

An Overview of Multi-Object Estimation via Labeled Random Finite Set

Ba-Ngu Vo, Ba-Tuong Vo, Tran Thien Dat Nguyen, and Changbeom Shim

Abstract—This article presents the Labeled Random Finite Set (LRFS) framework for multi-object systems—systems in which the number of objects and their states are unknown and vary randomly with time. In particular, we focus on state and trajectory estimation via a multi-object State Space Model (SSM) that admits principled tractable multi-object tracking filters/smoothers. Unlike the single-object counterpart, a time sequence of states does not necessarily represent the trajectory of a multi-object system. The LRFS formulation enables a time sequence of multi-object states to represent the multi-object trajectory that accommodates trajectory crossings and fragmentations. We present the basics of LRFS, covering a suite of commonly used models and mathematical apparatus (including the latest results not published elsewhere). Building on this, we outline the fundamentals of multi-object state space modeling and estimation using LRFS, which formally address object identities/trajectories, ancestries for spawning objects, and characterization of the uncertainty on the ensemble of objects (and their trajectories). Numerical solutions to multi-object SSM problems are inherently far more challenging than those in standard SSM. To bridge the gap between theory and practice, we discuss state-of-the-art implementations that address key computational bottlenecks in the number of objects, measurements, sensors, and scans.

Index Terms—State estimation, Filtering, Labeled random finite sets, Multi-Object tracking, Multi-Object system.

I. INTRODUCTION

A. State Space Model

State Space Model (SSM), also known as Hidden Markov Model (HMM), is a fundamental concept in dynamical systems theory. The state of the hidden object at sampling instant k is characterized by the state vector x_k , in some finite dimensional state space \mathbb{X} . This state generates an observation vector z_k in an (finite dimensional) observation space \mathbb{Z} , as depicted in Fig. 1. *State estimation*, *system identification*, and *control* are three interrelated fundamental problems in SSM.




Fig. 1: State Space Model of a single-object system.

Acknowledgment: This work was supported by the Australian Research Council under grants LP200301507 and FT210100506.

The authors are with the School of Electrical Engineering, Computing and Mathematical Sciences, Curtin University, Bentley, WA 6102, Australia (email: {ba-ngu.vo, ba-tuong.vo, t.nguyen1, changbeom.shim}@curtin.edu.au).

State estimation entails *estimating the state trajectory* $x_{0:k} \triangleq (x_0, \dots, x_k)$, and is the most basic problem upon which the other fundamental SSM problems are formulated. The trajectory can be estimated via *smoothing*, i.e., jointly estimating the states in batches, or via *filtering*, i.e., sequentially estimating the state at each time [1], [2], see Fig. 2. In practice, jointly estimating $x_{0:k}$ is intractable because the dimension of the variables (and computational load) per time step increases with k [2]. Hence, it is imperative to smooth over short windows to ensure the computational complexity per time step does not grow with time [2]. Filtering—the special case with a window length of one—is the most widely used [1], [2]. State estimation is an active research area popularized by the Kalman filter [1], [3] and particle filter [4], [5], which has far-reaching impact in many fields of study [1]–[3], [6].




Fig. 2: States and trajectory of a 1-D system. The state history $x_{0:k}$ defines the trajectory. Thus, a trajectory estimate can be obtained from a history of individual state estimates.

B. Multi-Object System

A *multi-object system* is a generalized dynamical system arising from a host of applications where, instead of a single object, the number of objects and their states are unknown and vary randomly with time, as depicted in Fig. 3. The *multi-object state* is the *set of states of individual objects*, while the set of their trajectories is the *multi-object trajectory*. The multi-object state $X_k \subset \mathbb{X}$ at time k generates an observation set $Z_k \subset \mathbb{Z}$. Each existing object may or may not generate an observation, while there could be false observations not generated by any object. This is compounded by *data association uncertainty*, i.e., it is not known which observations originated from which objects [7], [8].

In traditional SSM, the terminologies “state estimation” and “trajectory estimation” are used interchangeably because they both entail estimation of the trajectory (for the only




Fig. 3: Multi-object system. The number of states and observations vary with time. Existing objects may not be detected, false observations may occur, and it is not known which observations originated from which objects.

object). In line with this terminology, “*multi-object state estimation*” and “*multi-object trajectory estimation*” should both refer to estimation of the multi-object trajectory, and hence are abbreviated as “*multi-object estimation*” herein. To differentiate the task of estimating only the multi-object states (without trajectory information), we use the term “*multi-object localization*”. It is important to note from Fig. 4 that unlike single-object systems, the multi-object state history $X_{0:k}$ does not necessarily represent the multi-object trajectory. Consequently, *multi-object estimation* may not be possible via filtering nor smoothing over moving windows.

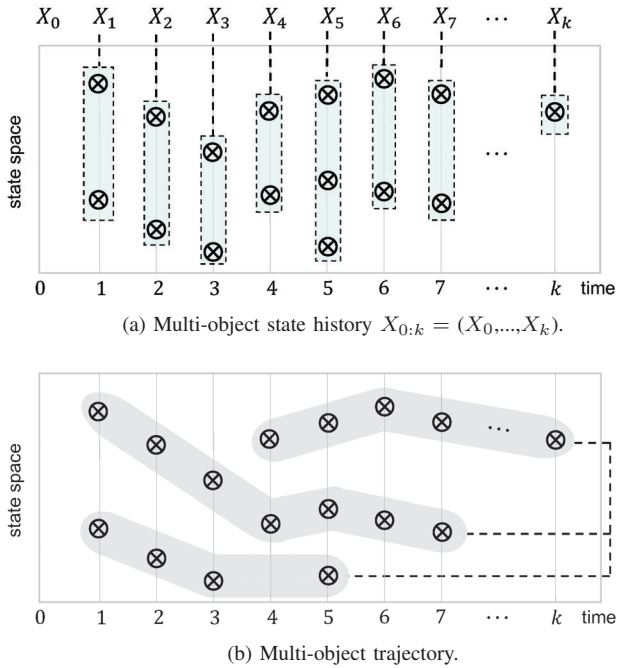


Fig. 4: States and trajectory of a 1-D multi-object system (note $X_0 = \{\}$ in this example). Objects may enter, exit, or reenter the state space. The multi-object state history $X_{0:k}$ does not necessarily represent (nor contains sufficient information to construct) the multi-object trajectory.

For a versatile multi-object state representation emulating single-object systems, the state history must be equivalent to the trajectory. Fundamentally, this is accomplished by augmenting distinct labels or provisional identities to each object state [9, pp. 135, 196-197], as illustrated in Fig. 5. This labeled multi-object representation enables multi-object




Fig. 5: 1-D labeled multi-object states and trajectory. The three objects are augmented with red, green, and white labels. Grouping the elements of the labeled multi-object state history $X_{0:k} = (X_0, \dots, X_k)$ according to labels gives the multi-object trajectory. A labeled multi-object trajectory estimate can be obtained from a history of labeled multi-object state estimates.

(trajectory) estimation to be done via filtering or smoothing over moving windows [10], [11]. Without the mechanism for linking trajectories from one window to another, even if segments of the trajectories can be estimated over short moving windows, multi-object (trajectory) estimation can only be performed over a window that grows with time, thus, computationally infeasible even for a single trajectory [2].

