A Random Finite Set Sensor Control Approach for Vision-based Multi-object Search-While-Tracking

Keith A. LeGrand, Pingping Zhu, and Silvia Ferrari

***Abstract*—Through automatic control, intelligent sensors can be manipulated to obtain the most informative measurements about objects in their environment. In object tracking applications, sensor actions are chosen based on the predicted improvement in estimation accuracy, or information gain. Although random finite set theory provides a formalism for measuring information gain for multi-object tracking problems, predicting the information gain remains computationally challenging. This paper presents a new tractable approximation of the random finite set expected information gain applicable to multi-object search and tracking. The approximation presented in this paper accounts for noisy measurements, missed detections, false alarms, and object appearance/disappearance. The effectiveness of the approach is demonstrated through a ground vehicle tracking problem using real video data from a remote optical sensor.**

***Index Terms*—sensor control, information gain, multi-object tracking, random finite set, cell multi-Bernoulli, bounded field-of-view, Kullback-Leibler divergence**

I. Introduction

MANY modern multi-object tracking applications involve mobile and reconfigurable sensors able to control the position and orientation of their field-of-view (FoV) in order to expand their operational tracking capacity and improve state estimation accuracy when compared to fixed sensor systems. By incorporating active sensor control in these dynamic tracking systems, the sensor can autonomously make decisions that produce observations with the highest information content based on prior knowledge and sensor measurements [1]–[3]. However, as the sensor FoV moves and covers extensive regions of interest, potentially for prolonged periods of time, several difficulties are introduced. The number of objects inside the FoV changes over time and is unknown *a priori*, as are the individual object states, which may also be time-varying and subject to significant measurement errors. As a result, existing tracking algorithms and information gain functions, such as those in [1]–[3], which assume a known number of objects and known data association, are either inapplicable or significantly degrade in performance due to measurement noise, object maneuvers, missed/spurious detections, and unknown measurement origin.

Keith LeGrand and Silvia Ferrari are with the Laboratory for Intelligent Systems and Controls (LISC), Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, New York, United States. Pingping Zhu is with the Department of Electrical Engineering at Marshall University, Huntington, West Virginia, United States. This work was supported in part by Office of Naval Research Grant N0014-19-1-2266, and by the Laboratory Directed Research and Development program at Sandia National Laboratories, a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Through the use of random finite set (RFS) theory, this paper formulates the multi-object information-driven control problem as a partially-observed Markov decision process (POMDP), wherein sensor actions are selected to maximize an expected reward conditioned on a probabilistic information state. Information-theoretic functionals known as *information gain*, such as expected entropy reduction (EER) [4], Cauchy-Schwarz Divergence (CSD) [5], [6], Kullback-Leibler divergence (KLD) [7], and Rényi divergence [8], [9], have been successfully used to represent sensing objectives, such as detection, classification, identification, and tracking, without an exhaustive enumeration of mission-specific contingencies. However, RFS-based information-theoretic sensor control remains computationally challenging. Most tractable solutions to date employ the so-called predicted ideal measurement set (PIMS) approximation [10], wherein sensor actions are selected based on ideal measurements with no measurement noise, false alarms, or missed detections. This paper presents a new computationally tractable higher-order approximation to the multi-object information gain using the cell multi-Bernoulli (cell-MB) approximation. Unlike previously proposed methods, the cell-MB approximation incorporates higher-order effects due to false alarms, missed detections, and non-Gaussian object probability distributions.

The cell-MB approximation and KLD information gain function presented in this paper also account for both detected and undetected objects by enabling the efficient computation of the RFS expectation operation. In particular, a partially piecewise homogeneous Poisson process is used to model undetected objects efficiently over space and time, including in challenging settings in which objects are diffusely distributed over a large geographic region. Prior work in [11] established a multi-agent probability hypothesis density (PHD)-based path planning algorithm aimed at maximizing the detection of relatively static objects. In [12], the exploration/exploitation problem was addressed by establishing an information theoretic uncertainty threshold for triggering pre-planned search modalities. The occupancy grid approach in [13] was successfully implemented for tracking and discovering objects with identity-tagged observations. A unified search and track solution was also proposed in [14] based on Poisson multi-Bernoulli mixture (PMBM) priors and a non-information-theoretic reward. However, these previous methods all rely on the PIMS approximation and, therefore, neglect the contribution of non-ideal measurements in the prediction of information gain.

The RFS information-driven approach presented in this paper, on the other hand, introduces a cell-MB approximation of the RFS reward expectation that accounts for non-ideal measurements. A new KLD reward is employed to measure information gain for detected and undetected objects, the efficient expectation of which is enabled by cell-MB approximation. The effectiveness of this approach is demonstrated using real video data in a challenging vehicle tracking application involving multiple closely-spaced and maneuvering objects in a cluttered environment. The proposed approach is shown to effectively track and maintain discovered objects while simultaneously searching and discovering new objects as they enter the surveillance region.

II. Problem Formulation

This paper considers an online search-while-tracking (SWT) problem involving a single sensor whose FoV can be manipulated through sensing platform motion and/or sensor configuration. The sensor objective is to discover and track multiple unidentified moving objects with partially hidden states and subject to unknown random inputs. The sensor control inputs are to be optimized at every time step in order to maximize the expected reduction in track uncertainty, as well as the overall state estimation performance.

The number of objects in the scene is unknown and changes over time, as objects enter and exit the surveillance region as well as, potentially, the sensor FoV, $\mathcal{S}_k$. Throughout this paper, single-object states are represented by lowercase letters (e.g. $\mathbf{x}$, $\mathbf{s}$), while multi-object states are represented by finite sets and denoted by italic uppercase letters (e.g. X, $\mathring{X}$). Bold lowercase letters are used to denote vectors (e.g. $\mathbf{x}$, $\mathbf{z}$). The accent "$\circ$" is used to distinguish labeled states and functions (e.g. $\mathring{f}$, $\mathring{x}$, $\mathring{X}$) from their unlabeled equivalents, where a state's label is simply a unique number or tuple to distinguish it from the states of other objects and associate track estimates over time. Spaces are represented by blackboard bold symbols (e.g. $\mathbb{X}$, $\mathbb{L}$), where $\mathbb{N}_\ell$ denotes the set of natural numbers

$$\mathbb{N}_\ell \triangleq \{1, \ldots, \ell\} \tag{1}$$

For brevity, the multi-object exponential notation,

$$h^A \triangleq \prod_{a \in A} h(a) \tag{2}$$

where $h^\emptyset \triangleq 1$, is adopted throughout. For multivariate functions, the dot "$\cdot$" denotes the argument of the multi-object exponential, e.g.:

$$[g(a, \cdot, c)]^B \triangleq \prod_{b \in B} g(a, b, c) \tag{3}$$

Let N_k denote the number of objects present in the surveillance region $\mathcal{W}$ at time t_k. The multi-object state X_k is the collection of N_k single-object states at time t_k and is expressed as the finite set

$$X_k = \{\mathbf{x}_{k,1}, \ldots, \mathbf{x}_{k,N_k}\} \in \mathcal{F}(\mathbb{X}) \tag{4}$$

where $\mathcal{F}(\mathbb{X})$ denotes the collection of all finite subsets of the object state space $\mathbb{X}$. Similarly, the multi-object measurement is the collection of M_k single-object measurements at time t_k and is expressed as the set

$$Z_k = \{\mathbf{z}_{k,1}, \ldots, \mathbf{z}_{k,M_k}\} \in \mathcal{F}(\mathbb{Z}) \tag{5}$$

where $\mathbb{Z}$ denotes the measurement space. The sensor resolution is such that single-object detections $\mathbf{z}_{k,i}$ are represented by points, e.g., a centroidal pixel, with no additional classification-quality information. Because detections contain no identifying information, the association between tracked objects and incoming measurement data is unknown.

