

Distributed Sensor Fusion for Object Position Estimation by Multi-Robot Systems

Ashley W. Stroupe, Martin C. Martin, and Tucker Balch

The Robotics Institute, Carnegie Mellon University

5000 Forbes Avenue

Pittsburgh, PA 15213

{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 of the canonical two-dimensional Gaussian distribution that corresponds more naturally to the observation space of a robot. 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, when using our approach, more observers achieve more accurate estimations of an object's position. The method is tested in three application areas, including object location, object tracking, and ball position estimation for robotic soccer. Quantitative evaluations of the technique in use on mobile robots are provided.

Index Terms—distributed sensing, multi-robot coordination

I. 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 between robots can increase the effective instantaneous visibility of the environment, allowing for more accurate modeling (at whatever level) and more appropriate response. If it is 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), the information transmitted between robots must be minimal and the computational demands to combine their 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. These are combined independently on each robot to produce improved estimates of target locations.

A more detailed version of this work is presented in a Carnegie Mellon University technical report [21].

II. BACKGROUND AND RELATED WORK

Two areas of robotics research are related to this work: localization and distributed sensing.

Two common probabilistic methods used for localization in robotics are Kalman filters [20] and Monte Carlo Localization [9]. Kalman filters have been used to localize robots within mapped environments. Typically, this approach is used to fuse position estimates from multiple synchronous or asynchronous sensors on a single robot [10][13][15][19]. Monte Carlo localization has been specifically applied to the multiple robot localization problem [9]. While most of these approaches rely on previously obtained maps of the environment, some recent research has been ongoing in simultaneously mapping the environment while localizing the robot within it [3][8]. Kalman filters have also been used for object tracking in several methods. A few examples are: estimation of error between different coordinate frames to determine the relative locations of objects [20], identification and prediction of future object locations [18][22], real-time tracking and camera pointing [23], and tracking for grasp using proximity sensors [17].

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]) and others share information to fill in blank areas in the world model (CS Freiburg [10][11], RMIT [4], 5dpo [7]). Sharing and merging overlapping information for object localization has received little attention. Two of the few examples are the use of multiple sequential robot observations as input to Kalman filter tracking [5][6] and updating of grid cell occupancy probabilities using sequential observations from 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. Our approach does not require discretization, working in the continuous spatial domain.

III. FUSING GAUSSIAN DISTRIBUTIONS

A. 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 any location, given the observation.

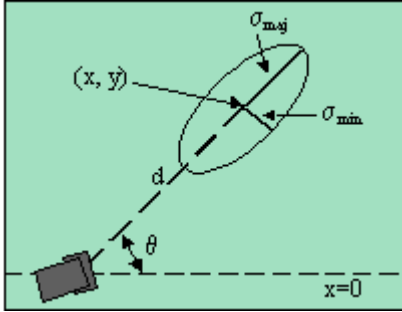


Figure 1. Gaussian distribution parameter definition: mean (x, y) , angle of major axis θ , standard deviations along major and minor axes σ_{maj} and σ_{min} , and distance to mean d .

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, first introduced by Duffin [1], based on Bayes' Rule and Kalman filters. 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, then, is to collect observations of multiple robots (as in Figure 2), then merge the corresponding Gaussian distributions to yield a better estimate of the location and uncertainty of the observed object.

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

$$p(X) = \frac{1}{2\pi\sqrt{|C|}} \exp\left(-\frac{1}{2}(X - \bar{X})^T C^{-1}(X - \bar{X})\right) \quad (1)$$

where

$$C = \begin{bmatrix} \sigma_x^2 & \rho\sigma_x\sigma_y \\ \rho\sigma_x\sigma_y & \sigma_y^2 \end{bmatrix} \quad (1b)$$

The parameterization of the Gaussian in this representation does not correspond to 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, distributions can be merged using matrix operations. After the observations are merged, we extract the mean and standard deviations from the merged result (these correspond to the estimated location and uncertainty of the observed object).

