

INFORMATION-BASED SENSOR MANAGEMENT FOR SIMULTANEOUS MULTITARGET TRACKING AND IDENTIFICATION

Christopher M. Kreucher^{a,b}, Alfred O. Hero III^b, Keith D. Kastella^a, and Ben Shapo^a

^aGeneral Dynamics Michigan Research & Development Facility
1200 Joe Hall Drive
Ypsilanti, MI 48197

Keith.Kastella@gd-ais.com, Ben.Shapo@gd-ais.com

^bThe University of Michigan, Dept. of EECS
1301 Beal Avenue, Ann Arbor, MI 48109-2122
ckreuche@umich.edu, hero@eecs.umich.edu

ABSTRACT

This paper addresses the problem of sensor management for simultaneous multitarget tracking and identification. In this context, sensor management refers to the process wherein an agile sensor is tasked, for example by choosing which mode the sensor should use and where the sensor should point. Simultaneous tracking and identification is a challenging sensor management problem because there are two criteria for optimization that must be simultaneously addressed.

The approach advocated here is based on information theory, where the value of a particular action is measured by the amount of information that is expected to be gained by its execution. Using information theory as a metric has the desirable property that the tracking and identification missions can be simultaneously addressed through a single metric without making any ad-hoc assumptions on the relative utility of the missions.

As the sensor management strategy is based on taking actions that maximize information gain, we require the probability density which describes the current uncertainty about the situation. In our application, the system state is characterized by the number of targets and the states (positions and velocities) of each. We address estimation of the density via particle filtering, which is a method of approximating the posterior by sampling the multitarget state space at discrete points (“particles”). Because particles represent a discrete sampling of the entire state space, they implicitly contain estimates of both the

number of targets and their individual states. By concentrating samples near likely multitarget states, the tracker provides a tractable Bayesian solution to the state estimation problem.

We illustrate the method in a scenario involving a simulated sensor with agility in pointing direction and mode, and real target trajectories collected as part of a training exercise at NTC. The performance of the information-based scheduling algorithm is compared to two other (non-information based) scheduling algorithms and a random strategy. The information-based approach is shown empirically to provide superior performance.

1. INTRODUCTION

In this paper, we address the problem of scheduling the resources of an agile sensor. This may include choosing the pointing angle and the mode of the sensor, or even choosing the emitted waveform. We advocate an information-based approach, where sensor tasking decisions are made based on the principle that actions should be chosen to maximize the information expected to be extracted from the scene. This approach provides a single metric able to automatically capture the complex tradeoffs involved when choosing between possible sensor allocations.

We apply this principle to the problem of tracking and identifying multiple moving ground targets from an airborne sensor. The aim is to task the sensor to most efficiently estimate both the number of targets and the state of each target simultaneously. The state of a target includes the kinematic quantities of position and velocity and also target class. This is a challenging sensor management problem as there are multiple criteria (estimation of target location and target class) that must be simultaneously addressed. The exper-

This work was supported under the USAF contract F33615-02-C-1199, AFRL contract SPO900-96-D-0080 and by ARO-DARPA MURI Grant DAAD19-02-1-0262. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force.

iments presented herein use high fidelity models of sensors and real collected target motion data.

The information-based approach to sensor management involves the development of two interrelated elements.

First, we form the joint multitarget probability density (JMPD), which is the fundamental entity capturing knowledge about the number of targets and the states of the individual targets. Unlike traditional methods, the JMPD does not assume any independence, but instead explicitly models coupling in uncertainty between target states, between targets, and between target state and the number of targets. Furthermore, the JMPD is not assumed to be of some parametric form (e.g., Gaussian). Because of this generality, the JMPD must be estimated using sophisticated numerical techniques. Our representation of the JMPD is via a novel multitarget particle filter with an adaptive sampling scheme.

Second, we use the estimate of the JMPD to perform (myopic) sensor resource allocation. The philosophy is to choose actions that are expected to maximize information extracted from the scene. This metric trades automatically between allocations that provide different types of information (e.g., actions that provide information about position versus actions that provide information about target class) without adhoc assumptions as to the relative utility of each. We employ a particular type of information measure, known as the Rényi (alpha) Divergence which is related asymptotically to the more widely used Kullback-Leibler Divergence.