Historically driven by the aerospace industry’s interests in *multi-object tracking* (MOT), multi-object system problems are found in many diverse application areas, including surveillance, situational awareness, oceanography, autonomous vehicles/drones, field robotics, remote sensing, computer vision, and cell biology [7], [8], [12]–[19]. MOT is a well-established field, with three main approaches: *Multiple Hypothesis Tracking* (MHT), *Joint Probabilistic Data Association* (JPDA), and *Random Finite Set* (RFS); we refer the reader to the texts [7], [20], [17] for more details, or to [21] for a brief overview. Due to false alarms, misdetections, and data association uncertainty, multi-object system problems are far more challenging than their single-object counterparts [17].

C. Aims and Scope

It has been over two decades since the introduction of the RFS framework for multi-object filtering [9]. Prior to this, related approaches based on point process theory have been developed in [22]–[25] using random measure theoretic formulations. The RFS approach gained significant traction due to its intuitive appeal as a geometric formulation, and timely emergence in an era when computing technology as well as numerical filtering techniques were sufficiently advanced. However, the original RFS filtering formulation only considers multi-object localization, and additional heuristics are needed to construct trajectory estimates [16, pp. 505-508]. Indeed, early RFS-based multi-object localization filters were so popular that their inability to accommodate trajectories for MOT was mistakenly attributed to the framework itself. On the contrary, the idea of augmenting distinct labels to individual states to represent trajectories (illustrated in Fig. 5) has been discussed in [9, pp. 135, 196-197]. Due to its simplicity, Occam’s Razor suggests that this labeled multi-

object representation is the proper approach to multi-object systems, but it was deemed computationally impractical at the time and abandoned in favor of the unlabeled representation. Nevertheless, it inspired the development of Labeled RFS [10], [11], culminating in a suite of analytical and numerical tools for multi-object SSMs as well as the growth in applications and commercial interests.

This article provides an overview of the Labeled RFS (LRFS) approach to multi-object estimation, covering topics from the elements of RFS to the Generalized Labeled Multi-Bernoulli (GLMB) filter [26] capable of handling over a million trajectories [27]. The LRFS formulation not only ensures that the multi-object trajectory is indeed given by the history of the multi-object states, accommodating trajectory crossings, fragmentations, and ancestries (for spawning objects), but also enables characterization of uncertainty on the multi-object trajectory ensemble and alleviation of critical computational bottlenecks. Existing overviews such as the reviews [28], [29], surveys [30], [31], and texts [16], [17], [32], mainly focus on RFS multi-object localization (which does not address MOT), though some early developments in LRFS filters have been discussed in the text [17], survey [31], and particle filtering overview [33]. Keeping in mind the balance between scientific rigor and utility, we present an up-to-date coverage of LRFS analytical and computational tools, as well as the ensuing multi-object estimation solutions that address key computational bottlenecks in the number of objects, measurements, sensors, and scans. This coverage also includes results on closed-form information divergences for a versatile class of LRFSSs, and estimators based on the notion of joint existence probability unique to LRFS, not previously published.

The rest of the article is organized as follows. Section II presents some background on Bayesian state estimation and the labeled set representation of the multi-object state. To model set-valued random variables, Section III presents the fundamentals of RFS theory and some classical RFS models. Section IV introduces LRFS and the mathematical apparatus for multi-object state/trajectory modeling. Building on this, Section V extends Bayesian state estimation to multi-object SSM along with solutions and discussions of related SSM problems. Closing remarks are given in Section VI.

II. BACKGROUND

We begin with an outline of Bayesian state estimation for discrete-time SSMs in Subsection II-A. Subsection II-B formalizes the set representation for the multi-object state/trajectory that emulates the single-object state/trajectory, while Subsection II-C discusses the challenges of working with sets.

A. Bayesian State Estimation

In an SSM, the system state vector $x_k \in \mathbb{X}$ at time k evolves from its previous value x_{k-1} , and generates an observation $z_k \in \mathbb{Z}$ according to the *state and observation equations*

$$x_k = F_k(x_{k-1}, u_{k-1}, \nu_{k-1}), \quad (1)$$

$$z_k = G_k(x_k, u_k, \mu_k), \quad (2)$$

TABLE I: Common notations from Section II onwards.

Notation	Description
$x_{m:n}$	x_m, x_{m+1}, \dots, x_n
\mathbb{X}	Finite dimensional state space
x_k	State vector at time k
\mathbb{Z}	Finite dimensional observation space
z_k	Observation vector at time k
$f_k(\cdot x_{k-1})$	Markov transition density to time k given x_{k-1}
$g_k(z_k x_k)$	Likelihood of observing z_k given x_k
$p_{0:k}(x_{0:k})$	Posterior density at $x_{0:k}$
$p_k(x_k)$	Filtering density at x_k
$\mathcal{D}(f)$	Domain of function f
ℓ	Label of an object
\mathbb{L}	(Discrete) space of labels
\mathbf{x}, \mathbf{x}_k	Labeled state (x_k, ℓ) of an object at time k
\mathbf{X}, \mathbf{X}_k	Labeled multi-object state at time k

where F_k and G_k are non-linear mappings, u_k (and u_{k-1}) is the control or input signal, ν_{k-1} and μ_k are, respectively, process and measurement noise.

State estimation, i.e., estimation of the state trajectory, together with *system identification* and *control* are the three fundamental problems in dynamical system. In system identification, the goal is to estimate the system parameters from the system input and output, while in control, the goal is to use the input signals to drive the system state/trajectory to prescribed regions of the state space. In this article, we focus on state estimation, and hence the control signal is omitted.

The Bayesian estimation paradigm models the state and observation as random vectors. The state equation is characterized by the Markov *transition density* $f_k(x_k|x_{k-1})$, i.e., the probability density of a transition to state x_k at time k given the previous state x_{k-1} . The observation equation is characterized by the measurement *likelihood function* $g_k(z_k|x_k)$, i.e., the probability density of observing z_k given state x_k . Further, it is assumed that measurements are conditionally independent, i.e., the probability density of the observation history $z_{1:k}$ condition on $x_{1:k}$ is given by

$$p_{1:k}(z_{1:k}|x_{1:k}) = g_k(z_k|x_k)g_{k-1}(z_{k-1}|x_{k-1})\dots g_1(z_1|x_1).$$

All information about the state history to time k is encapsulated in the *posterior* density $p_{0:k}(x_{0:k}) \triangleq p_{0:k}(x_{0:k}|z_{1:k})$, which can be computed recursively for any $k \geq 1$, starting from an initial prior p_0 , via the Bayes *posterior recursion*:

$$p_{0:k}(x_{0:k}) = \frac{g_k(z_k|x_k)f_k(x_k|x_{k-1})p_{0:k-1}(x_{0:k-1})}{\int g_k(z_k|s_k)f_k(s_k|s_{k-1})p_{0:k-1}(s_{0:k-1})ds_{0:k}}. \quad (3)$$

Note that the dependence on $z_{1:k}$ is omitted for notational compactness. The above recursion, also known as smoothing-while-filtering [34] (or simply smoothing in this paper), is not suitable for real-time operations. Since the dimension of the trajectory probability density grows with time, computational/memory requirement increases at each time step and quickly becomes intractable. Real-time applications require algorithms with computational complexity per time step that does not grow with time [2, pp. 53-54].




Fig. 6: Propagation of the filtering density p_k .

