2011 American Control Conference
on O'Farrell Street, San Francisco, CA, USA
June 29 - July 01, 2011

Information-Based Adaptive Sensor Management for Sensor Networks

Karen L. Jenkins and David A. Castañón

Abstract— **Consider the problem of controlling a network of surveillance sensors that are capable of selecting which areas to observe and which modes to observe these areas. In this paper, we study the problem of controlling the observations of these sensors adaptively in order to classify accurately a collection of objects using information on their observed features. Our proposed approach is modeling objects as templates of 3-D features, and modeling sensors as observing features of individual objects, subject to degradation by noise, obscuration, missed detections and background clutter. We exploit a statistical framework based on random sets similar to those used in multi-target tracking to model the statistical relationship between observed features and object types to compute information-theoretic estimates of the probability of error in classification. We present a novel approach for computation of these distances between distributions of random sets using k-best assignment algorithms. These estimates are combined with real time information to generate predictions of the information value of individual measurements for sensor management. Using these predictions, we develop assignment algorithms to compute sensor management strategies to minimize this bound. The resulting sensor management algorithms are capable of solving problems involving a large numbers of objects in real-time. We show simulations of the resulting algorithms for classifying 3-dimensional objects from 2-dimensional noisy projections that illustrate how the algorithms select complementary views to overcome obscuration and provide accurate classification. Our results show that our information-based sensor management algorithms achieve comparable classification accuracy to adaptive simulation-based approaches that evaluate the value of information, while requiring nearly five orders of magnitude less computation.**

I. INTRODUCTION

Intelligent sensors of multiple modalities are increasingly available in diverse applications, ranging from building security and defense to transportation and medicine. In turn, this has created a need for automated processing and reduction of sensor information, and the opportunity for controlling the nature of information collected so that the resulting systems can accomplish their task faster. One of the key tasks that these systems perform is is object recognition, such as the identification of explosives inside luggage, determining the type of vehicle in surveillance, or identifying the persons in a given location. These are complex tasks due to the large number of variations in ambient illumination as well as object appearance and pose in the sensed information. Robust object recognition is often based on extracting features of objects provide discrimination information, and which can be reliably observed under the range of sensed conditions. The set of observable features of an object by a specific sensor depends on factors such as environmental conditions, relative sensor-object position and orientation, positions of other objects and sensing modality. Thus, accurate recognition often requires fusion of information from multiple sensors using different views to obtain a sufficient collection of discriminating features on each object.

In this paper, we consider the problem of adaptive management of a team of distributed sensors in order to classify accurately a set of spatially distributed objects by observing errored subsets of their features. These sensors can be oriented dynamically to collect information on selected objects from different points of view in order to complement best the available information to achieve rapid classification.

Sensor management has received increased attention in recent years, as summarized in the recent book [1]. Most of this work focuses on finding and tracking objects, and not on identifying them [2]–[5]. A common approach in sensor management is to use information theory metrics as the basis for evaluating alternative sensing actions [6]–[11]. Alternative approaches use expected task performance such as as tracking error or Bayesian classification costs, leading to stochastic control formulations and corresponding algorithms [12]–[16]. The work in [17] compares task-driven sensor management schemes with information-driven schemes, and determines that the performance difference between the two classes of approaches was small.

The problem of sensor management for classification objectives was studied in [18] using information theoretic criteria. Alternative approaches based on stochastic control techniques and partially observed Markov decision processes were developed in [12], [15], [19]. In these papers, the sensor model for classification assumes that sensors observe a conditionally independent estimate of object type, which is a highly unrealistic assumption. In actuality, processing of information generated by sensors with a similar point of view will result in highly correlated tentative classification decisions because they will be based on the same set of features. Instead, a more accurate model would represent sensor measurements as providing conditionally independent information about object features, and determine how observation of additional features yields improved object classification decisions.

In our previous work [20], we proposed a model for feature-based sensor management representing objecdts as spatially related collections of features, characterized by object type and pose. Such models have been proposed for object recognition by several authors [21]–[23]. In this model, sensors obtain noisy observations of subsets of object

This work was supported by AFOSR grants FA9550-07-1-0361 and by ODDR&E MURI Grants FA9550-06-1-0324 and FA9550-07-1-0528

D. Castañón is with the Electrical and Computer Engineering Department at Boston University. K. Jenkins is with Livevol, Inc

features, where the observed subset of features depend on object type, and relative sensor/object pose, and includes missed features due to obscuration as well as missed detections, and clutter features, and are modeled as random point sets [24] derived from the object models. The results of [20] used as the objective for collected information the reduction in the expected probability of classification error. Real-time evaluation of these metrics requires the use of time-consuming simulations of future measurements. [20] developed a novel approach that combines off-line computation of a priori value of measurements using Bhattacharyya distances between models of random sets of point features, and real-time estimates of object type and pose generated from past measurements, to generate a prediction of the value of information that a sensor could collect on an object. This value was used in assignment algorithms to compute sensor management strategies to maximize the expected information collected.