The paper proceeds as follows. First, in Section 2, we describe how the JMPD captures the uncertainty in the multitarget tracking problem and is used in a recursive Bayes filtering scheme to track multiple targets. Second, in Section 3, we give a multitarget particle filter implementation of the JMPD. Third, in Section 4 we show uncertainty reduction in the JMPD is used to drive the sensor management algorithm. Next, in Section 5 we show simulation results involving a set of real targets and a simulated sensor that shows the benefit of the information-based strategy. We finish with a summary and some conclusions in Section 6.

2. THE JOINT MULTITARGET PROBABILITY DENSITY (JMPD)

Estimating the joint multitarget probability density (JMPD) provides a means for tracking an unknown number of targets in a Bayesian setting. Many others have studied Bayesian methods for tracking multiple targets, e.g., [1][2][3][4].

The statistical model uses the joint multitarget conditional probability density

$$p(\mathbf{x}_1^k, \mathbf{x}_2^k, \dots, \mathbf{x}_{T-1}^k, \mathbf{x}_T^k, T^k | \mathbf{Z}^k) = p(\mathbf{x}_1^k, \mathbf{x}_2^k, \dots, \mathbf{x}_{T-1}^k, \mathbf{x}_T^k | T^k, \mathbf{Z}^k) p(T^k | \mathbf{Z}^k) \quad (1)$$

as the probability density for exactly T targets with state vectors $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{T-1}, \mathbf{x}_T$ at time k based on a set of observations \mathbf{Z}^k . In this paper, the state vector \mathbf{x} is taken to include position, velocity and identification components, i.e., $\mathbf{x} = [x \ \dot{x} \ y \ \dot{y} \ c]'$.

The number of targets T is a variable to be estimated simultaneously with the states, although for simplicity the simulations presented in this paper treat the special case where the number of targets is known and fixed. More general discussion and related simulations are considered in [5].

The observation set \mathbf{Z}^k is the collection of measurements up to and including time k , i.e. $\mathbf{Z}^k = \{\mathbf{z}^1, \mathbf{z}^2, \dots, \mathbf{z}^k\}$, where each of the \mathbf{z}^i may be a single measurement or a vector of measurements made at time i .

For convenience, we will denote the multitarget state vector by \mathbf{X} , i.e. $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T]$, where \mathbf{X} is defined for $T = 1 \dots \infty$. Therefore, the JMPD at time k will be written in shorthand notation as $p(\mathbf{X}^k, T^k | \mathbf{Z}^k)$.

The likelihood $p(\mathbf{z} | \mathbf{X}, T)$ and the joint multitarget probability density $p(\mathbf{X}, T | \mathbf{Z})$ are conventional Bayesian objects manipulated by the usual rules of probability and statistics. Thus, a multitarget system has state $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_T)$ with probability distribution $p(\mathbf{x}_1, \dots, \mathbf{x}_T, T | \mathbf{Z})$. This can be viewed as a hybrid stochastic system where the discrete random variable T governs the dimensionality of \mathbf{X} . Therefore the normalization condition that the JMPD must satisfy is

$$\sum_{T=0}^{\infty} \int d\mathbf{x}_1 \dots d\mathbf{x}_T p(\mathbf{x}_1, \dots, \mathbf{x}_T, T | \mathbf{Z}) = 1 \quad (2)$$

where the single integral sign is used to denote the T integrations required.

The temporal update of the posterior likelihood proceeds according to the usual rules of Bayesian filtering. Time evolution of the JMPD is modeled by $p(\mathbf{X}^k, T^k | \mathbf{X}^{k-1}, T^{k-1})$ which will be referred to as the model of multitarget kinematics. This transition density includes models of how individual targets move, models of target birth and death, and any additional prior information that may exist such as terrain and roadway maps. Different kinematic models may be used for different target types.

Filtering proceeds in a manner completely analogous manner to the well known two step prediction-correction Kalman Filter recursions. However, since the JMPD is a non-linear filter, the framework is more general admitting arbitrary (e.g., non-Gaussian) densities, arbitrary (e.g., non-linear) models of temporal evolution and arbitrary (e.g., non-Gaussian) noise processes.