The *filtering* density $p_k(x_k) \triangleq \int p_{0:k}(x_{0:k}) dx_{0:k-1}$, i.e., the marginal of the posterior density at the current time, encapsulates all information about the current state, and can be computed recursively using the Bayes *filtering recursion*:

$$p_k(x_k) = \frac{g_k(z_k|x_k) \int f_k(x_k|y) p_{k-1}(y) dy}{\int g_k(z_k|x) \int f_k(x|y) p_{k-1}(y) dy dx}. \quad (4)$$

The above recursion propagates a function on a fixed dimensional space, see Fig. 6, and hence has a fixed computational complexity per time step [2]. The *smoothing* density $p_{k-l|k}(\cdot|z_{1:k})$, i.e., the marginal of the posterior at time $k-l$, can also be computed recursively (e.g., forward-backward or two-filter smoothing), but is not widely used, and will not be covered here. Interested readers are referred to [2], [34].

The state can be estimated from appropriate probability densities via the *expected a posteriori* (conditional mean) or *maximum a posteriori* (conditional mode) estimators. These estimators are Bayes optimal with respect to certain Bayes risks (or penalties), and statistically *consistent*, i.e., converge (almost surely) to the true state as more data accumulate [6], [35]. The trajectory can be estimated from the posterior jointly as a sequence of states, or as a sequence of individual state estimates from filtering (or smoothing). Smoothing refines state estimates from data that arrives later, and is expected to yield better estimates than filtering [34], [36], [37].

Numerical methods for Bayesian state estimation is an active area of research [1]–[3], [38]. The posterior recursion admits the Kalman smoother as an analytic solution for linear Gaussian models [1], while for general non-linear models, Sequential Monte Carlo (SMC) approximations have been proposed [39]. To maintain a fixed computational complexity per time step, smoothing is performed over short moving windows. The most widely used is filtering, the special case with a window length of one. Under linear Gaussian models, the filtering recursion admits the Kalman filter as an analytic solution [1]. In general, analytic solutions are not possible. Popular approximate solutions include the extended Kalman filter [1], [3], unscented Kalman filter [40], Gaussian sum filter [41], cubature Kalman filter [42], [43], SMC/particle filter [4], [5], [37], [38], [44] and particle flow filter [45].

B. Multi-Object Representation

Regardless of the differences in MOT approaches, mathematically, a *trajectory* in a state space \mathbb{X} and discrete-time window \mathbb{T} is a mapping $\tau : \mathbb{T} \rightarrow \mathbb{X}$ [27]. The *domain*, $\mathcal{D}(\tau) \subseteq \mathbb{T}$, corresponds to the set of instants when the object exists. This definition covers the so-called *fragmented trajectories/tracks*, i.e., trajectories with non-contiguous domains, to accommodate objects physically exiting and reentering the

state space, as well as trajectory estimates produced by most MOT algorithms, as depicted in Fig. 4(b).

For versatility and applicability to a wide range of problems, we require a multi-object representation that allows:

- (R.1) the multi-object trajectory to be determined from the multi-object state history (analogous to single-object systems), ensuring multi-object estimation to be accomplished with a complexity per time step that does not grow with time, for computational tractability;
- (R.2) fragmented trajectories (see Fig. 4(b)) and multiple objects occupying the same attribute state, to model scenarios unique to challenging multi-object estimation problems such as reappearance/reidentification and merging/occlusion [46]–[48].

Neither of these requirements could be met with the unlabeled representation, as illustrated in Fig. 4, and the fact that a set cannot contain repeated elements. The labeled representation [9, pp. 135, 196–197], [10], [11] described in the following is most suitable, meeting both requirements.

The *labeled state* of an object at a particular instant (if it exists) is represented by $\mathbf{x} = (x, \ell)$, where $x \in \mathbb{X}$ is its *attribute* state (e.g., kinematics, visual features, etc.), and ℓ is its *label*, a provisional distinct identity in some discrete space \mathbb{L} called a *label space*. A common practice is to use $\ell = (s, \iota)$, where s is the time of birth and ι is an index to distinguish objects born at the same time [10], [11].

A *labeled multi-object state* \mathbf{X} is a finite subset of the product space $\mathbb{X} \times \mathbb{L}$ with *distinct labels*, i.e., no two elements of \mathbf{X} share the same label. Fundamentally, macroscopic objects (whose extents are significantly larger than the de Broglie wavelengths) can be labeled with distinct identifiers [49]. In some applications, distinct identifiers are even explicitly included in the intrinsic states of objects, e.g., in [50] the state of each object contains a unique RFID signal (from a finite set of RFID signals). In practice, track labeling is a required functionality of a MOT system, see e.g., [7], [49].

Labeling enables a sequence $\mathbf{X}_{j:k}$ of multi-object states to completely characterize the *multi-object trajectory*—the set of trajectories of the objects—on the interval $\{j:k\}$ analogous to single-object systems, via the grouping of states according to labels (illustrated in Fig. 5). Formally, a trajectory in $\mathbf{X}_{j:k}$ is a time-stamped sequence $[(x_s, \ell), \dots, (x_t, \ell)]$ of labeled states in $\mathbf{X}_{j:k}$ with a common label, which indeed defines a mapping $\tau_\ell : i \mapsto x_i, i \in \mathcal{D}(\tau_\ell)$. The labeled multi-object representation allows: multi-object estimation (filtering or smoothing on moving windows, see Figs. 5 and 7(a)) with computational complexity per time step that does not grow with time; multiple objects simultaneously occupying the same attribute state, e.g., $\{(x, \ell_1), (x, \ell_2)\}$, which inevitably occurs in a finite state space (as often is the case in the HMM literature), representing mergings/collisions, and occlusions in a wide range of applications, e.g., [47], [48], [51], [52]; and fragmented trajectories, arising from objects exiting and reentering the state space (especially when it is bounded), intrinsic to the reappearance/reidentification problem in computer vision [46], [53], [54].

Pragmatically, labeling is unavoidable in real multi-object systems. Trajectory estimation over a growing window is

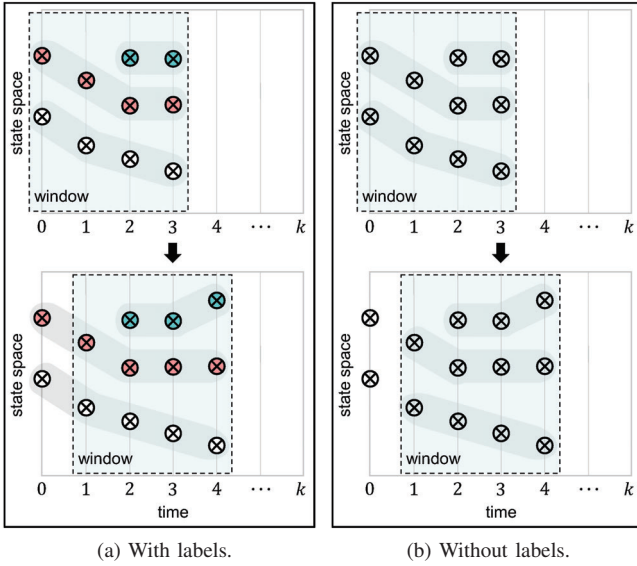


Fig. 7: Multi-object trajectories in moving windows. In (a) labels enable linking the trajectories between the windows $\{0:3\}$ and $\{1:4\}$ to represent trajectories on $\{0:4\}$. Without labels (b), there is no mechanism to link these trajectories.

numerically intractable as the computational complexity per time step grows with time [2]. Thus, it is imperative to use shorter moving windows, and link the trajectories between the windows together. This is straightforward with a labeled multi-object representation as illustrated in Fig. 7(a). In an unlabeled representation, there is no means to link the trajectories as illustrated in Fig. 7(b), making multi-object trajectory estimation infeasible without resorting to heuristic post-processing. Note that heuristically matching trajectories between windows via optimal assignment still requires labeling them.

