

Risk-Based Sensor Resource Management for Field of View Constrained Sensors

Sean Martin
The Johns Hopkins University
Applied Physics Laboratory
Laurel, MD
Email: Sean.Martin@jhuapl.edu

Abstract—This paper introduces a statistical risk based metric for field of view constrained sensor management in a target tracking scenario. This metric is based on a Bayesian estimate of both the target position and the target classification. Due to field of view restrictions it is assumed that more targets exist than the given sensor is capable of tracking simultaneously. It is also assumed that initially all target classifications are unknown and that a cost exists for incorrectly classifying a target track. This cost is higher for certain classes of targets than it is for others. To account for uncertainty in both the kinematic and classification state estimates, the proposed metric treats the cost as a random variable and uses a hierarchical statistical model to calculate the expected value of this cost when conditioned on the event of losing a target track. This metric is then applied to a simulated radar sensor manager to maintain an acceptable level of kinematic accuracy on targets of high cost. It is shown through empirical statistical tests that this sensor manager maintains track on high priority targets significantly better than other common methods.

I. INTRODUCTION

A common problem in target tracking applications is the ability to maintain a good estimate, or track, of a target's true state (i.e. classification, position, and velocity) in the presence of sensor imperfections. This is usually accomplished by maintaining an estimate of the target's state over time using an algorithm such as a Kalman filter [1]. However, it is often the case that sensors are incapable of maintaining a good estimate of all target states in a given area over a period of time due to field of view (FOV) restrictions. If the target's estimated location is poor enough such that its location falls outside this field of view, then no state measurements will be collected on the target. Essentially, the target will be lost.

As an example, consider a pan tilt zoom (PTZ) camera attempting to autonomously locate a person of interest in a crowd of people. Research has found that a larger number of camera pixels on a target substantially improves results during target classification [2]. Thus, the PTZ camera needs to zoom in on potential targets to classify them, at the sacrifice of collecting state measurements on other targets. This creates a field of view constraint and has led to a significant amount of research on how to autonomously manage a PTZ sensor resource to classify and track targets of interest [3] [4] [5] [6].

An analogous sensor management problem exists for radar sensors. Common radar systems use a 'pencil beam' mode with sufficient resolution to both classify and track targets of interest.

However, this mode typically requires a field of view of only several degrees in size and thus also leads to FOV constraints [7]. This has also led to a significant amount of research on how to autonomously manage radar sensors for applications ranging from satellite tracking to ground target tracking [8] [9] [10] [11].

Since both of these cases often involve tracking a large number of targets, a form of autonomous sensor management is typically employed to track as many targets as possible. This usually involves scheduling the sensor to measure the target track estimate with the largest uncertainty in the true target state. To measure the reduction in uncertainty, a metric of information gain is commonly employed. This can take the form of Kullback Liebler divergence, Fisher information gain, or Renyi divergence [12] [10] [9]. Information gain based metrics can be used for measuring both kinematic and classification based uncertainty during sensor management [13]. However, there is no prioritization of target tracks which is essential when only a subset of the total targets can be successfully tracked.

To circumvent the limitations on existing metrics, the idea of using a statistical risk model to calculate an expected cost as a metric has gained ground in recent research. Papageorgiou et. al. have used an expected cost for the problem of missile defense [14]. Wang et. al. have used Bayesian risk during sensor management for micro air vehicles and satellite tracking [8] [9]. DeSena et. al. have used heuristic approximations to an expected cost for surveillance problems [15]. This paper differs from these works in that its focus is on field of view constraints when calculating expected costs.

This paper considers autonomous sensor management for field of view constrained sensors when there are too many highly maneuverable targets for a single sensor to track. It is assumed that only a small number of targets need to be tracked (i.e. persons of interest in a crowd) and that initially all the target kinematic states are known, but their classification states are unknown. It is also assumed that a cost exists for making an incorrect decision on a target's true classification. The key observation in this paper is that the job of a sensor manager is not to make classification decisions on target tracks, nor is it to reduce all of the targets' state estimate uncertainties. Instead, the sensor manager needs to decide what allocation of sensors to what targets will result in the greatest reduction in expected cost of making an incorrect classification decision.

By framing the problem as one of reducing expected cost, the cost can then be conditioned on the event of losing a target track. Thus, the expected cost incorporates both the track kinematic estimate and the track classification estimate. Section II discusses the modeling of these estimates. The calculation of the expected cost and the amount of reduction in this cost from new measurements is discussed in Section III, and experiments showing the effectiveness of this metric compared to others are detailed in Section IV.

II. KINEMATIC AND CLASSIFICATION STATE ESTIMATION

For each target, a track:

$$\vec{X} = \left[\vec{X}_{kinematic} \quad \vec{X}_{classification} \right]$$

is maintained to estimate the true state of the target and provide needed data in calculation of the risk metrics provided in Section III. The specific values of this track \vec{X} are discussed in the following subsections.

A. Kinematic State Estimation

For simplicity, in this paper both the targets and the sensor exist in a two dimensional plane. Thus the true kinematic state consists of a two dimensional position (x and y) and velocity (\dot{x} and \dot{y}):

$$\vec{X}_{kinematic\ truth} = \begin{bmatrix} x \\ \dot{x} \\ y \\ \dot{y} \end{bmatrix}$$

The estimates of these true target states, referred to as target tracks, are modeled by four dimensional Gaussian distributions:

$$\vec{X}_{kinematic} \sim N(\hat{x}_k, \hat{P}_k)$$

Where \hat{x}_k is the mean state estimate and \hat{P}_k is the state estimate error covariance at time k . Both of these values are updated by the standard extended Kalman filter equations, with linearization around the measurement, which are summarized in the appendix. The parameters to these updates are detailed as follows.

