

Cell Multi-Bernoulli (Cell-MB) Sensor Control for Multi-Object Search-While-Tracking (SWT)

Keith A. LeGrand¹, Member, IEEE, Pingping Zhu², Member, IEEE, and
Silvia Ferrari³, Senior Member, IEEE

Abstract—Information-driven control can be used to develop intelligent sensors that can optimize their measurement value based on environmental feedback. In object tracking applications, sensor actions are chosen based on the expected reduction in uncertainty also known as information gain. Random finite set (RFS) theory provides a formalism for quantifying and estimating information gain in multi-object tracking problems. However, estimating information gain in these applications remains computationally challenging. This paper presents a new tractable approximation of the RFS expected information gain applicable to sensor control for multi-object search and tracking. Unlike existing RFS approaches, the information gain approximation presented in this paper considers the contributions of non-ideal noisy measurements, missed detections, false alarms, and object appearance/disappearance. The effectiveness of the information-driven sensor control is demonstrated through two multi-vehicle search-while-tracking experiments using real video data from remote terrestrial and satellite sensors.

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

1 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], [2], [3]. Also, the sensor FoV is able to move and cover large regions of interest, potentially for prolonged periods of time. By expanding the autonomy and

operability of sensors, however, several new challenges are introduced. As the sensor moves and reconfigures itself, the number of objects inside the FoV changes over time. Also, both the number of objects and the objects' states are unknown, time-varying, and subject to significant measurement errors. As a result, existing tracking algorithms and information gain functions (e.g., [1], [2], [3]) that 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.

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). Sensor actions can then be decided to maximize the expected information gain conditioned on a probabilistic information state. Information-theoretic functionals, such as expected entropy reduction (EER) [4], [5], Cauchy-Schwarz Divergence (CSD) [6], [7], Kullback-Leibler divergence (KLD) [8], and Rényi divergence [9], [10], have been successfully used to represent sensing objectives, such as detection, classification, identification, and tracking, circumventing exhaustive enumeration of all possible outcomes. However, RFS-based information-theoretic sensor control policies remain computationally challenging. Alternatively, they require simplifying assumptions that limit their applicability to vision-based search-while-tracking (SWT) systems. Tractable solutions to date employ the so-called predicted ideal measurement set (PIMS) approximation [11], by which 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 called the cell multi-Bernoulli (cell-MB) approximation for a restricted

- Keith A. LeGrand is with the School of Aeronautics and Astronautics, Purdue University, West Lafayette, IN 47907 USA. E-mail: klegrand@purdue.edu.
- Pingping Zhu is with the Department of Electrical Engineering, Marshall University, Huntington, WV 25755 USA. E-mail: zhup@marshall.edu.
- Silvia Ferrari is with the Laboratory for Intelligent Systems and Controls (LISC), Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY 14850 USA. E-mail: ferrari@cornell.edu.

Manuscript received 16 August 2021; revised 18 September 2022; accepted 16 November 2022. Date of publication 23 November 2022; date of current version 5 May 2023.

This work was supported in part by the Office of Naval Research under Grant N0014-19-1-2266, and in part by the Laboratory Directed Research and Development program, 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 Grant DE-NA0003525.

(Corresponding author: Keith A. LeGrand.)

Recommended for acceptance by R. Babuska.

This article has supplementary downloadable material available at <https://doi.org/10.1109/TPAMI.2022.3223856>, provided by the authors.

Digital Object Identifier no. 10.1109/TPAMI.2022.3223856

class of multi-object information gain functions satisfying cell-additivity constraints. Unlike existing approximation methods, the cell-MB approximation accounts for 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 discovered and undiscovered objects by enabling the efficient computation of the RFS expectation operation. In particular, a partially piecewise homogeneous Poisson process is used to model undiscovered objects efficiently over space and time, including in challenging settings in which objects are diffusely distributed over a large geographic region. Prior work in [12] established a multi-agent probability hypothesis density (PHD)-based path planning algorithm aimed at maximizing the detection of relatively static objects. In [13], the exploration/exploitation problem was addressed by establishing an information-theoretic uncertainty threshold for triggering pre-planned search modalities. The occupancy grid approach in [14] was successfully implemented for tracking and discovering objects with identity-tagged observations. A unified search and track solution was also proposed in [15] based on Poisson multi-Bernoulli mixture (PMBM) priors and a non-information-theoretic reward. However, these existing methods all rely on the PIMS approximation and, therefore, neglect the contribution of non-ideal measurements in the prediction of information gain. Preliminary results of this work were reported in [16].

The new RFS information-driven approach presented in this paper derives a cell-MB approximation of the RFS information gain expectation that accounts for non-ideal measurements. A new KLD function is shown to be cell-additive and employed to represent information gain for discovered and undiscovered objects and, subsequently, is approximated efficiently using the cell-MB decomposition. The effectiveness of this new approach is demonstrated using real video data in two distinct and challenging tracking applications involving multiple closely-spaced ground and marine vehicles maneuvering in a cluttered and remote environment. The proposed approach is demonstrated by tracking and maintaining discovered vehicles using an optical sensor with a bounded FoV, while simultaneously searching for and discovering new vehicles as they enter the surveillance region.

2 PROBLEM FORMULATION

This paper considers an online SWT problem involving a single sensor with a bounded and mobile FoV that can be manipulated by an automatic controller, as illustrated in Fig. 1. The sensor objective is to discover and track multiple unidentified moving objects in a region of interest (ROI) that far exceeds the size of the FoV. The objects are characterized by partially hidden states and are subject to unknown random inputs, such as driver commands, and may leave and enter the ROI at any time. 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 is unknown *a priori* and changes over time because objects enter and exit the surveillance region as well as, potentially, the sensor FoV. Let N_k denote

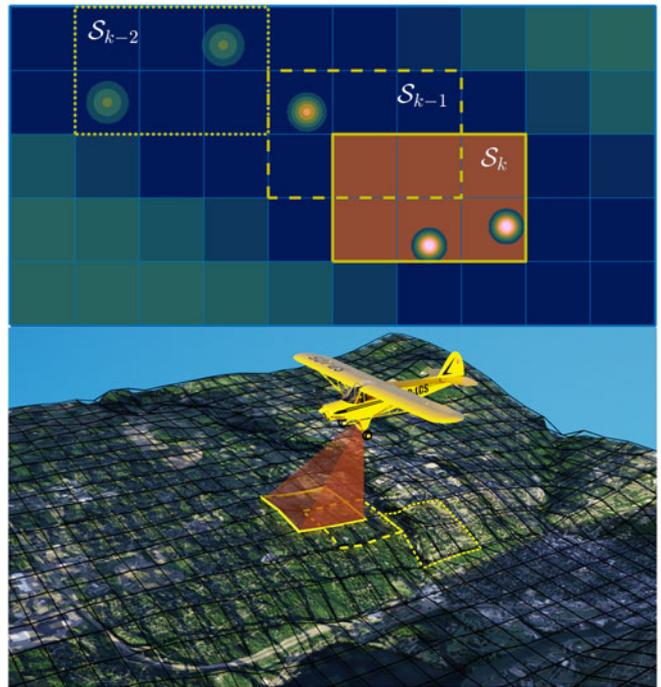


Fig. 1. Conceptual image of multi-object search-while-tracking, wherein the sensor field-of-view S is controlled to maximize the cell multi-Bernoulli approximated information gain.

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

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

where $\mathbf{x}_{k,i}$ is the i th element of X_k and $\mathcal{F}(\mathbb{X})$ denotes the collection of all finite subsets of the object state space \mathbb{X} . Throughout this paper, single-object states are represented by lowercase letters, while multi-object states are represented by finite sets and denoted by italic uppercase letters. Bold lowercase letters are used to denote vectors. Spaces are represented by blackboard bold symbols, where \mathbb{N}_ℓ denotes the set of natural numbers

$$\mathbb{N}_\ell \triangleq \{1, \dots, \ell\} \quad (2)$$

The multi-object measurement is the collection of M_k single-object measurements at time t_k and is expressed as the set

$$Z_k = \{\mathbf{z}_{k,1}, \dots, \mathbf{z}_{k,M_k}\} \in \mathcal{F}(\mathbb{Z}) \quad (3)$$

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 labels or features, the association between tracked objects and incoming measurement data is unknown.

Often in tracking, object detection may depend only a partial state $\mathbf{s} \in \mathbb{X}_s \subseteq \mathbb{R}^{n_s}$, where $\mathbb{X}_s \times \mathbb{X}_v = \mathbb{X} \subseteq \mathbb{R}^{n_x}$ forms the full object state space. For example, the instantaneous ability of a sensor to detect an object may depend only on the object's relative position. In that case, \mathbb{X}_s is the position space, and \mathbb{X}_v is composed of non-position states, such as object velocity. This nomenclature is adopted throughout

the paper while noting that the approach is applicable to other state definitions.

The sensor FoV is defined as a compact subset $\mathcal{S}_k \subset \mathbb{X}_s$. Then, object detection is assumed to be random and characterized by the probability function

$$p_{D,k}(\mathbf{x}_k; \mathcal{S}_k) = 1_{\mathcal{S}_k}(\mathbf{s}_k) \cdot p_{D,k}(\mathbf{s}_k) \quad (4)$$

where $1_{\mathcal{S}_k}(\mathbf{s}_k)$ is the indicator function and where the single-argument function $p_{D,k}(\mathbf{s}_k)$ is the probability of object detection for an unbounded FoV. When an object is detected, a noisy measurement of its state \mathbf{x}_k is produced according to the likelihood function

$$\mathbf{z}_k \sim g_k(\mathbf{z}_k | \mathbf{x}_k) \quad (5)$$

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 image plane is equally likely to produce a false alarm, and thus, the clutter process is modeled as a Poisson RFS process with PHD $\kappa_{c,k}(\mathbf{z})$ [17]. Further discussion on Poisson RFSs and the PHD function can be found in Section 3.1.

Let $\mathbf{u}_k \in \mathcal{U}_k$ denote the sensor control inputs that, through actuation, determine the position of the sensor FoV at time t_k , \mathcal{S}_k , where \mathcal{U}_k is the set of all admissible controls. The control \mathbf{u}_k influences both the FoV geometry, \mathcal{S}_k , and the sensor measurements, Z_k , 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. Therefore, given the admissible control inputs \mathcal{U}_k , let the field-of-regard (FoR) be defined as

$$\mathcal{T}_k \triangleq \bigcup_{\mathbf{u}_k \in \mathcal{U}_k} \mathcal{S}_k(\mathbf{u}_k) \quad (6)$$

and represent the composite of regions that the sensor can potentially cover (although not simultaneously) at the next time step.