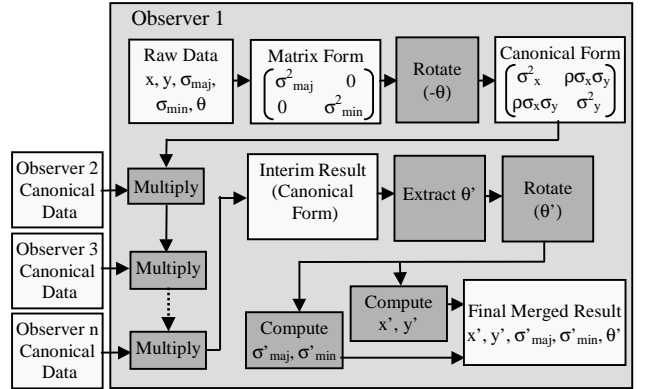


Figure 2. Block diagram representing the multi-distribution merging process. The multiplication step is conducted using the mathematical formulation described above. Each subsequent distribution is merged with the previous merging result.

B. Mathematical Details

We wish to determine the mean, standard deviations, and angle of the merged distribution to estimate object position and characterize the quality of the estimate. We compute these same parameters from sensor readings (mean and angle) and models of sensor error (deviations). Thus, we require a method of determining combined parameters from those of individual distributions.

The matrix formulation we use, adopted from Smith and Cheeseman [20], makes this computation relatively simple. The mean, standard deviations, and orientation of the major axis are independent of scaling; they can be extracted from the merged covariance matrices without the need to consider absolute probabilities.

The covariance matrix of an observation, C , is initially determined from the major and minor axis standard deviations in the local coordinate frame (designed L).

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

Since observations may be oriented arbitrarily with respect to the global coordinate frame, they be transformed to this frame. Rotation of X in equation 1 by leads to the following relationship:

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

where θ is the angle of the distribution's principal axis with respect to the global x -axis. This rotation accomplishes the transformation from observation parameters to the canonical form. Once the observation is in the canonical form, we continue merging.

Two covariance matrices are combined into a single covariance matrix representing the combined distribution:

$$C' = C_1 - C_1 [C_1 + C_2]^{-1} C_1 \quad (4)$$

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

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

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

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

where A , B , and D represent the components of the covariance matrix (upper left, upper right/lower left, and lower right, respectively).

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

$$C_L' = R(\theta')^T C' R(\theta') \quad (7)$$

and the resulting major and minor axis standard deviations can be directly computed from the covariance matrix by reversing Equation 2.

C. Simulated Example

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

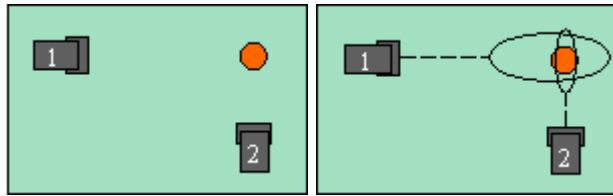


Figure 3. *Left:* Two robots see a target from different positions. *Right:* The robots' observations generate Gaussian distributions; uncertainty (1- σ ovals shown) increases with distance.

For this example, the robots are positioned with relative headings 90 degrees apart and looking directly at the target. The target is located at a position (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 the top of

Figure 4, the distributions resulting from the individual measurements by robot 1 and robot 2 are shown.

The result of merging distributions is shown in the bottom of Figure 4. It is easily observed that the implied uncertainty is reduced, and the resulting distribution is centered more accurately 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).