First, the posterior at time $k-1$, $p(\mathbf{X}^{k-1}, T^{k-1} | \mathbf{Z}^{k-1})$ is time-updated (predicted) using the model of multitarget

kinematics to form the prior at time k , $p(\mathbf{X}^k, T^k | \mathbf{Z}^{k-1})$.

$$p(\mathbf{X}^k, T^k | \mathbf{Z}^{k-1}) = \sum_{T^{k-1}=0}^{\infty} \int_{\mathbf{X}^{k-1}} d\mathbf{X}^{k-1} \times \quad (3)$$

$$p(\mathbf{X}^k, T^k | \mathbf{X}^{k-1}, T^{k-1}) p(\mathbf{X}^{k-1}, T^{k-1} | \mathbf{Z}^{k-1}) .$$

Second, the most recent measurements are used to measurement update (correct) the the prior. The prior at time k , $p(\mathbf{X}^k, T^k | \mathbf{Z}^{k-1})$ is used along with Bayes' rule and the most recent measurement \mathbf{z}^k to form the posterior at time k , $p(\mathbf{X}^k, T^k | \mathbf{Z}^k)$.

$$p(\mathbf{X}^k, T^k | \mathbf{Z}^k) = \frac{p(\mathbf{z}^k | \mathbf{X}^k, T^k) p(\mathbf{X}^k, T^k | \mathbf{Z}^{k-1})}{p(\mathbf{z}^k | \mathbf{Z}^{k-1})} . \quad (4)$$

This formulation allows JMPD to avoid altogether the problem of measurement-to-track association which is the fundamental computational issue in conventional multitarget tracking algorithms such as MHT and JPDA. There is no need to identify which target is associated with which measurement because the Bayesian framework keeps track of the entire joint multitarget density. This property, of course, introduces a different but related computational challenge which will be addressed later. Furthermore, there is no need for thresholded measurements (detections). A tractable sensor model merely requires the ability to compute the likelihood $p(\mathbf{z} | \mathbf{X}, T)$ for each measurement \mathbf{z} received. This property allows the JMPD technique to generalize and outperform other multitarget tracking algorithms particularly in low SNR environments.

3. THE PARTICLE FILTER REPRESENTATION OF THE JMPD

The sample space of \mathbf{X} is very large. It includes all configurations of state vectors for all values of T . Discretization on a grid has computational burden exponential in the number of targets and grid cells allotted to each state. A particle filter based implementation with carefully designed importance density allows for computational tractability [6].

To implement JMPD recursions via a particle filter, we approximate $p(\mathbf{X}, T | \mathbf{Z})$ by a set of N_{part} weighted samples (particles). Let the multitarget state vector be written as $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{T-1}, \mathbf{x}_T]$ and be defined for all T , $T = 0 \dots \infty$. Next, let the particle state vector be written $\mathbf{X}_p = [\mathbf{x}_{p,1}, \mathbf{x}_{p,2}, \dots, \mathbf{x}_{p,T_p}]$ where T_p is the estimate particle p has for the number of targets in the region. Letting δ_D denote the Dirac delta where it is understood that it is defined on the domain of its argument (i.e. finite dimensional real or complex vector), we define

$$\delta(\mathbf{X} - \mathbf{X}_p) = \begin{cases} 0 & T \neq T_p \\ \delta_D(\mathbf{X} - \mathbf{X}_p) & \text{otherwise} \end{cases} \quad (5)$$

Then the particle filter approximation to the JMPD is given by

$$p(\mathbf{X}, T | \mathbf{Z}) \approx \sum_{p=1}^{N_{part}} w_p \delta(\mathbf{X} - \mathbf{X}_p) \quad (6)$$

Different particles in the approximation may have different estimates target number, T_p . In practice, the maximum number of targets a particle may track is truncated at some large finite number T_{max} .

Particle filtering is a method of approximately solving the prediction and update equations by simulation [7]. Samples are used to represent the density and to propagate it through time. The prediction equation (eq. 3) is implemented by proposing new particles from the existing set of particles using a model of state dynamics and the measurements. The update equation (eq. 4) is implemented by assigning a weight to each of the particles that have been proposed using the measurements and the model of state dynamics.

This method differs from other PF algorithms where a single particle corresponds to a single target, as it explicitly enforces the multitarget nature of the problem by encoding in each particle an estimate of the number of targets and the states of those targets. Representing the full joint density rather than merely a factorized version provides the advantage that correlations between targets are explicitly modelled. However, due to the dramatic increase in dimensionality, a simplistic implementation leads to greatly increased computational burden. The key to tractability of the particle filter algorithm presented here is an adaptive sampling scheme for particle proposal that automatically factorizes the JMPD when targets or groups of targets are acting independently from the others (i.e. when there is no measurement to target association ambiguity), while maintaining the couplings when necessary. The importance density design is described in detail in [6].