State representation goes hand in hand with (mathematical) metrics for multi-object estimation error [55, Sec. 2.4]. A popular metric in the literature is the Optimal Sub-Pattern Assignment (OSPA) metric for multi-object states [56], which has been extended to multi-object trajectories in [27]. This extension, called the OSPA⁽²⁾ metric, penalizes errors in localization, number of trajectories, fragmentation and identity switching. The OSPA and OSPA⁽²⁾ metrics, respectively, gauge localization and tracking errors. The reader is referred to [57] for a comprehensive study of multi-object estimation performance evaluations, which includes such metrics.

C. Why Working with Sets is Tricky?

The Bayesian framework models the state (and observation) as a random variable, and consequently, in a multi-object SSM, a finite-set-valued random variable or a random finite set (RFS) is needed to model the multi-object state (and observation). Probability densities of RFSs are needed for the Markov transition density, likelihood function and posterior/filtering density characterizing the multi-object SSM.

The probability density of an RFS cannot be treated like that of a random vector as illustrated by the following example.




Fig. 8: Distribution of landing positions. Position x_1 is 3 times less likely than x_2 , and also 3 times less likely than x_3 .

Example 1. (Borrowed from [58]) Suppose that the number of apples falling in a day is uniformly distributed between 0 and 9, and that conditional on the number of fallen apples, the landing positions are independently and identically distributed (i.i.d.) according to the probability density p_f , shown in Fig. 8. Treating the set of landing positions as a random vector (see e.g., [59]), and noting that the probability of m apples falling ($m < 10$) is $\frac{1}{10}$, we have $p(\{x_{1:m}\}) = \frac{1}{10} \prod_{i=1}^m p_f(x_i)$. Consider the landing sets $\{x_1\}$ and $\{x_2, x_3\}$, where x_1, x_2 , and x_3 are shown in Fig. 8. Which of the sets is more likely? Since $p(\{x_1\}) = 2 \times 10^{-2}$ and $p(\{x_2, x_3\}) = 3.6 \times 10^{-2}$, it would appear that $\{x_2, x_3\}$ is more likely. However, had we measured distance in centimeters, then the probability density p_f is scaled by 10^{-2} (because it is non-zero on the interval $(-100, 100)$ and must integrate to 1), which results in $p(\{x_1\}) = 2 \times 10^{-4} > p(\{x_2, x_3\}) = 3.6 \times 10^{-6}$, contradicting the previous conclusion!

The notion of (probability) density goes hand in hand with integration. The density p_X of a random vector X , is defined such that $\Pr(X \in S) = \int_S p_X(x) dx$, for any (measurable) S . This means $p_X(x)$ is the instantaneous probability per unit hyper-volume (or formally, the Radon-Nikodým derivative of the probability distribution with respect to hyper-volume) at x , and is interpreted as the likelihood of instantiation x . Instantiations with higher density values are more likely (or probable) than those with lower ones, and their relative likelihoods are invariant to the unit of hyper-volume. The multi-object state space is the class of all finite subsets of \mathcal{X} , and does not inherit the usual Euclidean notion of density/integration. As noted in [28], the inconsistency in the above example arises because $p(\{x_1\})$ is measured in “m⁻¹” or “cm⁻¹” whereas $p(\{x_2, x_3\})$ is measured in “m⁻²” or “cm⁻²”, and hence cannot be meaningfully compared.

TABLE II: Common notations from Section III onwards.

Notation	Description
\mathcal{X}	Finite dimensional Euclidean space
$\mathcal{F}(\mathcal{X})$	Class of all finite subsets of \mathcal{X}
$\mathbf{1}_{\mathcal{T}}(\cdot)$	Indicator function for a set \mathcal{T}
π_{Σ}	Belief (or FISST) density of (the RFS) Σ
$ X $	Cardinality (number of elements) of X
h^X	Multi-object exponential $\prod_{x \in X} h(x)$, with $h^{\emptyset} = 1$
ρ_{Σ}	Cardinality distribution of Σ
v_{Σ}	Probability hypothesis density (PHD) of Σ
$\langle f, g \rangle$	Inner product $\int f(x)g(x)dx$ of functions f and g
$\delta_Y[X]$	Kronecker- δ , $\delta_Y[X] = 1$ if $X = Y$, and 0 otherwise

Another “tricky” issue is the *truncation of set-functions* (functions of sets), made up of intractably large sums in most multi-object system problems. Any function of a set $\{x_{1:n}\}$ must be symmetric in x_1, \dots, x_n , but truncation does not necessarily preserve symmetry, rendering the truncated sum an invalid set-function. Consider, e.g., the set-function $f(\{x, y\}) = x + y + xy$, which is clearly symmetric. However, truncating y gives the approximation $\hat{f}(x, y) = x + xy$, which is non-symmetric because $\hat{f}(x, y) \neq \hat{f}(y, x) = y + xy$, and thus not a valid function of the set $\{x, y\}$. In general, care must be taken to ensure validity of the results when truncating set-functions, see also Subsections IV-D, IV-F, and V-B.

III. RANDOM FINITE SETS FUNDAMENTALS

A *random finite set (RFS)* of \mathcal{X} is a *random variable taking values in the class of all finite subsets of \mathcal{X}* , hereon denoted as $\mathcal{F}(\mathcal{X})$. The cardinality and elements of an RFS are random variables, and the elements are unordered. A widely known example is the Poisson RFS, where the number of points is Poisson distributed and the points are i.i.d. according to some probability law on \mathcal{X} [60]. In this article, we restrict ourselves to a finite dimensional Euclidean space \mathcal{X} so that RFSs of \mathcal{X} have well-defined probability densities.

In Subsection III-A, we summarize some of the fundamental descriptors for RFSs. RFS statistics relevant to multi-object systems such as the cardinality and Probability Hypothesis Density (PHD) are presented in Subsection III-B, while multi-object estimators are discussed in Subsection III-C. Subsection III-D summarizes a number of popular RFS models.

A. Fundamental Descriptors

RFSs fall under the broader class of random (closed) sets in stochastic geometry [61], which also enables Bayesian inference with uncertainty models such as fuzzy set, Dempster-Shafer theory, and rule-based expert system [9]. Stochastic geometry has some overlaps with point process theory, and an RFS can be treated as a *simple finite point process*, i.e., a point process whose realizations are finite and have no repeated points [61], [62]. While the probability density of a point process may not exist, RFSs of \mathbb{R}^n do admit probability densities [62, Prop. 5.4.V, p. 138], [9, Prop. 19, p. 159], [60].