Then, the sensor control problem can be formulated as an RFS POMDP [9], [18], [19], that includes a partially- and noisily-observed state X_k , a known initial distribution of the state $f_0(X_0)$, a probabilistic transition model $f_{k|k-1}(X_k | X_{k-1})$, a set of admissible control actions \mathcal{U}_k , and a reward \mathcal{R}_k associated with each control action. At every time k , an RFS multi-object tracker provides the prior $f_{k|k-1}(X_k | Z_{0:k-1})$ and the sensor control input is chosen so as to maximize the expected information gain, or,

$$\mathbf{u}_k^* = \arg \max_{\mathbf{u}_k \in \mathcal{U}_k} \{E[\mathcal{R}_k(Z_k; \mathcal{S}_k, f_{k|k-1}(X_k | Z_{0:k-1}), \mathbf{u}_{0:k-1})]\} \quad (7)$$

where $E[\cdot]$ is the expectation operator and the functional dependence of Z_k and \mathcal{S}_k on \mathbf{u}_k is omitted for brevity here but is described in [20]. In this paper, \mathcal{R}_k is taken to be an information gain function, while noting that the presented results are more broadly applicable to any integrable reward function satisfying the cell-additivity constraint defined in Section 4.

A computationally tractable approximation of the expected information gain in (7) is derived using the new cell-MB approximation presented in Section 4. Based on this approximation, a

new sensor control policy for SWT applications is obtained in Section 5 using a dual information gain function. The dual information gain formulation treats discovered and undiscovered objects as separate processes, modeling undiscovered objects as a partially piecewise homogeneous Poisson process. By this approach, a computationally efficient sensor controller is developed for SWT over potentially large geographic regions.

3 BACKGROUND ON RANDOM FINITE SETS

RFS theory is a powerful framework for solving multi-sensor multi-object information fusion problems. In essence, RFS theory establishes multi-object analogs to random variables, density functions, moments, and other statistics, such that multi-object problems can be solved in a top-down fashion and with theoretic guarantees. For readers unfamiliar with RFS theory, [21] provides a gentle introduction to the subject, while [20], [22] provide a comprehensive treatment.

An RFS X is a random variable that takes values on $\mathcal{F}(\mathbb{X})$. A labeled random finite set (LRFS) \tilde{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) [20], [22]. This section reviews key RFS concepts and notation for the Poisson RFS, multi-Bernoulli (MB) RFS, and generalized labeled multi-Bernoulli (GLMB) LRFS distributions used in this paper.

3.1 Poisson RFS

The Poisson RFS is fundamental to RFS multi-object tracking due to its desirable mathematical properties and its usage in modeling false alarm and birth processes. For example, the popular PHD filter is derived from the assumption that the multi-object state is governed by a Poisson RFS process, which, in turn, leads to a computationally efficient tracking algorithm [23], [24], [25].

The density of a Poisson-distributed RFS X is

$$f(X) = e^{-N_X} [D]^X \quad (8)$$

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} . For brevity, the multi-object exponential notation,

$$h^A \triangleq \prod_{a \in A} h(a) \quad (9)$$

where $h^0 \triangleq 1$, is adopted throughout. For multivariate functions, the dot “ \cdot ” denotes the argument of the multi-object exponential, e.g.,:

$$[g(a, \cdot, c)]^B \triangleq \prod_{b \in B} g(a, b, c) \quad (10)$$

The PHD is an important statistic in RFS theory as its integral over a set $T \subseteq \mathbb{X}$ gives the expected number of objects in that set:

$$E[|X \cap T|] = \int_T D(\mathbf{x}) d\mathbf{x} \quad (11)$$

The PHD of a general RFS X is given in terms of its set density $f(X)$ as [23]

$$D(\mathbf{x}) = \int f(\{\mathbf{x}\} \cup X') \delta X' \quad (12)$$

The integral in (12) is a set integral, defined as

$$\int f(X) \delta X \triangleq \sum_{n=0}^{\infty} \frac{1}{n!} \int f(\{\mathbf{x}_1, \dots, \mathbf{x}_n\}) d\mathbf{x}_1 \cdots d\mathbf{x}_n \quad (13)$$

The set integral is a fundamental construct of RFS theory and enables the direct translation of the Bayes' filter recursion to the multi-object setting, as shown in [22] and discussed in Section 3.4. Set integration via (13) also presents practical challenges, as exact computation is rarely possible due to the infinite summation of nested multivariate integrals required. This challenge is a key motivation of the tractable cell-MB approximation introduced in Section 4.

3.2 Multi-Bernoulli RFS

In an MB distribution, a given object's existence is modeled as a Bernoulli random variable and specified by a probability of existence. As such, the MB RFS can accurately model a variety of multi-object processes when the true existence of objects is unknown and subject to change. The density of an MB distribution is [20, p. 102]

$$f(X) = \left[1 - r^{(\cdot)}\right]^{N_M} \sum_{1 \leq i_1 \neq \dots \neq i_n \leq M} \left[\frac{r^{i(\cdot)} p^{i(\cdot)}(\mathbf{x}_{(i)})}{1 - r^{i(\cdot)}} \right]^{N_n} \quad (14)$$

where $n = |X|$, 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. Given an MB distribution with density (14), its PHD is given by

$$D(\mathbf{x}) = \sum_{j=1}^M r^j p^j(\mathbf{x}) \quad (15)$$

3.3 GLMB RFS

The density of a GLMB distribution proposed in [26] is given by

$$\dot{f}(\dot{X}) = \Delta(\dot{X}) \sum_{\xi \in \Xi} w^{(\xi)}(\mathcal{L}(\dot{X})) [p^{(\xi)}]^{\dot{X}}, \quad (16)$$

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

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

The label of a labeled state \dot{x} is recovered by $\mathcal{L}(\dot{x})$, 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, $\mathcal{L}(\dot{X}) \triangleq \{\mathcal{L}(\dot{x}) : \dot{x} \in \dot{X}\}$. The distinct label indicator $\Delta(\dot{X}) = \delta_{(|\dot{X}|)}(|\mathcal{L}(\dot{X})|)$ ensures that only sets with distinct labels are considered.

3.4 Multi-Object Filtering

Online estimation of the multi-object state is performed using the data-driven GLMB filter, which provides the

Bayes-optimal solution of the measurement-driven Bayes filter recursion [27]:

$$\dot{f}_p(\dot{X}_{p,k} | Z_{0:k-1}) = \int \dot{f}(\dot{X}_{p,k} | \dot{X}_{k-1}) \dot{f}(\dot{X}_{k-1} | Z_{0:k-1}) \delta \dot{X}_{k-1} \quad (17)$$

$$\dot{f}(\dot{X}_k | Z_{0:k}) = \frac{g(Z_k | \dot{X}_k) \dot{f}_p(\dot{X}_{p,k} | Z_{0:k-1}) \dot{f}_b(\dot{X}_{b,k})}{\int g(Z_k | \dot{X}) \dot{f}_p(\dot{X}_{p,k} | Z_{0:k-1}) \dot{f}_b(\dot{X}_{b,k}) \delta \dot{X}} \quad (18)$$

The function time indices have been suppressed for brevity, and $\dot{f}_{p,k}(\dot{X}_{p,k})$ and $\dot{f}_{b,k}(\dot{X}_{b,k})$ denote the density of persisting and birth objects, respectively, where the joint state $\dot{X}_k = \dot{X}_{p,k} \cup \dot{X}_{b,k}$. The accent "˙" is used to distinguish labeled states and functions 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. $\dot{f}_{k|k-1}(\dot{X}_{p,k} | \dot{X}_{k-1})$ is the multi-object transition density, $g_k(Z_k | \dot{X}_k)$ is the multi-object measurement likelihood function, and g_k is used to denote both the single-object and multi-object measurement likelihood function. The nature of the likelihood function can be easily determined from its arguments.

3.5 Kullback-Leibler Divergence

Like single-object distributions, the similarity of RFS distributions may be measured by the KLD. Let f_1 and f_0 be integrable set densities where f_1 is absolutely continuous with respect to f_0 . Then, the KLD is [8, p. 206]

$$I(f_1; f_0) = \int f_1(Y) \log \left(\frac{f_1(Y)}{f_0(Y)} \right) \delta Y \quad (19)$$

Further simplification is possible if f_0 and f_1 are Poisson with respective PHDs D_0 and D_1 , in which case

$$I(f_1; f_0) = N_0 - N_1 + \int D_1(\mathbf{y}) \cdot \log \left(\frac{D_1(\mathbf{y})}{D_0(\mathbf{y})} \right) d\mathbf{y} \quad (20)$$

where $N_0 = \int D_0(\mathbf{y}) d\mathbf{y}$ and $N_1 = \int D_1(\mathbf{y}) d\mathbf{y}$. Importantly, when f_0 and f_1 represent prior and posterior densities, respectively, the KLD is a measure of information gain. By formulating control policies based on RFS divergence measures, the complicated effects of spatial uncertainty, false alarms, missed detections, existence uncertainty, and object appearance/disappearance are elegantly captured in a compact and abstract objective, as depicted in Fig. 2 and described in the following section.

4 INFORMATION-DRIVEN CONTROL

The objective of information-driven control is to maximize the value of the information gained by future measurements before they are known to the sensor. The expected information gain, therefore, can be obtained by marginalizing over the set Z_k , using an available measurement model. Then, the expected information gain obtained at the *next* time step can be obtained from the set integral

$$E[\mathcal{R}_k] = \int \mathcal{R}_k(Z_k; \cdot) f(Z_k) \delta Z_k \quad (21)$$

where $f(Z_k)$ is the predicted measurement density conditioned on past measurements. In general, direct evaluation of



Fig. 2. Graphical representation of important SWT considerations that are elegantly encapsulated in a compact and abstract RFS information gain function.

(21) is computationally intractable due to the infinite summation of nested single-object integrals (see (13)). 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 expected information gain.