4. INFORMATION BASED SENSOR MANAGEMENT

In this section, we detail our information-based myopic sensor management algorithm. At each instance when a sensor is available, we use an information-based method to compute the best sensing action to take. This is done by first enumerating all possible sensing actions. A sensing action may consist of choosing a particular mode (e.g., HRR mode or GMTI mode), a particular dwell point/pointing angle, or a combination of the two. Next, the *expected* information gain is calculated for each possible action, and the action that yields the maximum expected information gain is taken. The received measurement is then used to update the JMPD, which then drives the choice of the next sensing action.

Calculation of information gain between two densities f_1 and f_0 is done via the Rényi information divergence [8][9]:

$$D_\alpha(f_1||f_0) = \frac{1}{\alpha-1} \ln \int f_1^\alpha(x) f_0^{1-\alpha}(x) dx \quad (7)$$

We compute the divergence between the predicted and the updated density after a measurement is made. The PF approximation of the density simplifies eq. (7) to

$$D_\alpha(p(\cdot|\mathbf{Z}^{k+1})||p(\cdot|\mathbf{Z}^k)) \propto \ln \frac{1}{p(\mathbf{z})^\alpha} \sum_{p=1}^{N_p} w_p p(\mathbf{z}|\mathbf{X}_p)^\alpha \quad (8)$$

The sensor model $p(\mathbf{z}|\mathbf{X}_p)$ incorporates everything about the sensor, including signal to noise ratio, detection probabilities, and whether the locations are visible.

We wish to perform the measurement that makes the divergence between the current density and the density after a new measurement largest. This indicates the action has maximally increased information content of the measurement updated density with respect to the density before a measurement was made. To this end, we calculate the expected value of eq. (8) for each of the N possible sensing actions and choose the action that maximizes the expectation. Let a_i , $i = 1 \dots N$ to refer to the possible sensing actions under consideration, including but not limited to sensor mode selection and sensor beam positioning.

The expected value of eq. (8) is an integral over all possible outcomes z_{a_i} when performing action a_i :

$$\|D_\alpha\|_{a_i} = \int d\mathbf{z}_{a_i} p(\mathbf{z}_{a_i}|\mathbf{Z}^k) D_\alpha(p(\cdot|\mathbf{Z}^k, z_{a_i})||p(\cdot|\mathbf{Z}^k)) \quad (9)$$

In the special case of thresholded measurements (e.g., detections or non-detections), we have

$$\|D_\alpha\|_{a_i} \propto \sum_{z_{a_i}=0}^1 p(z_{a_i}) \ln \frac{1}{p(z_{a_i})^\alpha} \sum_{p=1}^{N_p} w_p p(z_{a_i}|\mathbf{X}_p)^\alpha \quad (10)$$

5. SIMULATION RESULTS

One of the principal benefits of the information based sensor management approach is that the complex tradeoffs between different sensing actions are automatically taken into account. The tradeoffs become even more complex in the case of a sensor that is able to decide between several modes of operation. In this section, we investigate via simulation studies a situation involving a sensor that has three modes available for use:

- A moving target indicator (MTI), which is a mode that able to detect the position of targets only when they are moving,
- A fixed target indicator (FTI), which is a mode that is able to detect the position of targets only when they are stopped, and
- An identification (ID) sensor, which is able to determine the type (e.g., jeep or tank) of a target

To choose the action, we compute the gain in information for each of the possible sensing actions. This includes each possible measurement when using the MTI sensor, each possible measurement when using the FTI sensor, and each possible measurement when using the ID sensor.

The MTI and FTI sensor modes model a GMTI and a SAR processing chain, respectively. When measuring a cell, the imager returns either a 0 (no detection) or a 1 (detection) which is governed by a probability of detection (P_d) and a per-cell false alarm rate (P_f). When a target is moving, the FTI sensor always returns the false alarm rate. Likewise, when a target is stationary, the MTI sensor always returns the false alarm rate. The signal to noise ratio (SNR) links these values together. In this illustration, we take $P_d = 0.5$, and $P_f = P_d^{(1+SNR)}$, which is a standard model for thresholded detection of Rayleigh returns. When there are T targets in the same cell, the detection probability increases according to $P_d(T) = P_d^{\frac{1+SNR}{1+T*SNR}}$. This model is known by the filter and used to evaluate (4).

