Submitted, 2005 Int'l Conference on Robotics and Automation
http://www.cds.caltech.edu/~murray/papers/2004t_cbm05-icra.html

On Decentralized Classification using a Network of Mobile Sensors

Timothy H. Chung*, Joel W. Burdick, Richard M. Murray
California Institute of Technology
Division of Engineering and Applied Sciences
Pasadena, CA 91125, USA
{*timothyc, jwb, murray*}*@caltech.edu*

***Abstract*— This paper considers how a team of mobile sensors should cooperatively move so as to optimally categorize a single moving target from their noisy sensor readings. The cooperative control procedure is based on the development of a cost function that quantifies the team's classification error. The robots' motions are then chosen to minimize this function. We particularly investigate the case where the sensor noise and class distributions are Gaussian. In this case, we can derive a *duality principle* which states that optimal classification will be realized when the covariance of the target estimate is minimized. That is, in this case, optimal estimation leads naturally to optimal classification. We extend previous work to develop a distributed discrete-gradient search algorithm that guides the team's location motions for purposes of optimal estimation and classification. The concepts developed are validated through numerical studies.**

***Index Terms*— Multiple robots, cooperating robots, distributed sensing, mobile sensors, classification**

I. Introduction

This paper considers the problem of how a team of mobile robots, endowed with noisy sensors, should cooperatively move and fuse their sensory data so as to maximize their collective ability to correctly classify a moving target. Many applications require a single robot or a team of robots to correctly categorize the target(s) of their observations, such that further actions can be made based on the classification. Distinguishing between rocks, bushes, and shadows, for example, can affect the choice in trajectory taken by unmanned off-road vehicles [1], [2]. As another example, directed exploration where agents are tasked with seeking and observing a particular class of entities, such as jellyfish [3], would greatly benefit from real-time classification of objects. Similarly, there are many applications where classification of teammates (friendly) or opponents (foe) would be useful in determining a course of action [4], [5].

The study of sensor networks, particularly ones in which sensor nodes can maneuver, has become an active research area, including work on distributed localization and mapping [6], optimal sensor placement [7], [8], and target tracking [9], [10]. The notion of active sensing, where sensors are configurable or actuated, has also been examined extensively [11], [12], [13]. In general, however, these works address the issue of estimating the state of objects, and do not examine the use of this information for the purposes of classifying the objects. As sensor networks become more "intelligent," there is a natural need to advance from sensor fusion to sensor interpretation for the purpose of intelligent decision making and the task of generating new network objectives. The classification of network data is a necessary step towards more intelligent sensor networks.

There is an ample literature in the fields of machine learning [14] and computer vision [15] that deals with classification of objects [16]. Further, detection theory examines the presence of types of signals, leading to results seen in, e.g. [17], including work done on distributed detection [18], [19], [20], [21]. These investigations and results do not, however, discuss either the improvement of measurements that can be obtained by sensor mobility, or the connection between estimation and classification tasks.

The main focus of this paper is to consider how cooperative teams of mobile sensors can improve their ability to classify a (possibly moving) target via the use of mobility and sensor fusion. We formulate a cost function whose minimization by cooperative team motions leads to improved classification.

We also investigate the relationship between distributed estimation and optimal classification. In particular, for the case of Gaussian sensor noise and Gaussian class distributions, we develop a *duality principle*, which states that optimal target estimation leads to optimal classification. We then extend some recent work [22] to develop a distributed discrete gradient search algorithm whose output drives the local robot motions so as to optimize their estimation and classification abilities. A simple simulation illustrates the method.

II. Problem Formulation

This section describes the framework used to formulate the problem of maneuvering mobile sensing agents cooperatively for the purpose of target classification.

A. Distributed Sensing

We are given a team of M mobile sensors that move in the plane (e.g., M holonomic point robots equipped with sensors). We assume that a target, $\mathcal{T}$, enters into the domain of the

*This work is supported in part by the Engineering Research Centers Program of the National Science Foundation under Award Number EEC-9402726 and also a grant from NASA.

robots' interest. The j-th sensor takes measurements $y_j \in \mathbb{R}^2$ of the target state $x \in \mathbb{R}^2$ according to the observation model:

$$y_j[k] = x[k] + T(\theta_j) v_j[k], \quad (1)$$

where $v_j[k] \in \mathbb{R}^2$ is zero-mean Gaussian measurement noise (in range and bearing) with covariance matrix $R_j[k] \in \mathbb{R}^{2\times 2}$. $T(\theta_j) \in SO(2)$ is the rotation matrix that transforms the noise from the robot's body fixed coordinates to a global (cartesian) coordinate system, under the assumption that the sensors have complete and perfect knowledge of their positions. Note that $T(\theta_j)$ introduces a dependence of the observations on the relative bearing between sensor j and the target. The most common sensor measurement will be range to the target, but a wide variety of sensing modalities can be incorporated into this framework.

