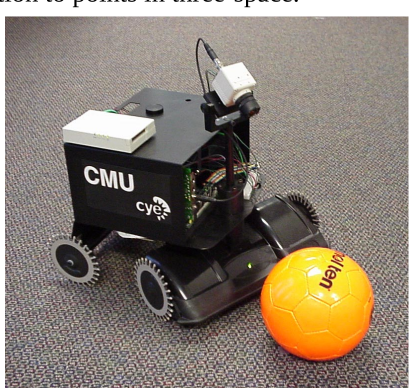

Figure 5. Enhanced Cye Robot platform.

A second calibration step characterizes systemic errors. Targets are placed at a set of fixed distances and angles relative to the robot and the distance and angle calculated by the vision system is recorded. Comparing measured distance versus actual distance provides a mean bias as a function of measurement distance. After correcting measurements for this bias, proportional errors are determined.

A histogram of these sensing errors determined that the corrected distances are distributed about actual distance approximately normally. From these errors, a standard deviation in percent distance can be directly determined. A similar process was completed for angle, though no bias correction was conducted. These deviation functions are used to compute parameters of the observation distributions.

4.2. Assumptions

Several assumptions are implicit in this approach. In addition to assuming independent, Gaussian sensor errors, the robot coordinate frames are assumed to be coincident. Without this, data are incompatible and the merging is meaningless.

We do not take into account robot positional uncertainty in the generation of target location distributions; our localization is not fully implemented at this time. Once determined, robot positional uncertainty can be incorporated by encoding it in object positional uncertainty. Merging the measurement uncertainty distribution with the robot's position distribution translated to the same mean spreads the object's position to accommodate both measurement and localization uncertainty.

Several additional assumptions were introduced for simplicity in experimentation and camera calibration. First, the camera parameter calibration assumes that objects are at a known height from the ground plane; unknown objects are therefore assumed to be on the ground plane. This reduces the transformation from three dimensions to two. This is not a highly restrictive limitation, as common obstacles, agents, landmarks, (etc) in environments are generally on the ground plane. Second, objects of interest are assumed to be unique in order to avoid the necessity of solving the association problem.

4.3. Experimental Setup

An experiment was devised to directly test this approach to distributed sensing and merged distributions without complications due to motion and positional uncertainty. Three stationary robots sequentially locate a stationary object at several pre-determined points (Figure 6). Robots share observations and compute resulting merged position estimates for each point. Computations were separately conducted on pairs of data points. Thus, accuracy of single-robot measurements can be directly compared to the accuracy obtained by combining data from two and three robots.

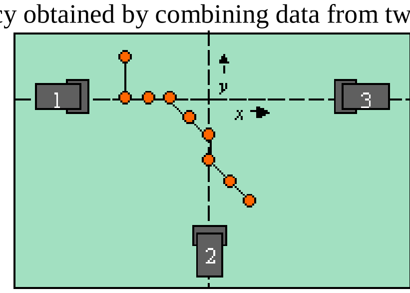

Figure 6. Three robots observe a ball at several fixed locations (shown as circles on line) and combine observations into single position estimates.

4.4. Experimental Results

Example experimental results are shown graphically. Figure 7 shows the results of successive merging compared to the actual trajectory. The top left compares robot 3's observations (largest errors at the greatest distances from the robot). The top right similarly compares the result of merging two robots' observations. The estimate resulting from merging three robots' measurements is shown at the bottom.