The ID sensor models a complete automatic target recognition (ATR) system that involves a high range resolution radar and a signal processing algorithm. There are 3 possible target types in this simulation. We model the performance by a confusion matrix, which describes the probability that algorithm will return a particular classification when it is pointed at a particular target type. The model is given in Table 1 and says that when a single target occupies a detection cell, the probability of correctly identifying the target is 0.6, with the misclassifications spread evenly about the other two classes. Also, when multiple targets occupy the same cell or no targets are in the cell, the ATR algorithm returns a random classification. This model is reasonable for common ATR systems, which rely on the geometry of scattering centers for targets to provide classification calls. When multiple targets are contributing to an energy return, this geometry is corrupted and ATR performs very poorly.

The expected myopic gain in information from using the ID sensor follows directly from (9), where the number of possible outcomes is now 3:

$$\langle D_\alpha \rangle_m = \frac{1}{\alpha-1} \sum_{z=1}^3 p(z) \ln \frac{1}{p(z)^\alpha} \sum_{p=1}^{N_{part}} w_p p(z|\mathbf{X}_p)^\alpha \quad (11)$$

Classification Probability	Actual Cell Status			Empty or Multiply Occupied
	Type 1	Type 2	Type 3	
Type 1	0.60	0.20	0.20	0.33
Type 2	0.20	0.60	0.20	0.33
Type 3	0.20	0.20	0.60	0.33

Table 1: The identification sensor is modeled using a confusion matrix. The confusion matrix given here, used for the simulations presented in this section, says each measurement of a single target is independent and provides the correct identification 60% of the time. Measurements of empty cells or cells containing multiple targets return a random classification call.

The goal is to use the sensor to simultaneously determine the position and target types of a group of maneuvering targets. The target motion is taken from real targets that are performing combat maneuvers. At each time step, the sensor must choose from among M different sensing actions (choosing both mode and pointing angle). Initially, the positions of the targets are known (with some covariance) and the identification is unknown.

We present in Figure 1 a comparison between the performance of the algorithm using the information based method, periodic scan, and two other non-divergence based methods. We compare the performance of the various managed strategies and the periodic scheme by looking at RMS error versus number of sensor dwells ("looks") and the number of targets correctly detected.

Sensor management algorithm "A" manages the sensor by pointing it at or near the estimated location of the targets. Specifically, algorithm "A" performs a gating procedure to restrict the portion of the surveillance area that the sensor will consider measuring. The particle filter approximation of the time updated JMPD (3) is used to predict the location of each of the targets at the current time. The set of cells that the sensor manager considers is then restricted to those cells containing targets plus the surrounding cells, for a total of 9 cells in consideration per target. The dwells are then allocated randomly among the gated cells.

Sensor management algorithm "B" tasks the sensor based on the estimated number of targets in each sensor cell. Specifically, the particle approximation of the time updated JMPD is projected into sensor space to determine the filter's estimate of the number of targets in each sensor cell. The cell to measure is then selected probabilistically, favoring cells that are estimated to contain more targets. In the single target case, this method reduces to measuring the cell that is most likely to contain the target.

Both algorithms "A" and "B" switch between the three sensor modes sequentially. Notice that both methods "A" and "B" introduce new assumptions that were not present in the information-based algorithm. In particular, both give value only to looking near where targets are predicted to be. There seems to be no natural extension of this heuristic to include detection of new targets. Furthermore, algorithm "A" assumes equal weight to all areas predicted to contain targets and algorithm "B" gives precedence to areas predicted to contain multiple targets.