In this paper, we extend the results of [20] to incorporate more accurate estimates of the reduction in probability of classification error. In particular, we develop estimates based on Chernoff coefficients instead of Bhattacharyya distances, which can lead to the use of more accurate Chernoff bounds on the expected probability of error between two random sets. We derive a new set of bounds on the expected reduction in probability of error based on the Chernoff coefficients, and integrate them into a real-time adaptive sensor management algorithm that scales well to scenarios with large numbers of objects. We also describe an efficient way for computation of Chernoff coefficients that avoids a combinatorial enumeration of hypotheses using k-best data association algorithms and importance sampling techniques.

We evaluate the proposed sensor management algorithms using data generated from synthetic 3-D models of object classes. We also evaluate the algorithms on synthetic aperture radar data generated under DARPA's MSTAR program and simulate sensors as extracting features from 2-D projections of these models. We compare the performance of our real-time sensor management algorithm with other information-theory approaches that use measurement simulations. Our real-time algorithms achieve comparable classification accuracy, while requiring nearly three orders of magnitude less computation. Our results establish the feasibility of a practical, scalable and accurate approach for the real-time management of a team of sensors.

An overview of the remainder of this paper is as follows: Section II, presents an overview of the mathematical formulation of feature-based sensor management using random set observations introduced in [20]. Section III describes two approaches for adaptive sensor management based on information-theory criteria, including our solution for combining off-line bounds with on-line estimates. Section IV discusses simulation results that compare the two algorithms in terms of classification performance and computation time. Section V discusses our results and presents directions for future work.

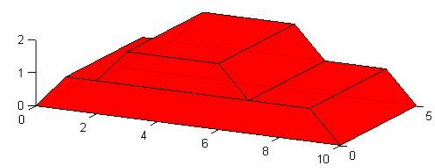

Fig. 1: Example of a three-dimensional model

II. Problem Formulation

Assume that there are M known locations that contain objects, with centers at locations $z_m \in \mathbb{R}^3$. Each location m contains an object of type $c_m \in \{1, \ldots, k\}$, that is oriented with a pose consisting of azimuth angle θ_m^{az} and elevation (or similarly, inclination) angle θ_m^{el}. We assume that the center locations z_m are known, and that objects are stationary. The unknown state of an object at location m is

$$x_m = (c_m, \theta_m^{az}, \theta_m^{el})$$

and the overall unknown system state is

$$X = (x_1, \ldots, x_M)$$

Assume for the purposes of inferencing that the set of possible azimuth and elevation orientations is finite, and that the states are independent random variables across location, with marginal prior distributions specified by $p_m(x_m), m = 1, \ldots, M$.

Associated with each object type $k = 1, \ldots, K$ is a collection of features, described by their spatial locations in a relative coordinate frame centered at the object centroid. Extending our formulation to adding additional information to features such as type is straightforward (e.g. [25]). Examples of such features can include extreme points of models such as corners, as illustrated in Fig. 1. Denote the relative locations of features for object type k as $M^k = \{M_1^k, M_2^k, ..., M_{n_k}^k\}$, where $M_i^k \in \mathbb{R}^3$. Given an object state x_m at known location m, the model M^{c_m} specifies unambiguously the number of features and their 3-D positions through application of the appropriate rotation matrices $R_z(\theta_m^{az}), R_y(\theta_m^{el})$, where $R_z(\theta)$ represent a clockwise rotation through an angle θ about the z axis and and $R_y(\theta)$ is a clockwise rotation through an angle θ about the y axis.

The true feature locations F^m of object m with state x_m are generated by translating and rotating model features from class c_m as:

$$F_i^m = z_o^m + R_z(\theta_m^{az}) * R_y(\theta_m^{el}) * M_i^{c_m}, \quad i = 1, \ldots, n_{c_m} \quad (1)$$

In this manner, three-dimensional objects can be generated as collections of points, spatially distributed about the known location of the object centroid.

Sensors collect measurements on object locations to estimate the states X. A single sensor observes only one location at a time and generates a set of projected noisy feature positions for that location, that misses obscured features, and includes random missed feature detections as clutter features. The statistical model for the observation is a random finite point set [24]. The measurement model is described constructively as follows: Denote the random set corresponding to measurement of location m by sensor j as Y^{mj}.

For each feature location M_i^k for object type k, we assume that there is a visibility map $I_i^k(\theta, \delta)$ which is a function of relative azimuth θ and elevation angle δ to a sensor location, which takes value 1 whenever this feature is visible from a relative direction corresponding to (θ, δ), and 0 otherwise because of self-obscuration. Without loss of generality, assume that the sensor is in the far field, so this obscuration depends only on the relative angle and not on the range to the sensor. Thus, given sensor location s_j and object state x_m at location m, the relative azimuth and orientation angles can be computed, and the subset of visible features for object type c_m are readily identified.