We assume that the motion of the target is not *a priori* known. Thus, for purposes of planning, we use a simple random walk model to predict the motion of the target in the plane:

$$x[k+1] = x[k] + w[k], \quad (2)$$

where $w[k] \in \mathbb{R}^2$ represents the process noise, which is assumed zero-mean, Gaussian and white with covariance matrix $Q[k] \in \mathbb{R}^{2\times 2}$.

Following the construction of [22], we assume that the range-measuring sensors, following standard sonar models [23], make measurements corrupted by noise in both range and bearing. The uncertainty of the measurement in the range direction is allowed to depend on the distance from the sensor to the target. Such a spatial dependence of the measurement noise on sensor position can be generalized for models of other types of sensors. Furthermore, the measurement noise of different sensors is assumed to be mutually independent.

The covariance matrix R_j, in keeping with the usual sonar models (see, e.g., [23]), has the diagonal form

$$R_j = \begin{bmatrix} \left(\sigma^j_{\text{range}}\right)^2 & 0 \\ 0 & \left(\sigma^j_{\text{bearing}}\right)^2 \end{bmatrix}. \quad (3)$$

$\left(\sigma^j_{\text{range}}\right)^2$ is the range measurement noise variance. The spatial dependence of the range uncertainty is represented by a function $f(r_j)$ of the distance r_j from sensor j to the target. As in [22], for the simulations in Section V we use a quadratic-in-range function for $f(r_j)$, such that the best measurements are attained at a "sweet spot" of the sensor. The bearing noise variance $\left(\sigma^j_{\text{bearing}}\right)^2$ is often modelled (e.g. see [24]) as a fixed multiple α of the range noise variance. Thus we obtain a covariance model of the form:

$$R_j[k] = \begin{bmatrix} f(r_j) & 0 \\ 0 & \alpha f(r_j) \end{bmatrix}. \quad (4)$$

Given these measurements from the sensing nodes, a sensor fusion algorithm is used to combine the local estimates $\hat{x}$ and estimate error covariance matrices P, which are generated by Kalman filters processing local measurements at each node. Addressing the issue of the cross-covariance components discussed in [22], [25], [26], we apply the inverse covariance form of the Kalman filter [27], [28], which is already decentralized in nature, for the fusion process. First, defining the inverse covariance forms of the estimate and covariance [10]

$$\hat{z}[k] \equiv P^{-1}[k]\ \hat{x}[k], \quad Z[k] \equiv P^{-1}[k],$$

the decentralized form of the Kalman filter for our system is summarized by the following sets of equations [10]. For the propagation step of the Kalman filter, we have:

$$\begin{aligned} \hat{z}[k]^- &= \left(I - Z[k-1]\left(Z[k-1] + Q[k]^{-1}\right)^{-1}\right)\hat{z}[k-1] \\ Z[k]^- &= \left(I - Z[k-1]\left(Z[k-1] + Q[k]^{-1}\right)^{-1}\right) Z[k-1]. \end{aligned}$$

The measurement update step, fusing all new measurements $y_j[k]$, is given by

$$\begin{aligned} \hat{z}[k] &= \hat{z}[k]^- + \sum_{j=1}^{M} R_j[k]^{-1} y_j[k] \\ Z[k] &= Z[k]^- + \sum_{j=1}^{M} R_j[k]^{-1} \equiv \Sigma[k]. \end{aligned}$$

The matrix Σ represents the covariance of the error in the global estimate obtained from the fused data and hence is an indicator of the quality of the fused estimate. Since the sensor noise covariance matrix is a function of the distance between the sensor and the target, the quality of the estimate depends on the distances between the various sensors and the targets. Thus by varying the positions of the sensors, we can vary the error covariance.

The challenge is: *Given these measurements of the state of target, how should the sensors move in order to improve their overall ability to correctly classify a target?* To address this question, we first develop a cost function which evaluates the quality of the classification. The intuitive choice is the probability of error in classification, denoted p_e, which describes the likelihood that the categorization made from measurements is not the true class of the target.

B. Classification Probability

First we formulate the expression for p_e generally, but will then concentrate on the binary hypothesis problem, where there are two classes (e.g. heads or tails, 0 or 1, friend or foe). Much literature already exists for this class of problems (e.g. see [29], [30]), and further, multiple hypothesis problems are a relatively straight-forward extension of the two-class scenario [17].

Let $\mathcal{C}$ denote the set of N classes, i.e. $\mathcal{C} = \{c_1, \ldots, c_N\}$. For example, in the binary hypothesis scenario, we would have $N = 2$ such that c_1 and c_2 represent "friend" and "foe," respectively.