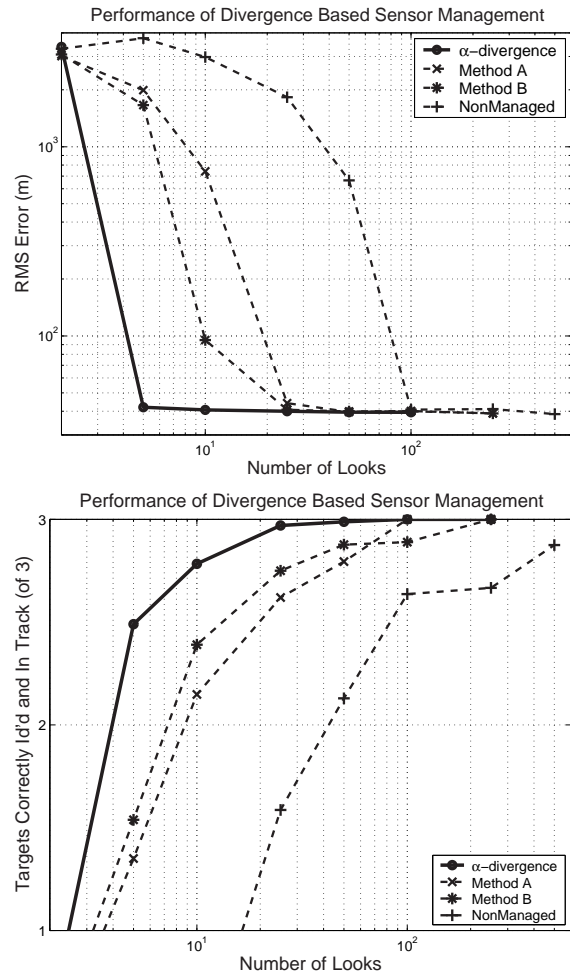


Figure 1: A comparison of the information-based method to periodic scan and two other methods for simultaneous track and ID. The performance is measured in terms of the (median) RMS error and the (average) number of targets correctly identified and in track. The α -divergence strategy outperforms the other strategies for a fixed amount of sensor resources.

6. CONCLUSION

This paper has illustrated a method of sensor management based on choosing actions that are expected to gain the most information. The compelling thing about information theoretic scheduling is that different mission goals such as detection, tracking, and identification can be simultaneously addressed with a single metric. Indeed, it can be shown [10] that information gain can be interpreted as an approximation to any task related measure such as minimizing tracking error or maximizing classification performance.

The main practical difficulty in applying information theoretic approaches is that a good representation of the posterior density is required. We have addressed this via a multitarget particle filter, which provides a computationally efficient and robust approach to implementing the nonlinear filter.

The benefit of the information theoretic strategy was shown on a model problem involving real data and a sensor with a agility in mode and pointing angle. In particular, the information-theoretic approach is shown to outperform two other algorithms in terms of tracking error and identification performance for a fixed number of sensor resources.

7. REFERENCES

- [1] L. D. Stone, T. L. Corwin, and C. A. Barlow, *Bayesian Multiple Target Tracking*, Artech House, 1999.
- [2] M. I. Miller, A. Srivastava, and U. Grenander, "Conditional mean estimation via jump-diffusion processes in multiple target tracking/recognition," *IEEE Transactions on Signal Processing* **43**(11), pp. 2678–2690, 1995.
- [3] R. E. Bethel and G. J. Paras, "A PDF multisensor multitarget tracker," *IEEE Transactions on Aerospace and Electronic Systems* **34**, pp. 153–168, January 1998.
- [4] R. Mahler, "Global optimal sensor allocation," *Proceedings of the Ninth National Symposium on Sensor Fusion* **1**, pp. 167–172, 1996.
- [5] C. M. Kreucher, M. Morelande, A. O. Hero III, and K. Kastella, "Particle filtering for multitarget detection and tracking," *The 2005 IEEE Aerospace Conference*, 2005.
- [6] C. M. Kreucher, K. Kastella, and A. O. Hero III, "Multitarget tracking using a particle filter representation of the joint multitarget probability density," *to appear in IEEE Transactions on Aerospace and Electronic Systems*, 2005.
- [7] A. Doucet, N. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*, Springer Publishing, 2001.
- [8] A. Rényi, "On measures of entropy and information," *Proceedings of the 4th Berkeley Symposium on Mathematics, Statistics, and Probability* **1**, pp. 547–561, 1961.
- [9] A. O. Hero III, B. Ma, O. Michel, and J. D. Gorman, "Alpha divergence for classification, indexing and retrieval," *Technical Report 328, Comm. and Sig. Proc. Lab. (CSPL), Dept. EECS, The University of Michigan*, 2001.
- [10] C. M. Kreucher and A. O. Hero III, "A comparison of task driven and information driven sensor management for target tracking," *The 2005 IEEE Conference on Decision and Control*, 2005.