1) *Probability Density*: Since the notion of hyper-volume needed for density/integration on \mathcal{X} does not extend to $\mathcal{F}(\mathcal{X})$, we adopt an equivalent construct—the uniform (probability) distribution. In point process theory, the unit Poisson RFS exhibits complete spatial randomness analogous to the uniform distribution on \mathcal{X} (see e.g., [60]–[62]). Specifically, the number of points is Poisson distributed with unit rate, and the points are uniformly i.i.d. in \mathcal{X} . The *Poisson measure*—the unnormalized probability distribution of a unit Poisson RFS—is defined for each (measurable) $\mathcal{T} \subseteq \mathcal{F}(\mathcal{X})$ by

$$\mu(\mathcal{T}) = \sum_{i=0}^{\infty} \frac{1}{i!U^i} \int \mathbf{1}_{\mathcal{T}}(\{x_{1:i}\}) dx_{1:i},$$

where $\mathbf{1}_{\mathcal{T}}(\cdot)$ is the indicator function for \mathcal{T} , and U is the unit of hyper-volume in \mathcal{X} , with the convention that the integral for $i = 0$ is the integrand evaluated at \emptyset , i.e., $\mathbf{1}_{\mathcal{T}}(\emptyset)$. Note that $1/U^i$ cancels out the unit U^i of $dx_{1:i}$. Further, integration of a unitless (or dimensionless) function f on $\mathcal{F}(\mathcal{X})$ is realized via the Lebesgue integral with respect to (w.r.t.) the Poisson measure μ [60], [63]:

$$\int f(X)\mu(dX) = \sum_{i=0}^{\infty} \frac{1}{i!U^i} \int f(\{x_{1:i}\})dx_{1:i}.$$

Analogous to random vector, the *probability density* of an RFS Σ is defined as a non-negative function p_{Σ} on $\mathcal{F}(\mathcal{X})$ such that for any (measurable) $\mathcal{T} \subseteq \mathcal{F}(\mathcal{X})$,

$$\Pr(\Sigma \in \mathcal{T}) = \int_{\mathcal{T}} p_{\Sigma}(X)\mu(dX) = \int \mathbf{1}_{\mathcal{T}}(X)p_{\Sigma}(X)\mu(dX),$$

i.e., integrating the probability density yields the probability distribution. This means the instantaneous probability per unit Poisson measure at $X \in \mathcal{F}(\mathcal{X})$ is

$$p_{\Sigma}(X) = \frac{\Pr(\Sigma \in dX)}{\mu(dX)},$$

the (unit-less) Radon-Nikodým derivative of the probability distribution w.r.t. μ . If $p_{\Sigma}(X) > 0$ implies $p_{\Sigma}(Y) > 0$ for all $Y \subseteq X$, Σ is said to be *hereditary* [64]. Unlike a random vector, the probability density $p_{\Sigma}(X)$ does not have the usual likelihood interpretation, see [58] for further details.

2) *Belief Density: Finite Set Statistics (FISST)* offers an alternative notion of density/integration on $\mathcal{F}(\mathcal{X})$, which bypasses measure theoretic constructs [9], [16], [17], [65]. In FISST, the *set integral* of a function f on $\mathcal{F}(\mathcal{X})$ over a (compact) region $S \subseteq \mathcal{X}$ is defined as:

$$\int_S f(X)\delta X \triangleq \sum_{i=0}^{\infty} \frac{1}{i!} \int_{S^i} f(\{x_{1:i}\})dx_{1:i}, \quad (5)$$

where S^i denotes the i -fold Cartesian product of S , with $S^0 = \emptyset$, and the integral for $i = 0$ is $f(\emptyset)$ by convention. Note that $f(\{x_{1:i}\})$ has unit of U^{-i} , which cancels the unit U^i of $dx_{1:i}$, hence $\int f(\{x_{1:i}\})dx_{1:i}$ and the set integral are unitless.

The *belief (or FISST) density* of an RFS Σ is defined as a non-negative function π_{Σ} on $\mathcal{F}(\mathcal{X})$ whose set integral over any region $S \subseteq \mathcal{X}$ gives the *belief functional* at S , i.e.,

$$\Pr(\Sigma \subseteq S) = \int_S \pi_{\Sigma}(X)\delta X.$$

The belief functional $\Pr(\Sigma \subseteq S)$, as a function of S , is not a probability distribution, but nonetheless, is a fundamental descriptor of Σ [9], [16], [17], from which the belief density can be obtained by taking *set derivatives* (interested readers are referred to [16], [17] for more details). A related fundamental descriptor is the *void probability functional*, defined as $\Pr(\Sigma \cap S = \emptyset)$, the probability that S contains no points of Σ , i.e., Σ is contained in the complement of S . This is the belief functional at the complement of S [61], [62]. Similarly, the *capacity functional*, defined as the probability that S contains at least one point of Σ , is also a fundamental descriptor because $\Pr(\Sigma \cap S \neq \emptyset) = 1 - \Pr(\Sigma \cap S = \emptyset)$ [61].

It is important to note that the probability density w.r.t. the Poisson measure and the belief density are equivalent [65]:

$$p_{\Sigma}(X) = U^{|X|} \pi_{\Sigma}(X),$$

where $|X|$ denotes the cardinality (number of elements) of X . Hence, they are collectively referred to as *multi-object densities*. For the purpose of introducing RFS algorithms, it is more convenient to use the belief density.

3) *Probability Generating Functional (PGFL)*: Another RFS fundamental descriptor pertinent to multi-object system is the *PGFL*, defined for any unitless *test function* $h: \mathcal{X} \rightarrow [0, 1]$ as the expectation [16], [17], [61], [62]

$$G_{\Sigma}[h] \triangleq \mathbb{E}[h^{\Sigma}] = \int h^X \pi_{\Sigma}(X) \delta X,$$

where the *multi-object exponential* $h^X \triangleq \prod_{x \in X} h(x)$, with the convention $h^{\emptyset} = 1$. It is clear from the definition that $G_{\Sigma}[h] \in [0, 1]$, and $G_{\Sigma}[\mathbf{1}_S] = \Pr(\Sigma \subseteq S)$. PGFLs are analogous to probability generating functions.

The convolution of multi-object densities translates to the multiplication of PGFLs. Suppose that Σ is the union of disjoint and statistically independent RFSs $\Sigma_1, \dots, \Sigma_n$. Then,

$$G_{\Sigma}[h] = G_{\Sigma_1}[h] \dots G_{\Sigma_n}[h], \quad (6)$$

$$\pi_{\Sigma}(X) = \sum_{W_1 \uplus \dots \uplus W_n = X} \pi_{\Sigma_1}(W_1) \dots \pi_{\Sigma_n}(W_n), \quad (7)$$

where the sum is taken over all mutually disjoint $W_1, \dots, W_n \subseteq X$ (including empty sets) that cover X [16], [17, pp. 85-86].