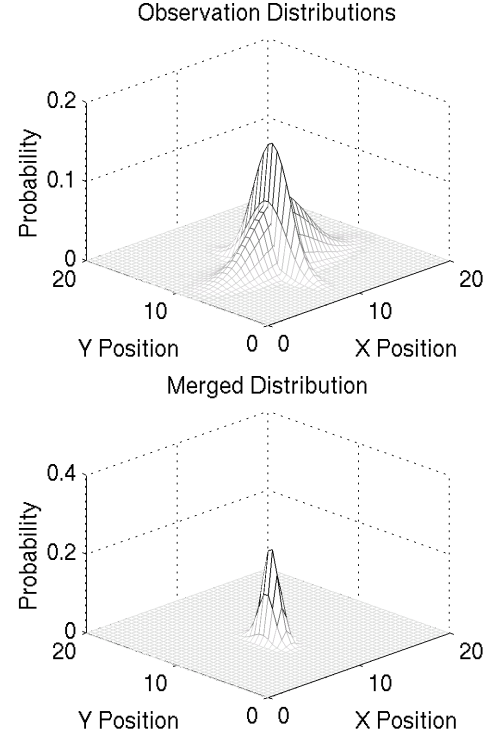


Figure 4. *Top:* The individual distributions, from observations, to be merged. *Bottom:* Resulting merged distribution, with reduced error and higher accuracy in the mean.

IV. ASSUMPTIONS

The primary assumption upon which the mathematics is abased is that sensor errors are independent and distributed normally. As demonstrated by the camera calibration procedure (see Hardware Platform section), this assumption approximately holds for this system.

The robot coordinate frames are assumed to be coincident. Without this, data are incompatible and the merging is meaningless.

We assume robots to be perfectly localized. Robot positional uncertainty is not taken into account in the generation of target location distributions, as at the current time our method of localization does not provide uncertainty estimates. If the errors in localization can be modeled as a Gaussian, the mathematics to take this uncertainty into account is similar to that of the merging of measurement distributions. The errors in robot configuration (x , y , θ) are geometrically mapped to standard deviations along axes oriented parallel to the major and minor axes of the measurement distribution. This creates a Gaussian approximation to localization

error. The standard deviations of the observations can then be “grown” to accommodate positional error by merging it with the position distribution (re-centered at the measurement distribution’s mean). Preliminary calibration of the Cyre odometry indicates that a Gaussian approximation is likely reasonable.

Several additional assumptions were introduced for simplicity in experimentation and camera calibration. Measurements are assumed simultaneous. High frame rates and low object speeds make this reasonable; largely different readings are not merged. 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. Last, objects are assumed to be unique to avoid the necessity of solving the association problem.

V. VALIDATION ON ROBOTS

A. Hardware Platform

The hardware platform used is a modified Cyre robot, an inexpensive and commercially available platform (Figure 5). This platform consists of a drive section and a trailer. The drive section uses two-wheel differential drive and is equipped with a front bump sensor; the trailer is passive. On board the Cyre 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.

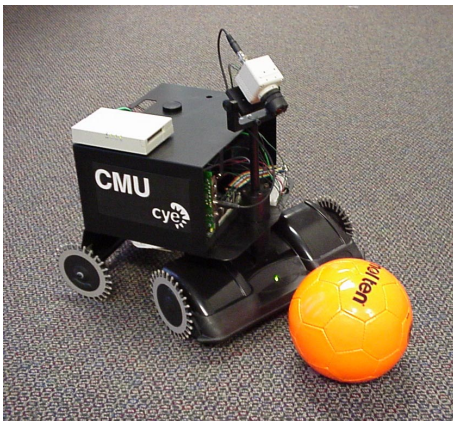


Figure 5. The modified Cyre robot platform.

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.

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 bias, proportional errors are determined.

A histogram of these sensing errors determined that the corrected distances are distributed about actual distance approximately normally (see Figure 6). 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.

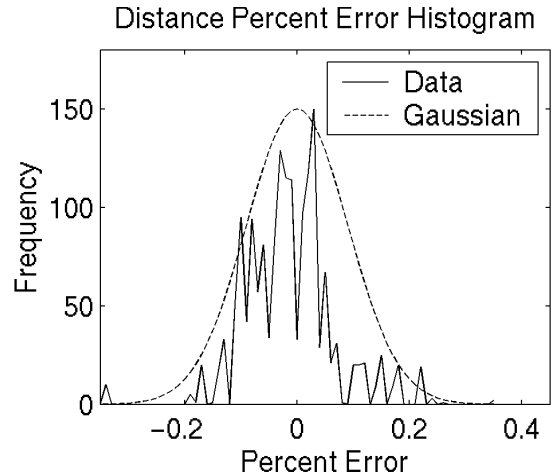


Figure 6. The histogram of distance error as a percent of distance is approximately Gaussian.