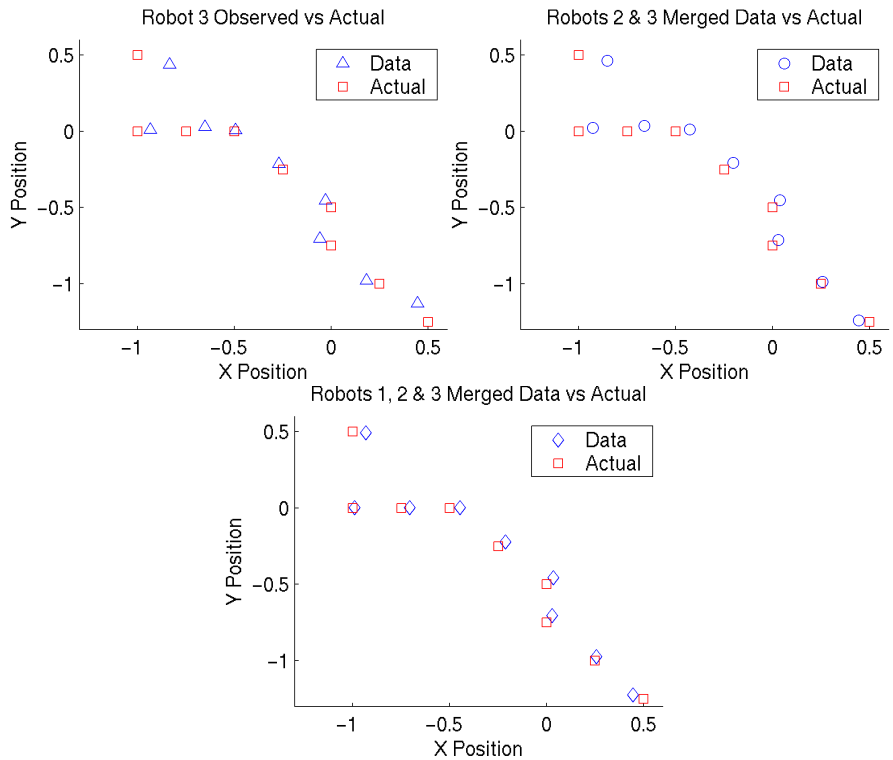

Figure 7. Reported data versus actual ball position for robot 3 (top left), robots 2 and 3 merged estimates (top right), and robots 1, 2, and 3 merged estimates (bottom).

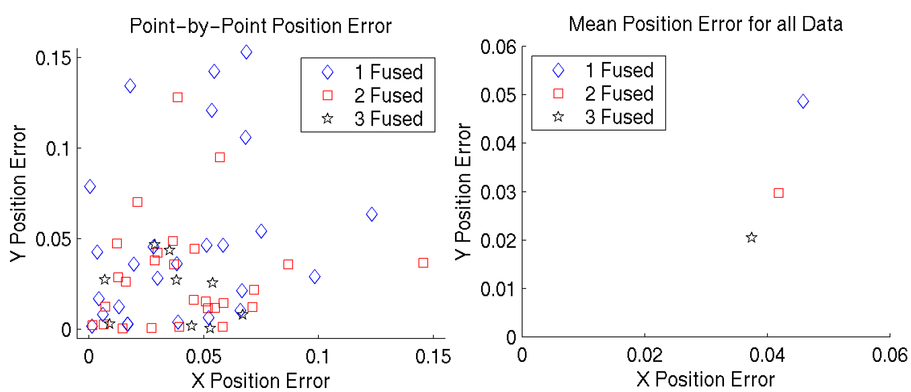

Figure 8. Left: Position error in x and y for each measurement. Each merging lowers position error bounds and reduces outlier frequency. Right: Mean position error in x and y over all single observations, all 2-observation merges, and all 3-observation merges.

Point-by-point trajectory errors are compared for all single-robot, two-robot, and three-robot measurements in Figure 8 (left). While individual trajectories are sometimes accurate at single points (in fact, occasionally more accurate than combined data), the consistency of accuracy in combined results is absent in the single-robot trajectories. Additionally, outlier frequency is reduced. This is best characterized by plotting the mean error of all single-robot observations, all two-robot observations, and all three-robot observations, as shown in Figure 8 (right).

5. Test Applications and Results

5.1. Location and Retrieval of Unseen Targets

This test exhibits the agents' increased effective field of view and ability to function with merged target positions. Initially, a robot is positioned so that it can see, but not reach, the target. Another cannot see the target, even with camera panning, but the path to the object is clear (Figure 9). By sharing information, robot 1 immediately obtains a target position without random search and successfully locates the object using only information provided by robot 2. Once the object is located, robot 1 reaches and manipulates it using the merged position provided by both robots. Due to the small distances traveled from known starting positions, assumptions on localization and coordinate frames hold.

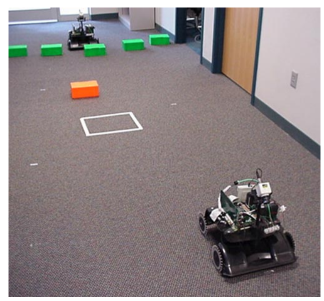

Figure 9. Robot 1 initially cannot see the target. Robot 2 provides initial target location.