Measurements consist of an angle $\theta = \arctan\left(\frac{y}{x}\right)$ and a range $r = \sqrt{x^2 + y^2}$:

$$\vec{Z}_{truth} = \begin{bmatrix} \theta \\ r \end{bmatrix}$$

To avoid singularities in the linearization process, measurements are converted to Cartesian coordinates and the linearization is performed on the noise [16]. This is reflected in the construction of the measurement error covariance matrix R as shown below. Measurements are thus represented as:

$$\vec{Z}_{truth}^{Cartesian} = \begin{bmatrix} r \cos(\theta) \\ r \sin(\theta) \end{bmatrix}$$

The linear state estimate to measurement transformation H is given by:

$$H = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}$$

Note that this simply extracts the x and y position estimates, since the sensor does not measure angular or range rates. Thus, the measurement error covariance matrix R using the linearization in [16] is given by:

$$R = \begin{pmatrix} r^2 \sigma_\theta^2 \sin^2 \theta + \sigma_r^2 \cos^2 \theta & r^2 \sigma_\theta^2 \cos^2 \theta + \sigma_r^2 \sin^2 \theta \\ r^2 \sigma_\theta^2 \cos^2 \theta + \sigma_r^2 \sin^2 \theta & (\sigma_r^2 - r^2 \sigma_\theta^2) \sin \theta \cos \theta \end{pmatrix}$$

Where $var(r) = \sigma_r^2$ and $var(\theta) = \sigma_\theta^2$. In this paper, $\sigma_r^2 = 1$ squared meter and $\sigma_\theta^2 = 3.0462e - 04$ squared radians (which corresponds to a standard deviation of 1 degree).

The process noise covariance matrix Q for all target tracks is given by the first order derivation (i.e. position and velocity) as derived in [17]:

$$Q = \Phi_s \begin{pmatrix} \frac{T_s^3}{2} & \frac{T_s^2}{2} & 0 & 0 \\ \frac{T_s^3}{2} & T_s & 0 & 0 \\ 0 & 0 & \frac{T_s^3}{2} & \frac{T_s^2}{2} \\ 0 & 0 & \frac{T_s^2}{2} & T_s \end{pmatrix}$$

Where $T_s = 1$ and $\Phi_s = 5$ in this paper.

Finally, the state transition model F is a basic constant velocity model:

$$F = \begin{pmatrix} 1 & dt & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & dt \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

B. Classification State Estimation

The next estimate is the classification state. The true classification is represented by the categorical random variable J with support $\{j \mid j \in [1, n]\}$ for n total possible classifications. Thus:

$$\vec{X}_{classification} = \begin{bmatrix} P(J = 1) \\ \vdots \\ P(J = n) \end{bmatrix}$$

the classification measurement is then represented by the discrete random variable M with support $\{m \mid m \in [1, n]\}$.

The classification probability is updated through direct application of Bayes' Theorem:

$$\begin{aligned} P'(J = i) &\equiv P(J = i | M = m) \\ &= \frac{P(M = m | J = i) P(J = i)}{P(M = m)} \\ &= \frac{P(M = m | J = i) P(J = i)}{\sum_{r=1}^n P(M = m | J = r) P(J = r)} \end{aligned}$$

Where P' indicates the posterior probability. Typically the measurement likelihoods $P(M = m | J = i)$ are represented by a normalized confusion matrix. This paper uses a normalized confusion matrix CC with the columns representing truth, the rows representing the classifier output, and the sum of

each column being 1 by definition of a conditional probability distribution:

$$CC = \begin{matrix} & \begin{matrix} 1 & \dots & n \end{matrix} \\ \begin{matrix} 1 \\ \vdots \\ n \end{matrix} & \begin{pmatrix} P(M=1|J=1) & \dots & P(M=1|J=n) \\ \vdots & \ddots & \vdots \\ P(M=n|J=1) & \dots & P(M=n|J=n) \end{pmatrix} \end{matrix}$$

III. STATISTICAL RISK

A. Type I Error Cost

As mentioned in the previous Section, there is a decision to be made as to what a target's classification is. If an incorrect decision is made, this results in a cost (i.e. lost sensor resources, lost coverage of a criminal act, etc.). In this paper, an incorrect decision is treated as falsely rejecting a null hypothesis in the context of statistical testing. Thus, the costs for making a wrong decision correspond to the type I error during statistical hypothesis testing. This is an approach that has been used before in economics when deciding whether or not to ban various genetically modified crops [18][19].

The matrix CM_1 as defined below contains the cost of committing a type I error when making a decision on a target's classification. In this cost matrix, the rows designate the decision and the columns designate the true classification. The diagonals are all zero since there is no cost for making the correct decision. Also note that the null hypothesis is the column and the alternate hypothesis is the row (decision is a false reject). Thus, with the exception of the diagonal entry of zero, each column entry is identical since they all represent the type I error cost:

$$CM_1 = \begin{matrix} & \begin{matrix} 1 & 2 & \dots & n \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ \vdots \\ n \end{matrix} & \begin{pmatrix} 0 & c_{12} & \dots & c_{1n} \\ c_{11} & 0 & \dots & c_{1n} \\ \vdots & \vdots & \ddots & \vdots \\ c_{11} & c_{12} & \dots & 0 \end{pmatrix} \end{matrix}$$

As an example, consider the decision that a target of true classification 1 is a different classification. This corresponds to the first column and any row other than the first row in CM_1 . The corresponding test with the decision underlined is:

$$H_0 : \text{Target is } \underline{1} \quad H_a : \text{Target is not } \underline{1}$$

A decision of $\neg 1$ falsely rejects H_0 resulting in a cost of c_{11} .