Depending on the sensor, the detectability of an object may depend only on the partial state $\mathbf{s}_k \in \mathbb{X}_s \subseteq \mathbb{X}$. For example, the instantaneous ability of a sensor to detect an object may depend only on the object's position. In that case, $\mathbb{X}_s$ is the position space, and the complement space $\mathbb{X} \setminus \mathbb{X}_s$ is comprised of non-position states, such as object velocity. This nomenclature is adopted throughout while noting that the presented approach is applicable to other non-spatial state definitions.

Object detection is assumed to be random and characterized by the probability function,

$$p_{D,k}(\mathbf{x}_k; \mathcal{S}_k) = 1_{\mathcal{S}_k}(\mathbf{s}_k) \cdot p_{D,k}(\mathbf{s}_k) \tag{6}$$

where the FoV $\mathcal{S}_k \subset \mathbb{X}_s$ and $p_{D,k}(\mathbf{s}_k)$ is the probability of object detection conditioned on the object's presence in $\mathcal{S}_k$. When an object is detected, a noisy measurement of its state $\mathbf{x}_k$ is produced according to the likelihood function

$$\mathbf{z}_k \sim g_k(\mathbf{z}|\mathbf{x}_k) \tag{7}$$

where $\mathbf{z}_k \in \mathbb{Z}$. In addition to detections originating from true objects, the sensor produces extraneous measurements due to random phenomena, which are referred to as clutter or false alarms. Each resolution cell (e.g., a pixel) of the sensor is equally likely to produce a false alarm, and thus, the clutter process is modeled as a Poisson RFS process with PHD $\kappa_c(\mathbf{z})$.

Let $\mathbf{u}_k \in \mathbb{U}_k$ denote the sensor control inputs that, through translation and/or rotation at time t_k, determine the position of the sensor FoV at time t_{k+1}, namely $\mathcal{S}_{k+1}$; and, let $\mathbb{U}_k$ denote the set of all admissible controls. Decisions about $\mathbf{u}_k$ influence both the FoV geometry, $\mathcal{S}_{k+1}$, and the sensor measurements, Z_{k+1} due to varying object visibility. Because in many modern applications the surveillance region $\mathcal{W}$ is much larger than the sensor FoV, only a fraction of the total object population can be observed at any given time. Given the admissible control inputs $\mathbb{U}_k$, the field-of-regard (FoR)

$$\mathcal{T}_{k+1} = \bigcup_{\mathbf{u}_k \in \mathbb{U}_k} \mathcal{S}_{k+1}(\mathbf{u}_k) \tag{8}$$

is the composite of regions that the sensor *could* (although not simultaneously) cover at the next time step.

The presented sensor control problem can be formulated as a POMDP [15], [16]. Elements of an RFS POMDP include a partially- and noisily-observed state X_k, a known initial distribution of the state $f_0(X_0)$, a probabilistic transition model $f_{k|k-1}(X_k|X_{k-1})$ that describes the stochastic evolution of the state, a set of admissible control actions $\mathbb{U}_k$, and a reward $\mathcal{R}_k$

associated with each control action. At every time step k, an RFS multi-object tracker processes incoming measurements to produce the posterior information state $f_k(X_k|Z_{0:k})$. Then, the sensor control input is chosen so as to maximize the expected information gain, or,

$$\mathbf{u}_k^* = \underset{\mathbf{u}_k \in \mathbb{U}_k}{\arg\max} \{\mathrm{E}\left[\mathcal{R}_k(Z_{k+1}; \mathcal{S}_{k+1},\, f_k(X_k|Z_{0:k}),\, \mathbf{u}_{0:k-1})\right]\} \tag{9}$$

where the functional dependence of Z_{k+1} and $\mathcal{S}_{k+1}$ on $\mathbf{u}_k$ is omitted for brevity here but is described in [17].

A computationally tractable approximation of the expected reward in (9) is found using the new cell-MB approximation presented in Section IV. A new sensor control policy for SWT applications is formulated in Section V based on dual detected/undetected information gain. The dual reward formulation treats detected and undetected objects as separate processes. A new model proposed in Section V models undetected objects as a partially piecewise homogeneous Poisson process, which enables computationally efficient SWT over potentially large geographic regions.

III. Background on Random Finite Sets

An RFS X is a random variable that takes values on $\mathcal{F}(\mathbb{X})$. A labeled random finite set (LRFS) $\mathring{X}$ is a random variable that takes values on $\mathcal{F}(\mathbb{X} \times \mathbb{L})$, where $\mathbb{L}$ is a discrete label space. Both RFS and LRFS distributions can be described by set density functions, as established by Mahler's finite set statistics (FISST) [17], [18]. Three distributions important to this paper are the Poisson RFS, multi-Bernoulli (MB) RFS, and generalized labeled multi-Bernoulli (GLMB) LRFS distributions.

A. Poisson RFS

The density of a Poisson-distributed RFS X is

$$f(X) = e^{-N_X}[D]^X \tag{10}$$

where N_X is the object cardinality mean, and $D(\mathbf{x})$ is the PHD, or intensity function, of X, which is defined on the single-object space $\mathbb{X}$. The PHD is an important statistic in RFS theory as its integral over a set gives the expected number of objects in that set. The PHD of a general RFS X is given in terms of its set density $f(X)$ as [19]

$$D(\mathbf{x}) = \int f(\{\mathbf{x}\} \cup X')\delta X' \tag{11}$$

The integral in (11) is a set integral, defined as

$$\int f(Y)\delta Y \triangleq \sum_{n=0}^{\infty} \frac{1}{n!} \int f(\{\mathbf{y}_1, \ldots, \mathbf{y}_n\})\mathrm{d}\mathbf{y}_1 \cdots \mathrm{d}\mathbf{y}_n \tag{12}$$

B. Multi-Bernoulli RFS

The density of an MB distribution is [17, p. 102]

$$f(X) = \left[1 - r^{(\cdot)}\right]^{\mathbb{N}_M} \sum_{1 \le i_1 \ne \cdots \ne i_n \le M} \left[\frac{r^{i_{(\cdot)}} p^{i_{(\cdot)}}(\mathbf{x}_{(\cdot)})}{1 - r^{i_{(\cdot)}}}\right]^{\mathbb{N}_n} \tag{13}$$

where $n = |X|$, $|\cdot|$ denotes the cardinality operator, M is the number of MB components and maximum possible object cardinality, r^i is the probability that the i^{th} object exists, and $p^i(\mathbf{x})$ is the single-object state probability density of the i^{th} object if it exists.