Following [17], the different hypotheses are represented as:

$$c_i : p(y|x_i) = \mathcal{N}(x_i, \Sigma). \quad (5)$$

Here, y represents the single resulting fused estimate arising from the sensor fusion process, such as the inverse covariance Kalman filter described in the previous section. Describing (5) in words, if the target is a member of class c_i, the measurements are distributed normally with covariance Σ and centered about the true state x_i. In addition, x_i itself is also a random variable whose distribution may depend on the class c_i. We will consider in particular the case where the probability distributions of members of a class are also random variables:

$$p(x_i) \equiv p(x|c_i) = \mathcal{N}(\bar{x}_i, \Delta_i). \quad (6)$$

where $\bar{x}_i$ and Δ_i represent the mean and covariance of the distribution of members in class c_i.

We assume that classification decision rules are known in advance. The rules defining the boundaries between classes can be learned, or determined from first principles. In either case, we denote the region in measurement space where c_i is the correct class by $\Omega_{y,i}$ and its complement by $\bar{\Omega}_{y,i}$.

With the above definitions in mind, following [31] the probability of error in classification can be expressed as:

$$\begin{aligned} p_e &= \sum_{i=1}^{N} p(error|c_i)p(c_i) \\ &= \sum_{i=1}^{N} p(c_i)\left(1 - \int_{\Omega_{y,i}} p(y|c_i)\ dy\right) \\ &= 1 - \sum_{i=1}^{N} \int_{\Omega_{y,i}} p(y|c_i)p(c_i)\ dy, \end{aligned}$$

where $p(c_i)$ represents the *a priori* probability of the target being in class c_i, and $\sum_{i=1}^{N} p(c_i) = 1$.

Note that the probability of correct classification, p_c, is related to p_e by $p_e = 1 - p_c$. In the case of many classes, it is easier to investigate the probability of correct classification. When the *a priori* class distributions are independent of the measurements, the utility function to be maximized is:

$$p_c = \sum_{i=1}^{N} p(c_i) \int_{\Omega_{y,i}} p(y|c_i)\ dy. \quad (7)$$

We seek an expression $p(y|c_i)$, which relates the probability distribution of the measurements with a particular class. By definition of marginal probabilities, we find that

$$p(y|c_i) = \int p(y, x|c_i)\ dx = \int p(y|x, c_i)p(x|c_i)\ dx, \quad (8)$$

where the second equation comes from simple application of Bayes' rule.

For a particular class c_i, we can rewrite (8) to incorporate the relationship between target state x and the i-th class c_i:

$$p(y|c_i) = \int p(y|x_i)p(x_i)\ dx_i. \quad (9)$$

where x_i is introduced as the nuisance parameter [17]. Thus (7) and (9), which depend implicity upon the robots' positions, define a utility function whose maximization over the set of possible robot motions leads to the best classification performance of a team of cooperating mobile sensors.

Gaussian noise and class distributions. In the particular case where the sensing noise is Gaussian and the class distributions are Gaussian, from (5) and (6) respectively we obtain:

$$p(y|x_i) = \frac{1}{(2\pi)^{\frac{n}{2}}|\Sigma|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(y - x_i)^T \Sigma^{-1}(y - x_i)\right)$$

$$p(x_i) = \frac{1}{(2\pi)^{\frac{n}{2}}|\Delta_i|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(x_i - \bar{x}_i)^T \Delta_i^{-1}(x_i - \bar{x}_i)\right).$$

such that (8) becomes:

$$\begin{aligned} p(y|c_i) &= \int \mathcal{N}(x_i, R)\mathcal{N}(\bar{x}_i, \Delta_i)\ dx_i \\ &= \mathcal{N}_y\left(\bar{x}_i, (\Sigma + \Delta_i)\right). \end{aligned}$$

We see that the integral of the product of Gaussians is itself a Gaussian distribution [32].

Thus in this specialized case, the utility function to be maximized, the probability of correct classification, is given by:

$$p_c = \sum_{i=1}^{N} p(c_i) \int_{\Omega_{y,i}} \mathcal{N}_y\left(\bar{x}_i, (\Sigma + \Delta_i)\right)\ dy. \quad (10)$$

III. Relationship between Estimation and Classification in the Gaussian Case

In general, the computation of the probability of correct classification is challenging, even when done numerically [32]. Optimization over sensor positions in order to determine the best trajectories of the sensors further complexifies the computation.

However, we may gain an understanding of the behavior of (10) by investigating its dependence on the sensor positions. Note that the covariance of the target estimate is affected by sensor motion, due to the spatially-dependent measurement noise. Before proceeding, we now state the following theorem, which will form the basis for a *duality principle* that simplifies the task of maximizing the utility function.