4.1 The Cell-MB Distribution

A new approximation of RFS density functions is presented in this section and then used to obtain the information gain expectation. This approach, referred to hereon as the cell-MB approach, 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 decomposition of the space \mathbb{Y} into P disjoint subspaces, or cells, as

$$\mathbb{Y} = \mathbb{Y}^1 \uplus \dots \uplus \mathbb{Y}^P \quad (22)$$

Given the cell decomposition (22), the RFS

$$Y = \{\mathbf{y}_1, \dots, \mathbf{y}_n\}$$

is considered to be cell-MB if it is distributed according to the density

$$f(Y) = \Delta(Y, \mathbb{Y}) \left[1 - r^{(\cdot)} \right]^{N_P} \cdot \sum_{1 \leq j_1 \neq \dots \neq j_n \leq P} \left[\frac{r^{j(\cdot)} p^{j(\cdot)}(\mathbf{y}_{(\cdot)})}{1 - r^{j(\cdot)}} \right]^{N_n} \quad (23)$$

where

$$\Delta(Y, \mathbb{Y}) \triangleq \begin{cases} 1 & |Y \cap \mathbb{Y}^j| \leq 1 \forall j \in \{1, \dots, P\} \\ 0 & \text{otherwise} \end{cases} \quad (24)$$

and

$$\int_{\mathbb{Y}} p^j(\mathbf{y}) d\mathbf{y} = 1, \quad j = 1, \dots, P \quad (25)$$

Note that 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.

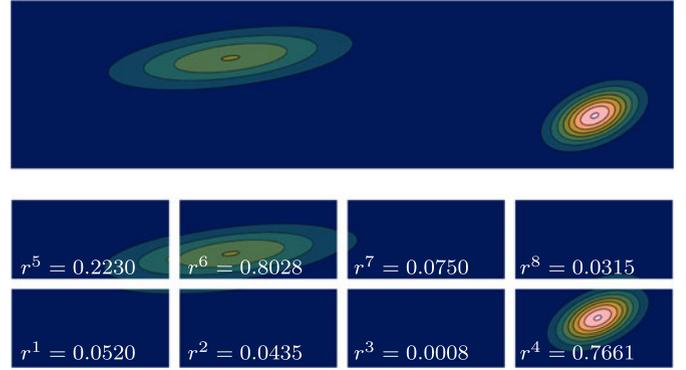


Fig. 3. The PHD of an RFS distribution $f(Y)$ (top) and its cell-MB approximation (bottom), where each cell j has a corresponding probability of existence r^j and spatial density $p^j(\mathbf{y})$.

In [28], a collection of Bernoulli distributions was defined over an occupancy grid by integration of the PHD for dynamic map estimation applications. Inspired by [28], in this paper, a general cell-MB approximation is developed for an arbitrary density and appropriate cell decomposition. The following proposition shows that 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.

Proposition 1. Let $f(Y)$ be an arbitrary set density with PHD $D(\mathbf{y})$ and $\mathbb{Y}^1 \uplus \dots \uplus \mathbb{Y}^P$ be a cell decomposition of space \mathbb{Y} such that

$$\int_{\mathbb{Y}^j} D(\mathbf{y}) d\mathbf{y} \leq 1, \quad j = 1, \dots, P \quad (26)$$

If $\bar{f}(Y)$ is a cell-MB distribution 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

$$r^j = \int_{\mathbb{Y}^j} 1_{\mathbb{Y}^j}(\mathbf{y}) D(\mathbf{y}) d\mathbf{y} \quad (27)$$

$$p^j(\mathbf{y}) = \frac{1}{r^j} 1_{\mathbb{Y}^j}(\mathbf{y}) D(\mathbf{y}) \quad (28)$$

The proof is provided in Appendix A, available online. An example 2×4 cell decomposition and corresponding cell-MB approximation is shown in Fig. 3 for illustration. As shown, the cell-MB approximation has a matching PHD surface, but within each cell j , the spatial density $p^j(\mathbf{y})$ is confined to the cell support \mathbb{Y}^j . When applied to the predicted measurement density, the cell-MB approximation results in a simplified multi-object expectation for a restricted class of information gain functions, as described in the following subsection.

4.2 Information Gain Expectation: Cell-MB

In order to reduce the computational complexity associated with the set integral in (21), this subsection shows that the multi-object information gain expectation simplifies to a finite sum involving only single-object integrals, assuming the measurement is cell-MB distributed and the information gain function in (21) is cell-additive, as defined in this subsection.

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

$$\overset{j}{\mathcal{S}} \triangleq \mathcal{S} \cap \overset{j}{\mathbb{X}}_s \quad (29)$$

Furthermore, 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} :

$$\overset{j}{\mathbb{X}}_s \setminus \overset{j}{\mathcal{S}} = \emptyset \quad \forall \overset{j}{\mathcal{S}} \neq \emptyset \quad (30)$$

This assumption is without loss of generality, as any violation is resolved by subdividing a cell $\overset{j}{\mathbb{X}}_s$ into two new cells $\overset{j}{\mathbb{X}}_s \cap \overset{j}{\mathcal{S}}$ and $\overset{j}{\mathbb{X}}_s \setminus \overset{j}{\mathcal{S}}$. Then, the cell-additivity condition can be defined as follows.

Definition 2. Given a decomposition $\overset{1}{\mathbb{Z}} \uplus \dots \uplus \overset{P}{\mathbb{Z}}$ of space \mathbb{Z} , the information gain function $\mathcal{R}_k(\cdot)$ is cell-additive if

$$\mathcal{R}_k(\mathbb{Z}_k; \mathcal{S}_k) = \sum_{j=1}^P \mathcal{R}_k(\mathbb{Z}_k \cap \overset{j}{\mathbb{Z}}; \overset{j}{\mathcal{S}}_k) \quad (31)$$

Theorem 1. Let Z_k be distributed according to the cell-MB density $f(Z_k)$ with parameters $\{r^j, p^j\}_{j=1}^P$ and the cell decomposition

$$\mathbb{Z} = \overset{1}{\mathbb{Z}} \uplus \dots \uplus \overset{P}{\mathbb{Z}} \quad (32)$$

If the information gain function $\mathcal{R}_k(\cdot)$ is integrable and cell-additive (Def. 2), then the expected information gain is

$$\mathbb{E}[\mathcal{R}_k] = \sum_{j=1}^P \mathcal{R}_k(\emptyset; \mathcal{S}_k) (1 - r^j) + \hat{\mathcal{R}}_{z,k}^j \cdot r^j \quad (33)$$

where

$$\hat{\mathcal{R}}_{z,k}^j \triangleq \int_{\overset{j}{\mathbb{Z}}} \mathcal{R}_k(\{\mathbf{z}\}; \overset{j}{\mathcal{S}}_k) p^j(\mathbf{z}) d\mathbf{z} \quad (34)$$

Proof of Theorem 1 is given in Appendix B, available in the online supplemental material.

Remark: In (31), (33), and (34), the auxiliary information gain arguments are suppressed for brevity and to highlight the structure of the cell-MB approximation.

The remainder of this paper considers information gain functions satisfying the cell-additivity constraint of (31), such as the PHD filter update based KLD information gain. 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 filter update based (hereon abbreviated as ‘‘PHD-based’’) KLD information gain is

$$\begin{aligned} \mathcal{R}_k(\mathbb{Z}; \mathcal{S}, D_{k|k-1}) &= \int_{\mathbb{X}} D_{k|k-1}(\mathbf{x}) \\ &\times \{1 - L_{\mathbb{Z}}(\mathbf{x}; \mathcal{S}) + L_{\mathbb{Z}}(\mathbf{x}; \mathcal{S}) \log [L_{\mathbb{Z}}(\mathbf{x}; \mathcal{S})]\} d\mathbf{x} \end{aligned} \quad (35)$$

where the pseudo-likelihood function

$$\begin{aligned} L_{\mathbb{Z}}(\mathbf{x}; \mathcal{S}) &= 1 - p_D(\mathbf{x}; \mathcal{S}) \\ &+ \sum_{z \in \mathbb{Z}} \frac{p_D(\mathbf{x}; \mathcal{S}) \cdot g(\mathbf{z}|\mathbf{x})}{\kappa_c(\mathbf{z}) + \int p_D(\mathbf{x}; \mathcal{S}) g(\mathbf{z}|\mathbf{x}) D_{k|k-1}(\mathbf{x}) d\mathbf{x}} \end{aligned} \quad (36)$$

is adopted from [20, p. 193]. The following proposition establishes that (35) is cell-additive for appropriate cell decompositions.

Proposition 2. Assume there exists a joint decomposition

$$\mathbb{Z} = \overset{1}{\mathbb{Z}} \uplus \dots \uplus \overset{P}{\mathbb{Z}}, \quad \mathbb{X} = \overset{1}{\mathbb{X}} \uplus \dots \uplus \overset{P}{\mathbb{X}} \quad (37)$$

such that (30) is satisfied, and assume that an object in cell $\overset{j}{\mathbb{X}}$ can only generate measurements within its corresponding measurement cell $\overset{j}{\mathbb{Z}}$; i.e.,:

$$D_{k|k-1}(\mathbf{x}) g_k(\mathbf{z}|\mathbf{x}) = 0 \quad \forall \mathbf{x} \in \overset{j}{\mathbb{X}}, \mathbf{z} \in \overset{j'}{\mathbb{Z}}, j \neq j' \quad (38)$$

Then, the PHD-based KLD is cell-additive:

$$\mathcal{R}_k(\mathbb{Z}; \mathcal{S}, D_{k|k-1}) = \sum_{j=1}^P \mathcal{R}_k(\mathbb{Z} \cap \overset{j}{\mathbb{Z}}; \overset{j}{\mathcal{S}}, D_{k|k-1}) \quad (39)$$

Proof of Proposition 2 is provided in Appendix C, available in the online supplemental material. Proposition 2 establishes that for appropriate cell decompositions, the PHD-based KLD for a given FoV is equivalent to the sum of PHD-based KLD information gains for smaller ‘‘virtual’’ FoVs. Perfect cell-additivity requires satisfying (38), which, in turn, implies that an object in cell $\overset{j}{\mathbb{X}}$ does not generate a measurement in $\overset{i}{\mathbb{Z}}$ for $i \neq j$. In general, violations of (38) 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 information gain 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 undiscovered objects, as is shown in the following section.

