

A Comparison of Task Driven and Information Driven Sensor Management for Target Tracking

Chris Kreucher^{a,b}, Alfred O. Hero III^a, and Keith Kastella^b

Abstract—Several authors have proposed sensor scheduling methods that are driven by information theoretic measures. In the information driven approach, the relative merit of different sensing actions is measured by the corresponding expected gain in information. Information driven approaches stand in stark contrast to task driven methods, i.e., methods that select some physical performance criteria and explicitly manage the sensor based on that criteria. This paper investigates the difference between a particular information driven approach, one that maximizes an alpha-Rényi measure of information gain, and task driven methods with a combination of theory and simulation. First, we give a mathematical relation that shows that when the decision error depends only weakly on the target state a certain type of marginalized information gain is a close approximation to the Bayes risk associated with performing a specific task. Second, we perform an empirical comparison between information driven and task driven approaches that maximize information gain or minimize risk, respectively. In particular, we give a task driven method that uses the sensor in a manner that is expected to maximize the probability the target is correctly located after the next measurement. We find as expected that the task driven method outperforms the information driven method when the performance is measured by risk, i.e., probability of localization error. However, the performance difference between the two methods is very small, suggesting that the information gain is a good surrogate for risk for this application.

I. INTRODUCTION

The problem of sensor management is to determine the best way to task a sensor or group of sensors when each sensor may have many modes and search patterns. Typically, the sensors are used to gain information about the kinematic state (e.g. position and velocity) and identification of a group of targets. There are many tasks that the sensor manager may be tuned to meet, e.g., minimization of track loss, probability of target detection, and identification accuracy. One can design a “task-driven” sensor management strategy that addresses one of these objectives but it may be very poor in addressing the others. Ad hoc methods that address a number of objectives by assigning relative weight to each task can also be used. However, this requires one to enumerate all of the possible objectives and assign relative values to each.

This work was supported by the USAF Contract No. F33615-02-C-1199, AFRL contract SPO900-96-D-0080, and 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.

^aThe University of Michigan Department of EECS, Ann Arbor, MI ckreuche@umich.edu, hero@eeecs.umich.edu

^bGeneral Dynamics Advanced Information Systems, Ypsilanti, MI Keith.Kastella@gd-ais.com

To deal with a multitude of performance criteria in a direct manner simultaneously, some researchers have proposed using information theoretic measures as a means of sensor management. In the context of Bayesian estimation, a good measure of the quality of a sensing action is the reduction in entropy of the posterior distribution that is expected to be induced by the measurement. Therefore, information theoretic methodologies strive to take the sensing action that maximizes the expected gain in information. The possible sensing actions are enumerated, the expected gain for each measurement is calculated, and the action that yields the maximum expected gain is chosen. Information as the criteria for driving sensing actions has the desirable property that different actions which gain different types of information can be compared using a single metric.

Hintz et. al. [1][2] focus on using the expected change in Shannon entropy when tracking a single target moving in one dimension with Kalman Filters. A related approach uses discrimination gain based on a measure of relative entropy, the Kullback-Leibler (KL) divergence. Schmaedeke and Kastella [3] use the KL divergence to determine optimal sensor-to-target tasking. Kastella [4][5] uses KL divergence to manage a sensor between tracking and identification mode in the multitarget scenario. Mahler [6] uses the KL divergence as a metric for optimal multisensor multitarget sensor allocation. Kreucher [7] uses the Rényi Divergence, a generalization of the KL divergence for multitarget detection, tracking, and identification. Zhao [8] compares several approaches, including simple heuristics, entropy, and relative entropy.

As a means of demonstrating the utility of information theoretic methods, this paper investigates the relationship between information driven and task driven methods. First, we derive an explicit mathematical relationship between an information driven sensor management algorithm based on the Rényi Divergence and task driven methods. This result indicates that such information gain methods can be expected to be nearly optimal when the decision error is only weakly dependent on the state and the incremental information gain from each measurement is small. Second, we perform an empirical comparison of information driven sensor scheduling and task driven scheduling for target tracking. This investigation is done by designing a sensor scheduling method that chooses the action so as to maximize the probability the target is correctly located. Here the “task” that we are specializing to is target tracking. We then compare this method to an information driven sensor scheduling algorithm that takes actions to maximize the gain in information. We look at two variants of the information driven approach, the

first where information is maximized and the second where marginalized information is maximized.