The multi-object density (and other statistics) can be obtained by differentiating the PGFL [16], [17], [61], [62]. For a functional G on the space of test functions, let

$$G^{(1)}[h; \zeta_1] \triangleq (dG)_h[\zeta] = \lim_{\epsilon \rightarrow 0} \frac{G[h + \epsilon \zeta] - G[h]}{\epsilon},$$

denote its *Gâteaux differential* at h in the direction ζ (if it exists). The n -th *Gâteaux differential* at h is a multi-linear form in the directions ζ_1, \dots, ζ_n , given recursively by

$$G^{(n)}[h; \zeta_{1:n}] = (dG^{(n-1)}[; \zeta_{1:n-1}])_h[\zeta_n].$$

Further, note that a multi-linear form can be expressed as $F[\zeta_{1:n}] = \int \zeta_1(y_1) \dots \zeta_n(y_n) f(y_{1:n}) dy_{1:n}$, and is completely characterized by the function f , which can be rewritten in the Dirac delta notation $F[\delta_{x_1}, \dots, \delta_{x_n}] \triangleq f(x_{1:n})$. This is suggestive of evaluating f at x_1, \dots, x_n via substituting the Dirac deltas $\delta_{x_1}, \dots, \delta_{x_n}$ into the integral that defines F , i.e., $f(x_{1:n})$ can be treated as the value of multi-linear form F at




Fig. 9: PHD or intensity function on a 1-D state space.

$\delta_{x_1}, \dots, \delta_{x_n}$. Using the *Volterra functional derivative* w.r.t. a finite set [17, pp. 66], defined by

$$\frac{\delta G}{\delta \{x_{1:n}\}}[h] \triangleq G^{(n)}[h; \delta_{x_1}, \dots, \delta_{x_n}],$$

with $\frac{\delta G}{\delta \emptyset}[h] \triangleq G[h]$, we have [17, pp. 95]

$$\frac{\delta G}{\delta X}[0] = \pi_{\Sigma}(X).$$

Volterra functional set derivatives can be calculated using a suite of differentiation rules [16, pp. 383–395], including a powerful generalized Faà di Bruno's chain rule [66]. Further properties and applications of PGFLs in multi-object systems can be found in the texts [16], [17].

B. Cardinality and Probability Hypothesis Density

In multi-object systems, relevant statistics of an RFS Σ often involve its *cardinality* (number of elements). The probability generating function (PGF) of the cardinality $|\Sigma|$, evaluated at $z \in [0, 1]$, is the PGFL evaluated at the test function $h(x) = z$, i.e., $G_{|\Sigma|}(z) = G_{\Sigma}[h = z]$. The cardinality distribution $\rho_{\Sigma}(n) \triangleq \Pr(|\Sigma| = n)$ can be computed via [16], [61], [67]

$$\rho_{\Sigma}(n) = \frac{1}{n!} \int \pi_{\Sigma}(\{x_{1:n}\}) dx_{1:n} = \frac{1}{n!} G_{|\Sigma|}^{(n)}(0),$$

where $G_{|\Sigma|}^{(n)}$ is the n -th derivative of the PGF. Statistics of the cardinality can be computed from ρ_{Σ} or $G_{|\Sigma|}$.

Another well-known cardinality-based statistic of an RFS is the *intensity function* [61], [62], also known as the *Probability Hypothesis Density (PHD)* in the MOT literature [16], [68]. As illustrated in Fig. 9, it is defined as a non-negative function v_{Σ} (on \mathcal{X}) whose integral over any region $S \subseteq \mathcal{X}$ gives the expected cardinality in S , i.e.,

$$\mathbb{E}[|\Sigma \cap S|] = \int_S v_{\Sigma}(x) dx. \quad (8)$$

The PHD can be computed from the multi-object density or PGFL by [17, pp. 93]

$$v_{\Sigma}(x) = \int \pi_{\Sigma}(\{x\} \cup W) \delta W = \frac{\delta G}{\delta \{x\}}[1],$$

and is the 1st of the *factorial moment densities* [61], [62]:

$$v_{\Sigma}(X) \triangleq \int \pi_{\Sigma}(X \cup W) \delta W = \frac{\delta G}{\delta X}[1].$$

Interestingly, knowledge of the PHD v_{Σ} is sufficient to calculate the expectation of random sums of a (measurable) real function f over Σ , via Campbell's Theorem [61], [64]

$$\mathbb{E}[\sum_{x \in \Sigma} f(x)] = \int f(x) v_{\Sigma}(x) dx.$$

Definition (8) means that the PHD is the density of the expected cardinality w.r.t. hyper-volume. This physically intuitive interpretation is one of the factors behind the appeal of the PHD filter. The local maxima of the PHD are points in \mathcal{X} with the highest local concentration of expected number of objects (per unit hyper-volume), suggesting that they are the most likely states for the underlying objects, see Fig. 9. Based on such interpretation, the *PHD estimate* and *cardinalized PHD (CPHD) estimate*, respectively, use $\hat{n} = \text{round}(\mathbb{E}[|\Sigma|])$ and $\hat{n} = \arg \max_n \rho_\Sigma(n)$ as the estimated number of objects, and the \hat{n} highest local maxima of the PHD as the estimated states, see also [16, pp. 595] for more details.

For a hereditary RFS, the intensity function can be extended to the *conditional intensity* at a point x , defined as [64]

$$v_\Sigma(x|Y) = \frac{p_\Sigma(Y \cup \{x\})}{p_\Sigma(Y)}.$$

Note that the intensity function is given by the expectation $\int v_\Sigma(x|Y) \pi_\Sigma(Y) \delta Y$. Moreover, the probability density p_Σ (and belief density π_Σ) is completely determined by its conditional intensity [61], [64]. Working with the conditional intensity eliminates the computation of the normalizing constant, but requires certain consistency conditions.

C. Multi-Object Estimators

Given a probability density, estimators are needed to determine estimates of the underlying random variable. Popular estimates such as the mean and mode are not directly applicable to RFSs, because there is no average for sets nor a likelihood interpretation for multi-object density [58]. Nonetheless, as the 1st moment, the PHD can be regarded as the expectation of an RFS, and hence the PHD estimate (Subsection III-B) can be treated as the mean estimate. The cardinality of the PHD estimate has a high variance for a large number of objects since the variance of a Poisson is equal to the mean. The CPHD estimate is an extension to improve the cardinality estimate.

While the multi-object density does not have a likelihood interpretation over $\mathcal{F}(\mathcal{X})$, when restricted to a cardinality it can be interpreted as a likelihood. This observation leads to an extension of the mode estimate, called the *Marginal Multi-object (MaM) estimate*, defined as *the most probable multi-object state given the most probable cardinality* [16, pp. 497-498] (which requires computing: the most probable cardinality from the cardinality distribution; and the supremum of the multi-object density conditioned on this cardinality). Another estimate that emulates the mode as an optimal Bayes estimator is the *Joint Multi-object (JoM) estimate* [16, pp. 498-500]

$$\hat{X} = \arg \sup_{X \in \mathcal{F}(\mathcal{X})} \frac{c^{|X|}}{|X|!} \pi_\Sigma(X),$$

where $c > 0$ is a constant with magnitude in the order of the desired accuracy, measured in units of hyper-volume in \mathcal{X} .

D. Classical RFS Models

The following RFSs are popular models in point process and multi-object SSM, especially for multi-object observations and multi-object localization. Hereon, we denote the *inner product*

$\int f(\zeta)g(\zeta)d\zeta$ of two functions f, g (or $\sum_{\ell=0}^{\infty} f(\ell)g(\ell)$ when they are sequences) as $\langle f, g \rangle$, and the generalized Kronecker delta that takes arbitrary arguments as

$$\delta_Y[X] \triangleq \begin{cases} 1, & \text{if } X = Y \\ 0, & \text{otherwise} \end{cases}.$$

1) *Poisson*: A *Poisson RFS* (or Poisson point process) of \mathcal{X} is completely characterized by its PHD (intensity function) v_Σ [17], [61], [62, pp. 98], with multi-object density and PGFL:

$$\pi_\Sigma(X) = e^{-\langle v_\Sigma, 1 \rangle} v_\Sigma^X, \quad G_\Sigma[h] = e^{\langle v_\Sigma, h-1 \rangle}.$$

Its cardinality is Poisson distributed with mean $\langle v_\Sigma, 1 \rangle$, i.e., $\rho_\Sigma(n) = e^{-\langle v_\Sigma, 1 \rangle} \langle v_\Sigma, 1 \rangle^n / n!$, and conditioned on the cardinality (number of distinct elements), each (distinct) element is i.i.d. according to $v_\Sigma / \langle v_\Sigma, 1 \rangle$.