5.2. Blind Robot Target Tracking

In this experiment, three robots are positioned around a target area (for convenience, at relative headings of 90 degrees). A ball is moved throughout the area, and all robots track the ball using the position obtained by merging all three observations. The robots are able to track the target in most cases, even at higher speeds, and always quickly recover lost objects. Even when the target travels along the line of sight of a single robot (diminished accuracy in the depth dimension), the additional point of view make up for this accuracy. One robot is subsequently blindfolded by covering the camera with a box (Figure 10, left). The ability of the blinded robot to track the ball using the merged position from the other two is not observably diminished. Fixed robot positions are precisely known; localization and coordinate frame assumptions hold.

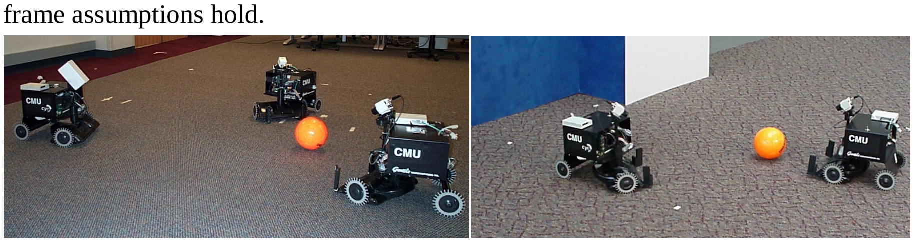

Figure 10. Left: A blind robot (left) can track targets by merging data provided by other robots. Right: A soccer test. Attacker (right) attempts to score on blue defender (left).

5.3. Robot Soccer

This approach to distributed sensing was applied to the CMU Hammerheads middle-sized robot soccer team (Figure 10, right). Robots transmit the position of the ball, if visible, so that it can be combined with all other current observations. This provides robots far from the ball with more accurate positions and allows robots to quickly locate an unseen ball. Conflicting observations (means differing by more than 2 standard deviations) were not merged to prevent false (or multiple) targets and data collected in incompatible coordinate frames from resulting in confusion.

The 2000 Hammerheads robots were localized entirely by odometry. Differences in coordinate frames arise from odometry drifts over time. As a result, the primary impact of distributed sensing was to provide a starting point for robots to locate lost balls. Despite coordinate frame discrepancy, frames were generally coherent enough that the camera's wide field of view allowed searching robots to immediately locate an unobstructed ball by looking at a provided target position.

6. Conclusions

We present a method to improve target position estimates by fusing data from two or more robot agents. This approach, based on Bayes' Rule and Kalman filter theory, implements real-time sensor data fusion on a reactive multi-robot system for many different applications. The successful ability to fuse these statistical measurements and the ability to receive position estimates of targets not visible allows our robots to quickly acquire targets and to more accurately estimate object position. While this work uses only vision for sensing, the approach can be applied to any sensor or suite of sensors which can be modeled by approximately Gaussian distributions.

This approach to distributed sensing and information sharing is very promising based on the applications presented here: unseen target location, accurate target acquisition and manipulation, and robot soccer. However, several extensions of this work are necessary for practical implementation. Even in well-localized systems, disparity between coordinate frames can arise and must be accommodated. Autonomously determining the relative transformation between coordinate frames using sensors will be investigated. Additionally, the accommodation of robot positional uncertainty will be incorporated into the target position distributions, as described previously. Lastly, it may be possible to implement a pixel-to-world coordinate transformation that does not assume that objects are at a known elevation, but this would require the development of a new method of camera calibration.

Acknowledgements

The authors wish to acknowledge the contributions and support of Jim Bruce, Scott Lenser, Kevin Sikorski, Hans Moravec, Manuela Veloso, and DARPA's Programs in Mobile Autonomous Robot Software and Control of Agent Based Systems.

References

[1] W. Anderson and R. Duffin. *J Mathematical Analysis Applications*, 26:576, 1969.
[2] T. Balch. TeamBots. http://www.teambots.org, 2000.

[3] R.G. Brown and B.R. Donald. Mobile robot self-localization without explicit landmarks. *Algorithmica*, 26(3/4):515-559, 2000.

[4] J. Brusey et al. RMIT Raiders. In: Veloso, Pagello, and Kitano (eds) *RoboCup-99: Robot Soccer World Cup III.* Springer-Verlag, Berlin, pages 741-744, 2000.

[5] A. Cai, T. Fukuda, and F. Arai. Information Sharing among Multiple Robots for Cooperation in Cellular Robotic System. In: *Proc IROS 97*, pages 1768-1773, 1997.