B. The Expected Cost of Committing a Type I Error

During the process of target tracking, several additional random variables (listed below) will influence the expected cost of making a decision. These variables are used in the following theorem, which is proved in the appendix.

C_1 := a discrete random variable representing the type I error cost. Support is the cost matrix CM_1 with entries $\{c_{1ij}\}$ defined below. The event space consists of all ways to make a

type I error.

c_{1ij} := Entry of an $n \times n$ dimensional cost matrix CM_1 incurred when a decision falsely rejects H_0 resulting in a type I error. Each row i of this matrix represents a decision on the classification. Each column j of this matrix represents the true target classification.

J := Categorical random variable denoting the actual classification of an object. Support is $\{j \mid j \in [1, n]\}$.

I := Categorical random variable denoting the decision on classification of an object. Support is $\{i \mid i \in [1, n]\}$.

I_{new} := Discrete uniformly distributed random variable denoting the classification decision on a reacquired object after its original track was lost. Support is $\{i \mid i \in [1, n]\}$.

L := Bernoulli random variable denoting whether or not the actual target is lost (i.e. outside the sensor's field of view when the sensor is centered on the track position component of the state estimate \hat{x}_k). The event space is $\{True, False\}$.

P_{lost} := The probability of the actual target being lost. This probability is assumed to be the portion of a multivariate normal distribution $N(\hat{x}_k, \hat{P}_k)$ not contained in the sensor field of view when the sensor aim-point is centered on $\hat{X}_{kinematic}$.

THEOREM 1:

The expected cost for making a type I error when deciding on the classification of a target track is given by:

$$\begin{aligned} E_{C_1}(C_1|I=i) &= \sum_{r \in I} c_{ri} P(J=i) P_{lost} \frac{n-1}{n} \forall r \neq i \\ &+ \sum_{r \in J} c_{ir} P(J=r) P_{lost} \frac{1}{n} \forall r \neq i \\ &+ \sum_{r \in J} c_{ir} P(J=r) (1 - P_{lost}) \forall r \neq i \end{aligned}$$

Proof follows by the law of total expectation as detailed in the appendix.

C. The Expected Risk Reduction

In calculating the expected risk reduction (ERR) the goal is always to maximize the reduction in risk. As mentioned in [14], risk always decreases with the addition of new measurements. In the context of the expected cost calculated in Theorem 1, these measurements reduce the probability of the target being lost or mis-classified. Assuming a sensor can only measure a single track at a time, the track with the *max* ERR is chosen for measurement.

When deciding what a track classification is, it is assumed that the decision of least cost is made. Thus, the *min* expected cost is chosen among all possible decisions for each track classification. Following the approach presented in [14], this

min cost is then used when calculating ERR. Start by denoting the min cost R before a measurement update as:

$$R_i \equiv E_{C_1} (C_1 | I = i)$$

$$R = \min_i \{R_i\}$$

Note that the probabilities in these expectations will change as measurements are accumulated by sensors being controlled by the sensor manager. It is assumed that probabilities change through a Bayesian update. In particular, both P_{lost} and $P(J = i)$ are updated by kinematic and classification measurements respectively. The posterior probabilities are denoted by P'_{lost} and $P'(J = i)$. It is assumed that P'_{lost} is the result of a Kalman filter covariance update using the measurement error covariance of the kinematic sensor. The classification probability is updated through direct application of Bayes' Theorem as discussed in Section II-B.

By Theorem 1 the resulting risk using these updated probabilities is:

$$R' = \min_i \{R'_i\}$$

$$= \min_i \left\{ \begin{aligned} & \sum_{r \in J} c_{ri} P'(J=i) P'_{lost} \frac{n-1}{n} \forall r \neq i \\ & \sum_{r \in J} c_{ir} P'(J=r) P'_{lost} \frac{1}{n} \forall r \neq i \\ & \sum_{r \in J} c_{ir} P'(J=r) (1 - P'_{lost}) \forall r \neq i \end{aligned} \right\}$$

$$= \min_i \left\{ \begin{aligned} & \sum_{r \in J} c_{ri} \frac{P(M=m|J=i)P(J=i)}{P(M=m)} P'_{lost} \frac{n-1}{n} \forall r \neq i \\ & \sum_{r \in J} c_{ir} \frac{P(M=m|J=r)P(J=r)}{P(M=m)} P'_{lost} \frac{1}{n} \forall r \neq i \\ & \sum_{r \in J} c_{ir} \frac{P(M=m|J=r)P(J=r)}{P(M=m)} (1 - P'_{lost}) \forall r \neq i \end{aligned} \right\}$$

Since any classification measurement is possible, it is necessary to take an additional expectation $\langle R' \rangle$ over all possible measurements. This is calculated over discrete classification measurements by:

$$\langle R' \rangle$$

$$= \sum_{m \in M} \min_i \{R'_i\} P(M=m)$$

$$= \sum_{m \in M} \min_i \left\{ \begin{aligned} & \sum_{r \in J} c_{ri} P(M=m|J=i) P(J=i) P'_{lost} \frac{n-1}{n} \forall r \neq i \\ & \sum_{r \in J} c_{ir} P(M=m|J=r) P(J=r) P'_{lost} \frac{1}{n} \forall r \neq i \\ & \sum_{r \in J} c_{ir} P(M=m|J=r) P(J=r) (1 - P'_{lost}) \forall r \neq i \end{aligned} \right\}$$

Note that this formulation of expected cost is guaranteed to decrease in value with the addition of measurements [14]. Thus, the expected risk reduction ΔR is given by:

$$\Delta R = R - \langle R' \rangle \quad (1)$$

The sensing action that yields the greatest reduction in this ERR value is chosen to take the actual measurement.