5 SEARCH-WHILE-TRACKING SENSOR CONTROL

This section presents a dual information gain function and associated sensor control policy that takes into account both discovered and undiscovered objects. The information gain function proposed in Section 5.1 balances the competing objectives of object search and tracking by means of a unified information-theoretic framework. Sections 5.2 and 5.3 derive the expected information gain functions for discovered and undiscovered objects, respectively, the combination of which is maximized by the sensor control policy in Section 5.4. Sections 5.5 and 5.6 describe multi-object filters for recursive estimation of the undiscovered and discovered object densities, respectively, and the overall SWT algorithm is summarized in Section 5.7.

5.1 Dual Information Gain Function

Separate density parameterizations for discovered and undiscovered 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 $0, \dots, k-1$ 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

information gain as

$$\mathbf{u}_k^* = \arg \max_{\mathbf{u}_k \in \mathcal{U}_k} \left\{ \begin{aligned} &E[\mathcal{R}_k^d(V_k; \mathcal{S}_k(\mathbf{u}_k))] \\ &+ E[\mathcal{R}_k^u(W_k; \mathcal{S}_k(\mathbf{u}_k))] \end{aligned} \right\} \quad (40)$$

where

$$\mathcal{R}_k^d(\cdot; \cdot) = \mathcal{R}_k(\cdot; \cdot, D_{d,k|k-1}) \quad (41)$$

$$\mathcal{R}_k^u(\cdot; \cdot) = \mathcal{R}_k(\cdot; \cdot, D_{u,k|k-1}) \quad (42)$$

are used for brevity, and $D_{d,k|k-1}$ and $D_{u,k|k-1}$ are the prior PHDs of discovered and undiscovered objects, respectively. The individual information gain expectations for discovered and undiscovered objects are derived in the following subsections.

5.2 Expected Information Gain of Discovered Objects

If $f_{k|k-1}(V_k)$ is cell-MB with parameters $\{r_v^j, p_v^j\}_{j=1}^P$, then from Theorem 1 it follows that

$$E[\mathcal{R}_k^d] = \sum_{j=1}^P \mathcal{R}_k^d(\emptyset; \mathcal{S}_k^j)(1 - r_v^j) + \hat{\mathcal{R}}_{v,k}^{d,j}(\mathcal{S}_k) \cdot r_v^j \quad (43)$$

where

$$\hat{\mathcal{R}}_{v,k}^{d,j}(\mathcal{S}) \triangleq \int_{\mathcal{Z}} \mathcal{R}_k^d(\{\mathbf{z}\}; \mathcal{S}) p_v^j(\mathbf{z}) d\mathbf{z} \quad (44)$$

$$r_v^j(\mathcal{S}) = \int_{\mathcal{Z}} \mathbf{1}_j(\mathbf{z}) D_{v,k|k-1}(\mathbf{z}; \mathcal{S}) d\mathbf{z} \quad (45)$$

$$p_v^j(\mathbf{z}; \mathcal{S}) = \frac{1}{r_v^j} \mathbf{1}_j(\mathbf{z}) D_{v,k|k-1}(\mathbf{z}; \mathcal{S}) \quad (46)$$

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

$$D_{d,k|k-1}(\mathbf{x}) = \sum_{(I,\xi) \in \mathcal{F}(L) \times \Xi} \sum_{\ell \in I} w^{(\xi)}(I) p^{(\xi)}(\mathbf{x}, \ell) \quad (47)$$

The PHD $D_{v,k|k-1}$ can be obtained from the predicted measurement density $f_{k|k-1}(V_k)$ through application of (12). From the prior GLMB density,

$$f_{k|k-1}(V_k) = \int g_k(V_k | \hat{X}) \hat{f}_{k|k-1}(\hat{X}) \delta \hat{X} \quad (48)$$

Given a GLMB prior, explicit computation of the predicted measurement density is computationally challenging because it requires summation over all possible object-to-measurement association hypotheses. Instead, the discovered object PHD is readily obtained from the GLMB prior, from which $D_{v,k|k-1}$ is approximated as

$$D_{v,k|k-1}(\mathbf{z}; \mathcal{S}) \approx \int D_{d,k|k-1}(\mathbf{x}) p_{D,k}(\mathbf{x}; \mathcal{S}) g_k(\mathbf{z}|\mathbf{x}) d\mathbf{x} + \kappa_{c,k}(\mathbf{z}) \quad (49)$$

Because an analytic solution of the integral in (44) is not available, a numerical quadrature rule is employed. In the proposed approach, a measurement cell is further tessellated into regions $\{\Omega_i^j\}_{i=1}^{R_j} \subset \mathcal{Z}$ based on the anticipated information value of measurements within each region, as illustrated in Fig. 4. Then, given a representative

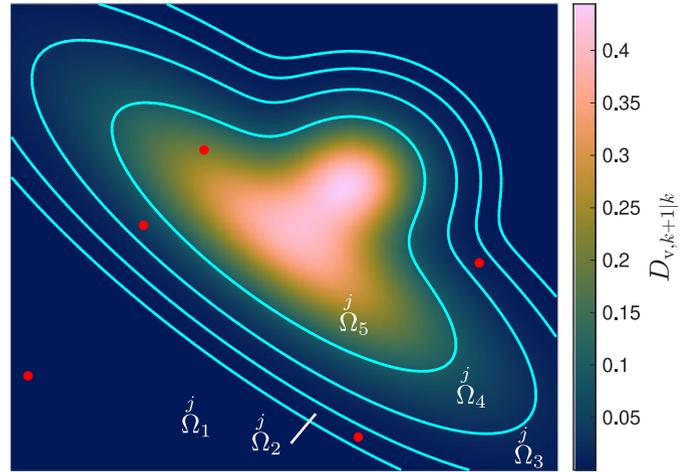


Fig. 4. Example quadrature of the single-measurement conditional expected information gain, where representative measurements $\mathbf{z}_{j,i}$ are denoted by red dots and quadrature regions are outlined in cyan.

measurement $\mathbf{z}_{j,i}$ for each region, the conditional information gain expectation is approximated as

$$\hat{\mathcal{R}}_{v,k}^{d,j}(\mathcal{S}) \approx \sum_{i=1}^{R_j} \mathcal{R}_k^d(\{\mathbf{z}_{j,i}\}; \mathcal{S}) p_v^j(\mathbf{z}_{j,i}) A_{j,i} \quad (50)$$

where $A_{j,i}$ is the volume of region Ω_i^j . By this approach, the PHD-based KLD information gain function is evaluated only R_j times. Further details regarding the computation of the quadrature regions and representative measurement points are provided in Appendix D, available in the online supplemental material.

5.3 Expected Information Gain of Undiscovered Objects

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

$$D_{u,k|k-1}(\mathbf{s}) = \sum_{j=1}^P \frac{\mathbf{1}_j(\mathbf{s})}{A(\hat{X}_s^j)} \cdot \lambda_{j,k|k-1} \quad (51)$$

where $\lambda_{j,k|k-1}$ is the expected number of undiscovered objects in \hat{X}_s^j at time step k and $A(\hat{X}_s^j)$ is the volume of cell \hat{X}_s^j . For ease of exposition, the undiscovered object PHD is modeled using the same cell decomposition employed in the cell-MB approximation. Modeling undiscovered objects as a Poisson point process is one of the core ideas of the PMBM filter, where discovered objects are modeled as a multi-Bernoulli mixture (MBM) RFS. While discovered objects are modeled as a GLMB distribution in this work, the cell-MB SWT framework is amenable to any discovered object RFS prior, including Poisson, independently and identically distributed cluster (i.i.d.c.), MB, MBM, labeled multi-Bernoulli (LMB), and GLMB distributions.

If $f_{k|k-1}(W_k)$ is cell-MB with parameters $\{r_w^j, p_w^j\}_{j=1}^P$, then by Theorem 1,

$$E[\mathcal{R}_k^u] = \sum_{j=1}^P \mathcal{R}_k^u(\emptyset; \mathcal{S}_k) (1 - r_w^j) + \hat{\mathcal{R}}_{w,k}^{u,j}(\mathcal{S}_k) \cdot r_w^j \quad (52)$$

where

$$\hat{\mathcal{R}}_{w,k}^{u,j}(\mathcal{S}) \triangleq \int_{\mathcal{Z}} \mathcal{R}_k^u(\{\mathbf{z}\}; \mathcal{S}) p_w^j(\mathbf{z}) d\mathbf{z} \quad (53)$$

$$r_w^j(\mathcal{S}) = \int_{\mathcal{Z}} 1_j(\mathbf{z}) D_{w,k|k-1}(\mathbf{z}; \mathcal{S}) d\mathbf{z} \quad (54)$$

$$\approx \frac{\lambda_{j,k|k-1}}{A(\mathbb{X}_s)} \int_{\mathbb{X}_s} p_D(\mathbf{s}; \mathcal{S}) d\mathbf{s} \quad (55)$$

$$p_w^j(\mathbf{z}; \mathcal{S}) = \frac{1}{r_w^j} 1_j(\mathbf{z}) D_{w,k|k-1}(\mathbf{z}; \mathcal{S}) \quad (56)$$

$$D_{w,k|k-1}(\mathbf{z}; \mathcal{S}) = \int D_{u,k|k-1}(\mathbf{x}) p_{D,k}(\mathbf{x}; \mathcal{S}) g_k(\mathbf{z}|\mathbf{x}) d\mathbf{x} + \kappa_{c,k}(\mathbf{z}) \quad (57)$$

Under a piecewise homogeneous PHD, the undiscovered object information gain simplifies drastically if the measurement likelihood is independent of non-position states: i.e., $g_k(\cdot|\mathbf{x}) = g_k(\cdot|\mathbf{s})$. Following (35),

$$\begin{aligned} \mathcal{R}_k^u(W_k; \mathcal{S}_k) &= \int_{\mathbb{X}_s} D_{u,k|k-1}(\mathbf{s}) \{1 - L_{W_k}(\mathbf{s}; \mathcal{S}_k) \\ &\quad + L_{W_k}(\mathbf{s}; \mathcal{S}_k) \log [L_{W_k}(\mathbf{s}; \mathcal{S}_k)]\} d\mathbf{s} \end{aligned} \quad (58)$$