B. Experimental Setup

To test the approach, an experiment was devised in which three robots track an object through several sequential points in the environment. In this way, the accuracy of single-robot measurements can be directly compared to the accuracy obtained by combining data from two and three robots. An illustration of this experimental setup is shown in Figure 7; each robot is 1.5 meters from the origin of the coordinate frame.

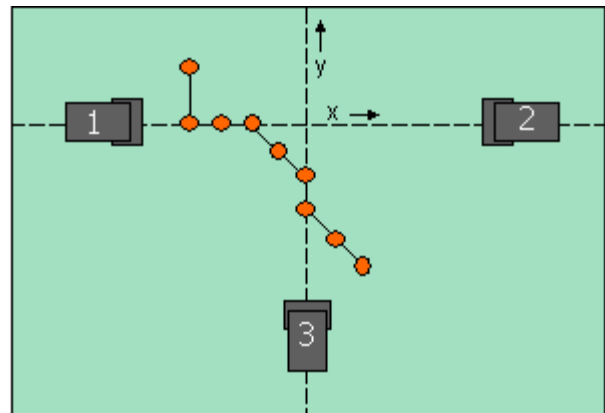


Figure 7. Experimental setup for validation. Ball locations marked as circles along trajectory. Dotted lines show the global coordinate frame.

The ball was placed at a series of known discrete points along a trajectory. At each point, the ball location measured by each robot was recorded. These individual observations were merged together and in pairs.

C. Experimental Results

The experimental results are shown graphically in the following figures. In Figure 8, example observed trajectories are shown. An example trajectory seen by an individual robot is compared to the actual path of the target is shown at the top. In the middle, a merged result from two robots is similarly compared. The merged result from all three robots is shown at the bottom.

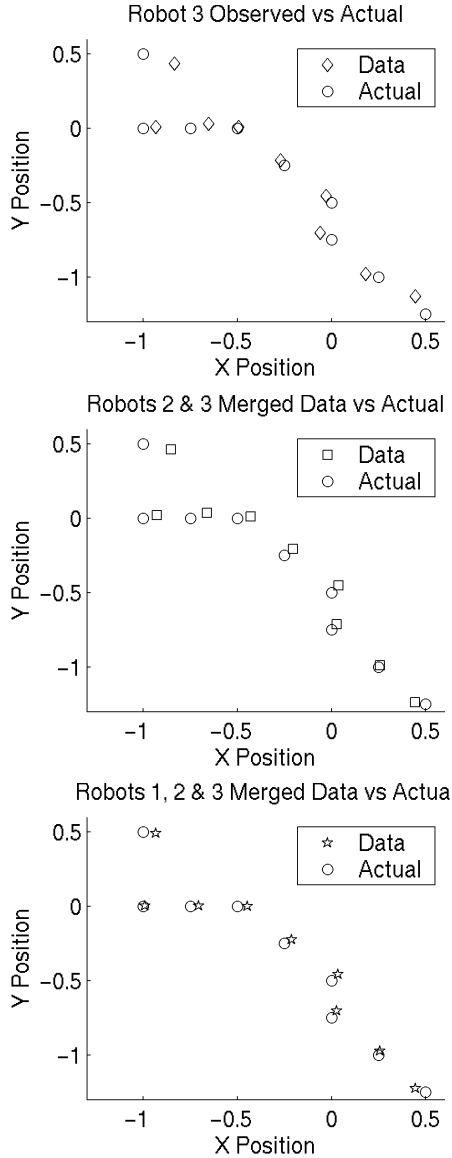


Figure 8. Reported data versus actual ball position. *Top:* Robot 3's estimate. *Middle:* Robots 2 and 3 merged estimates. *Bottom:* Robots 1, 2, and 3 merged estimates.

In Figure 9, trajectory errors point-by-point and mean error are shown and compared for all single-robot, two-robot, and three-robot measurements.

While individual trajectories are sometimes accurate at single points, (in fact, occasionally more accurate than the combined information) the consistency of accuracy shown in the combined results is absent in the single-robot trajectories. This is best characterized by plotting the mean error of single-robot observations, two-robot observations, and three-robot observations (Figure 10).