IV. EXPERIMENTS

It is worth re-iterating that all the metrics discussed up to this point depend on expected values and probabilities of various events occurring. Thus, there will always be situations

where targets will be mis-classified or they will maneuver in an unexpected way such that tracks on important targets are lost. The value of a sensor management algorithm is therefore better reflected by either average or median performance over time. For this reason, this section examines several tests over Monte Carlo runs of measurements generated from the ground truth for a single scenario consisting of 200 time steps of one second each. This scenario involves 10 maneuvering ground targets and 1 sensor. Of these 10 targets, 3 are targets of interest (targets 1, 3, and 10) and should be tracked. The ground truth over the entire 200 second scenario is shown in Figure 1. A single sensor capable of measuring: bearing, range, and target classification is located at position $x = 0$ and $y = 0$. The sensor field of view (FOV) is a 500 meter square region centered on the track position estimate location. If the target ground truth corresponding to the track position is outside this field of view, i.e. the state estimate is very poor, then the track is considered lost and no further measurements are generated on it. The sensor takes a measurement only on even time steps (i.e. every 2 seconds). Thus, every two seconds a sensor resource management algorithm based on the risk metrics introduced in previous sections runs to decide which target track estimate to measure using the sensor.

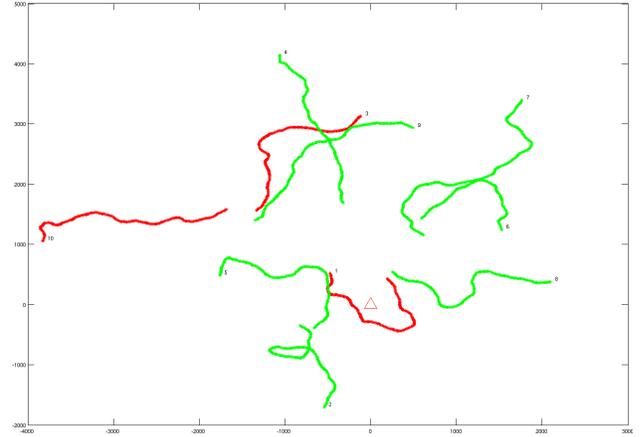


Fig. 1: Position and classification ground truth for targets of interest (in red), targets not of interest (in green), and the sensor location (red triangle). The numeric target identifiers are shown at the target position at the scenario end (200 seconds).

Each target starts with a high accuracy kinematic track, however, the classifications of all targets are unknown at the start of the scenario. Thus, the task of the single sensor is to: correctly classify, maintain track on, and allocate measurements to the targets of interest. As experiments will show, a single sensor cannot maintain track on all 10 maneuvering targets, thus the problem of allocating the sensor is non-trivial. The software used in this section is available at [20].

A. Two Classifications

In this section a binary classification state is considered where either the target being tracked is a target of interest

($J = 1$) or not a target of interest ($J = 2$). Thus:

$$\vec{X}_{classification} = \begin{bmatrix} P(J = 1) \\ P(J = 2) \end{bmatrix}$$

and the binary classification measurement M has support $\{m \mid m \in [1, 2]\}$. The confusion matrix CC and cost matrix CM_1 are:

$$CC = \frac{1}{2} \begin{pmatrix} .8 & .2 \\ .2 & .8 \end{pmatrix} \quad CM_1 = \frac{1}{2} \begin{pmatrix} 0 & 1 \\ 30 & 0 \end{pmatrix}$$

Figure 2 shows the scenario described at the beginning of this section. The sensor position is denoted by a red triangle. The ground truth for the targets of interest (1,3, and 10) is shown by the red trails. The green trails show ground truth for non targets of interest. The track position state estimate is shown by the diamonds. The color of the diamonds reflects the classification estimate. The color is black if $P(J = 1) = .5$, red if $P(J = 1) > .5$, and green if $P(J = 1) < .5$. Ellipses indicate a 95% confidence region for the location of the corresponding target. A line is drawn from the sensor to the track location showing line of sight, and a green box shows the sensor's field of view.

For comparison purposes, the performance of a sensor manager based on the expected risk reduction metric introduced in Section III-C was compared to several other common methods for sensor management. In particular: random assignment, round robin assignment where targets are repeatedly selected in order, and choosing the target that maximizes the Kullback-Leibler divergence in kinematic accuracy before and after a measurement. Fifty Monte Carlo runs were conducted using each method. The track error was calculated between estimated position and ground truth for all tracks at the end of each run and saved for analysis. Table I shows the resulting median error on each target for each sensor scheduling method. From this table it is evident that only the sensor manager using ERR maintains track on targets 1, 3, and 10 with error less than the sensor FOV radius (250m). The random method performs very poorly, the round robin method has poor performance on target 3, and the KL divergence method has poor performance on target 1. Observe from the ground truth shown in Figure 1 that both targets 1 and 3 have sharp turns which likely result in them becoming lost and consequently having large errors. When using ERR, more measurements can be assigned to these targets since they are targets of interest and thus tracking performance improves at the sacrifice of tracks on non targets of interest. To examine these results more rigorously, it is necessary to show with high confidence that the average track error on targets 1, 3, and 10 is less than 250 meters and thus still within the sensor field of view. For each analysis below, the highest and lowest 5% of the error measurements were discarded to remove outliers. If the average error on any of targets 1, 3, or 10 is greater than 250 meters then conclude that the scheduling method was ineffective on average. Performing a single sample one sided Student t-test (H_0 : true mean error is ≤ 250 meters)