Theorem 1: Let the scalar-valued function $f : \mathbb{R}^{n \times n} \to \mathbb{R}$ that operates on a matrix $A \in \mathbb{R}^{n \times n}$ be given as:

$$f(A) = \sum \kappa \int \mathcal{N}(A)\ dz,$$

where κ is a constant, $z \in \mathbb{R}^n$, and

$$\mathcal{N}(A) = \frac{1}{(2\pi)^{\frac{n}{2}}\ |A|^{\frac{1}{2}}} \exp\left(-\frac{1}{2} z^T A^{-1} z\right).$$

Then, given that A and $B \in \mathbb{R}^{n \times n}$ are positive definite,

$$f(A + B) < f(A).$$

Proof: Examining the function f applied to the perturbed matrix $A+B$, first we find that

$$\begin{aligned}\mathcal{N}(A{+}B) &= \frac{1}{(2\pi)^{\frac{n}{2}}\,|A{+}B|^{\frac{1}{2}}}\exp\left(-\frac{1}{2}z^T\,(A+B)^{-1}\,z\right)\\ &< \frac{1}{(2\pi)^{\frac{n}{2}}\,|A|^{\frac{1}{2}}}\exp\left(-\frac{1}{2}z^T\,(A+B)^{-1}\,z\right),\end{aligned}$$

noting that $|C+D|>|C|$ for C and D positive definite [33]. Further, use of the identity [34]

$$\left(C^{-1}+D^{-1}\right)^{-1}=C-C\,(C+D)^{-1}\,C,$$

and the fact that A and B are nonsingular allow further simplification:

$$\begin{aligned}\mathcal{N}(A{+}B) &< \frac{1}{(2\pi)^{\frac{n}{2}}\,|A|^{\frac{1}{2}}}e^{-\frac{1}{2}z^T\left(A^{-1}-A^{-1}\left(A^{-1}+B^{-1}\right)^{-1}A^{-1}\right)z}\\ &< \mathcal{N}(A)\;e^{\frac{1}{2}z^T\left(A^{-1}\left(A^{-1}+B^{-1}\right)^{-1}A^{-1}\right)z}\\ &< \mathcal{N}(A).\end{aligned}$$

The last inequality is seen by observing that a quadratic form is always positive, and hence, the exponential factor must be greater than or equal to identity.

Noting that f is the integral over positively-valued $\mathcal{N}$, we immediately arrive at the desired result. □

Theorem 1 leads to the conclusion that the probability of correct classification, p_c, is inversely related to the sum of Σ and Δ_i, which are both positive definite. Hence, we arrive at the following corollary.

Corollary 1: In the case of Gaussian sensor noise and Gaussian class distributions, maximization of the probability of correct classification, p_c, is achieved by minimization of the determinant of the estimate error covariance matrix $|\Sigma|$.

Proof: Note that the class distribution covariance Δ_i is independent of sensor positions, and so we can focus on changes in Σ due to the motion of the sensors. Define A greater than B for $A,B\in\mathbb{R}^{n\times n}$ and positive definite if and only if the matrix $A-B$ is also positive definite. Hence, for A,B positive definite matrices,

$$|A|>|B|\Leftrightarrow A>B,$$

or in other words, decreasing a positive definite matrix, such as Σ, decreases its determinant. Thus, we immediately see the inverse relationship between p_c and $|\Sigma|$ from application of Theorem 1. □

Note that the objective for distributed optimal estimation tasks is to choose the collective team motions to minimize the estimation cost function, commonly given by the determinant of the estimate error covariance matrix. For this reason, we term the statement of Corollary 1 the *estimation-classification duality principle*.

Thus, we see that for the special case of Gaussian noise and class distributions, there exists a direct connection between the classification problem and the task of estimating the target state. This result simplifies the distributed classification objective to one of distributed sensing, for which many possible approaches have been suggested in literature, such as [22], [35], [9]. These methods yield the optimal sensor motion paths and configurations for estimating or tracking targets, which are, for the given formulation, the same paths and configurations for optimal classification of targets.

IV. Discrete Gradient Search Algorithm

For completeness, we now discuss recent work [22] on a gradient-descent-based algorithm used to optimize sensor motions for estimation/tracking of target(s).

Assume the presence of a single target in the sensing field, and assume every sensor observes the target. Recall from Section II-A that at each time step, each sensor makes an observation, processes the measurement locally using a Kalman filter to update its estimate, and fuses this estimate with information from other sensors to obtain a global estimate $\hat{z}$ and its global error covariance matrix $Z\equiv\Sigma$.

The objective in optimal estimation and tracking is to determine the sensor motions which minimize the uncertainty present in the fused estimate. Furthermore, we require a distributed solution such that each sensor identifies its optimal location for the next time step. However, the complexity of the problem makes attaining such a solution for all sensors, whether analytic or numerical, rather intractable. Given that the gradient provides the locally optimal direction of movement, we use a gradient descent algorithm which defines the optimal control action as that which will position the sensors to minimize $|\Sigma|$ in the following time step. This approach for optimization is appropriate due to the fact that the calculation of the gradient is intrinsically decentralized [36].