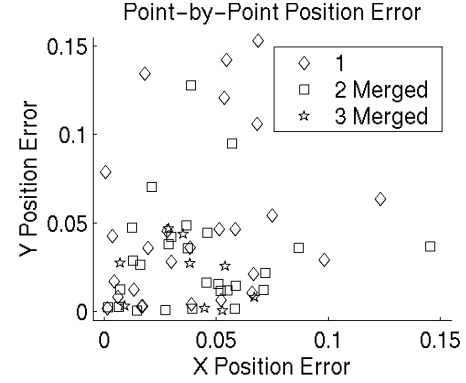


Figure 9. Position error in x and y for each measurement. Each merging lowers position error bounds and reduces outlier frequency.

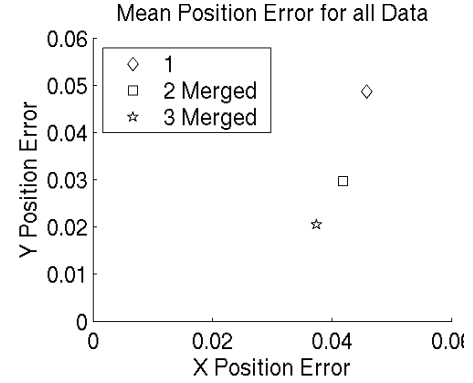


Figure 10. Mean position error in x and y over all single observations, all 2-observation merges, and all 3-observation merges.

It should be noted that improvement in x is of a smaller degree than that in y. While this one experiment may not be typical, this is likely a result of the specific experiment configuration. In this case, only one robot is aligned along the y-axis to improve accuracy along the x-axis, whereas two robots have better accuracy along the y-axis. This imbalance may be the cause of this discrepancy.

VI. TEST APPLICATIONS AND RESULTS

In each of these tests, robots begin in precisely known starting positions. Deviations from these positions are small or zero. Thus, assumptions on robot localization and coordinate frames hold. All targets are constrained to the ground plane (except in soccer) to remain within the capabilities of the vision system calibration.

A. Location and Retrieval of Unseen Targets

This test application for this method demonstrates the ability to increase the effective field of view of agents. In this experiment, one robot is positioned so that it can see,

but not reach, the target object. The other robot cannot initially see the target object, even with camera panning, but the path to the object is unobstructed (Figure 11).

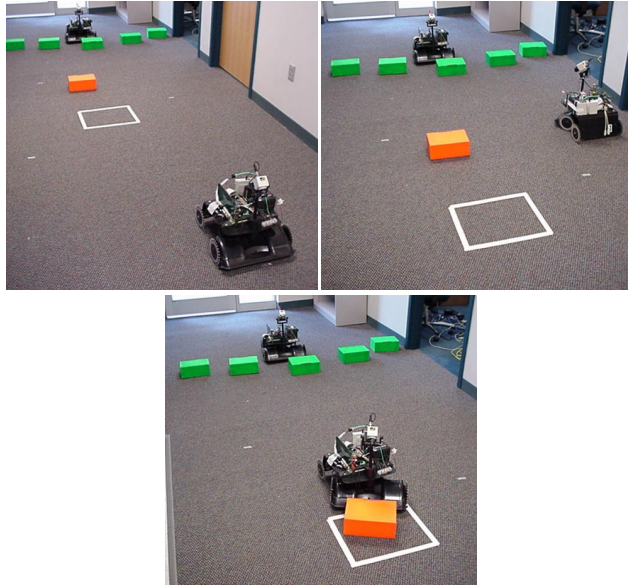


Figure 11. Unseen target location and pushing experiment setup. *Top left:* Robot 1 (closer) cannot see target from its starting location, even with camera panning. *Top Right:* Robot 1 proceeds directly to the target without searching, given the target location by Robot 2. *Below:* Robot 1 pushes target to goal using a merged target position.

By sharing visual information, the unobstructed robot immediately obtains a target location without requiring random search. The robot is able to successfully locate the object using information provided exclusively by the second robot. Once the object is located, it effectively reaches and manipulates the target using the merged position provided by both robots.