The sensor's location s_j defines a line-of-sight vector to the location z_m, and a projection plane perpendicular to this vector. The sensor measures the projections of the features which are visible according to the visibility map to the measurement plane. If feature i is visible, then its projected location is

$$Y_i^{mj} = (F_i^m - s_j) - \frac{(z_m - s_j)(F_i^m - s_j)^T(z_m - s_j)}{\| z_m - s_j \|^2}$$

These projected locations are represented as two-dimensional locations in a sensor-centered coordinate frame. Misdetection of visible features is modeled by a Bernoulli detection process which is independent across features, measurements and sensors, with probability of detection p_D. Let $d_i \in \{0, 1\}$ denote the indicator that feature i for object of type c_m is both visible and detected when viewed from sensor j, for $i = 1, \ldots, n_{c_m}$.

Detected projected locations as measured imprecisely, with Gaussian additive errors, independent across features and measurements, yielding noisy measurements for $\tilde{Y}_i^{mj} = Y_i^{mj} + v_{mji}$ where $v_{mji} \sim N(0, \sigma_{mj}^2)$ has covariance that depends on the sensor-object location geometry and is a 2-dimensional random error in the projection measurement plane of sensor j. The collection of these detected noisy features for sensor j of location m are denoted as Y_d^{mj}, defined as

$$Y_d^{mj} = \{\tilde{Y}_i^{mj} | d_i = 1, i = 1, \ldots, n_{c_m}\}$$

Clutter features in the measurement sets are modeled by a homogeneous spatial Poisson point process with total intensity λ distributed uniformly over the sensor's field of view, independent across sensors and measurement times. Denote this random set of points Y_c^{mj}. Then, the observation of location m by sensor j is defined as $Y^{mj} = Y_d^{mj} \cup Y_c^{mj}$ The resulting observation Y^{mj} is a random set of projected locations [24], [26], [27].

The above model defines the likelihood $p(Y^{mj}|x_m)$. Let $I_m(t)$ denote all the random set observations of location m collected for all times prior to and including time t, and the sensor actions $u_j(m) = s$ that collected the observations. The sufficient statistic or information state summarizing the observations at location m is the conditional probability $p(x_m|I_m(t))$. Let $I(t) = \cup_{m=1}^N I_m(t)$. Given the independence assumptions, it is straightforward to show that the sufficient statistic for the overall state X at time t is given by

$$p(X|I(t)) \equiv \Pi(t+1) = \prod_{m=1}^{M} p(x_m|I_m(t))$$

The evolution of the information state is governed by Bayes' rule, given the new observation $Y^{mj}(t)$ at time t, as

$$p(x_m|I_m(t)) = \frac{p(Y^{mj}(t)|x_m)p(x_m|I_m(t-1))}{\sum_{x'_m} p(Y^{mj}(t)|x'_m)p(x'_m|I_m(t-1))}$$

This inference involves averaging over possible data associations from the random set locations to the features of object m under state x_m [26], [27].

With the above setup, adaptive sensor management problem selects sensing actions $u_j(t) \in \{1, \ldots, M\}$ for each sensor, sensors $j = 1, \ldots, J$ identifying the location at for each time t that each sensor will observe, based on knowledge of the information state $\Pi(t)$. Sensor actions are constrained so that a sensor can only observe locations which are within its field of regard. In this paper, we pursue the approach of [20] to use surrogate performance measures based on information theory as the basis for selection of sensor actions, as discussed in the next section.

III. Information-Based Adaptive Sensor Management

The adaptive sensor management schemes we consider are based on generating stage-by-stage optimal policies. At time t, the sensor management system must select a sensing action $u_i(t)$ for each sensor i, that are functions of the information state $\Pi(t)$. We define utility functions $R(j, m)$ associated with the value of measuring location m with sensor j at time t, which depend on the current information state $\Pi(t)$ and the random set observation models described in Section II. Given these utility functions, computed for each sensor-location pair, the actions at time t will be selected to maximize the total collected utility by solving an inequality-constrained linear assignment problem of the form

$$\begin{aligned} \max \sum_{j=1}^{N} \sum_{m=1}^{M} & R_{j,m} a_{j,m} \quad \text{subject to} \qquad (2) \\ \sum_{m=1}^{M} & a_{j,m} = 1, j = 1, \ldots, N \\ \sum_{j=1}^{N} & a_{j,m} \leq 1, m = 1, \ldots, M \\ & a_{j,m} \in \{0, 1\}, j = 1, \ldots N, m = 1, \ldots, M \end{aligned}$$

where $a_{j,m} = 1$ if sensor i is assigned to measure location m at time t, and 0 otherwise. The solution of this problem translates to decision actions at time t, as $u_j(t) = m$ if $a_{j,m} = 1$. Note that this formulation restricts the actions of sensors so that no two sensors can observe the same location at the same time. This restriction could be relaxed, but the resulting problem would require a combinatorial enumeration of the predicted performance of combinations of sensors for each location, significantly increasing the computation. Given the presence of multiple observation times, the loss in performance by this restriction is minimal, as sensors that should observe a location that is also currently observed by other sensors can simply delay their observation to another time.