The paper proceeds as follows. First, in Section II we give an overview of a Bayesian method of capturing uncertainty about the number of targets, their states, and their identifications. A more detailed version is given elsewhere [9]. This probabilistic approach is fundamental to describe the entropy of the current estimate and to compute the expected information gain achievable from each subsequent action. Next, in Section III, we describe our information-driven approach to sensor tasking. We show how this method can be used to emphasize certain types of information over others (e.g., kinematic information over identification information) by marginalization of the relevant densities. Third, in Section IV, we give a bound that describes the relationship between information driven scheduling using the Rényi Divergence and task driven scheduling. Fourth, in Section V, we describe a task driven method that is designed to take actions that maximize the probability of correctly estimating the target position. Fifth, Section VI presents a simulation result where the different methods are empirically compared. Finally, in Section VII, we give some concluding remarks.

II. THE JOINT MULTITARGET PROBABILITY DENSITY (JMPD)

In this section, we give an overview of a Bayesian method of recursively estimating the uncertainty about the number of targets, their states, and their identifications. This method is based on recursive estimation of the joint multitarget probability density (JMPD) [5][9]. Among the other researchers that have studied Bayesian methods for tracking multiple targets are [10][11][12] and [13].

The estimated JMPD is used in the sensor management strategies discussed in Sections III and V. In the information driven case (Section III), actions are taken to maximally gain information about the JMPD as measured by the Rényi Divergence. In the task driven case (Section V), actions are taken to maximize the probability targets are correctly located.

Recursive estimation of the JMPD provides a means for tracking an unknown number of targets in a Bayesian setting. The statistical model employed 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 states $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{T-1}, \mathbf{x}_T$ at time k based on a set of past observations \mathbf{Z}^k , i.e., $\mathbf{Z}^k = \{\mathbf{z}^0, \mathbf{z}^1, \dots, \mathbf{z}^{k-1}, \mathbf{z}^k\}$ where \mathbf{z}^k denotes the measurement(s) collected at time k .

With a slight abuse of terminology, we call $p(\mathbf{x}_1^k, \mathbf{x}_2^k, \dots, \mathbf{x}_{T-1}^k, \mathbf{x}_T^k, T^k | \mathbf{Z}^k)$ a density since T is a discrete valued random variable. In fact, as (1) shows, the JMPD continuous discrete hybrid as it is a product of the probability mass function $p(T^k | \mathbf{Z}^k)$ and the probability density function $p(\mathbf{x}_1^k, \mathbf{x}_2^k, \dots, \mathbf{x}_{T-1}^k, \mathbf{x}_T^k | T^k, \mathbf{Z}^k)$.

The normalization condition is thus

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

Each of the state vectors \mathbf{x}_t in the density $p(\mathbf{x}_1^k, \mathbf{x}_2^k, \dots, \mathbf{x}_{T-1}^k, \mathbf{x}_T^k | T^k, \mathbf{Z}^k)$ is a vector quantity and may (for example) be of the form $[x, \dot{x}, y, \dot{y}, c]$, where c is the target class. For convenience, the density will be written more compactly in the traditional manner as $p(\mathbf{X}^k | T^k, \mathbf{Z}^k)$, which implies that the state-vector \mathbf{X} represents a variable number of targets each possessing their own state vector. We will drop the time superscript k for notational simplicity when no confusion will arise.

The temporal update of the posterior likelihood proceeds according to the usual rules of Bayesian filtering. The model of how the JMPD evolves over time is given by $p(\mathbf{X}^k, T^k | \mathbf{X}^{k-1}, T^{k-1})$ and will be referred to as the kinematic prior (KP). The kinematic prior includes models of target motion, target birth and death, and any additional prior information that may exist such as terrain and roadway maps. In the case where target identification is part of the state being estimated, different kinematic models may be used for different target types.

The time-updated (prediction) density is computed via the *model update* equation as

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

The *measurement update* equation uses Bayes' rule to update the posterior density with a new measurement \mathbf{z}^{k+1} as

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

For the purposes of the simulations included in this paper, we specialize to the case where T is known and fixed at $T = 1$. The more general case, where T is unknown, large, and time varying is treated in [9] and [14].