To further ameliorate the challenge of computation, we reduce the gradient descent algorithm to a discrete gradient search algorithm by restricting the possible control actions for each sensor to a finite, discrete set of motions. This algorithm is summarized by Table I.

In this manner, the responsibility for optimization of sensor motion is given to each sensor, in place of a central computation node. The only required interaction between sensor nodes is the communication of local information for data fusion. Note that each sensor obtains position information of the other sensor nodes implicitly from the transmitted local information. Further, the decentralized nature of the algorithm enables a sensor to simply disregard any non-communicating nodes in the sensor fusion step.

V. Simulation Results

To illustrate the ideas presented in the previous sections, we investigate a simplified but demonstrative example of a classification scenario. The goal of this example is to generate the optimal motion trajectories for multiple mobile sensing agents

TABLE I

DECENTRALIZED GRADIENT-SEARCH-BASED ALGORITHM

```
For k=1:sim_time

  % —– Local Observation —–
  Take local measurement;
  Update local estimate ẑ and error covariance matrix P;

  % —– Sensor Fusion —–
  Transmit local information to other sensors;
  Receive information from other sensors;
  Fuse all local information to get global estimate Σ;

  % — Optimization of sensor position —
  Assume other sensors do not move;
  Propagate P of other sensors by one time step;
    For all own allowable motion actions
        Propagate own error covariance matrix P;
        Fuse with propagated P of other sensors;
        Obtain cost function estimate;
    end
  Identify cost-minimizing action;

  % —– Update position —–
  Update position for next time step;
end
```

that lead to the best classification of a target as “teammate” or an “opponent.” Such friend-or-foe classification schemes arise often in team-based competitive game situations like RoboFlag [4].

Assume all sensing agents are equipped with identical sensors, which measure the position of the target according to (1). For simplicity’s sake, consider the case where the target’s class can be determined based on measurements of its position. Note that this case can be easily generalized to use other state information, such as target velocity, to classify the target.

Let the class distributions (6) be described by:

$$\bar{x}_1 = \begin{pmatrix} -15 \\ -12 \end{pmatrix}, \qquad \Delta_1 = \begin{pmatrix} 100 & 0 \\ 0 & 100 \end{pmatrix}$$
$$\bar{x}_2 = \begin{pmatrix} 15 \\ 12 \end{pmatrix}, \qquad \Delta_2 = \begin{pmatrix} 100 & 0 \\ 0 & 100 \end{pmatrix},$$

as illustrated in Fig. 1. The classification boundary is the x-axis, where the negative half of the simulation field is the “teammate” zone and the positive half-plane represents the “opponent” region.

Further, the target moves in a random walk (2) with process noise covariance

$$Q[k] = \begin{bmatrix} 0.01 & 0 \\ 0 & 0.01 \end{bmatrix}.$$

In addition, the scale parameter in the measurement noise covariance matrix described by (4) was set to $\alpha = 5$, and the quadratic range-dependent function

$$f(r_j) = a_2 r_j^2 + a_1 r_j + a_0,$$

with $a_0 = 0.3481,\ a_1 = -0.0250,$ and $a_2 = 0.0008$.

We utilize a modified version decentralized discrete gradient-search method proposed in [22], incorporating the inverse covariance Kalman filter for the fusion process, to generate the best motion trajectories of the sensors to achieve the best estimates of the target state. The resulting trajectories of the three sensors (numbered circles) observing the target (square) are depicted in Fig. 2. Note that the steady-state configuration corresponds to the results predicted from the analysis (see [22]) and also the intuitively optimal arrangement of sensors.

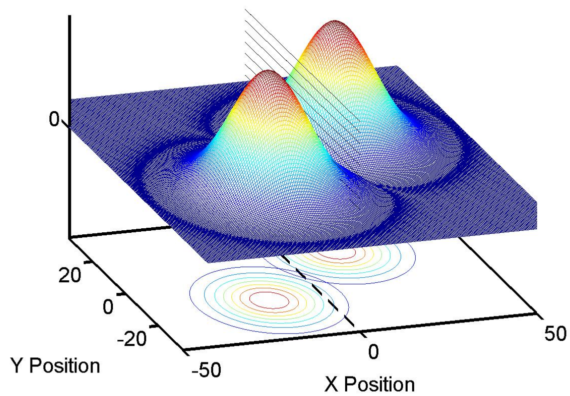

Fig. 1. Probability Distributions for Two Classes

Plotted in Fig. 3 and Fig. 4 are the resulting fused costs in estimation and classification, respectively. As the estimate error uncertainty Σ decreases as time evolves, we see that the correct classification probability p_c increases, verifying the inverse relationship between the two quantities discussed in the previous section.