The main issue in our approach is the choice of surrogate utility function that decomposes additively over sensor actions, as required by the optimization approach. One alternative is the discrimination gain, or equivalently, the expected reduction in sample conditional entropy of the underlying conditional probability distributions of the object states at each time [4], [10], [18]. The sample conditional entropy of the sufficient statistic $\Pi(t)$ as

$$\begin{aligned} H[\Pi(t))] &= -\sum_{m=1}^{M}\sum_{x_m} p(x_m|I_m(t-1)) \ \log p(x_m|I_m(t-1) \\ &= \sum_{m=1}^{M} H_m[p(x_m|I_m(t))] \end{aligned} \quad (3)$$

Note that we are not averaging over $I(t)$ as would be required for the standard definition of conditional entropy, since $H[\Pi(t)]$ is adapted to the sufficient statistic $\Pi(t)$.

As in [18], the expected discrimination gain at time t of a collection of sensing actions $U(t)$ is the sum of the expected sample conditional entropy reduction for each action. For action $u_j(t) = m$, the discrimination gain from location m is

$$\begin{aligned} D_{j,m}(p(x_m|I_m(t-1))) &= H_m[p(x_m|I_m(t-1))] \\ &- E_{Y^{mj}(t)}\left[H_m[p(x_m|I_m(t-1), u_j(t)=m, Y^{mj}(t))]\right] \end{aligned} \quad (4)$$

For sensor management, we define the stage t utility functions $R(j,m) = D_{j,m}(p(x_m|I_m(t-1)))$. Note that the overall cost (3) is additive in the individual location discrimination gains. Direct computation of the expectation in (4) using our random set model is difficult, so the functions $R(j,m)$ are typically estimated through simulation of the random sets $Y^{mj}(t)$ for each location m and sensor j.

A different metric for sensor management is reduction in the expected probability of object classification error based on maximum a posteriori classification for each location. For location m, given the information state $\Pi(t)$ at state t, the probability of making a classification error using the maximum a posteriori classifier before collecting additional measurements is

$$R_m(t) = 1 - \max_{x_m} p(x_m|I_m(t-1))$$

Similarly, after sensor j collects a measurement $Y^{mj}(t)$, the posterior probability of classification error is

$$R_m(t, Y^{mj}(t)) = 1 - \max_{x_m} p(x_m|I_m(t-1), Y^{mj}(t))$$

The expected reduction in probability of error from measuring location m with sensor j is

$$\begin{aligned} R_{j,m} &= E_{Y^{mj}}[\max_{x_m} p(x_m|I_m(t-1), Y^{mj}(t))] \\ &- \max_{x_m} p(x_m|I_m(t-1)) \end{aligned} \quad (5)$$

By summing over the set of possible sensing actions, the utility functions $R_{j,m}$ in (5) maximize the expected reduction in the number of classification errors over the set of locations. Again, direct computation of this objective requires averaging over the random sets $Y^{mj}(t)$ for each sensor-location pair using Monte Carlo techniques.

The approach proposed in [20] is to use bounds on the probability of error to estimate the expected reduction above. In particular, these bounds do not require on-line simulation of random sets, and use quantities based on information-theoretic distances that can be computed off-line. We extend the results of [20] to use bounds on the probabilities of error in confusing potential states x_m, x'_m in terms of Chernoff coefficients [28], as described next.

Assume we have two distributions for the same random variable Y, given by $p(Y|x)$ and $p(Y|x')$. The Chernoff coefficient for parameter $0 < \alpha < 1$ between these distributions is defined as

$$\rho^{\alpha}_{x,x'} = \int_Y p(Y|x)^{\alpha} p(Y|x')^{1-\alpha} dY \quad (6)$$

The special symmetric case for $\alpha = 1/2$ is the Bhattacharyya coefficient, which was used in [20], and is denoted without superscript as

$$\rho_{x,x'} = \int_Y \sqrt{p(Y|x)p(Y|x')} dY \quad (7)$$

The Chernoff coefficients in (6) depend on the distribution of the random set measurements conditioned on the discrete states x_m, x'_m, and not on the real-time information $I_m(t)$. Thus, these coefficients can be computed off-line, which we will exploit for fast real-time sensor management.