C. GLMB RFS

The density of a GLMB distribution is given by [20]

$$\mathring{f}(\mathring{X}) = \Delta(\mathring{X}) \sum_{\xi \in \Xi} w^{(\xi)}(\mathcal{L}(\mathring{X}))[p^{(\xi)}]^{\mathring{X}}\,, \tag{14}$$

where Ξ is a discrete space, and where each $\xi \in \Xi$ represents a history of measurement association maps. Each $p^{(\xi)}(\cdot, \ell)$ is a probability density on $\mathbb{X}$, and each weight $w^{(\xi)}$ is non-negative with

$$\sum_{(I,\xi) \in \mathcal{F}(\mathbb{L}) \times \Xi} w^{(\xi)}(I) = 1 \tag{15}$$

The label of a labeled state $\mathring{x}$ is recovered by $\mathcal{L}(\mathring{x})$, where $\mathcal{L} : \mathbb{X} \times \mathbb{L} \mapsto \mathbb{L}$ is the projection defined by $\mathcal{L}((\mathbf{x}, \ell)) \triangleq \ell$. Similarly, for LRFSs, $\mathcal{L}(\mathring{X}) \triangleq \{\mathcal{L}(\mathring{x}) : \mathring{x} \in \mathring{X}\}$. The distinct label indicator $\Delta(\mathring{X}) = \delta_{(|\mathring{X}|)}(|\mathcal{L}(\mathring{X})|)$ ensures that only sets with distinct labels are considered.

D. Multi-Object Filtering

Online estimation of the multi-object state is performed using the data-driven GLMB filter, which provides the recursive solution of the measurement-driven Bayes filter [21]:

$$\mathring{f}_{\mathrm{p}}(\mathring{X}_{\mathrm{p},k}|Z_{0:k-1}) = \int \mathring{f}(\mathring{X}_{\mathrm{p},k}|\mathring{X}_{k-1})\mathring{f}(\mathring{X}_{k-1}|Z_{0:k-1})\delta\mathring{X}_{k-1} \tag{16}$$

$$\mathring{f}(\mathring{X}_k|Z_{0:k}) = \frac{g(Z_k|\mathring{X}_k)\mathring{f}_{\mathrm{p}}(\mathring{X}_{\mathrm{p},k}|Z_{0:k-1})\mathring{f}_{\mathrm{b}}(\mathring{X}_{\mathrm{b},k})}{\int g(Z_k|\mathring{X})\mathring{f}_{\mathrm{p}}(\mathring{X}_{\mathrm{p},k}|Z_{0:k-1})\mathring{f}_{\mathrm{b}}(\mathring{X}_{\mathrm{b},k})\delta\mathring{X}} \tag{17}$$

where the function time indices have been suppressed for brevity, and where $\mathring{f}_{\mathrm{p},k}(\mathring{X}_{\mathrm{p},k})$ and $\mathring{f}_{\mathrm{b},k}(\mathring{X}_{\mathrm{b},k})$ denote the density of persisting and birth objects, respectively, $\mathring{X}_k = \mathring{X}_{\mathrm{p},k} \cup \mathring{X}_{\mathrm{b},k}$, $\mathring{f}_{k|k-1}(\mathring{X}_{\mathrm{p},k}|\mathring{X}_{k-1})$ is the multi-object transition density, $g_k(Z_k|\mathring{X}_k)$ is the multi-object likelihood function, and, as a slight abuse of notation, g_k is used to denote both the single-object and multi-object likelihood function. The correct function usage can be easily determined by the nature of its arguments.

IV. Information-Driven Control

The objective of information-driven control is to maximize the information gained through future measurements when they are still unknown to the sensor. The reward expectation over unknown measurements Z_{k+1} can be obtained using the set integral

$$\mathrm{E}[\mathcal{R}_k] = \int \mathcal{R}_k(Z_{k+1}; \cdot)f(Z_{k+1})\delta Z_{k+1} \tag{18}$$

where $f(Z_{k+1})$ is the predicted measurement density conditioned on past measurements. Unfortunately, direct evaluation of (18) is, in general, intractable due to the infinite summation of nested single-object integrals. Furthermore, each integrand evaluation encompasses a multi-object filter update and subsequent divergence computation. As such, principled approximations are needed for tractable computation of the reward expectation.

A. The Cell-MB Distribution

In this paper, a new approximation is established for computing the reward expectation. We refer to this as the cell-MB approach, which approximates an arbitrary measurement density as an MB density with existence probabilities and single-object densities derived from a cell decomposition of the measurement space.

Definition 1: Consider the tessellation of the space $\mathbb{Y}$ into P disjoint subspaces, or cells, as

$$\mathbb{Y} = \overset{1}{\mathbb{Y}} \uplus \cdots \uplus \overset{P}{\mathbb{Y}} \tag{19}$$

Given the cell-decomposition (19), the RFS $Y = \{\mathbf{y}_1, \ldots, \mathbf{y}_n\}$ is considered to be cell-MB if it is distributed according to the density

$$f(Y) = \Delta(Y, \mathbb{Y}) \left[1 - r^{(\cdot)}\right]^{\mathbb{N}_P} \sum_{1 \le j_1 \ne \cdots \ne j_n \le P} \left[\frac{r^{j_{(\cdot)}} p^{j_{(\cdot)}}(\mathbf{y}_{(\cdot)})}{1 - r^{j_{(\cdot)}}}\right]^{\mathbb{N}_n} \tag{20}$$

where

$$\Delta(Y, \mathbb{Y}) \triangleq \begin{cases} 1 & |Y \cap \overset{j}{\mathbb{Y}}| \le 1 \, \forall \, j \in \{1, \ldots, P\} \\ 0 & \text{otherwise} \end{cases} \tag{21}$$

and

$$\int_{\overset{j}{\mathbb{Y}}} p^j(\mathbf{y}) \mathrm{d}\mathbf{y} = 1, \qquad j = 1, \ldots, P \tag{22}$$

In essence, the cell-MB distribution is a special case of the MB distribution in which the probability of more than one object occupying the same cell is zero.

As shown in [22], a collection of Bernoulli distributions can be defined over an occupancy grid by integration of the PHD. Inspired by [22], in this paper, the cell-MB approximation for an arbitrary density and appropriate cell-decomposition is established. The best cell-MB approximation, as defined by KLD minimization, has a matching PHD and cell weights equal to the expected number of objects in each cell. This is established in the following proposition.

Proposition 1: Let $f(Y)$ be an arbitrary set density with PHD $D(\mathbf{y})$ and $\mathbb{Y} = \overset{1}{\mathbb{Y}} \uplus \cdots \uplus \overset{P}{\mathbb{Y}}$ be a cell decomposition such that

$$\int_{\overset{j}{\mathbb{Y}}} D(\mathbf{y}) \mathrm{d}\mathbf{y} \le 1, \quad j = 1, \ldots, P \tag{23}$$

If $\bar{f}(Y)$ is a cell-MB over the same cell-decomposition with parameters $\{r^j, p^j\}_{j=1}^P$, the KLD between $f(Y)$ and $\bar{f}(Y)$ is minimized by parameters

$$r^j = \int 1_{\overset{j}{\mathbb{Y}}}(\mathbf{y}) D(\mathbf{y}) \mathrm{d}\mathbf{y} \tag{24}$$

$$p^j(\mathbf{y}) = \frac{1}{r^j} 1_{\overset{j}{\mathbb{Y}}}(\mathbf{y}) D(\mathbf{y}) \tag{25}$$

The proof is provided in [23]. The cell-MB approximation, when applied to the predicted measurement density, results in a simplified multi-object expectation for a specific class of reward functions, as described in the following subsection.