\mathbf{X} has a very large sample space. It contains 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 implementation with appropriately designed importance density allows for computational tractability [9]. To implement JMPD via a particle filter, we approximate the joint multitarget probability density by a set of N_p weighted samples (particles), i.e. $p(\mathbf{X} | \mathbf{Z}) \approx \sum_{p=1}^{N_p} w_p \delta(\mathbf{X} - \mathbf{X}_p)$.

In summary, by recursive estimation of the JMPD, we capture all of the uncertainty about the number of targets and their states conditioned on all of the measurements made. Implementation of this is done via a particle filter with an adaptive importance density. The JMPD then drives the sensor management algorithms as described in the next sections.

III. THE RÉNYI DIVERGENCE FOR INFORMATION DRIVEN SENSOR MANAGEMENT

This section gives the details of our information driven sensor management algorithm. As others have realized [5][6][8], a good measure of the quality of a sensing action is the reduction in entropy of the posterior distribution that is induced by the measurement. Since we wish to determine the best sensing action to take before actually executing it, we require a measure of the *expected* reduction in entropy that a sensing action will produce. This is done by first enumerating all possible sensing actions. A sensing action may consist of choosing a particular sensor mode, a particular dwell point/pointing angle, or a combination of the two. Next, the *expected* information gain is calculated for each of the possible actions, and the action that yields the maximum expected information gain is taken. The measurement received is used to update the JMPD, which is in turn used to determine the next measurement to make.

A. The Rényi Divergence

In this work, the calculation of information gain between two densities p_1 and p_0 is done using the Rényi information divergence [15][16], also known as the α -divergence:

$$D_\alpha(p_1||p_0) = \frac{1}{\alpha-1} \ln \int p_1^\alpha(x)p_0^{1-\alpha}(x)dx . \quad (5)$$

The α parameter may be used to adjust how heavily one emphasizes the tails of the two distributions p_1 and p_0 . In the limiting case of $\alpha \rightarrow 1$ the Rényi divergence becomes the commonly utilized Kullback-Leibler discrimination (6).

$$\lim_{\alpha \rightarrow 1} D_\alpha(p_1||p_0) = \int p_0(x) \ln \frac{p_0(x)}{p_1(x)} dx . \quad (6)$$

If $\alpha = 0.5$, the Rényi information divergence becomes the Hellinger affinity $2 \ln \int \sqrt{p_1(x)p_0(x)}dx$, which is related to the Hellinger-Battacharya distance squared [17] via

$$D_{Hellinger}(p_1||p_0) = 2 \left(1 - \exp \left(.5 D_{\frac{1}{2}}(p_1||p_0) \right) \right) . \quad (7)$$

B. The Rényi Divergence In The JMPD Setting

The function D_α in (5) is a measure of the divergence between the densities p_0 and p_1 . In our application, we are interested in computing the divergence between the predicted density $p(\mathbf{X}^k, T^k|\mathbf{Z}^{k-1})$ and the updated density after a measurement is made, $p(\mathbf{X}^k, T^k|\mathbf{Z}^k)$. Therefore, the divergence of interest is given by (8).

$$D_\alpha(p(\cdot|\mathbf{Z}^k)||p(\cdot|\mathbf{Z}^{k-1})) = \frac{1}{\alpha-1} \times \ln \sum_{T^k} \int_{\mathbf{X}^k} p^\alpha(\mathbf{X}^k, T^k|\mathbf{Z}^k) p^{1-\alpha}(\mathbf{X}^k, T^k|\mathbf{Z}^{k-1}) d\mathbf{X}^k . \quad (8)$$

Using Bayes' rule applied to the JMPD and simple algebra, we can write D_α as (9).

$$D_\alpha(p(\cdot|\mathbf{Z}^k)||p(\cdot|\mathbf{Z}^{k-1})) = \frac{1}{\alpha-1} \ln \frac{1}{p^\alpha(\mathbf{z}^k|\mathbf{Z}^{k-1}, m)} \times \sum_{T^k} \int_{\mathbf{X}^k} p^\alpha(\mathbf{z}^k|\mathbf{X}^k, T^k, m) p(\mathbf{X}^k, T^k|\mathbf{Z}^{k-1}) d\mathbf{X}^k . \quad (9)$$

where m is the sensing action and \mathbf{z}^k is the most recently received measurement.