The Chernoff coefficient provides a bound on the probability of classification error in making a decision between a pair of hypotheses x_m and x'_m. If the prior probabilities on the classes are $p(x_m)$, $p(x'_m)$, and the maximum a posteriori (MAP) rule is used to select one of the two values x_m or x'_m, the average probability of error can be bounded by

$$\begin{aligned} P_{err}(x_m, x'_m) &= E_Y[\min(\frac{p(Y|x_m)P(x_m)}{P(Y)}, \frac{p(Y|x'_m)P(x'_m)}{P(Y)})] \\ &= \int_Y \min(p(x_m)p(Y|x_m), p(x'_m)p(Y|x'm))dY \\ &\le p(x_m)^{\alpha} p(x'_m)^{1-\alpha} \rho^{\alpha}_{x_m,x'_m} \end{aligned} \quad (8)$$

Note that the above bound combines the available information given by the prior probabilities, together with the off-line computation of the Chernoff coefficient. We can compute this bound for multiple values of the Chernoff coefficients α, and choose the minimum as the best bound on the probability of error.

The above bound on the average probability of error in selecting between two hypotheses can be extended to compute a bound for the probability of error when there are many discrete states $x_m \in \{X_m^1, \ldots, X_m^N\}$ assuming that the estimate is selected by the MAP rule.

Theorem 1: Assume there are N possible hypothesis x_m for location m, with prior probabilities $p(x_m)$, and that we observe a measurement Y with known conditional probabilities $P(Y|x_m), x_m = 1, \ldots, N$. Then, for any $\alpha \in (0, 1)$, the expected probability of error for the maximum likelihood decision rule, P_{err} is bounded by

$$P_{err} \leq \sum_{j=1}^{N} \sum_{j'=j+1}^{N} p(x_m = j)^\alpha p(x_m = j')^{1-\alpha} \rho_{j,j'}^\alpha$$

The proof appears in the appendix.

Theorem 1 gives a bound on the probability of classification error for a multi-hypothesis problem, in terms of the Chernoff coefficients between pairs of measurement likelihood models. For sensor management, we want to use this bound to approximate the expected reduction in probability of error in (5). The expected reduction $R(j, m)$ is

$$R(j, m) = R_m(t) - R_m(t, Y^{mj}(t))$$

The expected probability of error after collecting a measurement from sensor j to location m, denoted by $R_m(t, Y^{mj}(t))$ is bounded above from Theorem 1, by

$$R_m(t, Y^{mj}) \leq \sum_{k=1}^{N} \sum_{k'=k+1}^{N} p(x_m = k|I_m(t-1))^\alpha \\ p(x_m = k'|I_m(t-1))^{1-\alpha} \rho_{k,k'}^\alpha$$

We use this bound to define the utility function

$$R(j, m) = R_m(t) - \sum_{k=1}^{N} \sum_{k'=k+1}^{N} p(x_m = k|I_m(t-1))^\alpha \\ p(x_m = k'|I_m(t-1))^{1-\alpha} \rho_{k,k'}^\alpha \quad (9)$$

which is an estimate of the reduction in expected probability of error in estimating the state x_m given a new random set measurement $Y^{mj}(t)$ by sensor j. Note in particular that this utility is easily computed given the information states $\Pi(t)$ and the precomputed Chernoff coefficients.

A. Computation of Chernoff Coefficients

Computing Chernoff coefficients requires integration over random sets, a time-consuming task to perform exactly. Instead of (6), we propose to compute the Chernoff coefficient between two possible states x_m, x'_m using importance sampling, as:

$$\rho_{x_m,x'_m}^\alpha = \int_Y \Big[\frac{p(Y^{mj}|x_m)}{p(Y^{mj}|x'_m)}\Big]^\alpha p(Y^{mj}|x'_m) dY^{mj} \quad (10)$$

We use samples $\tilde{Y}_\ell^{mj}, \ell = 1, \ldots, L$ generated from the distribution $p(Y^{mj}|x'_m)$, along with the likelihood models of Section II to evaluate ρ_{x_m,x'_m} using Monte Carlo techniques to yield the result:

Lemma 1: The estimator

$$\hat{\rho}_{x_m,x'_m}^\alpha = \sum_{\ell=1}^{L} \Big[\frac{p(\tilde{Y}_\ell^{mj}|x_m)}{p(\tilde{Y}_\ell^{mj}|x'_m)}\Big]^\alpha$$

is an unbiased estimator of ρ_{x_m,x'_m}^α.
The proof is straightforward from the properties of independent samples.

The computation of Chernoff coefficients from Monte Carlo samples $\tilde{Y}_\ell^{mj}$ involves the computation of likelihoods $p(Y_\ell^{mj}|x_m)$. Note that Y_ℓ^{mj} is a random set, therefore computation of this likelihood requires summation over all possible data associations of points in Y_ℓ^{mj} with the predicted visible features for sensor j from the models at location m with state x_m [24]. From our model in Section II, these likelihoods become:

$$p(Y_\ell^{mj}|x_m) = \sum_M p(Y_\ell^{mj}|x_m, M) p(M|x_m) \quad (11)$$

where M is a data association hypothesis that explains uniquely the correspondence of points in Y_ℓ^{mj} to clutter points and visible features for sensor j for an object in state x_m at location m. The details of this computation can be found in [23], and other similar references.