Target	Random	Round Robin	KL Divergence	ERR
1	6535.93	10.32	7489.81	13.85
2	5992.96	56.51	40.28	1964.07
3	5664.38	4752.66	16.48	8.29
4	67.22	69.63	18.16	3427.44
5	1662.32	52.50	65.93	1609.45
6	3151.88	83.65	94.22	1662.93
7	44.39	36.03	29.71	32.79
8	1278.00	36.25	63.84	94.61
9	3559.79	23.97	58.72	5136.31
10	387.56	21.51	10.48	16.22

TABLE I: Median position error in meters per target for each metric after 50 Monte Carlo runs (targets of interest are 1, 3, 10)

Metric	Target 1	Target 3	Target 10
Random	1.65e-10	3.464e-11	1.673e-07
Round Robin	1	2.2e-16	1
KL Divergence	2.2e-16	1	1
ERR	0.0988	0.1438	0.07948

TABLE II: p-values for H_0 : residual error ≤ 250 m

yields the p values in Table II. Based on the results in this table observe that the scheduling method based on ERR is the only method that effectively maintains tracks on all of targets 1, 3, and 10 with 95% confidence that H_0 is not rejected. Additional tests were conducted with just the ERR metric to examine the median error under 100 Monte Carlo runs. Table III shows the first three quartiles of error values when targets 1, 3, and 10 are targets of interest, while Table IV shows error quartiles when targets 4, 6, and 8 are targets of interest. In both cases the targets of interest all have track error less than 250m even at the upper quartile.

B. Three Classifications

In this section a tertiary classification state is considered where either the target being tracked is a target of high interest ($J = 1$) or medium interest ($J = 2$) or low interest ($J = 3$). Thus:

$$\vec{X}_{classification} = \begin{bmatrix} P(J = 1) \\ P(J = 2) \\ P(J = 3) \end{bmatrix}$$

and the classification measurement M has support $\{m \mid m \in [1, 2, 3]\}$. The confusion matrix CC and cost matrix

Target	Lower Quartile	Median	Upper Quartile
1	7.18	13.28	28.72
2	1951.48	1962.23	1973.26
3	7.96	8.35	9.25
4	885.66	3417.70	3452.57
5	1605.71	1614.14	2332.90
6	1101.07	1649.18	1668.15
7	22.02	393.58	2114.64
8	52.44	594.52	1706.32
9	5136.31	5136.31	5136.31
10	11.45	15.27	39.58

TABLE III: First, second, and third quartiles for position error in meters per target (targets of interest are: 1, 3, 10) after 100 Monte Carlo runs

Target	Lower Quartile	Median	Upper Quartile
1	7483.12	7488.36	7494.21
2	1946.15	1961.79	1971.43
3	2597.55	2627.39	2649.93
4	26.18	28.73	40.18
5	1605.11	1611.91	1619.73
6	11.66	17.32	51.33
7	972.91	2115.00	2132.47
8	9.84	14.96	24.67
9	6.45	9.40	5136.31
10	115.86	573.32	1099.58

TABLE IV: First, second, and third quartiles for position error in meters per target (targets of interest are: 4, 6, 8) after 100 Monte Carlo runs

Target	Lower Quartile	Median	Upper Quartile
1	9.67	48.36	101.08
2	1953.16	1963.16	1971.08
3	8.18	8.81	22.32
4	18.00	28.94	66.53
5	1605.20	1613.06	1618.07
6	1635.69	1653.45	1666.71
7	33.81	939.76	2096.89
8	299.36	1391.61	1900.84
9	5136.31	5136.31	5136.31
10	10.56	14.77	25.76

TABLE V: First, second, and third quartiles for position error in meters per target (the high priority target is 10, medium priority targets are: 1, 3, 4) after 100 Monte Carlo runs

CM_1 are:

$$CC = \begin{matrix} & \begin{matrix} 1 & 2 & 3 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{pmatrix} .8 & .1 & .1 \\ .1 & .8 & .1 \\ .1 & .1 & .8 \end{pmatrix} \end{matrix} \quad CM_1 = \begin{matrix} & \begin{matrix} 1 & 2 & 3 \end{matrix} \\ \begin{matrix} 1 \\ 3 \\ 3 \end{matrix} & \begin{pmatrix} 0 & 20 & 1 \\ 30 & 0 & 1 \\ 30 & 20 & 0 \end{pmatrix} \end{matrix}$$

Since ERR is the only metric sensitive to prioritized targets, the results for the other metrics would be unchanged in this section. Thus, only the median error test with 100 Monte Carlo runs is performed with results in Table V. As seen in this table, both the high priority targets and the medium priority targets are tracked with error less than 250m even at the upper quartile.

V. CONCLUSION

This paper introduced a metric based on the expected cost (or risk) of making an incorrect decision on a target's classification under uncertainty. The metric was then conditioned on the event of losing a target track which allowed for the combination of classification and kinematic uncertainty in the same metric. Further, this metric was then applied in the context of a sensor resource manager which takes field of view constrained sensor actions based on the maximum expected risk reduction. It has been shown empirically that unlike other common methods of sensor management, a sensor manager based on maximizing ERR can maintain track on targets of interest when it is not possible for a single sensor to track all targets in the environment.