Given that at most one measurement may exist per cell, two cases need to be considered: the null measurement case and the singleton measurement case. Letting $W_k = \emptyset$, and after some algebraic manipulation, the undiscovered object information gain for a null measurement can be written as

$$\mathcal{R}_k^u(\emptyset; \mathcal{S}_k) = \sum_{j=1}^P \mathcal{R}_k^u(\emptyset; \mathcal{S}_k) \quad (59)$$

$$\mathcal{R}_k^u(\emptyset; \mathcal{S}_k) = \lambda_{j,k|k-1} \cdot d_j \cdot (1 - \delta_\emptyset(\mathcal{S}_k)) \quad (60)$$

$$d_j \triangleq \frac{1}{A(\mathbb{X}_s)} \int_{\mathbb{X}_s} p_D(\mathbf{s}) + (1 - p_D(\mathbf{s})) \log [1 - p_D(\mathbf{s})] d\mathbf{s} \quad (61)$$

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

$$p_D(\mathbf{s}) = p_{D,j} \quad \forall \mathbf{s} \in \mathbb{X}_s \quad (62)$$

then (61) simplifies to

$$d_j = p_{D,j} + (1 - p_{D,j}) \log (1 - p_{D,j}) \quad (63)$$

For the singleton measurement case, similar analytic simplifications of the conditional information gain (53) are limited. However, within a cell, the uniform position density of undiscovered objects is known *a priori* up to an unknown factor $\lambda_{j,k|k-1}$. Thus, the undiscovered object information gain can be pre-computed for efficiency and

$$\hat{\mathcal{R}}_{w,k}^{u,j}(\mathcal{S}_k) \approx \bar{\mathcal{R}}_w^{u,j}(\lambda_{j,k|k-1}) \quad (64)$$

where the function $\bar{\mathcal{R}}_w^{u,j}(\lambda_{j,k|k-1})$ returns interpolated information gain values over $\lambda_{j,k|k-1} \in [0, 1]$.

Remark: In the special case that $p_D(\mathbf{s}) = 1$, the term $\mathcal{R}_k^u(\emptyset; \mathcal{S}_k)$ is equivalent to the search objective term proposed in [15]. The cell-MB approach differs in that perfect detection is not assumed, and that the information gained from detecting an undiscovered object, captured in the term $\hat{\mathcal{R}}_{w,k}^{u,j}$, is also considered.

5.4 Field-of-View Optimization and Sensor Control

Prior to optimization of the FoV, the information gain associated with each cell in the FoR is computed, as described in Algorithm 1. The FoR cell information gains for discovered and undiscovered objects are stored as arrays $\{\mathcal{R}_k^d[j]\}_{j=1}^P$ and $\{\mathcal{R}_k^u[j]\}_{j=1}^P$, respectively. Then, the optimal FoV is found as the one composed of the cells with the highest composite information gain, without reevaluating the information gain. With this, the sensor control that produces the desired optimal FoV can be written as

$$\mathbf{u}_k^* = \arg \max_{\mathbf{u} \in \mathcal{U}_k} \left\{ \sum_{j \in \mathbb{N}_P, \mathbb{X}_s \subseteq \mathcal{S}_k(\mathbf{u})} (\mathcal{R}_k^u[j] + \mathcal{R}_k^d[j]) \right\} \quad (65)$$

where $\mathcal{T}_k^j \triangleq \mathcal{T}_k \cap \mathbb{X}_s^j$.

Algorithm 1. FoR Information Gain Pseudocode

Input: $\mathcal{T}_k, f_{k|k-1}(\hat{X}), D_{u,k|k-1}(\mathbf{x})$
 Compute $D_{d,k|k-1}(\mathbf{x})$ from $f_{k|k-1}(\hat{X})$ (47)

Compute $D_{v,k|k-1}(\mathbf{z}; \mathcal{T}_k)$ (49)

for $j = 1, \dots, P$ **for** j such that $\mathbb{X}_s^j \in \mathcal{T}_k$ **do**

$$r_v^j \leftarrow \int_{\mathcal{Z}} 1_j(\mathbf{z}) D_{v,k|k-1}(\mathbf{z}; \mathcal{T}_k) d\mathbf{z}$$

$$r_w^j \leftarrow \int_{\mathcal{Z}} 1_j(\mathbf{z}) D_{w,k|k-1}(\mathbf{z}; \mathcal{T}_k) d\mathbf{z}$$

$$\text{Compute } \hat{\mathcal{R}}_{v,k}^{d,j}(\mathcal{T}_k) \quad (50)$$

$$\text{Compute } \hat{\mathcal{R}}_{w,k}^{u,j}(\mathcal{T}_k) \quad (64)$$

$$\mathcal{R}_k^d[j] \leftarrow \mathcal{R}_k^d(\emptyset; \mathcal{T}_k) (1 - r_v^j) + \hat{\mathcal{R}}_{v,k}^{d,j}(\mathcal{T}_k) \cdot r_v^j$$

$$\mathcal{R}_k^u[j] \leftarrow \mathcal{R}_k^u(\emptyset; \mathcal{T}_k) (1 - r_w^j) + \hat{\mathcal{R}}_{w,k}^{u,j}(\mathcal{T}_k) \cdot r_w^j$$

end for

return $(\mathcal{R}_k^d[j])_{j=1}^P, (\mathcal{R}_k^u[j])_{j=1}^P$

Remark: Explicit computation of the cell-MB single-object densities p_w^j and p_w^j is not required. Instead, these densities are implicitly computed when evaluating the conditional information gain expectations $\hat{\mathcal{R}}_{v,k}^{d,j}$ and $\hat{\mathcal{R}}_{w,k}^{u,j}$.

5.5 Undiscovered Object Prediction and Update

The prediction and update of the undiscovered object PHD are accomplished using the cell-discretized PHD filter. The prediction step incorporates undiscovered object motion, birth, and death. The undiscovered object distribution parameters are predicted and updated as

$$\lambda_{j,k|k-1} = \lambda_{B,j,k} + \sum_{i=1}^P p_{S,i,k} \cdot P_{j|i} \cdot \lambda_{i,k-1} \quad (66)$$

$$\lambda_{j,k} = \left[1 - p_{D,j} \cdot (1 - \delta_\emptyset^j(\mathcal{S}_k)) \right] \cdot \lambda_{j,k|k-1} \quad (67)$$

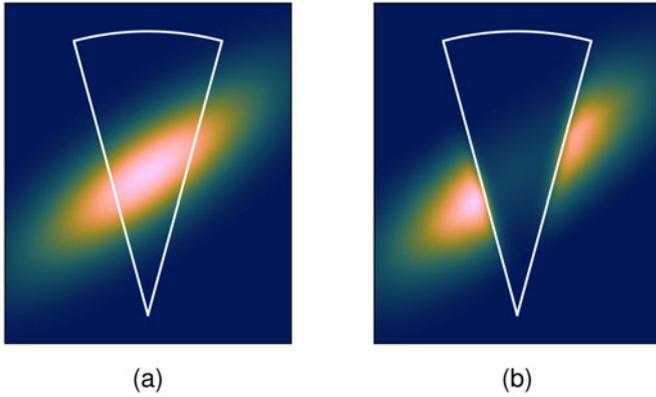


Fig. 5. Prior object density and FoV (a), and posterior object density after recursive split and non-detection (b).

where $\lambda_{B,j,k}$ is the expected number of newborn objects in cell j , $p_{S,i,k}$ is the probability that an undiscovered object in cell i survives, and P_{ji} is the probability that an undiscovered object moves to cell j given that it exists in cell i .

5.6 Discovered Object Tracking

Discovered object tracking is performed using the data-driven GLMB filter. While a detailed description of the data-driven GLMB filter is beyond the scope of this paper, one important consideration is highlighted involving the FoV-dependent nonlinear probability of detection. The data-driven GLMB is implemented in GM form, such that single-object densities are

$$p^{(\xi)}(\mathbf{x}, \ell) = \sum_{i=1}^{J^{(\xi)}(\ell)} w_i^{(\xi)}(\ell) \mathcal{N}(\mathbf{x}; \mathbf{m}_i^{(\xi)}(\ell), \mathbf{P}_i^{(\xi)}(\ell)) \quad (68)$$

It is through the FoV-dependent p_D that the filter probabilistically incorporates the knowledge of where objects were not observed.

In the filter, products of the form $p_D(\mathbf{x}; \mathcal{S})p^{(\xi)}(\mathbf{x})$ are expanded about the GM component means in a zeroth-order Taylor expansion. The accuracy of this approximation is dependent on the GM resolution near the FoV boundaries. Thus, a recursive splitting algorithm [29] is employed that identifies and splits Gaussian components that overlap the FoV boundaries into several “smaller” Gaussians. The resulting $J^{(\xi)}(\ell)$ component mixture replaces the original density, enabling the accurate approximation

$$p_D(\mathbf{x}; \mathcal{S})p^{(\xi)}(\mathbf{x}, \ell) \approx \sum_{i=1}^{J^{(\xi)}(\ell)} w_i^{(\xi, \ell)} p_D(\mathbf{m}_i^{(\xi, \ell)}; \mathcal{S}) \mathcal{N}(\mathbf{x}; \mathbf{m}_i^{(\xi, \ell)}, \mathbf{P}_i^{(\xi, \ell)}) \quad (69)$$

An example is provided in Fig. 5, wherein the prior density is split prior to a Bayes update, allowing for the accurate incorporation of negative information from a non-detection.

5.7 Numerical Implementation

This subsection summarizes the SWT algorithmic implementation. At each step k , a time-update (17), (66) of the previous posterior densities $f(\dot{X}_{k-1}|Z_{0:k-1})$ and $D_{u,k-1}$ yields predicted prior densities for the time of the next decision. The FoR is constructed from admissible control actions as shown in (6),

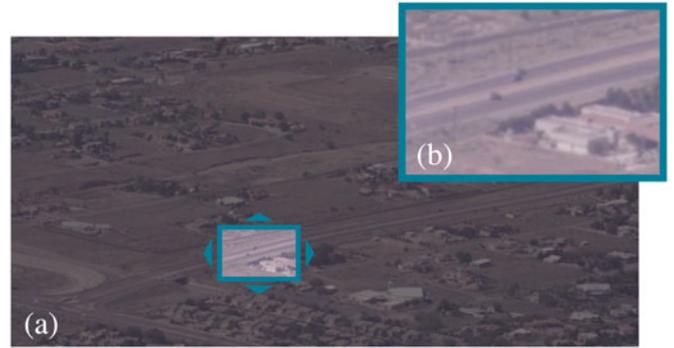


Fig. 6. Example video frame (a) from Albuquerque dataset, artificially windowed to emulate smaller, movable FoV, which is enlarged in (b) to show detail.

and the expected information gain for each cell within the FoR is computed. The candidate FoV that contains the maximizing sum of cell information gains is found, and the corresponding control (65) that yields that FoV is applied. The sensor collects a new multi-object measurement which is processed in the data-driven GLMB filter to update the multi-object density, giving the posterior density in (18). The algorithm is summarized in Algorithm 2.