[6] A. Cai, T. Fukuda, F. Arai, and H. Ishihara. Cooperative Path Planning and Navigation Based on Distributed Sensing. In: *Proc 1996 IEEE Int Conf Robotics and Automation*, pages 2079-2084, 1996.

[7] P. Costa et al. 5dpo-2000 Team Description. In: Veloso, Pagello, and Kitano (eds) *RoboCup-99: Robot Soccer World Cup III*, Springer-Verlag, Berlin, pages 754-757, 2000.

[8] G. Dissanayake, H. Durant-Whyte, and T. Bailey. A Computationally Efficient Solution to the Simultaneous Localisation and Map Building (SLAM) Problem. In: *Proc 2000 IEEE Int Conf on Robotics & Automation*, pages 1009-1014, 2000.

[9] D. Fox, W. Burgard, H. Kruppa, and S. Thrun. Collaborative Multi-Robot Localization. In: *Autonomous Robots on Heterogeneous Multi-Robot Systems, Special Issue*, 8(3), 2000.

[10] J.S. Gutmann et al. Reliable self-localization, multirobot sensor integration, accurate path-planning and basic soccer skills: playing an effective game of robotic soccer. In: *Proc Ninth Int Conference on Advanced Robotics*, pages 289-296, 1999.

[11] J.S. Gutmann et al. The CS Freiburg Robotic Soccer Team: Reliable Self-Localization, Multirobot Sensor Integration, and Basic Soccer Skills. In: Asada and Kitano (eds) *RoboCup-98: Robot Soccer World Cup II*, Springer-Verlag, Berlin, pages 93-108, 1999.

[12] R. Hanek, T. Schmitt, M. Klupsch, and S. Buck. From Multiple Images to a Consistent View. To appear in: Stone, Balch, and Kraetzschmar (eds) *RoboCup-2000: Robot Soccer World Cup IV*, Springer-Verlag, Berlin, 2001.

[13] T.D. Larsen, M. Bak, N.A. Andersen, and O. Raven. Design of Kalman filters for mobile robots; evaluation of the kinematic and odometric approach. In: *Proc 1999 IEEE Int Conf Control Applications, Vol 2*, pages 1021-1026, 1999.

[14] H. Moravec. Robust Navigation by Probabilistic Volumetric Sensing. http://www.ri.cmu.edu/~hpm/project.archive/robot.papers/2000/ARPA.MARS.reports.00/Report.0001.html, 2000.

[15] L. Moreno, J.M. Armingol, A. de la Escalera, and M.A. Salichs. Global integration of ultrasonic sensors information in mobile robot localization. In: *Proc Ninth Int Conf Advanced Robotics*, pages 283-288, 1999.

[16] D. Nardi et al. ART99–Azzurra Robot Team. In: Veloso M, Pagello E, Kitano H (eds) *RoboCup-99: Robot Soccer World Cup III*, Springer-Verlag, Berlin, pages 695-698, 2000.

[17] G. Petryk, and M. Buehler. Robust estimation of pre-contact object trajectories. *Robot Control 1997, Vol 2*, pages 793-799, 1997.

[18] D. Rembold, U. Zimmermann, T. Langle, and H. Worn. Detection and handling of moving objects. In: *Proc 24th Annual Conf IEEE Ind Electron Soc, IECON '98, Vol 3*, pages 1332-1337, 1998.

[19] J.Z. Sasiadek and P. Hartana. Sensor data fusion using Kalman filter. In: *Proc Third Int Conf Information Fusion, Vol 2*, pages 19-25, 2000.

[20] R. C. Smith and P. Cheeseman. "On the Representation and Estimation of Spatial Uncertainty." In: *The International Journal of Robotics Research*, 5(4):56-68, 1986.

[21] A.W. Stroupe, M.C. Martin, and T. Balch. Merging Probabilistic Observations for Mobile Distributed Sensing. Carnegie Mellon University Techical Report CMU-RI-TR-00-30, 2000.

[22] C.Y. Tang, Y.P. Hung, S.W. Shih, and Z. Chen Z. feature-based tracker for multiple object tracking. In: *Proc National Science Council, Republic of China, Part A*, 23(1):151-168, 1999.

[23] H. Wang, C.S. Chua, and C.T. Sim. Real-time object tracking from corners. *Robotica*, 16(1):109-116, 1999.