APPENDIX A PROOF OF THEOREM 1

Using the law of iterated expectation for each random variable on which the cost depends, observe:

$$\begin{aligned} & E_{C_1}(C_1|I=i) \\ &= E_L(E_{C_1}(C_1|I=i,L)) \\ &= E_J(E_L(E_{C_1}(C_1|I=i,L,J))) \\ &= E_J(E_{C_1}(C_1|I=i,L=True,J))P(L=True) \\ &\quad + E_J(E_{C_1}(C_1|I=i,L=False,J))P(L=False) \\ &= E_{C_1}(C_1|I=i,L=True,J=i)P(L=True)P(J=i) \\ &\quad + E_{C_1}(C_1|I=i,L=False,J=i)P(L=False)P(J=i) \\ &\quad + E_{C_1}(C_1|I=i,L=True,J \neq i)P(L=True)P(J \neq i) \\ &\quad + E_{C_1}(C_1|I=i,L=False,J \neq i)P(L=False)P(J \neq i) \\ &= E_{C_1}(C_1|I=i,L=True,J=i)P(L=True)P(J=i) \\ &\quad + 0 \\ &\quad + E_{C_1}(C_1|I=i,L=True,J \neq i)P(L=True)P(J \neq i) \\ &\quad + E_{C_1}(C_1|I=i,L=False,J \neq i)P(L=False)P(J \neq i) \end{aligned}$$

Note that for the second term in the above summation, the correct decision is made and the target is not lost. Thus, the cost is zero since it falls on the diagonal of the cost matrix CM_1 . Now at this point, if a target was lost, it is necessary to consider the case in which the target is re-acquired. Note that regardless of the classifier accuracy, it is possible that the target acquired is not the original target. Thus, I_{new} is modeled as a discrete uniform distribution. Also note that depending on the density of targets of interest in a given scenario, this may be a different distribution.

$$\begin{aligned} & E_{C_1}(C_1|I=i) \\ &= E_{I_{new}}(E_{C_1}(C_1|I=i,L=True,J=i,I_{new}))P(L=True)P(J=i) \\ &\quad + E_{I_{new}}(E_{C_1}(C_1|I=i,L=True,J \neq i,I_{new}))P(L=True)P(J \neq i) \\ &\quad + E_{C_1}(C_1|I=i,L=False,J \neq i)P(L=False)P(J \neq i) \\ &= E_{C_1}(C_1|I=i,L=True,J=i,I_{new}=i)P(L=True)P(J=i)P(I_{new}=i) \\ &\quad + E_{C_1}(C_1|I=i,L=True,J=i,I_{new} \neq i)P(L=True)P(J=i)P(I_{new} \neq i) \\ &\quad + E_{C_1}(C_1|I=i,L=True,J \neq i,I_{new}=i)P(L=True)P(J \neq i)P(I_{new}=i) \\ &\quad + E_{C_1}(C_1|I=i,L=True,J \neq i,I_{new} \neq i)P(L=True)P(J \neq i)P(I_{new} \neq i) \\ &\quad + E_{C_1}(C_1|I=i,L=False,J \neq i)P(L=False)P(J \neq i) \\ &= 0 \\ &\quad + E_{C_1}(C_1|I=i,L=True,J=i,I_{new} \neq i)P(L=True)P(J=i)P(I_{new} \neq i) \\ &\quad + E_{C_1}(C_1|I=i,L=True,J \neq i,I_{new}=i)P(L=True)P(J \neq i)P(I_{new}=i) \\ &\quad + 0 \\ &\quad + E_{C_1}(C_1|I=i,L=False,J \neq i)P(L=False)P(J \neq i) \end{aligned}$$

Observe that the first term in the above summation still has the correct decision even though the track was lost and later re-acquired. Thus, there is no cost. Also note that in the fourth term of the summation, $I_{new} \neq i$ and $J \neq i$. Thus, there is no cost with regard to the initial decision of $I = i$ prior to the track being lost. This leaves three terms to the summation. The

first term is the cost incurred if the target is lost and a decision on its classification is wrong after the target is re-acquired. The second term is the cost for a decision on the target classification never being correct even after being lost. The last term is the cost for a target classification decision being wrong even though the target is never lost. These terms are related to specific rows and columns of the cost matrix CM_1 as follows:

$$\begin{aligned}
& E_{C_1}(C_1|I=i) \\
&= E_{C_1}(C_1|I=i, L=True, J=i, I_{new} \neq i) P(L=True) P(J=i) P(I_{new} \neq i) \\
&\quad + E_{C_1}(C_1|I=i, L=True, J \neq i, I_{new} = i) P(L=True) P(J \neq i) P(I_{new} = i) \\
&\quad + E_{C_1}(C_1|I=i, L=False, J \neq i) P(L=False) P(J \neq i) \\
&= \sum_{c_{1ij}} c_{1ij} P(C_1=c_{1ij} | I=i, L=True, J=i, I_{new} \neq i) P(L=True) P(J=i) P(I_{new} \neq i) \\
&\quad + \sum_{c_{1ij}} c_{1ij} P(C_1=c_{1ij} | I=i, L=True, J \neq i, I_{new} = i) P(L=True) P(J \neq i) P(I_{new} = i) \\
&\quad + \sum_{c_{1ij}} c_{1ij} P(C_1=c_{1ij} | I=i, L=False, J \neq i) P(L=False) P(J \neq i) \\
&= \sum_{r \in I} c_{1_{ri}} P(J=i) P(L=True) P(I_{new} \neq i) \forall r \neq i \\
&\quad + \sum_{r \in J} c_{1_{ir}} P(J=r) P(L=True) P(I_{new} = i) \forall r \neq i \\
&\quad + \sum_{r \in J} c_{1_{ir}} P(J=r) P(L=False) \forall r \neq i
\end{aligned}$$