C. The Expected Rényi Divergence for a Sensing Action

Our real aim is to choose the sensing action to take *before actually receiving* the measurement \mathbf{z} . Specifically, we would like to choose to take the action that makes the divergence between the current density and the density after a new measurement as large as possible. This indicates that the sensing action has maximally increased the information content of the measurement updated density, $p(\mathbf{X}^k, T^k|\mathbf{Z}^k)$, with respect to the density before a measurement was made, $p(\mathbf{X}^k, T^k|\mathbf{Z}^{k-1})$. However, we cannot choose the action that maximizes the divergence as we do not know the outcome of the action before taking it.

We propose, then, as a method of sensor management to calculate the expected value of the Rényi Divergence for each of the M possible sensing actions and to choose to take the action that maximizes the expectation. In this notation m ($m = 1 \dots M$) will refer to any possible sensing action under consideration, including but not limited to sensor mode selection and sensor beam positioning. In this manner, we say that we are making the measurement that maximizes the expected gain in information as measured by the Rényi Divergence.

The expectation may be written as an integral over all possible outcomes \mathbf{z} when performing sensing action m as

$$\langle D_\alpha \rangle_m = \int_{\mathbf{z}^k} dz^k p(\mathbf{z}^k|\mathbf{Z}^{k-1}, m) D_\alpha(p(\cdot|\mathbf{Z}^k)||p(\cdot|\mathbf{Z}^{k-1})) . \quad (10)$$

Therefore, the information based scheduling strategy is to choose the action m_{opt} as

$$m_{opt} = \arg \max_m \langle D_\alpha \rangle_m . \quad (11)$$

D. Scheduling Using The Rényi Divergence Between Marginalized JMPDs

If is known *a priori* that certain types of information are more important than others, one can use a Rényi divergence between marginalized JMPDs to allocate the sensor.

Consider the case where one only wants to use kinematic information when scheduling the sensor. The state of the system \mathbf{X} contains both kinematic (x, \dot{x}, y, \dot{y}) and target class (c) information

$$\mathbf{X} = [x_1, \dot{x}_1, y_1, \dot{y}_1, c_1 \dots x_T, \dot{x}_T, y_T, \dot{y}_T, c_T] . \quad (12)$$

In this case, one can marginalize across target type and generate a new state variable \mathbf{Y} that only contains kinematic information

$$\mathbf{Y} = [x_1, \dot{x}_1, y_1, \dot{y}_1, \dots x_T, \dot{x}_T, y_T, \dot{y}_T] . \quad (13)$$

The marginalization is performed according to

$$p(x_1, \dot{x}_1, y_1, \dot{y}_1, \dots x_T, \dot{x}_T, y_T, \dot{y}_T|\mathbf{Z}) = \int_{c_1} \dots \int_{c_T} dc_1 \dots dc_T p(x_1, \dot{x}_1, y_1, \dot{y}_1, c_1, \dots|\mathbf{Z}) . \quad (14)$$

The density $p(\mathbf{Y}|\mathbf{Z})$ can then be used to calculate expected information gain on the reduced state space using the divergence between marginalized JMPDs

$$D_\alpha(p(\cdot|\mathbf{Z}^k)||p(\cdot|\mathbf{Z}^{k-1})) = \frac{1}{\alpha-1} \ln \int_{\mathbf{Y}} p^\alpha(\mathbf{Y}^k, T^k|\mathbf{Z}^k) p^{1-\alpha}(\mathbf{Y}^k, T^k|\mathbf{Z}^{k-1}) d\mathbf{Y}^k. \quad (15)$$

One may additionally envision treating kinematic and target type information separately (although this marginalization ignores the coupling in uncertainty between the two) by marginalizing out target type to get \mathbf{Y} and marginalizing out kinematics to get \mathbf{V} and using

$$D_\alpha(p(\cdot|\mathbf{Z}^k)||p(\cdot|\mathbf{Z}^{k-1})) = \frac{w_y}{\alpha-1} \ln \int_{\mathbf{Y}} d\mathbf{Y}^k p^\alpha(\mathbf{Y}^k, T^k|\mathbf{Z}^k) p^{1-\alpha}(\mathbf{Y}^k, T^k|\mathbf{Z}^{k-1}) + \frac{w_v}{\alpha-1} \ln \int_{\mathbf{V}} d\mathbf{V}^k p^\alpha(\mathbf{V}^k, T^k|\mathbf{Z}^k) p^{1-\alpha}(\mathbf{V}^k, T^k|\mathbf{Z}^{k-1}). \quad (16)$$

where w_y and w_v are the relative weight given to the different information types.