B. Reward Expectation: Cell-MB Approximation

In order to reduce the computational complexity associated with the set integral, this subsection shows that the multi-object reward expectation simplifies to a finite sum involving only single-object integrals, assuming the measurement is cell-MB distributed and the reward function is additive over FoV subsets.

Given the FoV $\mathcal{S} \subset \mathbb{X}_s$, define

$$\overset{j}{\mathcal{S}} \triangleq \mathcal{S} \cap \overset{j}{\mathbb{X}}_s \tag{26}$$

Further assume that position state cells do not overlap at the FoV bounds, such that each position state cell $\overset{j}{\mathbb{X}}_s$ is either wholly included in or wholly excluded by $\mathcal{S}$:

$$\overset{j}{\mathbb{X}}_s \setminus \overset{j}{\mathcal{S}} = \emptyset \quad \forall \overset{j}{\mathcal{S}} \ne \emptyset \tag{27}$$

Proposition 2: Let Z_{k+1} be distributed according to the cell-MB density $f(Z_{k+1})$ with parameters $\{r^j, p^j\}_{j=1}^P$ and the cell decomposition

$$\mathbb{Z} = \overset{1}{\mathbb{Z}} \uplus \cdots \uplus \overset{P}{\mathbb{Z}} \tag{28}$$

If the reward function $\mathcal{R}_k(\cdot)$ is integrable and additive over disjoint cells, i.e.,

$$\mathcal{R}_k(Z_k; \mathcal{S}_k) = \sum_{j=1}^{P} \mathcal{R}_k(Z_{k+1} \cap \overset{j}{\mathbb{Z}}; \overset{j}{\mathcal{S}}_k) \tag{29}$$

then the expected reward is

$$\mathrm{E}[\mathcal{R}_k] = \sum_{j=1}^{P} \mathcal{R}_k(\emptyset; \overset{j}{\mathcal{S}}_k) \left(1 - r^j\right) + \hat{\mathcal{R}}^j_{\mathrm{z},k} \cdot r^j \tag{30}$$

where

$$\hat{\mathcal{R}}^j_{\mathrm{z},k} \triangleq \int_{\overset{j}{\mathbb{Z}}} \mathcal{R}_k(\{\mathbf{z}\}; \overset{j}{\mathcal{S}}_k) p^j(\mathbf{z}) \mathrm{d}\mathbf{z} \tag{31}$$

The proof is provided in [23].

The remainder of this paper considers reward functions satisfying the cell-additivity constraint of (29), such as the PHD filter based KLD reward. Note that adopting the PHD filter for estimating the information gain does not require using it for multi-object tracking. Given an arbitrary RFS prior density $f_{k|k-1}(X)$ and its PHD $D_{k|k-1}(\mathbf{x})$, the PHD-based KLD reward is

$$\mathcal{R}_k(Z; \mathcal{S}, D_{k|k-1}) = \int_{\mathbb{X}} D_{k|k-1}(\mathbf{x}) \cdot \{1 - L_Z(\mathbf{x}; \mathcal{S}) + L_Z(\mathbf{x}; \mathcal{S}) \log[L_Z(\mathbf{x}; \mathcal{S})]\} \mathrm{d}\mathbf{x} \tag{32}$$

where the pseudo-likelihood function is [17, p. 193]

$$L_Z(\mathbf{x}; \mathcal{S}) = 1 - p_D(\mathbf{x}; \mathcal{S}) + \sum_{\mathbf{z} \in Z} \frac{p_D(\mathbf{x}; \mathcal{S}) \cdot g(\mathbf{z}|\mathbf{x})}{\kappa_c(\mathbf{z}) + \int p_D(\mathbf{x}; \mathcal{S}) g(\mathbf{z}|\mathbf{x}) D_{k+1|k}(\mathbf{x}) \mathrm{d}\mathbf{x}} \tag{33}$$

The following proposition establishes that (32) is cell-additive for appropriate cell decompositions.

Proposition 3: Assume there exists a decomposition

$$\mathbb{Z} = \overset{1}{\mathbb{Z}} \uplus \cdots \uplus \overset{P}{\mathbb{Z}}, \qquad \mathbb{X} = \overset{1}{\mathbb{X}} \uplus \cdots \uplus \overset{P}{\mathbb{X}} \tag{34}$$

such that (27) is satisfied and

$$D_{k+1|k}(\mathbf{x}) g_{k+1}(\mathbf{z}|\mathbf{x}) = 0 \qquad \forall\, \mathbf{x} \in \overset{j}{\mathbb{X}},\, \mathbf{z} \in \overset{j'}{\mathbb{Z}}, j \neq j' \tag{35}$$

Then, the PHD-based KLD is additive over cells:

$$\mathcal{R}_k(Z; \mathcal{S}, D_{k|k-1}) = \sum_{j=1}^{P} \mathcal{R}_k(Z \cap \overset{j}{\mathbb{Z}}; \overset{j}{\mathcal{S}}, D_{k|k-1}) \tag{36}$$

The proof of Proposition 3 is provided in [23]. Perfect cell-additivity requires satisfying (35), which, in turn, implies that an object in cell $\overset{i}{\mathbb{X}}$ does not generate a measurement in $\overset{j}{\mathbb{Z}}$ for $i \neq j$. In general, violations of (35) are tolerable and result in approximation errors that are negligible in comparison to the stochastic variations in the actual information gain. Furthermore, these simplifying assumptions need not be satisfied by the multi-object tracker.

The cell-MB approximation accounts for the potential reward of non-ideal measurements, which may include missed detections, clutter, and measurements originating from new objects. The latter case is particularly important for the search of undetected objects, as is shown in the following section.

V. Search-While-Tracking Methodology

This section presents a dual reward function that takes into account both detected and undetected objects. The proposed reward function balances the competing objectives of object search and tracking by means of a unified information-theoretic framework.