Note that the first term in this summation is a function of the rows of the cost matrix over column $J = i$. This reflects an incorrect decision after the target was re-acquired. Finally, plugging in the assumption of I_{new} being uniformly distributed and the previously mentioned calculation of P_{lost} :

$$\begin{aligned}
& E_{C_1}(C_1|I=i) \\
&= \sum_{r \in I} c_{1_{ri}} P(J=i) P(L=True) P(I_{new} \neq i) \forall r \neq i \\
&\quad + \sum_{r \in J} c_{1_{ir}} P(J=r) P(L=True) P(I_{new} = i) \forall r \neq i \\
&\quad + \sum_{r \in J} c_{1_{ir}} P(J=r) P(L=False) \forall r \neq i \\
&= \sum_{r \in I} c_{1_{ri}} P(J=i) P_{lost} \frac{n-1}{n} \forall r \neq i \\
&\quad + \sum_{r \in J} c_{1_{ir}} P(J=r) P_{lost} \frac{1}{n} \forall r \neq i \\
&\quad + \sum_{r \in J} c_{1_{ir}} P(J=r) (1 - P_{lost}) \forall r \neq i
\end{aligned}$$

APPENDIX B KALMAN FILTER EQUATIONS

The equations for the recursive Kalman filter estimators are summarized below. The notation $\hat{x}_{n|m}$ denotes the estimate of x at time n given measurements up to and including time $m \leq n$. Prediction for the state estimate $\hat{x}_{k|k-1}$ and state estimate error covariance $\hat{P}_{k|k-1}$ are given by:

$$\begin{aligned}
\hat{x}_{k|k-1} &= F_k \hat{x}_{k-1|k-1} \\
\hat{P}_{k|k-1} &= F_k \hat{P}_{k-1|k-1} F_k^T + Q_k
\end{aligned}$$

An update to these estimates given the measurement z_k at time k is given by:

$$\begin{aligned}
\tilde{y}_k &= z_k - H_k \hat{x}_{k|k-1} \\
S_k &= H_k \hat{P}_{k|k-1} H_k^T + R_k \\
K_k &= \hat{P}_{k|k-1} H_k^T S_k^{-1} \\
\hat{x}_{k|k} &= \hat{x}_{k|k-1} + K_k \tilde{y}_k \\
\hat{P}_{k|k} &= (I - K_k H_k) \hat{P}_{k|k-1}
\end{aligned}$$

Where: F_k is the state transition model, H_k is the measurement model, Q_k is the process noise covariance, \tilde{y}_k is the measurement residual, S_k is the residual covariance, R_k is the measurement noise covariance, and K_k is the optimal Kalman gain.

ACKNOWLEDGMENT

The author would like to thank Anshu Saksena, Christopher Boswell, Michael Williams, Jonathan DeSena, and Andrew Newman for highly constructive input on this work. This work was supported by an internal research and development grant at the Johns Hopkins University Applied Physics Laboratory.

REFERENCES

- [1] R. Kalman, "A new approach to linear filtering and prediction problems," *Journal of basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.
- [2] P. Phillips, W. Scruggs, A. OToole, P. Flynn, K. Bowyer, C. Schott, and M. Sharpe, "Frvt 2006 and ice 2006 large-scale results," Tech. Rep. 7408, National Institute of Standards and Technology, 2007.
- [3] E. Sommerlade and I. Reid, "Information-theoretic active scene exploration," in *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pp. 1–7, IEEE, 2008.
- [4] Y. Cai, G. Medioni, and T. Dinh, "Towards a practical ptz face detection and tracking system," in *Applications of Computer Vision (WACV), 2013 IEEE Workshop on*, pp. 31–38, IEEE, 2013.
- [5] S. Lim, L. Davis, and A. Elgammal, "Scalable image-based multi-camera visual surveillance system," in *Advanced Video and Signal Based Surveillance, 2003. Proceedings. IEEE Conference on*, pp. 205–212, IEEE, 2003.
- [6] C. J. Costello, C. P. Diehl, A. Banerjee, and H. Fisher, "Scheduling an active camera to observe people," in *Proceedings of the ACM 2nd international workshop on Video surveillance & sensor networks*, pp. 39–45, ACM, 2004.
- [7] P. Lacomme, J. Marchais, J. Hardange, and E. Normant, *Air and space-borne radar systems: An introduction*. William Andrew, 2001.
- [8] Y. Wang and I. Hussein, *Risk-Based Sequential Decision-Making Strategy*, vol. 427 of *Lecture Notes in Control and Information Sciences*. Springer London, 2012.
- [9] Y. Wang, I. Hussein, and R. Erwin, "Risk-based sensor management for integrated detection and estimation," *Journal of Guidance, Control, and Dynamics*, vol. 34, no. 6, pp. 1767–1778, 2011.
- [10] A. Sinha, T. Kirubarajan, and Y. Bar-Shalom, "Autonomous ground target tracking by multiple cooperative uavs," in *Aerospace Conference, 2005 IEEE*, pp. 1–9, IEEE, 2005.
- [11] P. Williams, D. Spencer, and R. Erwin, "Coupling of estimation and sensor tasking applied to satellite tracking," *Journal of Guidance, Control, and Dynamics*, vol. 36, no. 4, pp. 993–1007, 2013.
- [12] S. Kullback and R. Leibler, "On information and sufficiency," *The Annals of Mathematical Statistics*, pp. 79–86, 1951.
- [13] K. Kastella, "Discrimination gain to optimize detection and classification," *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, vol. 27, no. 1, pp. 112–116, 1997.
- [14] D. Papageorgiou and M. Raykin, "A risk-based approach to sensor resource management," *Advances in Cooperative Control and Optimization*, vol. 369, pp. 129–144, 2007.