The number of terms in the summation in (11) grows exponentially with the number of features. Furthermore, most of the terms $p(Y^{mj}|x_m, M)$ tend to be vanishingly small, so the sum is dominated by few terms. A common approach in multiobject tracking is to compute only the largest term in the sum, using a generalized likelihood method, for which efficient solutions are available [23]. The idea is to approximate

$$M^{max} = argmax_M p(Y^{mj_\ell}|x_m, M) p(M|x_m) \\ p(Y_\ell^{mj}|x_m) \approx p(Y_\ell^{mj}|x_m, M^{max})$$

However, this generalized likelihood method can be very inaccurate. Instead, we use a K-best approximation, where we compute the probabilities of the K most-likely matchings from the simulated measurements to the model for state x_m to approximate the sum in (11). The best K matches can be found efficiently in time that scales polynomially with the number of points in the random sets using variations of assignment algorithms [29].

IV. Experiments

The experiments in this section consist of scenarios with 7 sensors and 20 objects corresponding to two different classes of cars, shown in Figure 2. Sensors correspond to cameras that image objects that are within a sensing radius, but are constrained to be capable of viewing just one object at a time. In our simulations, we start with the a priori information that each car type is equally likely. Each object is presented at an unknown azimuth angle which is uniformly distributed between 0 and 2π, with respect to the system reference axis. We model a non-flat terrain by a distribution on object elevation angle, which is uniformly distributed on $-\pi/6$ to $+\pi/6$. Object templates are three-dimensional models of cars, with features defined as the locations of corner points. Objects are opaque, so features can be obscured based on observation orientation. Visible features are projected noisily onto the sensor viewing plane. Measurements of object features are conditionally independent and Gaussian distributed with a measurement noise variance of $\sigma_m = 0.5$. On average, each sensor has 4 objects within its sensing radius. Features are detected with a probability $p_D = 0.9$ and false alarms are generated at a rate of $\lambda = 1$ uniformly on the measured image.

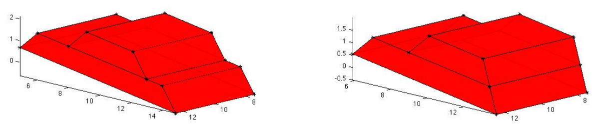

Fig. 2: Two object models at the same pose

Figure 3 shows the performance of our real-time approach to sensor management, in terms of average probability of error versus number of decision times. The results compare an on-line discrimination gain maximization algorithm with our real-time algorithms using two different approaches at computing off-line bounds: the Bhattacharyya indices proposed in [20], and the minimum bound obtained by selecting the optimal index for Chernoff coefficient α in the bound of (6). The performance curves represent Monte Carlo averages in classification error over 100 simulations, where each simulation varied the location and orientation of the objects present as well as the measurement errors. Note that the performances of the algorithm using Bhattacharyya indices and Chernoff indices are nearly identical, which suggests that the choice of best sensor/object geometry is not very sensitive to the quality of the bound used. The main point to note is that the performance achieved by our on-line/off-line approach is comparable to that achieved by the discrimination gain algorithm, but reducing computation in our MATLAB implementations by three orders of magnitude.

V. Conclusions

In this paper, we presented algorithms for computing adaptive sensor management strategies for a group of sensors seeking to classify a set of spatially distributed objects. Our approach used a novel model based on random set observations in terms of collections of extracted features from measured imagery, which are angle-dependent.

We presented a novel approach that is scalable and suitable for large numbers of objects based on off-line computation of Chernoff coefficients between measurement likelihoods under pairs of hypotheses, together with real-time estimates of conditional probabilities of the state hypotheses given past

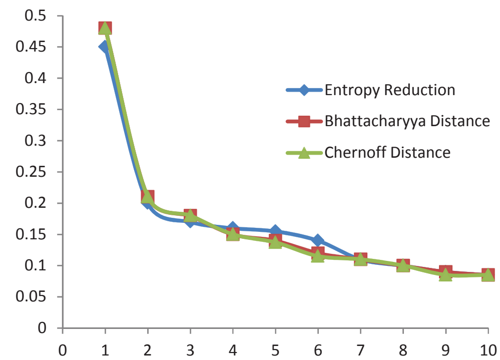

Fig. 3: Comparison between Discrimination Gain and our algorithms using Bhattacharyya and Chernoff coefficients

measurements. We also proposed a new algorithm for computation of the off-line Chernoff coefficients using a k-best assignment algorithm. Our simulations show that our off-line/on-line algorithms achieve performance comparable to that of algorithms that use real-time Monte Carlo predictions of performance, while reducing computation by three orders of magnitude.