Algorithm 2. SWT Sensor Control Pseudocode

Input: $f_0(\dot{X}), D_{u,0}(\mathbf{x})$
for $k = 1, \dots, K$ **do**
 $\hat{f}_{k|k-1}(\dot{X}), D_{u,k|k-1}(\mathbf{x}) \leftarrow \text{filter_prediction}(\hat{f}_{k-1}(\dot{X}), D_{u,k-1}(\mathbf{x}))$ (17), (66)
 $(\mathcal{R}_k^d[j])_{j=1}^P, (\mathcal{R}_k^u[j])_{j=1}^P \leftarrow \text{FoR_information_gain}(\mathcal{T}_k, \hat{f}_{k|k-1}(\dot{X}), D_{u,k|k-1}(\mathbf{x}))$ (Algorithm 1)
 $\mathbf{u}_k^* \leftarrow \text{maximize_expected_reward}((\mathcal{R}_k^d[j])_{j=1}^P, (\mathcal{R}_k^u[j])_{j=1}^P)$ (65)
 $\mathcal{S}_k(\mathbf{u}_k^*) \leftarrow \text{apply sensor control}$
 $Z_k \leftarrow \text{obtain measurement}$
 $\hat{f}_{k|k-1}(\dot{X}) \leftarrow \text{split_for_FoV}(\hat{f}_{k|k-1}(\dot{X}), \mathcal{S}_k)$ [29]
 $\hat{f}_{k|k}(\dot{X}), D_{u,k|k} \leftarrow \text{filter_update}(\hat{f}_{k|k-1}(\dot{X}), D_{u,k|k-1}(\mathbf{x}), Z_k, \mathcal{S}_k)$ (18), (67)
end for

6 APPLICATION TO REMOTE MULTI-VEHICLE SWT

The cell-MB SWT framework is demonstrated in two distinct vehicle tracking problems using real video data. The first experiment, hereon referred to as the “Albuquerque” experiment, is based on a video recorded by a fixed camera pointed at a remote location where multiple mobile ground vehicles are observed. The second “Sydney” experiment involves tracking multiple mobile maritime surface vehicles using real satellite video¹ taken of Sydney, Australia from the Chinese low Earth-orbiting satellite, Jilin-1. In both experiments, real-time FoV controlled motion is simulated by windowing the data over a small fraction of the available frame, as illustrated in Fig. 6. These datasets present 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.

1. Video publicly available at <https://mall.charmingglobe.com>

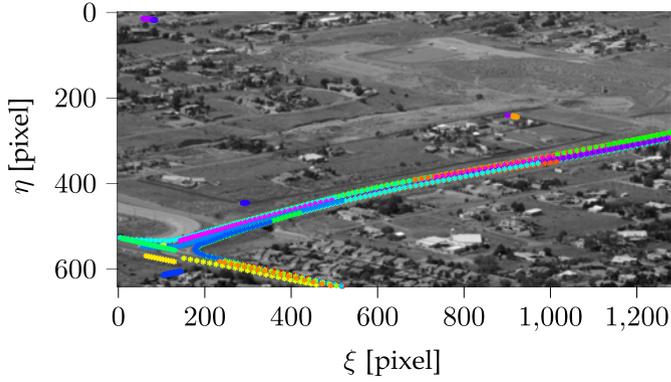


Fig. 7. True trajectories of moving objects with an example image frame as background.

6.1 Vehicle Dynamics

Vehicle dynamics are modeled directly in the image frame. While vehicle dynamics are more naturally expressed in the terrestrial frame, the cameras' precise location and orientation are unknown. Thus, the transformation between image and terrestrial coordinates could not be readily established for the Albuquerque experiment. The Sydney video is georegistered such that world coordinate motion maps directly to scaled image coordinate motion.

The object state is modeled as

$$\mathbf{x}_k = [\mathbf{s}_k^T \quad \boldsymbol{\zeta}_k^T]^T \quad (70)$$

$$\mathbf{s}_k = [\xi_k \quad \eta_k]^T, \quad \boldsymbol{\zeta}_k = [\dot{\xi}_k \quad \dot{\eta}_k \quad \Omega_k]^T \quad (71)$$

where ξ_k and η_k are the horizontal and vertical coordinates, respectively, of the vehicle position with respect to the full-frame origin, $\dot{\xi}_k$ and $\dot{\eta}_k$ are the corresponding rates, and Ω_k is the vehicle turn rate.

Vehicle motion is modeled using the nearly coordinated turn model with directional process noise [17], [30] as

$$\mathbf{x}_{k+1} = \mathbf{f}_k(\mathbf{x}_k) + \boldsymbol{\Gamma}_k \mathbf{v}_k(\mathbf{s}_k) \quad (72)$$

where \mathbf{f}_k is defined in [31, Ch. 11] and

$$\boldsymbol{\Gamma}_k = \begin{bmatrix} \frac{1}{2}(\Delta t)^2 \mathbf{I}_{2 \times 2} & \mathbf{0}_{2 \times 1} \\ (\Delta t) \mathbf{I}_{2 \times 2} & \mathbf{0}_{2 \times 1} \\ \mathbf{0}_{1 \times 2} & \Delta t \end{bmatrix} \quad (73)$$

where $\Delta t = 1$ [sec] 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

$$\mathbb{E}[\mathbf{v}_k \mathbf{v}_k^T] = \mathbf{Q}_k(\mathbf{s}) = \begin{bmatrix} \mathbf{D}^T(\mathbf{s}) \mathbf{Q}_d \mathbf{D}(\mathbf{s}) & \mathbf{0}_{2 \times 1} \\ \mathbf{0}_{1 \times 2} & \sigma_{\Omega, < \text{ABQ, SYD} >}^2 \end{bmatrix} \quad (74)$$

$$\mathbf{Q}_d = \begin{bmatrix} \sigma_{t, < \text{ABQ, SYD} >}^2 & 0 \\ 0 & \sigma_{n, < \text{ABQ, SYD} >}^2 \end{bmatrix} \quad (75)$$

$$\mathbf{D}(\mathbf{s}) = \begin{bmatrix} \cos \Psi(\mathbf{s}) & \sin \Psi(\mathbf{s}) \\ -\sin \Psi(\mathbf{s}) & \cos \Psi(\mathbf{s}) \end{bmatrix} \quad (76)$$

where $\sigma_{\Omega, \text{ABQ}} = 180$ [arcmin/sec] and $\sigma_{\Omega, \text{SYD}} = 30$ [arcmin/sec] are the turn rate process noise standard deviations, $\sigma_{t, \text{ABQ}} = 5$ [pixel/sec²] and $\sigma_{n, \text{ABQ}} = 0.01$ [pixel/sec²] are the standard

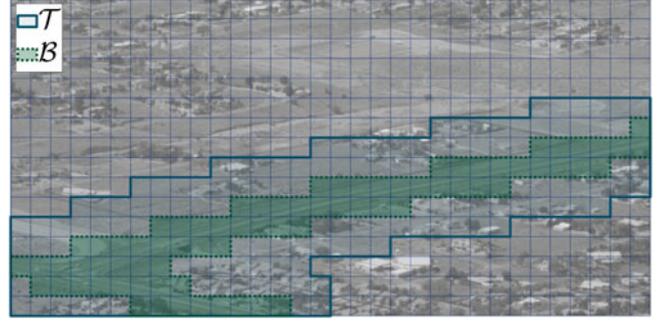


Fig. 8. Albuquerque field-of-regard, \mathcal{T} , and primary road region \mathcal{B} , with an example image frame as background.

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. Information-driven sensor control efficacy fundamentally depends on the accuracy of motion prediction and the rate of uncertainty growth in the absence of observations. Evaluation of multiple candidate motion models revealed that the nearly coordinated turn model with directional process noise offered better motion prediction and more precise uncertainty growth over the simpler constant velocity models. Road geometry is not applicable to the Sydney experiment, and thus an isotropic linear process noise is used with $\sigma_{t, \text{SYD}} = \sigma_{n, \text{SYD}} = 1$ [pixel/sec²]. The true trajectories of all moving objects in the Albuquerque experiment are shown in Fig. 7.

6.2 Sensor and Scene Model

Object detections are generated from raw frame data using normalized difference change detection [32] and fast approximate power iteration subspace tracking [33] for temporal background estimation. The single-object measurement function is linear-Gaussian with corresponding likelihood

$$g_k(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mathbf{z}; \mathbf{H}\mathbf{x}, \mathbf{R}), \quad (77)$$

$$\mathbf{H} = [\mathbf{I}_{2 \times 2} \quad \mathbf{0}_{2 \times 3}], \quad \mathbf{R} = \sigma_{z, < \text{ABQ, SYD} >}^2 \cdot \mathbf{I}_{2 \times 2} \quad (78)$$

where $\sigma_{z, \text{ABQ}}^2 = 9$ [pixel²] and $\sigma_{z, \text{SYD}}^2 = 100$ [pixel²]. The Albuquerque experiment sensor FoV is a rectangular region that is 240 pixels wide and 160 pixels tall. Rectangular FoV geometry is also emulated in the Sydney experiment, where the FoV dimensions are 1024 pixels wide and 640 pixels tall. Moving objects within the FoV are assumed to be detectable with probability $p_{D,k}(\mathbf{s}_k) = 0.9$. The mean false alarm rates are assumed to be five and thirty false detections per scene frame in the Albuquerque and Sydney experiments, respectively.