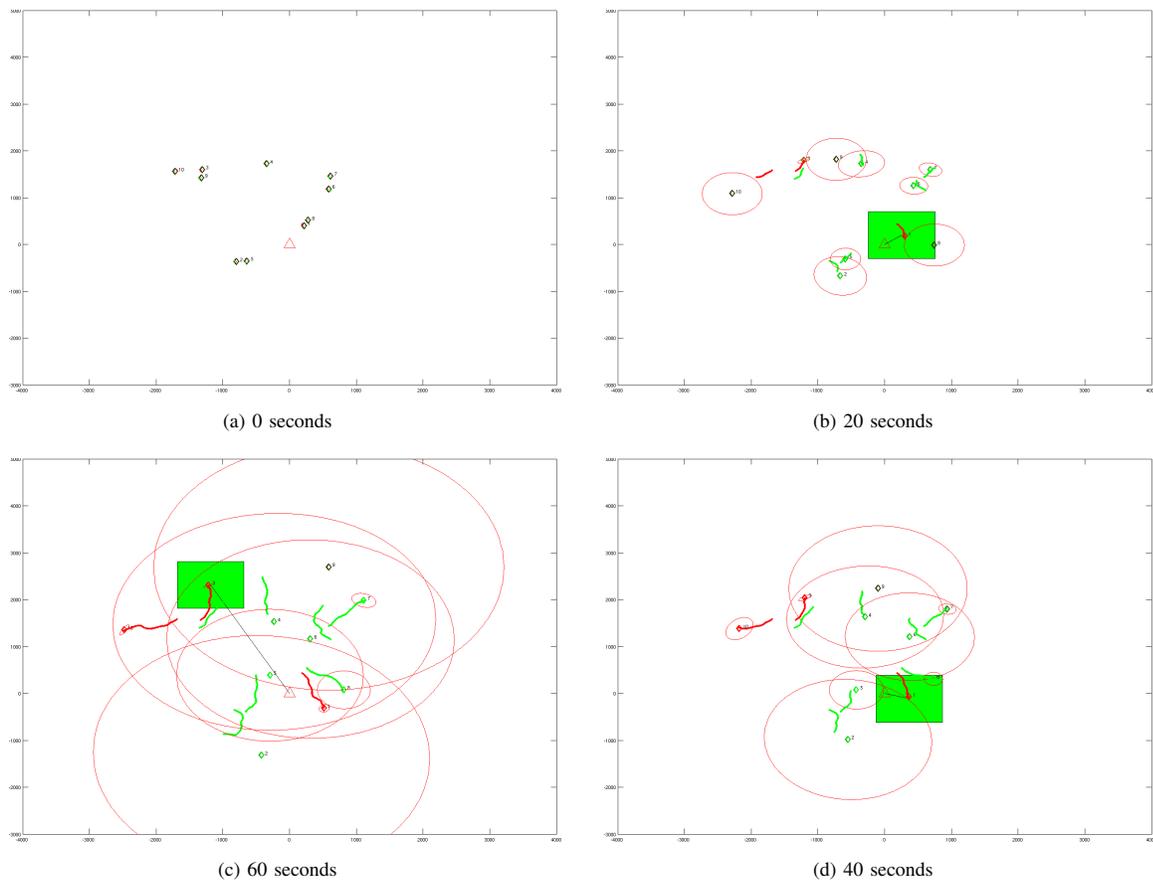


Fig. 2: Selected time steps for a scenario containing a sensor manager making decisions that maximize the ERR for type I error. The diamonds indicate track position state estimate means and are color coded to represent unknown class (black), targets not of interest (green), and targets of interest (red). The red ellipses represent 95% confidence ellipsoids on target positions. The green box is the sensor's field of view and the sensor is located at the red triangle. Scenario times are shown in increasing order clockwise from the upper left. Observe that by 40 seconds the targets of interest have been identified and by 60 seconds the position uncertainty of the targets of interest have been substantially reduced while most of the targets that are not of interest are ignored.

[15] J. DeSena, S. Martin, J. Clarke, D. Dutrow, and A. Newman, "Dynamic optimization of isr sensors using a risk-based reward function applied to ground and space surveillance scenarios," in *SPIE Defense, Security, and Sensing*, pp. 83920B–83920B, International Society for Optics and Photonics, 2012.

[16] Y. Bar-Shalom and X. Li, *Multitarget-Multisensor Tracking: Principles and Techniques*. YBS Publishing, 1995.

[17] P. Zarchan and H. Musoff, *Fundamentals of Kalman filtering: a practical approach*, vol. 208. AIAA, 2005.

[18] D. Hennessy and M. GianCarlo, "Regulatory actions under adjustment costs and the resolution of scientific uncertainty," *American Journal of Agricultural Economics*, vol. 88, no. 2, pp. 308–323, 2006.

[19] E. Ansink and J. Wesseler, "Quantifying type i and type ii errors in decision-making under uncertainty: the case of gm crops," *Letters in Spatial and Resource Sciences*, vol. 2, no. 1, pp. 61–66, 2009.

[20] "Risk-based sensor resource management for field of view constrained sensors." <https://www.mathworks.com/matlabcentral/fileexchange/50996-risk-based-sensor-resource-management>. 2015.