A. Dual Reward

Separate density parameterizations for detected and undetected objects are employed such that their unique characteristics may be leveraged for computational efficiency. Let $X_{u,k} \in \mathcal{F}(\mathbb{X})$ be the state of objects that were not detected during steps $0, \ldots, k-1$ and $X_{d,k} \in \mathcal{F}(\mathbb{X})$ be the state of objects detected prior to k. Denote by $Z_{u,k}$, $Z_{d,k}$, and $Z_{c,k}$ the detections generated by $X_{u,k}$, $X_{d,k}$, and clutter, respectively. Let $V_k \triangleq Z_{d,k} \cup Z_{c,k}$ and $W_k \triangleq Z_{u,k} \cup Z_{c,k}$. Then, the sensor control policy is defined in terms of the dual reward as

$$\begin{aligned}\mathbf{u}_k^* = \underset{\mathbf{u}_k \in \mathbb{U}_k}{\arg\max} \Big\{ &\mathrm{E}[\mathcal{R}_k^d(V_{k+1}; \mathcal{S}_{k+1}(\mathbf{u}_k))] \\ &+ \mathrm{E}[\mathcal{R}_k^u(W_{k+1}; \mathcal{S}_{k+1}(\mathbf{u}_k))] \Big\}\end{aligned} \tag{37}$$

where

$$\mathcal{R}_k^d(\cdot;\cdot) = \mathcal{R}_k(\cdot;\cdot, D_{d,k+1|k}) \tag{38}$$

$$\mathcal{R}_k^u(\cdot;\cdot) = \mathcal{R}_k(\cdot;\cdot, D_{u,k+1|k}) \tag{39}$$

are used for brevity, and $D_{d,k+1|k}$ and $D_{u,k+1|k}$ are the prior PHDs of detected and undetected objects, respectively. The individual reward expectations for detected and undetected objects are derived in the following subsections.

B. Detected Object Reward Expectation

If $f(V_{k+1})$ is cell-MB with parameters $\{r_{\mathrm{v}}^j, p_{\mathrm{v}}^j\}_{j=1}^P$, then by Proposition 2

$$\mathrm{E}[\mathcal{R}_k^d] = \sum_{j=1}^{P} \mathcal{R}_k^d(\emptyset; \overset{j}{\mathcal{S}}_{k+1}) \left(1 - r_{\mathrm{v}}^j\right) + \hat{\mathcal{R}}_{\mathrm{v},k}^{d,j}(\overset{j}{\mathcal{S}}_{k+1}) \cdot r_{\mathrm{v}}^j \tag{40}$$

where

$$\hat{\mathcal{R}}_{\mathrm{v},k}^{d,j}(\overset{j}{\mathcal{S}}) \triangleq \int_{\overset{j}{\mathbb{Z}}} \mathcal{R}_k^d(\{\mathbf{z}\}; \overset{j}{\mathcal{S}}) p_{\mathrm{v}}^j(\mathbf{z}) \mathrm{d}\mathbf{z} \tag{41}$$

$$r_{\mathrm{v}}^j(\mathcal{S}) = \int 1_{\overset{j}{\mathbb{Z}}}(\mathbf{z}) D_{\mathrm{v},k+1|k}(\mathbf{z}; \mathcal{S}) \mathrm{d}\mathbf{z} \tag{42}$$

$$p_{\mathrm{v}}^j(\mathbf{z}; \mathcal{S}) = \frac{1}{r_{\mathrm{v}}^j} 1_{\overset{j}{\mathbb{Z}}}(\mathbf{z}) D_{\mathrm{v},k+1|k}(\mathbf{z}; \mathcal{S}) \tag{43}$$

The multi-object tracker provides the prior GLMB density $\mathring{f}_{\mathrm{p},k+1|k}(\mathring{X}_{\mathrm{p},k+1}|Z_{0:k})$, from which the detected object PHD is obtained as

$$D_{d,k+1|k}(\mathbf{x}) = \sum_{(I,\xi) \in \mathcal{F}(\mathbb{L}) \times \Xi} \sum_{\ell \in I} w^{(\xi)}(I) p^{(\xi)}(\mathbf{x}, \ell) \tag{44}$$

The PHD $D_{\mathrm{v},k+1|k}$ can be obtained from the predicted measurement density $f_{k+1|k}(V_{k+1})$ through application of (11). From the prior GLMB density,

$$f_{k+1|k}(V_{k+1}) = \int g_{k+1}(V_{k+1}|\mathring{X}) \mathring{f}_{k|k-1}(\mathring{X}) \delta \mathring{X} \tag{45}$$

Given a GLMB prior, explicit computation of the predicted measurement density is computationally challenging. Thus, $D_{\mathrm{v},k+1|k}$ is directly computed from the detected object PHD using the approximation

$$\begin{aligned}D_{\mathrm{v},k+1|k}(\mathbf{z}; \mathcal{S}) \approx \int &D_{d,k+1|k}(\mathbf{x}) p_{D,k+1}(\mathbf{x}; \mathcal{S}) g_{k+1}(\mathbf{z}|\mathbf{x}) \mathrm{d}\mathbf{x} \\ &+ \kappa_{c,k+1}(\mathbf{z})\end{aligned} \tag{46}$$

C. Undetected Object Reward Expectation

This subsection presents a new approach to efficiently model the undetected object distribution, which may be diffuse over a large region. Although Gaussian mixtures (GMs) and particle representations can be used to model undetected objects, they are highly inefficient at representing diffuse distributions. Thus, in this paper, the position-marginal density of undetected objects is taken to be piecewise homogeneous with PHD

$$D_{u,k+1|k}(\mathbf{s}) = \sum_{j=1}^{P} \frac{1_{\overset{j}{\mathbb{X}}_s}(\mathbf{s})}{A(\overset{j}{\mathbb{X}}_s)} \cdot \lambda_{j,k+1|k} \tag{47}$$

where $\lambda_{j,k+1|k}$ is the expected number of undetected objects in $\overset{j}{\mathbb{X}}_s$ at time step $k+1$ and $A(\overset{j}{\mathbb{X}}_s)$ is the volume of cell $\overset{j}{\mathbb{X}}_s$. For ease of exposition, the undetected object PHD is modeled using the same cell decomposition employed in the cell-MB approximation.

If $f(W_{k+1})$ is cell-MB with parameters $\{r_{\rm w}^j, p_{\rm w}^j\}_{j=1}^P$, then by Proposition 2,

$$\mathrm{E}[\mathcal{R}_k^u] = \sum_{j=1}^{P} \mathcal{R}_k^u(\emptyset; \overset{j}{\mathcal{S}}_{k+1}) \left(1 - r_{\rm w}^j\right) + \hat{\mathcal{R}}_{{\rm w},k}^{u,j}(\overset{j}{\mathcal{S}}_{k+1}) \cdot r_{\rm w}^j \tag{48}$$

where

$$\hat{\mathcal{R}}_{{\rm w},k}^{u,j}(\overset{j}{\mathcal{S}}) \triangleq \int_{\overset{j}{\mathbb{Z}}} \mathcal{R}_k^u(\{\mathbf{z}\}; \overset{j}{\mathcal{S}}) p_{\rm w}^j(\mathbf{z}) \mathrm{d}\mathbf{z} \tag{49}$$

$$r_{\rm w}^j(\mathcal{S}) = \int 1_{\overset{j}{\mathbb{Z}}}(\mathbf{z}) D_{{\rm w},k+1|k}(\mathbf{z}; \mathcal{S}) \mathrm{d}\mathbf{z} \tag{50}$$

$$\approx \frac{\lambda_{j,k+1|k}}{A(\overset{j}{\mathbb{X}}_s)} \int_{\overset{j}{\mathbb{X}}_s} p_D(\mathbf{s}; \mathcal{S}) \mathrm{d}\mathbf{s} \tag{51}$$

$$p_{\rm w}^j(\mathbf{z}; \mathcal{S}) = \frac{1}{r_{\rm w}^j} 1_{\overset{j}{\mathbb{Z}}}(\mathbf{z}) D_{{\rm w},k+1|k}(\mathbf{z}; \mathcal{S}) \tag{52}$$

$$D_{{\rm w},k+1|k}(\mathbf{z}; \mathcal{S}) = \int D_{u,k+1|k}(\mathbf{x}) p_{D,k+1}(\mathbf{x}; \mathcal{S}) g_{k+1}(\mathbf{z}|\mathbf{x}) \mathrm{d}\mathbf{x} + \kappa_{c,k+1}(\mathbf{z}) \tag{53}$$