The task of classifying the target can be carried out once a certain threshold for the probability of correct classification has been crossed [32]. Alternatively, once the steady-state of the system has been reached, the sensors may collectively categorize the target, since the sensors have achieved the configuration which yields the best classification.

VI. CONCLUSION

Increasingly, sensor networks will be applied to higher-level tasks such as classification and decision-making, rather than simple data gathering. This paper introduced new methods to improve the ability of a mobile sensor team to classify a target. The method was based on the use of a utility function that quantifies collective classification behavior, and a decentralized optimization algorithm to distribute the algorithm across vehicles.

For the particular case of Gaussian noise and class distributions, we proved that as the uncertainty in the collective estimate of the target state is reduced, the probability that the target will be correctly classified increases. Our discovery of this connection between estimation and classification for this class of categorization problems is this paper’s main contribution. For the cases to which this result applies, it significantly reduces the computational cost of determining

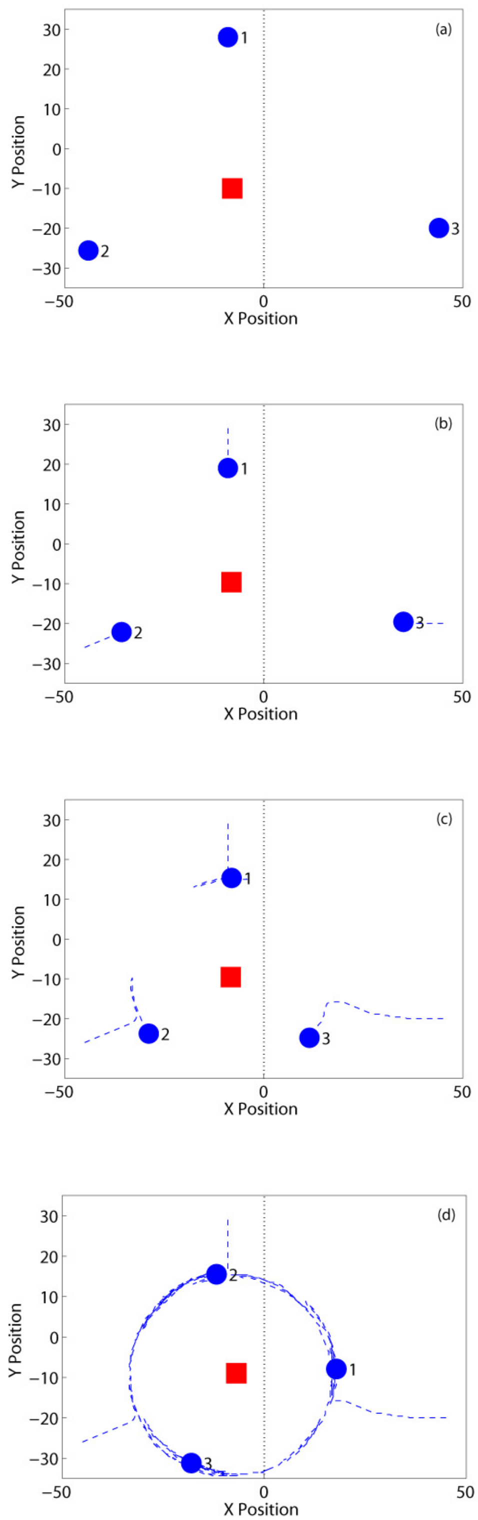

Fig. 2. Classification example: (a) Initial positions of three sensing agents (numbered circles) and randomly-walking target (square); (b) Optimal sensor motion for improved estimates and classification under the proposed discrete gradient-search algorithm; (c) Optimal configuration of sensors; (d) Evolution of sensor trajectories.

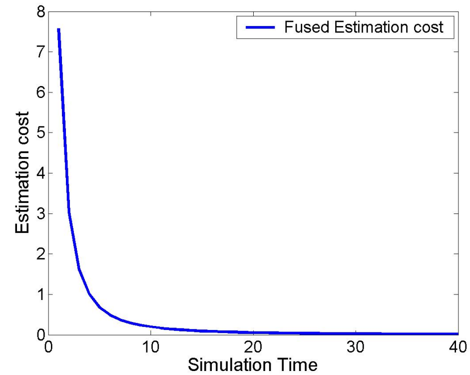

Fig. 3. Evolution of the Estimation Error Cost

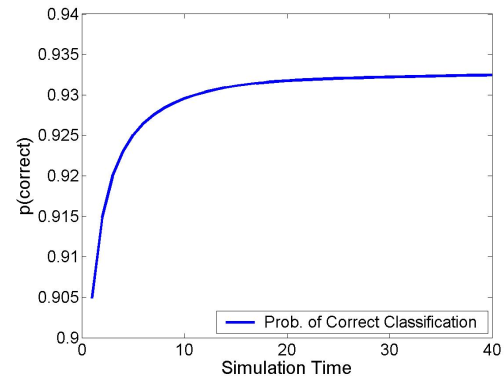

Fig. 4. Evolution of the Probability of Correct Classification

the best mobile sensor motions. Additionally, we showed that a distributed gradient search algorithm can provide the necessary infrastructure for motion planning.