B. Blind Robot Target Tracking

This experiment is designed to test increased effective field of view, rapid target acquisition, and tracking using a merged position. In this experiment, three robots are positioned around a target area. A target (soccer ball) is moved throughout the target area, and all three 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. When one robot loses sight of the target, it immediately looks in the correct direction and quickly recovers the object. Even when the target travels within the line of sight of a single robot (with diminished accuracy in this dimension), the additional point of view typically makes up for this lack.

One of the robots is subsequently blindfolded by covering the camera (with a box, for example, as in Figure 12). For convenience, these robots are typically positioned at relative headings of 90 degrees. The ability of the blinded robot to track the ball using only the merged position from the other two is not substantially diminished.

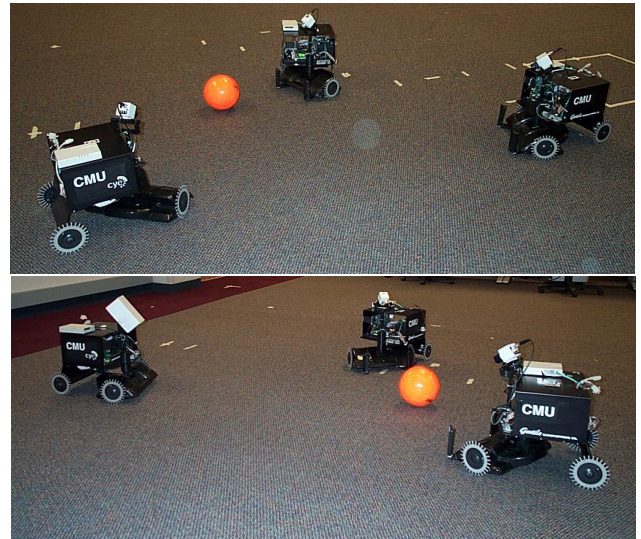


Figure 12. Blind robot ball tracking experiment setup. *Top:* In this experiment, three robots track the ball with a location generated by combining all three observations. *Bottom:* In this experiment, a robot is blinded (for example, a box covers the left robot's camera) and it still successfully tracks the ball with combined data from the remaining two.

C. Robot Soccer

This approach to distributed sensing was applied to the CMU Hammerheads RoboCup middle-sized team in 2000 (Figure 13). At each cycle, robots would transmit the position of the ball, if visible, so that it could be combined with all other current observations. Widely conflicting observations (for example, distributions with means differing by at least 2 standard deviations) were not merged. This eliminates confusion resulting from false targets (or more generally, multiple targets) and data collected in incompatible coordinate frames. This, in theory, would provide robots far from the ball with more accurate positions and allow robots that could not see the ball to quickly locate it.

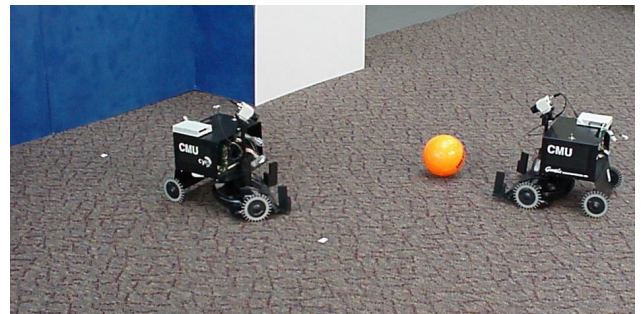


Figure 13. A test run of robot soccer. Attacker (right) attempts to score on blue goal defender (left).

The CMU Hammerheads 2000 team robots were localized entirely by on-board odometry. This odometry drifts over time, leading to differences in coordinate frames. As a result, the primary impact of distributed sensing was to provide a starting point for robots to locate balls that had become invisible. This use of shared information did allow robots to more quickly locate a lost ball during competition when it was not entirely obstructed. Despite

the discrepancy in coordinate frames, the coordinate frames were generally coherent enough that when a robot looked at a potential target position, the ball became visible within the camera's wide field of view.