Under a piecewise homogeneous PHD, the undetected object reward simplifies drastically if the measurement likelihood is independent of non-position states: i.e. $g(\cdot|\mathbf{x}) = g(\cdot|\mathbf{s})$. Following (32),

$$\begin{aligned} &\mathcal{R}_k^u(W_{k+1}; \mathcal{S}_{k+1}) \qquad (54)\\ &= \int_{\mathbb{X}_s} D_{u,k+1|k}(\mathbf{s}) \{1 - L_{W_{k+1}}(\mathbf{s}; \mathcal{S}_{k+1}) \\ &\quad + L_{W_{k+1}}(\mathbf{s}; \mathcal{S}_{k+1}) \log[L_{W_{k+1}}(\mathbf{s}; \mathcal{S}_{k+1})]\} \mathrm{d}\mathbf{s} \end{aligned}$$

Given that at most one measurement may exist per cell, two cases need be considered: the null measurement case and the singleton measurement case. Substitution of $W_{k+1} = \emptyset$ and some algebraic manipulation yields

$$\mathcal{R}_k^u(\emptyset; \mathcal{S}_{k+1}) = \sum_{j=1}^{P} \mathcal{R}_k^u(\emptyset; \overset{j}{\mathcal{S}}_{k+1}) \tag{55}$$

$$\mathcal{R}_k^u(\emptyset; \overset{j}{\mathcal{S}}_{k+1}) = \lambda_{j,k+1|k} \cdot d_j \cdot (1 - \delta_\emptyset(\overset{j}{\mathcal{S}}_{k+1})) \tag{56}$$

$$d_j \triangleq \frac{1}{A(\overset{j}{\mathbb{X}}_s)} \int_{\overset{j}{\mathbb{X}}_s} p_D(\mathbf{s}) + (1 - p_D(\mathbf{s})) \log[1 - p_D(\mathbf{s})] \mathrm{d}\mathbf{s} \tag{57}$$

Furthermore, if the probability of detection is homogeneous within cells such that

$$p_D(\mathbf{s}) = p_{D,j} \quad \forall \mathbf{s} \in \overset{j}{\mathbb{X}}_s \tag{58}$$

then (57) simplifies to

$$d_j = p_{D,j} + (1 - p_{D,j}) \log(1 - p_{D,j}) \tag{59}$$

Within a cell, the uniform position density of undetected objects is known *a priori* up to an unknown factor $\lambda_{j,k+1|k}$. Thus, the undetected object reward can be pre-computed for efficiency and

$$\hat{\mathcal{R}}_{{\rm w},k}^{u,j}(\overset{j}{\mathcal{S}}_{k+1}) \approx \bar{\mathcal{R}}_{\rm w}^{u,j}(\lambda_{j,k+1|k}) \tag{60}$$

where the function $\bar{\mathcal{R}}_{\rm w}^{u,j}(\lambda_{j,k+1|k})$ returns interpolated reward values over $\lambda_{j,k+1|k} \in [0, 1]$.

VI. Application to Vehicle Tracking

The cell-MB SWT framework is demonstrated in a vehicle tracking problem using real video data. The video was recorded using a fixed camera with a large FoV (Fig. 1.a), and real-time FoV controlled motion was simulated by windowing the data over a small fraction of the image, as illustrated in Fig. 1.b. This dataset presents significant tracking challenges, including jitter-induced noise and clutter, unknown measurement origin, merged detections from closely-spaced vehicles, and most significantly, temporal sparsity of detections.

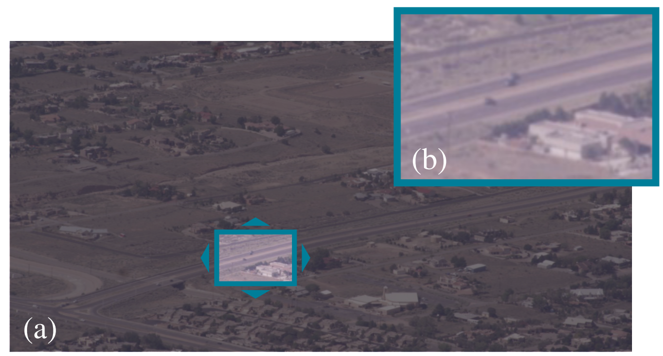

Fig. 1. Example video frame (a), artificially windowed to emulate smaller, movable FoV, which is enlarged in (b) to show detail.

A. Vehicle Dynamics

Vehicle dynamics are modeled directly in the image frame. While vehicle dynamics are more naturally expressed in the terrestrial frame, the camera's precise location and orientation is unknown. Thus, transformation between image and terrestrial coordinates could not be readily established.

The object state is modeled as

$$\mathbf{x} = [\mathbf{s}^T \quad \boldsymbol{\zeta}^T]^T \tag{61}$$

$$\mathbf{s} = [\xi \quad \eta]^T, \qquad \boldsymbol{\zeta} = [\dot{\xi} \quad \dot{\eta} \quad \Omega]^T \tag{62}$$

where ξ and η are the horizontal and vertical coordinates, respectively, of the vehicle position with respect to the full frame origin, $\dot{\xi}$ and $\dot{\eta}$ are the corresponding rates, and Ω is the vehicle turn rate.

Vehicle motion is modeled using the nearly coordinated turn model with directional process noise [24], [25] as

$$\mathbf{x}_{k+1} = \mathbf{f}_k(\mathbf{x}_k) + \boldsymbol{\Gamma}_k \boldsymbol{v}_k(\mathbf{s}_k) \tag{63}$$

where

$$\boldsymbol{\Gamma}_k = \begin{bmatrix} \frac{1}{2}(\Delta t)^2 \mathbf{I}_{2\times2} & \mathbf{0}_{2\times1} \\ (\Delta t)\mathbf{I}_{2\times2} & \mathbf{0}_{2\times1} \\ \mathbf{0}_{1\times2} & \Delta t \end{bmatrix} \tag{64}$$

where Δt is the discrete time step interval, $\mathbf{I}_{n\times n}$ denotes the $n \times n$ identity matrix, and $\mathbf{0}_{m\times n}$ denotes the $m \times n$ matrix whose elements are zero. The covariance of the process noise is

$$\mathrm{E}[\boldsymbol{v}_k \boldsymbol{v}_k^T] = \mathbf{Q}_k(\mathbf{s}) = \begin{bmatrix} \mathbf{D}^T(\mathbf{s})\mathbf{Q}_d\mathbf{D}(\mathbf{s}) & 0 \\ \mathbf{0}_{1\times 2} & \sigma_\Omega^2 \end{bmatrix} \tag{65}$$

$$\mathbf{Q}_d = \begin{bmatrix} \sigma_t^2 & 0 \\ 0 & \sigma_n^2 \end{bmatrix}, \quad \mathbf{D}(\mathbf{s}) = \begin{bmatrix} \cos\Psi(\mathbf{s}) & \sin\Psi(\mathbf{s}) \\ -\sin\Psi(\mathbf{s}) & \cos\Psi(\mathbf{s}) \end{bmatrix} \tag{66}$$

where σ_Ω is the turn rate process noise standard deviation, σ_t and σ_n are the standard deviation of process noise tangential and normal to the road, respectively, and $\Psi(\mathbf{s})$ is the angle of the road segment nearest $\mathbf{s}$, measured from the horizontal axis to the tangent direction.