There are many directions in which this work can be extended. The results can be readily extended to moving objects and moving sensors where sensor-object geometry would be changing dynamically. In addition, our simple feature model can be extended to add categorical feature types in addition to feature locations. Another extension would allow model uncertainty, so that the templates could have variations over where features are located. These directions remain as subjects for future investigation.

Appendix

Proof of Theorem 1:
Given hypotheses $x \in \{1, \ldots, N\}$ with prior probabilities $p(x)$, and observations Y with conditional distributions $p(Y|x)$, the maximum aposteriori estimate

$$\hat{x} = \max_{x \in \{1,\ldots,N\}} \left\{ \frac{p(Y|x)p(x)}{p(Y)} \right\}$$

and the probability of error in making this decision is:

$$P_{err} = \sum_{j=1}^{N} p(j)p(\hat{x} \neq j|j)$$

Define $A_{j'}$ to denote the event that the classification decision $\hat{x}_m = j'$. This corresponds to a collection of observations Y for which $p(x = j'|Y) \geq p(x = j|Y)$ for all $j \neq j'$. Then,

$$P_{err} = \sum_{j=1}^{N} p(j)p(\bigcup_{j' \neq j} A_{j'}|x = j)$$

We can assume the events $A_{j'}$ are disjoint by using a tie-breaking rule for the maximum a posteriori estimate. Thus,

$$p(\bigcup_{j' \neq j} A_{j'}|x = j) = \sum_{j'=1, j' \neq j}^{N} p(A_{j'}|x = j)$$

so we can write

$$P_{err} = \sum_{j=1}^{K} p(j) \sum_{j'=1, j' \neq j}^{N} p(A_{j'}|x=j)$$
$$= \sum_{j=1}^{N} p(j) \sum_{j'=1, j' \neq j}^{N} p(\hat{x}=j'|x=j)$$

from the definition of $A_{j'}$. We can rearrange symmetric terms in the summation to obtain:

$$P_{err} = \sum_{j=1}^{N} \sum_{j'=j+1}^{N} [p(\hat{x}=j'|x=j)p(x=j)+ p(\hat{x}=j|x=j')p(x=j')]$$

We define Y^j to denote those values of observation Y for which event A_j is true, that is, the classification decision is as $\hat{x}=j$. Then,

$$P_{err} = \sum_{j=1}^{N} \sum_{j'=j+1}^{N} \int_{Y^{j'}} p(j)p(Y|x=j)dY + \int_{Y^j} p(xj')p(Y|x=j')dY$$
$$= \sum_{j=1}^{N} \sum_{j'=j+1}^{N} \int_{Y^j \cup Y^{j'}} \min \left\{ p(j)p(Y|j), p(j')p(Y|j') \right\} dY$$
$$\leq \sum_{j=1}^{N} \sum_{j'=j+1}^{N} \int_{Y} \min \left\{ p(j)p(Y|j), p(j')p(Y|j') \right\} dY$$
$$\leq \sum_{j=1}^{N} \sum_{j'=j+1}^{N} \int_{Y} p(j)^{\alpha} p(Y|j)^{\alpha} p(j')^{1-\alpha} p(Y|j')^{1-\alpha} dY$$

where the first equality follows because $p(j)p(Y|j) \leq p(j')p(Y|j')$ under $Y \in Y^{j'}$ and the last inequality holds for $0 < \alpha < 1$. This gives us the result:

$$P_{err} \leq \sum_{j=1}^{N} \sum_{j'=j+1}^{N} p(j)^{\alpha} p(j')^{1-\alpha} \int_{Y} p(Y|j)^{\alpha} p(Y|j')^{1-\alpha} dY$$
$$= \sum_{j=1}^{N} \sum_{j'=j+1}^{N} p(j)^{\alpha} p(j')^{1-\alpha} \rho_{j,j'}^{\alpha}$$

References

[1] A. Hero, D. A. Castañón, D. Cochran, and K. Kastella, *Foundations and Applications of Sensor Management*. Springer, 2008.

[2] Y. Li, L. W. Krakow, E. K. P. Chong, and K. N. Groom, "Dynamic sensor management for multisensor multitarget tracking," in *Proceedings of the 40th Annual Conference on Information Sciences and Systems*, 2006, pp. 1397–1402.

[3] H. Wang, K. Yao, and D. Estrin, "Information-theoretic approaches for sensor selection and placement in sensor networks for target localization and tracking," *Journal Of Communication and Networks*, vol. 7, pp. 438–449, 2005.

[4] C. Kreucher, K. Kastella, and A. Hero, "Sensor management using an active sensing approach," *Signal Processing*, vol. 3, no. 85, pp. 607–624, March 2005.

[5] J. L. Williams, J. F. III, and A. Willsky, "An approximate dynamic programming approach to a communication constrained sensor management problem," *Proc. Eighth International Conference of Information Fusion*, July 2005.