IV. MATHEMATICAL RELATION BETWEEN THE MARGINALIZED RÉNYI DIVERGENCE AND RISK

In this section we show that there is a simple mathematical relation between marginalized Rényi α divergence and task based performance, i.e., risk. This result establishes that, with the right weighting, a certain function of the Rényi divergence can be viewed as an approximation to the risk.

We start by proving a simple bound linking the expectation of a non-negative random variable to weighted divergence. Let U be an arbitrary r.v., let p and q be densities for U , and define $E_p[g(U)] = \int g(u)p(u)du$. Assume that q dominates p , i.e. $q(u) = 0$ implies $p(u) = 0$, and $\alpha > 0$. Then for any bounded non-negative function g , with $w = \inf_u g(u)$, $W = \sup_u g(u)$:

$$E_p[g] \geq w E_q^{1/\alpha} \left[\left(\frac{p}{q} \right)^\alpha \right], \quad \alpha \leq 1$$

$$E_p[g] \leq W E_q^{1/\alpha} \left[\left(\frac{p}{q} \right)^\alpha \right], \quad \alpha > 1, \quad (17)$$

with equality when either $\alpha = 1$ or when g is constant and $p = q$.

Proof: For $\alpha \leq 1$ ($\alpha \geq 1$) concavity (convexity) of the function $g(z) \stackrel{\text{def}}{=} z^\alpha$ and Jensen's inequality imply that $E_q[Z^\alpha] \leq E_q^\alpha[Z]$ ($E_q[Z^\alpha] \geq E_q^\alpha[Z]$) for any non-negative r.v. Z . Combining this inequality with the fact that $E_q[gp/q] = E_p[g]$ yields the result (17).

To apply the above bound to the JMPD tracking application we define $U = [\mathbf{X}^k, \mathbf{z}^k, m]^T$, define $E[g(U)]$ as the conditional Bayes risk associated with cost function $g(U)$ under action m when using the optimal estimate of state \mathbf{X}^k after observing \mathbf{z}^k , and define $p(u) = p(\mathbf{X}^k|\mathbf{Z}^k, m)$ and $q(u) = p(\mathbf{X}^k|\mathbf{Z}^{k-1}, m)$ the posterior densities of the state \mathbf{X}^k given the m -th action and observation sets \mathbf{Z}^k

and \mathbf{Z}^{k-1} , respectively. For example, for a single target, if the state $\mathbf{X}^k = [x^k, y^k, c]^T$ is composed of target position $\xi^k = [x^k, y^k]^T$ and target class c then one can define $g(U) = (\xi^k - \hat{\xi}^k(\mathbf{z}^k, m))^2$ or $g(U) = 1 - I(c - \hat{c}(\mathbf{z}^k, m))$ when the cost function is associated with the tracking MSE of the optimal tracker or the classification error probability of the optimal classifier, respectively.

More generally, let ξ denote the components of the state \mathbf{X}^k over which the risk $g(U) = g(\xi^k, \mathbf{z}^k, m)$ is not constant. When the above bound is applied to the conditional expectation $E[g(U)|\mathbf{z}^k, m]$ we obtain for $\alpha \leq 1$ and bounded cost function,

$$E[g(U)|\mathbf{z}^k, m] = \int g(\xi, \mathbf{z}^k, m) p(\xi|\mathbf{z}^k, m) d\xi$$

$$\leq W \left(\int \left(\frac{p(\xi|\mathbf{Z}^k, m)}{p(\xi|\mathbf{Z}^{k-1}, m)} \right)^\alpha d\xi \right)^{1/\alpha}$$

$$= W \exp \left(-\frac{1-\alpha}{\alpha} D_\alpha(p(\cdot|\mathbf{Z}^k)||p(\cdot|\mathbf{Z}^{k-1})) \right), \quad (18)$$

To obtain this bound, we have set $q = p(\xi|\mathbf{Z}^{k-1}, m)$. Thus we conclude the following:

- The RHS of the bound (18) involves the Rényi divergence of the marginals of the posterior state given \mathbf{Z}^k and \mathbf{Z}^{k-1} , respectively.
- When the cost function g is independent of the state ξ , the RHS of (18) can be more tightly bounded by a weighted Rényi divergence of the marginals.
- The RHS of (18) implies it is better to work with the divergences between marginalized posteriors than the divergence between the posteriors and the full state \mathbf{X}^k .
- One can sandwich the conditional Bayes risk by selecting $\alpha_1 \in [0, 1]$ and $\alpha_2 > 1$ in the bounds (17)

$$\left(w \int \left(\frac{p(\xi|\mathbf{Z}^k, m)}{p(\xi|\mathbf{Z}^{k-1}, m)} \right)^{\alpha_2} d\xi \right)^{1/\alpha_2}$$

$$\leq E[g(\xi, \mathbf{z}^k, m)|\mathbf{z}^k, m] \leq$$

$$\left(W \int \left(\frac{p(\xi|\mathbf{Z}^k, m)}{p(\xi|\mathbf{Z}^{k-1}, m)} \right)^{\alpha_1} d\xi \right)^{1/\alpha_1}$$

- The above bound becomes tighter when the cost function g only weakly depends on ξ and $p(\xi|\mathbf{Z}^k, m)$ is close to $p(\xi|\mathbf{Z}^{k-1}, m)$.

V. TASK DRIVEN SENSOR MANAGEMENT

In this section, we derive a task driven sensor scheduling algorithm which chooses actions to minimize the expected tracking error [18].

The details of this algorithm are as follows. The task driven sensor management algorithm is designed to choose the action that results in maximal certainty about target position after the action is taken. Mathematically, this criteria can be stated as taking the action that leads to the largest *maximum a posteriori* (MAP) probability of target location. As in the information driven tasking algorithm, we do not know the outcome of a sensing action before taking the

action. Therefore, we again choose to take the action that maximizes the expected certainty about target position.

We specialize here to the case where we are concerned only with the location in terms of the sensor grid, although the method is general. Let c denote a sensor cell. Furthermore, let $p(c|\mathbf{Z}^{k-1})$ denote the probability that the target is in cell c given the measurements up to time $k-1$, \mathbf{Z}^{k-1} . Then the sensor allocation strategy that maximizes location certainty is to choose the action m that is expected to maximize the *a posteriori* probability the target is in cell c for the most likely cell c , i.e.,

$$m_{opt} = \arg \max_m E_z \left(\max_c p(c|\mathbf{Z}^{k-1}, \mathbf{z}^k, m) \right) \quad (19)$$

$$\arg \max_m \int_{\mathbf{z}^k} p(\mathbf{z}^k|\mathbf{Z}^{k-1}, m) \left(\max_c p(c|\mathbf{Z}^{k-1}, \mathbf{z}^k, m) \right) d\mathbf{z}^k .$$

The important probability density $p(c|\mathbf{Z}^{k-1}, \mathbf{z}^k, m)$ is computed in two steps. First, the posterior JMPD $p(\mathbf{X}^k|\mathbf{Z}^{k-1}, \mathbf{z}^k, m)$ is computed via Bayes' rule using the received measurement \mathbf{z}^k . In the case of interest in this paper, this is just the single target posterior $p(\mathbf{x}^k|\mathbf{Z}^{k-1}, \mathbf{z}^k, m)$. Then the probability $p(c|\mathbf{Z}^{k-1}, \mathbf{z}^k, m)$ is derived from this by integrating over sensor cell c , i.e.,

$$p(c|\mathbf{Z}^{k-1}, \mathbf{z}^k, m) = \int_{\mathbf{x} \in c} d\mathbf{x}^k p(\mathbf{x}^k|\mathbf{Z}^{k-1}, \mathbf{z}^k, m) . \quad (20)$$

Note that this could be generalized to any arbitrary region (not necessarily tied to the sensor grid) by appropriate choice of integration region corresponding to "cell" c .

VI. SIMULATION RESULT

We consider the following model problem to illustrate the performance of the two sensor scheduling algorithms developed here. There is a single target with unknown position and identification present in a surveillance region. An airborne sensor has the ability to interrogate any cell in the surveillance region, where a cell measures $100\text{m} \times 100\text{m}$. The sensor may choose to interrogate with one of two modes: an identification mode or a target indication mode.

The identification mode idealizes a synthetic aperture radar dwell and the associated signal processing algorithms. In our abstraction, the result of an identification dwell is a classification call. There are ten possible target classes. We model the performance of the identification mode by the confusion matrix given in Table I.