B. Sensor and Scene Model

Object detections are generated from raw frame data using normalized difference change detection [26] and fast approximate power iteration subspace tracking [27] for temporal background estimation. The single-object measurement function is linear-Gaussian with corresponding likelihood

$$g(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mathbf{z};\, \mathbf{H}\mathbf{x},\, \mathbf{R})\,, \tag{67}$$

$$\mathbf{H} = \begin{bmatrix} \mathbf{I}_{2\times 2} & \mathbf{0}_{2\times 3} \end{bmatrix}, \qquad \mathbf{R} = 9 \cdot \mathbf{I}_{2\times 2}\,[\text{pixel}^2] \tag{68}$$

The sensor FoV is a rectangular region that is 240 pixels wide and 160 pixels tall. Moving objects within the FoV are assumed to be detectable with probability $p_{D,k}(\mathbf{s}_k) = 0.9$.

The scene is tessellated by a 16×32 grid of uniformly sized rectangular cells as shown in Fig. 2. Within the scene, a region of interest is specified which contains the scene's two primary roads and is denoted by $\mathcal{T}$ due to its equivalence to the FoR for this problem. Within the region of interest, cells containing road pixels comprise the set $\mathcal{B}$, which is used to establish an initial uniform distribution of undetected objects.

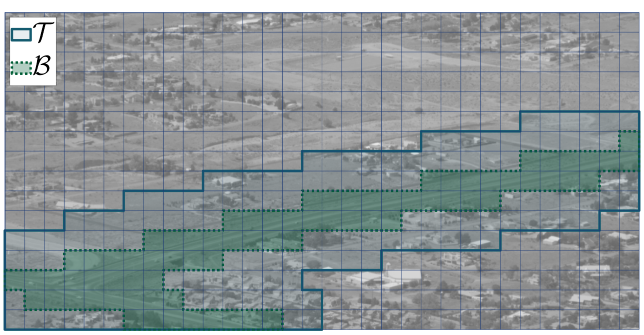

Fig. 2. Field-of-regard, $\mathcal{T}$, and primary road region $\mathcal{B}$, with example image frame as background.

C. Experiment Results

An experiment consisting of sixty time steps is performed. To emulate a pan/tilt camera from the wider available frame data, the FoV is assumed to be able to be move to any location within the scene in a single time step. This is a reasonable assumption as these adjustments would be less than a degree.

At each step, the incoming detections are processed by a GM implementation of the data-driven GLMB tracker [21] to compute the posterior multi-object set density. The negative information content from missed detections is leveraged to refine the multi-object density, incorporating the knowledge of where objects were not seen at a given instant. This is achieved by recursively splitting the density's GM components that overlap the FoV bounds, as shown in [28].

Some key frames of the experiment are shown in Fig. 3. In the early time steps, the FoV motion is dominated by the undetected object component of the reward. As more objects are discovered and tracked, the observed actions demonstrate a balance of revisiting established tracks to reduce state uncertainty and exploring new areas where undetected objects may exist.

The performance of the presented SWT approach is evaluated by the multi-object tracking accuracy, as measured using the generalized optimal sub-pattern assignment (GOSPA) metric [29]. The GOSPA metric along with its missed and false object components, are shown in Fig 4. Note that, unlike similar metrics, the GOSPA is unnormalized and may exceed the cutoff distance. The cell-MB approach effectively balances the competing objectives of new object discovery and maintenance of established tracks, as illustrated by the decline in missed objects and consistently low number of false tracks. An increase in GOSPA is observed in the final time steps of the experiment, which is caused by a sharp uptick in new object appearances.

VII. Conclusion

This paper presents a novel cell multi-Bernoulli (cell-MB) approximation that enables the tractable higher-order approximation of the expectation of set functions that are additive over disjoint measurable subsets. The cell-MB approximation is useful for approximating the expectation of computationally-expensive set functions, such as information-theoretic reward functions employed in sensor control applications. The approach is developed in the context of information-driven sensor control in which the objective is to discover and track an unknown time-varying number of non-cooperative objects with minimal estimation error. The problem is formulated as a partially-observed Markov decision process with a new Kullback-Leibler divergence based reward that incorporates both detected and undetected object information gain. In a demonstration using real sensor data, the approach is used to manipulate the sensor field-of-view to discover and track multiple moving ground objects from an aerial vantage point.

References

[1] S. Ferrari and T. A. Wettergren, *Information-driven Planning and Control*. MIT Press, 2021.

[2] H. Wei, P. Zhu, M. Liu, J. P. How, and S. Ferrari, "Automatic pan–tilt camera control for learning Dirichlet process Gaussian process (DPGP) mixture models of multiple moving targets," *IEEE Transactions on Automatic Control*, vol. 64, no. 1, pp. 159–173, 2019.

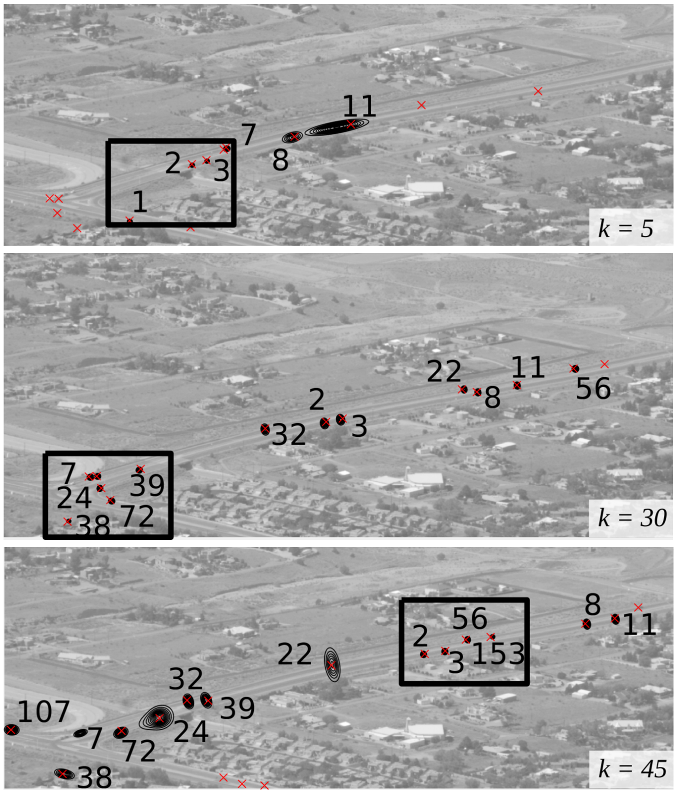

Fig. 3. FoV position and tracker estimates in the form of single-object density contours for objects with probabilities of existence greater than 0.5, shown at select time steps.