The Poisson RFS models “no interaction” or “complete spatial randomness”, and is one of the best-known and most tractable of point processes [60]–[62], [69]. In multi-object systems, it is a popular model for clutter or false alarms.

2) *I.I.D. Cluster*: An *i.i.d. cluster RFS* is a generalization of the Poisson RFS to accommodate cardinality distributions other than Poisson, and is completely characterized by its cardinality distribution ρ_Σ and PHD v_Σ . Specifically, the multi-object density and PGFL are [17], [61], [62, pp. 99]:

$$\pi_\Sigma(X) = \frac{\rho_\Sigma(|X|) |X|!}{\langle v_\Sigma, 1 \rangle^{|X|}} v_\Sigma^X, \quad G_\Sigma[h] = G_{|\Sigma|} \left(\frac{\langle v_\Sigma, h \rangle}{\langle v_\Sigma, 1 \rangle} \right).$$

The cardinality of an i.i.d. cluster RFS is distributed according to the prescribed ρ_Σ , and conditioned on the cardinality, each (distinct) element is i.i.d. according to $v_\Sigma / \langle v_\Sigma, 1 \rangle$. When ρ_Σ is Poisson, this reduces to the Poisson RFS [61], [62].

Remark 1. Multi-object densities for i.i.d. cluster (and Poisson) RFSs are defined for a fixed dimensional \mathcal{X} . They may not exist otherwise [62, Prop. 5.4.V, p. 138], [60].

3) *Multi-Bernoulli*: A *Bernoulli RFS*, parameterized by the pair (r, p) , has probability $1-r$ of being empty, and probability r of being a singleton, conditioned on which the element is distributed according to the probability density p (on \mathcal{X}). The multi-object density, PGFL, cardinality distribution, and PHD are given by [17, pp. 100]

$$\begin{aligned} \pi_\Sigma(X) &= (1-r)\delta_\emptyset[X] + rp(x)\delta_{\{x\}}[X], \\ G_\Sigma[h] &= 1-r + r\langle p, h \rangle, \\ \rho_\Sigma(n) &= (1-r)\delta_0[n] + r\delta_1[n], \\ v_\Sigma(x) &= rp(x). \end{aligned}$$

A *multi-Bernoulli RFS* (for compactness we drop the ‘RFS’) is a union of disjoint and independent Bernoulli RFSs with parameters $\{(r^{(\zeta)}, p^{(\zeta)})\}_{\zeta \in \Psi}$ [17, pp. 101]. Its PGFL is the product of the constituent Bernoulli RFSs’ PGFLs,

$$G_\Sigma[h] = \prod_{\zeta \in \Psi} (1-r^{(\zeta)} + r^{(\zeta)} \langle p^{(\zeta)}, h \rangle),$$

and its multi-object density is the convolution of the constituent Bernoulli RFSs’ densities [17, pp. 102]

$$\pi_\Sigma(\{x_{1:n}\}) = \pi_\Sigma(\emptyset) \sum_{1 \leq i_1 \neq \dots \neq i_n \leq |\Psi|} \prod_{j=1}^n \frac{r^{(\zeta_{i_j})} p^{(\zeta_{i_j})}(x_j)}{1-r^{(\zeta_{i_j})}}, \quad (9)$$

TABLE III: Common notations from Section IV onwards.

Notation	Description
\mathcal{A}	Attribute projection $(x, \ell) \mapsto x$
\mathcal{L}	Label projection $(x, \ell) \mapsto \ell$
$\Delta(\mathbf{X})$	Distinct label indicator $\delta_{ \mathbf{X} }[\ \mathcal{L}(\mathbf{X})\]$
$\langle f \rangle(\{\ell_{1:n}\})$	Label-marginal $\int f(\{(x_1, \ell_1), \dots, (x_n, \ell_n)\}) dx_{1:n}$
π_{Σ}	Multi-object density of (the LRFS) Σ
$w_{\Sigma}(L)$	Probability that Σ has label set L
ρ_{Σ}	Cardinality distribution of Σ
\mathbf{v}_{Σ}	PHD or intensity of Σ
$\mathbf{X}_{j:k}$	Multi-object sequence/trajectory on $\{j:k\}$
$\mathcal{L}(\mathbf{X}_{j:k})$	$(\mathcal{L}(\mathbf{X}_j), \dots, \mathcal{L}(\mathbf{X}_k))$
$T(\ell)$	Set of instants that $\mathbf{X}_{j:k}$ contains label ℓ
$\mathbf{x}_{T(\ell)}^{(\ell)}$	$[(x_i, \ell) \in \mathbf{X}_i]_{i \in T(\ell)}$, trajectory of ℓ in $\mathbf{X}_{j:k}$
$h^{\mathbf{X}_{j:k}}$	$\prod_{\ell \in \mathcal{L}(\mathbf{X}_{j:k})} h(\mathbf{x}_{T(\ell)}^{(\ell)})$,
$\prod_{i=j}^k S_i$	$S_j \times \dots \times S_k$
$\Delta(\mathbf{X}_{j:k})$	Multi-scan distinct label indicator $\prod_{i=j}^k \Delta(\mathbf{X}_i)$

where $\pi_{\Sigma}(\emptyset) = \prod_{\zeta \in \Psi} (1 - r^{(\zeta)})$, and $\zeta_1, \dots, \zeta_{|\Psi|}$ enumerate all the elements of Ψ . By convention, the sum reduces to 1 when $n = 0$, and zero when $n > |\Psi|$. It is implicitly assumed that $r^{(\zeta)} \in [0, 1)$. Statistics such as the cardinality distribution and PHD are given by [17, pp. 102],

$$\rho_{\Sigma}(n) = \pi_{\Sigma}(\emptyset) \sum_{1 \leq i_1 < \dots < i_n \leq |\Psi|} \prod_{j=1}^n \frac{r^{(\zeta_{i_j})}}{1 - r^{(\zeta_{i_j})}}, \quad (10)$$

$$\mathbf{v}_{\Sigma}(x) = \sum_{\zeta \in \Psi} r^{(\zeta)} p^{(\zeta)}(x). \quad (11)$$

Multi-Bernoullis are often used for modeling object survival/death, and detection uncertainty in observations.