The Albuquerque and Sydney scenes are tessellated by 16×32 and 24×32 grids, respectively, of uniformly sized rectangular cells as shown in Figs. 8 and 9. Within the Albuquerque scene, an ROI 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 Albuquerque ROI, cells containing road pixels comprise the set \mathcal{B} , which is used to establish an initial uniform distribution of undiscovered objects. In the Sydney experiment, the ROI is defined as the main water region, including the piers and wharves. Thus, following the assumptions established in Section 5.3, the initial undiscovered object position marginal

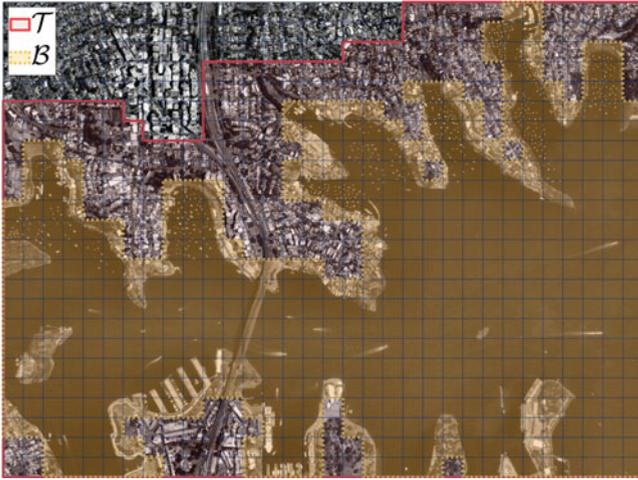


Fig. 9. Sydney field-of-regard, \mathcal{T} , and water region \mathcal{B} , with example image frame as background.

PHD is characterized by (51) with

$$\lambda_{j,0} = \begin{cases} \lambda_{\langle \text{ABQ}, \text{SYD} \rangle} & \text{if } \mathcal{X}^j \subseteq \mathcal{B} \\ 0 & \text{otherwise} \end{cases} \quad (79)$$

where $\lambda_{\text{ABQ}} = 0.137$ and $\lambda_{\text{SYD}} = 0.0593$ correspond to initial estimates of ten and thirty undiscovered objects in the scene, respectively.

6.3 Experiment Results

The Albuquerque and Sydney experiments consist of 60 and 64 time steps, respectively. To emulate a pan/tilt camera

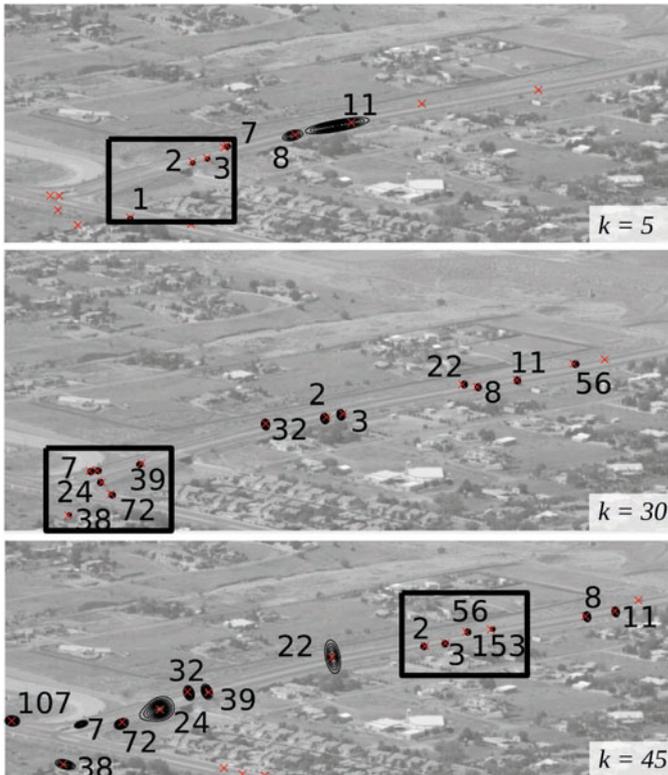


Fig. 10. FoV position and tracker estimates in the form of single-object density contours for objects with probabilities of existence greater than 0.5, shown at select time steps for the Albuquerque experiment.

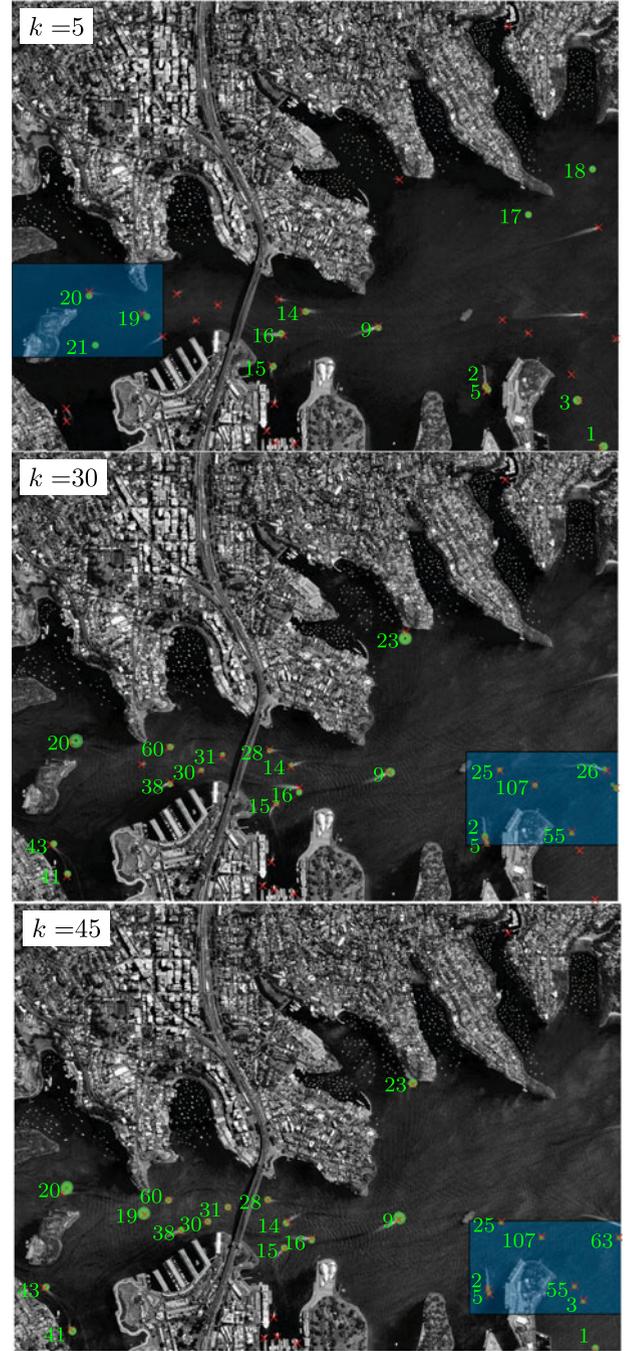


Fig. 11. FoV position and tracker estimates in the form of single-object density contours for objects with probabilities of existence greater than 0.5, shown at select time steps for the Sydney experiment.

from the wider available frame data, the FoV is assumed to be able to be moved 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.

Some key frames of the Albuquerque and Sydney experiments are shown in Figs. 10 and 11, respectively. In the early time steps, the FoV motion is dominated by the undiscovered object component of the information gain. 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 undiscovered objects may exist.

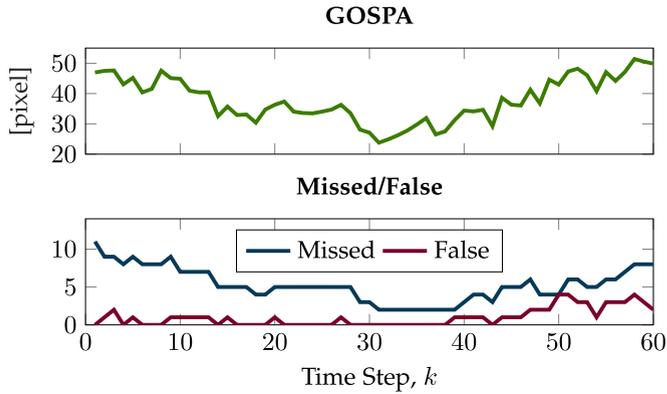


Fig. 12. Albuquerque dataset SWT performance in terms of GOSPA metric and component errors over time using cutoff distance $c = 20$ [pixel], order $p = 2$, and $\alpha = 2$.

Because the overall objective of the SWT sensor control is to reduce multi-object tracking uncertainty, SWT sensor control performance is most naturally quantified by the resulting multi-object tracking accuracy, as measured using the generalized optimal sub-pattern assignment (GOSPA) metric [34]. For the metric parameters selected in this work, the GOSPA metric is equal to the sum of localization errors for properly tracked objects and penalties for missed and false tracks. The GOSPA metric and the number of false and missed objects are shown over time for the Albuquerque and Sydney experiments in Figs. 12 and 13, respectively. The cell-MB SWT sensor control 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 Albuquerque experiment, which is caused by a sharp uptick in new object appearances.

The average GOSPA over the experiment is compared with the PIMS-based information driven control and random FoV motion in Table 1. The cell-MB sensor control achieves significant improvement with respect to other methods in the number of missed and false tracks, as well as the overall GOSPA metric, which encompasses cardinality errors and localization errors. While the PIMS approach exhibits degraded performance in these applications, it should still be considered as a viable method when using an information gain function that is not cell-additive.

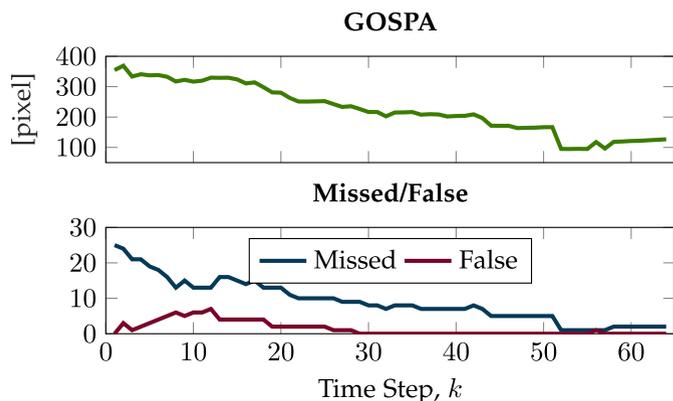


Fig. 13. Sydney dataset SWT performance in terms of GOSPA metric and component errors over time using cutoff distance $c = 100$ [pixel], order $p = 2$, and $\alpha = 2$.