The target indication mode idealizes a ground moving target indicator dwell and the associated signal processing algorithms. In our abstraction, the result of a target indication dwell is binary, corresponding to a no-detection or a detection. We model the performance of the target indication mode using a detection rate P_d which specifies the probability of detection if the cell is occupied and a false alarm rate P_f which specifies the probability of detection if the cell is not occupied. In this simulation we set $P_d = 0.5$ and $P_f = .001$, corresponding to an SNR of 10dB .

Classification Probability	Cell Status				
	Type 1	Type 2	...	Type 10	Empty
Type 1	0.60	0.04	...	0.04	0.05
Type 2	0.04	0.60	...	0.04	0.05
.
.
.
Type 10	0.04	0.04	...	0.60	0.05

TABLE I

THE MODEL FOR THE IDENTIFICATION SENSOR. EACH MEASUREMENT OF A SINGLE TARGET IS INDEPENDENT AND PROVIDES THE CORRECT IDENTIFICATION 60% OF THE TIME. THE CLASSIFICATION ERRORS ARE DISTRIBUTED UNIFORMLY OVER THE OTHER CLASSES. MEASUREMENTS OF EMPTY CELLS RETURN A RANDOM CLASSIFICATION CALL.

At each time step, the sensor manager must choose (a) which cell to interrogate, and (b) which sensor mode to use for the interrogation. The chosen cell is then interrogated and the measurement is used to update the probability density describing the scenario. This procedure repeats for an entire vignette.

We compare the information driven method, the information driven method between marginalized JMPDs (where only the position is important) and the task driven method as the means of choosing which cell to interrogate via a set of Monte Carlo runs. Motivated by the results of Section IV, we maximize the integral in (10) without taking the logarithm. We look at the performance of the algorithms in terms of two quantities: The tracking error and the identification error.

Figure 1 shows that the scheduler derived specifically to minimize tracking error is the best performing algorithm when only the tracking performance is considered, although only marginally so. The methods based on maximizing expected information gain perform only slightly poorer when considering tracking performance. On the other hand, using identification quality as the metric, the information gain criterion clearly outperforms the other two methods. This is due to the fact that the information gain criteria tries to learn about the target position and class simultaneously and so uses some sensor dwells to identify the target.

An interesting property of the Figure is the oscillatory behavior in tracking error. This property is due to the sensor model used in this investigation, and has nothing to do with the sensor management method used. The sensor model in tracking mode is that the sensor is capable of making measurements as to the presence or absence of a target in a cell. This capability manifests itself in the tracking as follows. When a target first enters a cell, its position is known very well. The fact that a target has transitioned from one cell to the next allows very accurate position prediction. This explains the minimum of the oscillation in the error curves. After a target remains in a cell for some time, its position becomes less certain to the tracker. As the tracker is only being updated with information about the presence of the target in the cell, it will eventually estimate uniform

probability of position over the entire cell. This explains the maximum in the oscillation in the error curve. Eventually, the target transitions to another cell, repeating the process. We see approximately 8 cell transitions in the simulation.

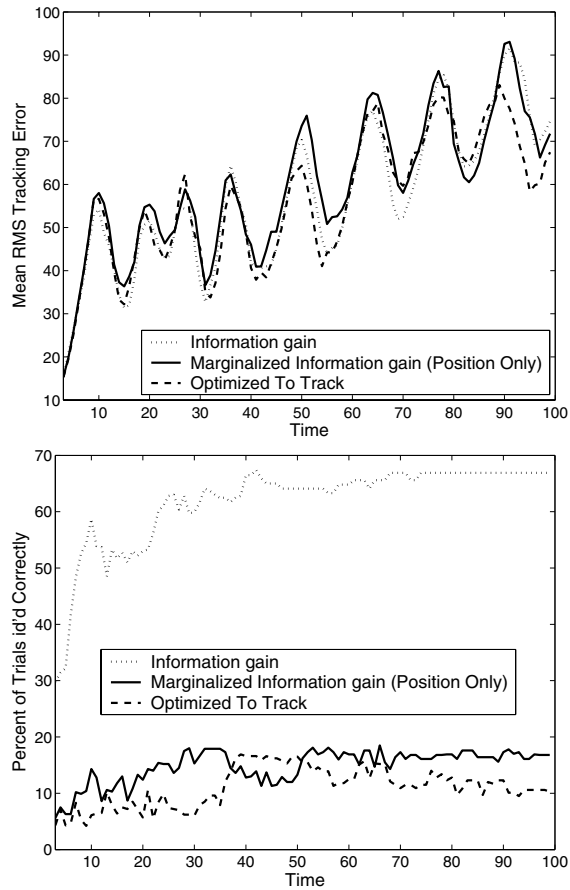


Fig. 1. Comparison between the Rényi divergence scheduler, a Rényi divergence between marginalized JMPDs, and a scheduler designed to minimize tracking error. The divergence criteria is able to do good identification with only a small degradation in tracking performance. As expected, the scheduler designed to minimize track error performs best in terms of tracking error, followed by the divergence between marginalized JMPDs and the divergence. However, in terms of identification performance, the divergence between full state JMPDs is superior.