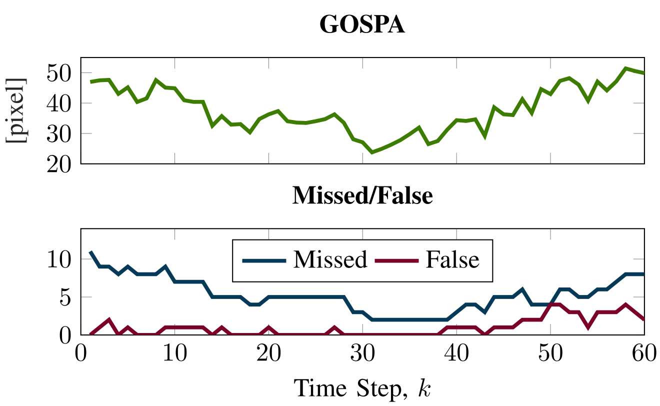

Fig. 4. GOSPA metric and component errors over time using cutoff distance $c = 20$ [pixel], order $p = 2$, and $\alpha = 2$.

[3] H. Wei, W. Lu, P. Zhu, S. Ferrari, M. Liu, R. H. Klein, S. Omidshafiei, and J. P. How, "Information value in nonparametric Dirichlet-process gaussian-process (DPGP) mixture models," *Automatica*, vol. 74, no. June, pp. 360–368, 2016.

[4] C. Cai and S. Ferrari, "Information-driven sensor path planning by approximate cell decomposition," *IEEE Transactions on Systems, Man, and Cybernetics - Part B*, vol. 39, no. 3, pp. 672–689, 2009.

[5] H. G. Hoang, B. T. N. Vo, B. T. N. Vo, and R. Mahler, "The Cauchy-Schwarz divergence for Poisson point processes," *IEEE Transactions on Information Theory*, vol. 61, no. 8, pp. 4475–4485, 2015.

[6] M. Beard, B.-T. Vo, B.-N. Vo, and S. Arulampalam, "Sensor control for multi-target tracking using Cauchy-Schwarz divergence," in *18th International Conference on Information Fusion*, Washington, DC, 2015, pp. 937–944.

[7] I. Goodman, R. P. Mahler, and H. T. Nguyen, *Mathematics of Data Fusion*. Kluwer Academic Publishers, 1997.

[8] B. Ristic and B.-n. Vo, "Sensor control for multi-object state-space estimation using random finite sets," *Automatica*, vol. 46, no. 11, pp. 1812–1818, 2010.

[9] B. Ristic, B.-N. Vo, and D. Clark, "A note on the reward function for PHD filters with sensor control," *IEEE Transactions on Aerospace and Electronic Systems*, pp. 1521 – 1529, 2011.

[10] R. P. S. Mahler and T. R. Zajic, "Probabilistic objective functions for sensor management," in *Signal Processing, Sensor Fusion, and Target Recognition XIII*, vol. 5429, no. August 2004, aug 2004, pp. 233–244.

[11] J. Olofsson, G. Hendeby, T. R. Lauknes, and T. A. Johansen, "Multi-agent informed path-planning using the probability hypothesis density," *Autonomous Robots*, 2020.

[12] S. Gehly, B. A. Jones, and P. Axelrad, "Search-detect-track sensor allocation for geosynchronous space objects," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 54, no. 6, pp. 2788–2808, 2018.

[13] H. V. Nguyen, H. Rezatofighi, B.-N. Vo, and D. C. Ranasinghe, "Multi-objective multi-agent planning for jointly discovering and tracking mobile objects," vol. 34, no. 05 SE - AAAI Technical Track: Multiagent Systems, apr 2020, pp. 7227–7235.

[14] P. Bostrom-Rost, D. Axehill, and G. Hendeby, "Sensor management for search and track using the Poisson multi-Bernoulli mixture filter," *IEEE Transactions on Aerospace and Electronic Systems*, 2021.

[15] V. Krishnamurthy, *Partially Observed Markov Decision Processes: From Filtering to Controlled Sensing*. Cambridge: Cambridge University Press, 2016.

[16] D. Castañón and L. Carin, "Stochastic control theory for sensor management," in *Foundations and Applications of Sensor Management*, A. O. Hero, D. Castañón, D. Cochran, and K. Kastella, Eds. Springer Science & Business Media, 2007, ch. 2.

[17] R. P. Mahler, *Advances in Statistical Multisource-Multitarget Information Fusion*. Artech House, 2014.

[18] ——, *Statistical Multisource-Multitarget Information Fusion*. Artech House Boston, 2007.

[19] ——, "Multitarget Bayes filtering via first-order multitarget moments," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 39, no. 4, pp. 1152–1178, 2003.

[20] B.-N. Vo, B.-T. Vo, and D. Phung, "Labeled random finite sets and the Bayes multi-target tracking filter," *IEEE Transactions on Signal Processing*, vol. 62, no. 24, pp. 6554–6567, 2014.

[21] K. A. LeGrand and K. J. DeMars, "The data-driven δ-generalized labeled multi-Bernoulli tracker for automatic birth initialization," in *Signal Processing, Sensor/Information Fusion, and Target Recognition XXVII*, vol. 10646. International Society for Optics and Photonics, 2018, p. 1064606.

[22] D. Nuss, S. Reuter, M. Thom, T. Yuan, G. Krehl, M. Maile, A. Gern, and K. Dietmayer, "A random finite set approach for dynamic occupancy grid maps with real-time application," *The International Journal of Robotics Research*, vol. 37, no. 8, pp. 841–866, jul 2018.

[23] K. A. LeGrand, P. Zhu, and S. Ferrari, "Cell multi-Bernoulli (cell-MB) sensor control for multi-object search-while-tracking (SWT)," 2021. [Online]. Available: https://arxiv.org/abs/2108.11236

[24] Y. Bar-Shalom, P. K. Willett, and X. Tian, *Tracking and Data Fusion*. YBS publishing, 2011.

[25] M. Yeddanapudi, Y. Bar-Shalom, and K. Pattipati, "IMM estimation for multitarget-multisensor air traffic surveillance," *Proceedings of the IEEE*, vol. 85, no. 1, pp. 80–96, 1997.

[26] K. M. Simonson and T. J. Ma, "Robust real-time change detection in high jitter," Tech. Rep., 8 2009.

[27] R. Badeau, B. David, and G. Richard, "Fast approximated power iteration subspace tracking," *IEEE Transactions on Signal Processing*, vol. 53, no. 8, pp. 2931–2941, 2005.

[28] K. A. LeGrand and S. Ferrari, "The role of bounded fields-of-view and negative information in finite set statistics (FISST)," in *2020 23rd International Conference on Information Fusion (FUSION 2020)*, July 2020, pp. 1–9.

[29] A. S. Rahmathullah, Á. F. García-Fernández, and L. Svensson, "Generalized optimal sub-pattern assignment metric," in *2017 20th International Conference on Information Fusion (Fusion)*, 2017, pp. 1–8.