VII. CONCLUSIONS

We present a method for using a simple representation of two-dimensional Gaussian distributions to fuse target position estimates from two or more robot agents. This approach is based on accepted Bayes' Rule and Kalman filter theory and implements real-time sensor data fusion on a reactive multi-robot system for several different applications. The successful ability to fuse these statistical measurements and the ability to receive position estimates on targets not visible allows our robots to quickly acquire targets and to more accurately estimate object position. While this work is used vision for sensing, the approach can be applied to any sensor or suite of sensors that 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 practically implementing this method of distributed sensing and information sharing. Even in well-localized systems, disparity between coordinate frames can arise. Such disparity must be accommodated and/or corrected. Autonomous re-merging of coordinate frames using sensors will be investigated. Additionally, well-localized robots still maintain some uncertainty in position, as computed by the localization. The accommodation of robot positional uncertainty will be incorporated into the target position distributions. Lastly, it may be possible to remove the ground plane for unknown objects; this involves a different method of transforming pixel location into world coordinates and requires further research.

VIII. ACKNOWLEDGEMENTS

The authors would like to acknowledge the contributions of Jim Bruce (vision), Scott Lenser (localization), Kevin Sikorski (camera calibration), Hans Moravec (statistics and camera calibration) and Manuela Veloso. The authors would also like to acknowledge DARPA's Mobile Autonomous Robot Software Program (MARS) and Control of Agent Based Systems Programs for support of this research.

IX. REFERENCES

- [1] W. Anderson and R. Duffin. *Journal of 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, A. Jennings, M. Makies, C. Keen, A. Kendall, L. Padgham, and D. Singh. "RMIT Raiders." In: Veloso, Pagello, and Kitano, eds. *RoboCup-99: Robot Soccer World Cup III*. Springer-Verlag, pages 741-744, 2000.
- [5] A. Cai, T. Fukuda, and F. Arai. "Information Sharing among Multiple Robots for Cooperation in Cellular Robotic System." In: *Proceedings of 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: *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, pages 2079-2084, 1996.
- [7] P. Costa, A. Moreira, A. Sousa, P. Marques, P. Costa, and A. Matos. "5dpo-2000 Team Description." In Veloso, Pagello, and Kitano, eds. *RoboCup-99: Robot Soccer World Cup III*. Springer-Verlag, 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: *Proceedings of the 2000 IEEE International Conference on Robotics and 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, W. Hatzack, I. Herrmann, B. Nebel, F. Rittinger, A. Topor, and T. Weigel. "Reliable self-localization, multirobot sensor integration, accurate path-planning and basic soccer skills: playing an effective game of robotic soccer." In: *Proceedings of the Ninth International Conference on Advanced Robotics*, pages 289-296, 1999.
- [11] J.-S. Gutmann, W. Hatzack, I. Herrmann, B. Nebel, F. Rittinger, A. Topor, T. Weigel, and B. Welsch. "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, 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, 2001.
- [13] T. D. Larsen, K. L. Hansen, N. A. Andersen, O. and Ravn. "Design of Kalman filters for mobile robots; evaluation of the kinematic and odometric approach." In: *Proceedings of the 1999 IEEE International Conference on 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>.
- [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: *Proceedings of the Ninth International Conference on Advanced Robotics*, pages 283-288, 1999.
- [16] D. Nardi, G. Adorni, A. Bonarini, A. Chella, G. Clemente, E. Pagello, and M. Piaggio. "ART99 - Azzurra Robot Team." In: Veloso, Pagello, and Kitano, eds. *RoboCup-99: Robot Soccer World Cup III*. Springer-Verlag, pages 695-698, 2000.
- [17] G. Petryk and M. Buehler. "Robust estimation of pre-contact object trajectories." In: *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: *Proceedings of the 24th Annual Conference of the 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: *Proceedings of the Third International Conference on 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 Technical Report CMU-RI-TR-00-30, 2000.
- [22] C. Y. Tang, Y. P. Hung, S. W. Shih, and Z. Chen. "A feature-based tracker for multiple object tracking." In: *Proceedings of the National Science Council, Republic of China*, Part A, Vol. 23, No. 1, pages 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.