VII. CONCLUSIONS

This paper has compared target tracking and identification performance under information driven and task driven sensor scheduling algorithms. As expected, when making the performance comparison by looking at the task of interest, the task driven scheduler slightly outperforms the information driven scheduler. However, the power of the information driven approach is evident when there are multiple competing performance criteria. In this case, target classification is used to show that the information driven approach is able to balance the two desires. We see that the information based approach sacrifices a small amount of tracking performance to yield a high level of classification performance.

VIII. ACKNOWLEDGEMENTS

The authors greatly acknowledge their interaction with Professor Jeffrey Fessler and Professor Susan Murphy which provided the impetus to investigate the relative performance of information driven and task driven sensor management.

REFERENCES

- [1] K. J. Hintz and E. S. McVey, "Multi-process constrained estimation," *IEEE Transactions on Man, Systems, and Cybernetics*, vol. 21, no. 1, pp. 434–442, January/February 1991.
- [2] K. J. Hintz, "A measure of the information gain attributable to cueing," *IEEE Transactions on Systems, Man and Cybernetics*, vol. 21, no. 2, pp. 237–244, 1991.
- [3] W. Schmaedeke and K. Kastella, "Event-averaged maximum likelihood estimation and information-based sensor management," *Proceedings of SPIE*, vol. 2232, pp. 91–96, June 1994.
- [4] K. Kastella, "Discrimination gain to optimize classification," *IEEE Transactions on Systems, Man and Cybernetics-Part A: Systems and Humans*, vol. 27, no. 1, January 1997.
- [5] —, "Discrimination gain for sensor management in multitarget detection and tracking," *IEEE-SMC and IMACS Multiconference CESA*, vol. 1, pp. 167–172, 1996.
- [6] R. Mahler, "Global optimal sensor allocation," *Proceedings of the Ninth National Symposium on Sensor Fusion*, vol. 1, pp. 167–172, 1996.
- [7] C. M. Kreucher, K. Kastella, and A. O. Hero III, "Information based sensor management for multitarget tracking," *Proceedings of SPIE Conference on Signal and Data Processing of Small Targets*, 2003.
- [8] F. Zhao, J. Shin, and J. Reich, "Information-driven dynamic sensor collaboration," *IEEE Signal Processing Magazine*, pp. 61–72, March 2002.
- [9] C. M. Kreucher, K. Kastella, and A. O. Hero III, "Tracking multiple targets using a particle filter representation of the joint multitarget probability density," *Proceedings of SPIE Conference on Signal and Data Processing of Small Targets*, 2003.
- [10] L. D. Stone, T. L. Corwin, and C. A. Barlow, *Bayesian Multiple Target Tracking*. Artech House, 1999.
- [11] 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*, vol. 43, no. 11, pp. 2678–2690, 1995.
- [12] R. E. Bethel and G. J. Paras, "A PDF multisensor multitarget tracker," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 34, no. 1, pp. 153–168, January 1998.
- [13] I. Goodman, R. Mahler, and H. Nguyen, *Mathematics of Data Fusion*. Kluwer Academic Publishers, 1997.
- [14] 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.
- [15] A. Rényi, "On measures of entropy and information," *Proceedings of the 4th Berkeley Symposium on Mathematics, Statistics, and Probability*, vol. 1, pp. 547–561, 1961.
- [16] A. O. Hero III, B. Ma, O. Michel, and J. Gorman, "Applications of entropic spanning graphs," *IEEE Signal Processing Magazine (Special Issue on Mathematics in Imaging)*, vol. 19, no. 5, pp. 85–95, 2002.
- [17] 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.
- [18] K. Kastella and S. Musick, "The search for optimal sensor management," *Proceedings of the SPIE Symposium on Aerospace/Defense Sensing and Controls*, vol. 2759, pp. 318–329, April 1996.