[6] J. L. Williams, J. W. Fisher, and A. S. Willsky, "An approximate dynamic programming approach for communication constrained inference," *IEEE Transactions on Signal Processing*, vol. 55, no. 8, pp. 3995–4003, August 2007.

[7] C. M. Kreucher, K. D. Kastella, and A. O. Hero III, "Information-based sensor management for multitarget tracking," in *Signal and Data Processing of Small Targets 2003*, O. E. Drummond, Ed., vol. 5204. SPIE, 2003, pp. 480–489.

[8] C. M. Kreucher, K. D. Kastella, and A. O. H. III, "Multi-platform information-based sensor management," in *Defense Transformation and Network-Centric Systems*, R. Suresh, Ed., vol. 5820. SPIE, 2005, pp. 141–151.

[9] W. W. Schmaedeke, "Information-based sensor management," in *Signal Processing, Sensor Fusion, and Target Recognition II*, I. Kadar and V. Libby, Eds., vol. 1955. SPIE, 1993, pp. 156–164.

[10] W. W. Schmaedeke and K. D. Kastella, "Event-averaged maximum likelihood estimation and information-based sensor management," in *Signal Processing, Sensor Fusion, and Target Recognition III*, I. Kadar and V. Libby, Eds., vol. 2232. SPIE, 1994, pp. 91–96.

[11] B. Grocholsky, A. Makarenko, and H. Durrant-Whyte, "Information-theoretic coordinated control of multiple sensor platforms," in *Proc. IEEE International Conference on Robotics and Automation, 2003*, Sep. 2003, pp. 1521–1526.

[12] D. A. Castañón, "Approximate dynamic programming for sensor management," in *Proceedings 36th IEEE Conference on Decision and Control*, 1997, pp. 1202–1207.

[13] V. Krishnamurthy and J. Evans, "Optimal sensor scheduling for hidden markov model state estimation," *International Journal of Control*, vol. 74, no. 18, pp. 1737–1742, Dec. 2001.

[14] V. Krishnamurthy, "Algorithms for optimal scheduling and management of hidden markov model sensors," *IEEE Transactions on Signal Processing*, vol. 50, no. 6, pp. 1382–1397, Jun. 2002.

[15] D. A. Castañón, "Stochastic control bounds on sensor network performance," in *Proceedings 44th IEEE Conference on Decision and Control*, Dec. 2005, pp. 4939–4944.

[16] J. Wintenby and V. Krishnamurthy, "Hierarchical resource management in adaptive airborne surveillance radars," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 42, no. 2, pp. 401–420, Apr. 2006.

[17] C. Kreucher, A. Hero, and K. Kastella, "A comparison of task driven and information driven sensor management for target tracking," *44th IEEE Conference on Decision and Control*, December 2005.

[18] K. Kastella, "Discrimination gain to optimize detection and classification," *IEEE Transactions on Systems, Man and Cybernetics, Part A*, vol. 27, no. 1, pp. 112–116, Jan. 1997.

[19] D. C. Hitchings and D. A. Castañón, "Receding horizon stochastic control algorithms for sensor management," in *Proc. American Control Conference)*, Baltimore, MD, June 2010.

[20] K. L. Jenkins and D. A. Castañón, "Adaptive sensor management for feature-based classification," in *Proc. of the Conference on Decision and Control*, December 2010.

[21] W. M. Wells, "Statistical approaches to feature-based object recognition," *International Journal of Computer Vision*, vol. 21, no. 1/2, pp. 63–98, 1997.

[22] H. Chiang and R. L. Moses, "Model-based bayesian feature matching with application to synthetic aperture radar target recognition," in *Special Issue on Data and Information Fusion in Image Processing and Computer Vision*, 2001, pp. 1539–1553.

[23] Z. Ying and D. Castañón, "Partially occluded object recognition using statistical models," *International Journal of Computer Vision*, vol. 49, no. 1, pp. 57–78, 2002.

[24] J. Goutsias, R. P. Mahler, and H. T. Nguyen, *Random Sets: Theory and Applications*, ser. IMA volumes in mathematics and its applications. Springer Verlag, 1997, vol. 97.

[25] H. Chiang, R. L. Moses, and L. C. Potter, "Model-based classification of radar images," *IEEE Transactions on Information Theory*, vol. 46, pp. 1842–1854, 2000.

[26] R. B. Washburn, "A random point process approach to multiobject tracking," in *Proc. American Control Conference*, 10–12 June 1987, pp. 1846–1852.

[27] R. Mahler, "Random-set approach to data fusion," in *SPIE*, vol. 2234, July 1994.

[28] H. Chernoff, "A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations," *Annals of Mathematical Statistics*, vol. 23, pp. 493–507, 1952.

[29] K. G. Murty, "An algorithm for ranking all the assignments in order of increasing cost," *Operations Research*, vol. 16, no. 3, pp. 682–687, 1968.