TABLE 1
GOSPA Performance on Albuquerque and Sydney Experiments, Averaged Over Experiment Duration, With Percentage Improvement Over Baseline Random Control Shown Parenthetically

	Algorithm	GOSPA [pixel]	Missed	False
ABQ	Cell-MB	37.84 (56%)	5.27 (168%)	0.97 (158%)
	PIMS	47.46 (24%)	9.95 (42%)	0.90 (178%)
	Random	59.07 (N/A)	14.10 (N/A)	2.50 (N/A)
SYD	Cell-MB	225.00 (41%)	9.03 (85%)	1.41 (97%)
	PIMS	267.44 (19%)	11.98 (39%)	2.31 (20%)
	Random	318.00 (N/A)	16.67 (N/A)	2.78 (N/A)

The computational efficiencies of the cell-MB and PIMS approaches are very similar and are on the order of seconds per decision. The PIMS approximation is a single-sample approximation, and thus requires fewer operations than the cell-MB in general. However, when applied to the undiscovered object information gain presented in this work, the cell-MB approximation is, in fact, cheaper due to the properties of the piecewise homogeneous Poisson point process, which allow the undiscovered object information gain to be pre-computed for each cell and interpolated at runtime. Although not pursued in this work, further computational improvements can be realized in both approaches through further parallelization and software optimization.

7 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 information gain that incorporates both discovered and undiscovered object information gain. In demonstrations using real terrestrial and satellite sensor data, the search-while-tracking sensor control is used to manipulate the sensor fields-of-view to discover and track multiple moving ground vehicles and boats from an aerial vantage point.

REFERENCES

- [1] S. Ferrari and T. A. Wettergren, *Information-Driven Planning and Control*. Cambridge, MA, USA: 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 Trans. Autom. Control*, vol. 64, no. 1, pp. 159–173, Jan. 2019.
- [3] H. Wei et al., "Information value in nonparametric Dirichlet-process Gaussian-process (DPGP) mixture models," *Automatica*, vol. 74, pp. 360–368, 2016.

- [4] C. Cai and S. Ferrari, "Information-driven sensor path planning by approximate cell decomposition," *IEEE Trans. Syst. Man Cybern. Part B*, vol. 39, no. 3, pp. 672–689, Jun. 2009.
- [5] G. Zhang, S. Ferrari, and C. Cai, "A comparison of information functions and search strategies for sensor planning in target classification," *IEEE Trans. Syst., Man, Cybern. Part B*, vol. 42, no. 1, pp. 2–16, Feb. 2012.
- [6] H. G. Hoang, B.-N. Vo, B.-T. Vo, and R. P. Mahler, "The Cauchy-Schwarz divergence for Poisson point processes," *IEEE Trans. Inf. Theory*, vol. 61, no. 8, pp. 4475–4485, Aug. 2015.
- [7] M. Beard, B.-T. Vo, B.-N. Vo, and S. Arulampalam, "Sensor control for multi-target tracking using Cauchy-Schwarz divergence," in *Proc. IEEE 18th Int. Conf. Inf. Fusion*, 2015, pp. 937–944.
- [8] I. Goodman, R. P. Mahler, and H. T. Nguyen, *Math. Data Fusion*. Norwell, MA, USA: Kluwer Academic Publishers, 1997.
- [9] 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.
- [10] B. Ristic, B.-N. Vo, and D. Clark, "A note on the reward function for PHD filters with sensor control," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 47, no. 2, pp. 1521–1529, Apr. 2011.
- [11] R. P. 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. Bellingham, WA, USA: SPIE, pp. 233–244.
- [12] J. Olofsson, G. Hendeby, T. R. Lauknes, and T. A. Johansen, "Multi-agent informed path-planning using the probability hypothesis density," *Auton. Robots*, vol. 44, pp. 913–925, 2020.
- [13] S. Gehly, B. A. Jones, and P. Axelrad, "Search-detect-track sensor allocation for geosynchronous space objects," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 54, no. 6, pp. 2788–2808, Dec. 2018.
- [14] H. V. Nguyen, H. Rezatofighi, B.-N. Vo, and D. C. Ranasinghe, "Multi-objective multi-agent planning for jointly discovering and tracking mobile objects," in *Proc. AAAI Conf. Artif. Intell.*, 2020, pp. 7227–7235.
- [15] P. Bostrom-Rost, D. Axehill, and G. Hendeby, "Sensor management for search and track using the Poisson multi-Bernoulli mixture filter," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 57, no. 5, pp. 2771–2783, Oct. 2021.
- [16] K. A. LeGrand, P. Zhu, and S. Ferrari, "A random finite set sensor control approach for vision-based multi-object search-while-tracking," in *Proc. IEEE 24th Int. Conf. Inf. Fusion*, 2021, pp. 1–8.
- [17] Y. Bar-Shalom, P. K. Willett, and X. Tian, *Tracking and Data Fusion*. England: YBS Publishing, 2011.
- [18] V. Krishnamurthy, *Partially Observed Markov Decision Processes: From Filtering to Controlled Sensing*. New York, NY, USA: Cambridge Univ. Press, 2016.
- [19] 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., Berlin, Germany: Springer, 2007.
- [20] R. P. Mahler, *Advances in Statistical Multisource-Multitarget Information Fusion*. Norwood, MA, USA: Artech House, 2014.
- [21] B.-N. Vo et al. *Multitarget Tracking*. Hoboken, NJ, USA: Wiley, 2015, pp. 1–15.
- [22] R. P. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Norwood, MA, USA: Artech House Boston, 2007.
- [23] R. P. Mahler, "Multitarget Bayes filtering via first-order multitarget moments," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 39, no. 4, pp. 1152–1178, Oct. 2003.
- [24] H. Sidenbladh, "Multi-target particle filtering for the probability hypothesis density," in *Proc. 6th Int. Conf. Inf. Fusion*, 2003, pp. 800–806.
- [25] B.-N. Vo and W.-K. Ma, "The Gaussian mixture probability hypothesis density filter," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4091–4104, Nov. 2006.
- [26] B.-N. Vo, B.-T. Vo, and D. Phung, "Labeled random finite sets and the Bayes multi-target tracking filter," *IEEE Trans. Signal Process.*, vol. 62, no. 24, pp. 6554–6567, Dec. 2014.
- [27] 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. Bellingham, WA, USA: SPIE, 2018, Art. no. 1064606.
- [28] D. Nuss et al., "A random finite set approach for dynamic occupancy grid maps with real-time application," *Int. J. Robot. Res.*, vol. 37, no. 8, pp. 841–866, Jul. 2018.
- [29] K. A. LeGrand and S. Ferrari, "The role of bounded fields-of-view and negative information in finite set statistics (FISST)," in *Proc. IEEE 23rd Int. Conf. Inf. Fusion*, 2020, pp. 1–9.

- [30] M. Yeddanapudi, Y. Bar-Shalom, and K. Pattipati, "IMM estimation for multitarget-multisensor air traffic surveillance," *Proc. IEEE*, vol. 85, no. 1, pp. 80–96, Jan. 1997.
- [31] Y. Bar-Shalom, X.-R. Li, and T. Kirubarajan, *Estimation With Applications To Tracking and Navigation*. Hoboken, NJ, USA: Wiley, 2001.
- [32] K. M. Simonson and T. J. Ma, "Robust real-time change detection in high jitter," Sandia National Laboratories, Tech. Rep. 8, 2009.
- [33] R. Badeau, B. David, and G. Richard, "Fast approximated power iteration subspace tracking," *IEEE Trans. Signal Process.*, vol. 53, no. 8, pp. 2931–2941, Aug. 2005.
- [34] A. S. Rahmathullah, A. F. García-Fernández, and L. Svensson, "Generalized optimal sub-pattern assignment metric," in *Proc. IEEE 20th Int. Conf. Inf. Fusion*, 2017, pp. 1–8.



Keith A. LeGrand (Member, IEEE) received the BS and MS degrees in aerospace engineering from the Missouri University of Science and Technology, and PhD degree in aerospace engineering from Cornell University, where he worked in the Laboratory for Intelligent Systems and Controls (LISC). He is currently an assistant professor of Aeronautics and Astronautics with Purdue University. Prior to that, he was a Senior Member of Technical Staff with Sandia National Laboratories in Albuquerque, New Mexico where he conducted research in inertial navigation, space systems, and multi-object tracking. He is the recipient of the U.S. DoD National Defense Science and Engineering Graduate (NDSEG) Fellowship (2021) and multiple best paper awards with the International FUSION Conference. His research interests include multi-sensor multi-object tracking; space domain awareness; and information theory.



Pingping Zhu (Member, IEEE) received the BS degree in electronics and information engineering and the MS degree from the Institute for Pattern Recognition and Artificial Intelligence, Huazhong University of Science and Technology, China, in 2006 and 2008, respectively, and PhD degrees in electrical and computer engineering from the University of Florida, USA, in 2013. He is currently an assistant professor with the Department of Computer Science and Electrical Engineering with Marshall University. Prior to that, he was a research associate with the Department of Mechanical and Aerospace Engineering, Cornell University. His research interests include approximate dynamic programming, reinforcement learning, signal processing, information theoretical learning, machine learning, artificial intelligence, and neural networks.



Silvia Ferrari (Senior Member, IEEE) received the BS degree from Embry–Riddle Aeronautical University and the MA and PhD degrees from Princeton University. She is John Brancaccio professor of Mechanical and Aerospace Engineering with Cornell University. Prior to that, she was professor of Engineering and Computer Science with Duke University, and Founder and Director of the NSF Integrative Graduate Education and Research Traineeship (IGERT) and Fellowship program on Wireless Intelligent Sensor Networks (WiSeNet). Currently, she is the Director of the Laboratory for Intelligent Systems and Controls (LISC) with Cornell University and the co-Director of the Cornell-UNIBO V'ho Institute on Vehicle Intelligence with Cornell Tech. Her principal research interests include active perception, robust adaptive control, learning, and approximate dynamic programming, and control of multiscale dynamical systems. She is the author of the book "Information-driven Path Planning and Control," MIT Press (2021), and of the TED talk "Do robots dream of electric sheep?." She is a Fellow of ASME and AIAA, and a Member of SPIE and SIAM. She is the recipient of the ONR young investigator Award (2004), the NSF CAREER Award (2005), and the Presidential Early Career Award for Scientists and Engineers (PECASE) Award (2006).

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.