Many interesting avenues for continued research exist in using mobile sensors to cooperatively classify objects. Generalization of the probability distributions to non-Gaussian, potentially multi-modal, distributions would allow for the results of this paper to address a larger class of classification problems, including the incorporation of multiple hypotheses testing. Similarly, scenarios where the parameters of the class distributions are not known beforehand can be determined real-time by means of methods such as sequential learning or particle filters [31], [37].

Additionally, further study of the effects such as sensor localization uncertainty and network communication constraints on the performance of the classification would be of interest. In particular, for purposes of practical implementation, an understanding of the effect of limiting the number of nodes each sensor communicates with is desirable.

In the longer-term, logical extension of this work include the the incorporation of multiple features (e.g. size, color, shape) in the classification process, as well as the geometry of the target. E.g., when the targets are not simple point objects, certain features may vary with the robot's relative position, thereby altering the optimal position and motion of the sensors. Relevant applications for the use of geometric targets might include, for example, human recognition and classification.

VII. Acknowledgements

The authors gratefully acknowledge the support of the National Science Foundation for this work. The authors would like to thank Vijay Gupta for useful discussions and input.

References

[1] S. Kristensen and H.I. Christensen. Decision-Theoretic Multisensor Planning and Integration for Mobile Robot Navigation. In *Proc. Intl. Conf. on Multisensor Fusion and Integration for Intelligent Systems*, pages 517–524, 1996.

[2] T.F. Bastos-Filho, M. Sarcinelli-Filho, and R.A.C. Freitas. A Multi-Sensorial Integration Scheme to Help Mobile Robot Navigation through Obstacle Recognition. In *Proc. Intl. Conf. on Emerging Technologies and Factory Automation*, volume 1, pages 549–558, 1999.

[3] Dirk Walther, Duane R.Edgington, and Christof Koch. Detection and Tracking of Objects in Underwater Video. In *Proc. Intl. Conf. on Computer Vision and Pattern Recognition*, volume I, pages 544–549, 2004.

[4] A. Chaudhry, R. D'Andrea, and M. Campbell. Roboflag - a framework for exploring control, planning, and human interface issues related to coordinating multiple robots in a realtime dynamic environment. In *Proc. of Intl. Conf. on Advanced Robotics*, 2003.

[5] R. Vidal, O. Shakernia, H. Kim, D. Shim, and S. Sastry. Probabilistic pursuit-evasion games: Theory, implementation and experimental evaluation. *IEEE Transactions on Robotics and Automation*, 18(5):662–669, Oct 2002.

[6] S. Roumeliotis and G. Bekey. Distributed multi-robot localization. *IEEE Transactions on Robotics and Automation*, 18(5):781–795, Oct 2002.

[7] S. Dhillon and K. Chakrabarty and S. Iyengar. Sensor Placement for Grid Coverage under Imprecise Detections. In *Proc. Intl. Conf. on Information Fusion*, pages 1571–1587, 2002.

[8] Dariusz Ucinski. *Measurement Optimization for Parameter Estimation in Distributed Systems*. Technical University Press, 1999.

[9] John R. Spletzer and Comillo J. Taylor. Dynamic Sensor Planning and Control for Optimally Tracking Targets. *International Journal of Robotics Research*, 22(1):7–20, January 2003.

[10] A. Makarenko, E. Nettleton, B. Grocholsky, S. Sukkarieh, and H. Durrant-Whyte. Building a decentralized active sensor networks. In *11th Intl Conf on Advanced Robotics*, Coimbra, Portugal, 2003.

[11] M. Sznaier and O.I. Camps. Control issues in active vision: open problems and some answers. In *Proc. IEEE Conf. on Decision and Control*, volume 3, pages 3238–3244, Dec 1998.

[12] P. Mowforth. Active sensing for mobile robots. In *Proc. Intl. Conf. on Control*, volume 2, pages 1141–1146, 1991.

[13] Lyudmila Mihaylova, Tine Lefebvre, Herman Bruyninckx, Klaas Gadeyne, , and Joris De Schutter. Active Sensing for Robotics - A Survey. In *Proc. of the 5th Intl. Conf. on Numerical Methods and Applications*, pages 316–324, Borovets, Bulgaria, August 2002.

[14] J. Kittler. Multi-Sensor Integration and Decision Level Fusion. In *IEE Workshop on Intelligent Sensor Processing*, Feb 2001.

[15] S. Shah, J.K. Aggarwal, J. Eledath, and J. Ghosh. Multisensor Integration for Scene Classification: An Experiment in Human Form Detection. In *Proc. Intl. Conf. on Image Processing*, volume 2, pages 199–202, 1997.

[16] Xiaoling Wang, Hairong Qi, and S.S. Iyengar. Collaborative Multi-Modality Target Classification in Distributed Sensor Networks. In *Proc. Intl. Conf. on Information Fusion*, volume 1, pages 285–290, 2002.

[17] Ralph D. Hippenstiel. *Detection Theory: Applications and Digitial Signal Processing*. CRC Press, 2002. §4.9.

[18] R. Viswanathan and P.K. Varshney. Distributed detection with multiple sensors i. fundamentals. *Proceedings of the IEEE*, 85(1):54–63, Jan 1997.

[19] R. S. Blum, S. A. Kassam, and H. V. Poor. Distributed detection with multiple sensors ii. advanced topics. *Proceedings of the IEEE*, 85(1):64–79, Jan 1997.

[20] S. Alhakeem and P.K. Varshney. A unified approach to the design of decentralized detection systems. *IEEE Transactions on Aerospace and Electronic Systems*, 31(1):9–20, Jan 1995.

[21] I.Y. Hoballah and P.K. Varshney. Distributed bayesian signal detection. *IEEE Transactions on Information Theory*, 35(5):995–1000, Sep 1989.

[22] Timothy H. Chung, Vijay Gupta, Joel W. Burdick, and Richard M. Murray. On a Decentralized Active Sensing Strategy using Mobile Sensor Platforms in a Network. In *Proc. of the IEEE Conf. on Decision and Control*, Paradise Island, Bahamas, Dec 2004.

[23] K.V. Ramachandra. *Kalman Filtering Techniques for Radar Tracking*. Marcel Dekker, Inc., New York, NY, 2000.

[24] K. Umeda, J. Ota, and H. Kimura. Fusion of multiple ultrasonic sensor data and imagery data for measuring moving obstacle's motion. In *Proc. of Intl. Conf. on Multisensor Fusion and Integration for Intelligent Systems*, pages 742–748, December 1996.

[25] Y. Bar-Shalom. On the Track-to-Track Correlation Problem. *IEEE Trans. on Automatic Control*, 26(2):571–572, April 1981.

[26] Y. Bar-Shalom and L. Campo. The effect of the common process noise on the two-sensor fused-track covariance. *IEEE Trans. on Aerospace and Electronic Systems*, 22(6):803–805, 1986.

[27] P.S. Maybeck. *Stochastic Models, Estimation and Control*, volume 1. Academic Press, 1979.

[28] B.S.Y. Rao, H.F. Durrant-Whyte, and J.A. Sheen. A fully decentralized multi-sensor system for tracking and surveillance. *Intl. Journal of Robotics Research*, 12:20–44, 1993.

[29] E. Drakopoulos C.-C. Lee. Optimum multisensor fusion of correlated local decisions. *IEEE Transactions on Aerospace and Electronic Systems*, 27(4):593–606, Jul 1991.

[30] Clayton Scott and Rob Nowak. Minimum Probability of Error Decision Rule. Online. http://cnx.rice.edu/content/m11534/latest/.

[31] Richard Duda, Peter Hart, and David Stork. *Pattern Classification*. John Wiley & Sons, Inc., second edition, 2001.

[32] Keinosuke Fukunaga. *Introduction to Statistical Pattern Recognition*. Academic Press, Inc., 1972. § 3.2.

[33] Jan R. Magnus and Heinz Neudecker. *Matrix Differential Calculus with Applications in Statistics and Econometrics*. John Wiley and Sons, 1st edition, 1988.

[34] Mike Brookes. Matrix Reference Manual. Online. http://www.ee.ic.ac.uk/hp/staff/dmb/matrix/calculus.html.

[35] B. Grocholsky, A. Makarenko, T. Kaupp, and H. Durrant-Whyte. Scalable control of decentralised sensor platforms. In *Proc. Of the 2nd International Workshop on Information Processing in Sensor Networks*, Palo Alto, CA, 2003.

[36] D. P. Bertsekas and J. N. Tsitsiklis. *Parallel and distributed computation : numerical methods*. Prentice Hall, Englewood Cliffs, NJ, 1989.

[37] K. Murphy and S. Russell. *Sequential Monte Carlo Methods in Practice*, chapter Rao-Blackwellised Particle Filtering for Dynamic Bayesian Networks. Springer-Verlag, 2001.