Remark 2. For an i.i.d. cluster (and Poisson) RFS, the joint distribution of the elements is conditioned on the cardinality (i.e., the number of distinct elements), which ensures distinctness of the elements. However, for a multi-Bernoulli, the distribution of each element is conditional on its existence, independent from others, and there is no mechanism to prevent two constituent (independent) Bernoulli RFS's from sharing the same point. For example, consider the Bernoulli RFSs $B^{(i)}$ with $r^{(i)} = 0.5$, $p^{(i)}(\cdot) = \mathcal{N}(\cdot; 0, 1)$, $i = 1, 2$. Noting that the likelihood of a realization x_1 from $B^{(1)}$ and x_2 from $B^{(2)}$ is $0.25\mathcal{N}(x_1; 0, 1)\mathcal{N}(x_2; 0, 1)$, the likelihood of $x_1 = 0$ and $x_2 = 0$ is not only positive, but the highest possible among all values of x_1, x_2 , and hence, not negligible. The disjoint condition between the constituent Bernoulli RFSs implies dependence. The expressions for the multi-Bernoulli's PGFL and multi-object density implicitly assume negligible dependence between constituent components. This assumption is reasonable for traditional detection processes based on thresholding, which return distinct observations.

IV. LABELED RANDOM FINITE SET

A labeled RFS is a special class of RFSs introduced in [10], [11] for modeling multi-object states/trajectories. It provides a versatile multi-object estimation framework that admits

characterization of uncertainty for the multi-object trajectory ensemble, and meaningful approximations with characterizable errors that are requisite for principled solutions.

Formally, a *labeled RFS* (LRFS) with *attribute space* \mathbb{X} and (discrete) *label space* $\mathbb{L} = \{\alpha_i : i \in \mathbb{N}\}$, is an RFS Σ of the product space $\mathbb{X} \times \mathbb{L}$ such that each realization has distinct labels. Defining the *attribute projection* $\mathcal{A} : (x, \ell) \mapsto x$, and *label projection* $\mathcal{L} : (x, \ell) \mapsto \ell$, so that $\mathcal{A}(\mathbf{X})$ and $\mathcal{L}(\mathbf{X})$ are, respectively, the sets of *attributes* and *labels* of \mathbf{X} , then \mathbf{X} has *distinct labels* if the *distinct label indicator*

$$\Delta(\mathbf{X}) \triangleq \delta_{|\mathbf{X}|}[\|\mathcal{L}(\mathbf{X})\|], \quad (12)$$

equals 1, i.e., $\mathcal{L}(\mathbf{X})$ and \mathbf{X} have the same cardinality.

An LRFS can be thought of as a simple finite *marked point process*¹ of $\mathbb{X} \times \mathbb{L}$, with distinct marks from the discrete mark space. Note the distinction between “simple finite marked” and “marked simple finite” point processes. The former is simple, but the point process formed by *unmarking* (discarding the marks) is not necessarily simple. The latter is the special case with a simple unmarked version, because it is constructed by marking an RFS, and hence, has the same cardinality as its unmarked version [11]. In general, an LRFS does not necessarily have the same cardinality as its unmarked version.

Subsection IV-A introduces the concept of joint existence probability and the ensuing multi-object estimators unique to LRFSs. Subsections IV-B to IV-D summarize popular LRFS models. Information divergences for LRFS are presented in Subsection IV-E, including closed-form expressions not previously published. LRFS approximations and spatio-temporal modeling are discussed in Subsections IV-F and IV-G.

Following [11], vectors are represented by lower case letters (e.g., x and \mathbf{x}), and finite sets are represented by upper case letters (e.g., X and \mathbf{X}), where the symbols for labeled entities and their distributions are bolded (e.g., \mathbf{x} , \mathbf{X} , $\boldsymbol{\pi}$, etc.) to distinguish them from unlabeled ones.

A. Joint Existence Probability

Pertinent to multi-object estimation and unique to LRFSs is the notion of joint existence probability. Let us denote the *label-marginal* of a function $f : \mathcal{F}(\mathbb{X} \times \mathbb{L}) \rightarrow \mathbb{R}$, by

$$\langle f \rangle(\{\ell_{1:n}\}) \triangleq \int f(\{(x_1, \ell_1), \dots, (x_n, \ell_n)\}) dx_{1:n}, \quad (13)$$

with $\langle f \rangle(\emptyset) \triangleq f(\emptyset)$. For an LRFS Σ with multi-object density $\boldsymbol{\pi}_{\Sigma}$, we define the *joint existence probability* of $L \subseteq \mathbb{L}$, and the *label-conditioned joint attribute (probability) density* for distinct labels $\ell_{1:n}$, respectively, as [70]

$$w_{\Sigma}(L) \triangleq \langle \boldsymbol{\pi}_{\Sigma} \rangle(L), \quad (14)$$

$$\pi_{\Sigma|\ell_{1:n}}(x_{1:n}) \triangleq \frac{\boldsymbol{\pi}_{\Sigma}(\{(x_1, \ell_1), \dots, (x_n, \ell_n)\})}{w_{\Sigma}(\{\ell_{1:n}\})}, \quad (15)$$

with the convention that $\pi_{\Sigma|\ell_{1:n}}(x_{1:n}) = 1$ whenever $w_{\Sigma}(\{\ell_{1:n}\}) = 0$. Then $w_{\Sigma}(\cdot)$ is a probability distribution on $\mathcal{F}(\mathbb{L})$, and $\pi_{\Sigma|\ell_{1:n}}(\cdot)$ is a probability density on \mathbb{X}^n [70]. The joint existence probability $w_{\Sigma}(L)$ is the probability that the LRFS Σ has label set L .

¹A marked point process Σ of $\mathbb{X} \times \mathbb{M}$, with mark space \mathbb{M} , satisfies $|\Sigma \cap (S \times \mathbb{M})| < \infty$ for any bounded $S \subseteq \mathbb{X}$ [60]–[62], [64].

In the context of estimators, the joint existence probability is more informative than the cardinality distribution [70]

$$\rho_{\Sigma}(n) = \sum_{L \subseteq \mathbb{L}} \delta_n[|L|] w_{\Sigma}(L), \quad (16)$$

in the following sense. The label sets of cardinality n , each with small $w_{\Sigma}(L)$, could accumulate up to a large $\rho_{\Sigma}(n)$. Thus, sets with the most probable cardinality, e.g., the MaM estimate in Subsection III-C, could have negligible joint existence probability compared to the highest joint existence probability achievable (by sets with other cardinalities).

A judicious alternative to the MaM estimate is the *label-MaM* estimate, defined as *the most probable multi-object state given the most probable label set* (which is meaningful because the label set is a discrete variable). This estimator seeks the label set L^* with highest joint existence probability, whose attributes are given by the mode of the corresponding label-conditioned joint attribute density. If the most probable label set is not unique, we select one with the most probable cardinality, and if this is still not unique, we select the one that yields the highest label-conditioned joint attribute density.

LRFS furnishes the PHD or intensity value $\mathbf{v}_{\Sigma}(x, \ell)$ with an additional interpretation as the attribute PHD at x , of (the object with) label ℓ . Recall the definition of the PHD from (8) that for any $S \subseteq \mathbb{X}$, the expected number of objects with attributes in S and label ℓ is given by

$$\mathbb{E}[|\Sigma \cap (S \times \{\ell\})|] = \int_S \mathbf{v}_{\Sigma}(x, \ell) dx.$$

Thus, for a given label ℓ , the function $\mathbf{v}_{\Sigma}(\cdot, \ell)$ is its attribute PHD. Further, the distinct label property means that the cardinality $|\Sigma \cap (\mathbb{X} \times \{\ell\})|$ is either 1, if Σ has a member with label ℓ , or 0 otherwise. Consequently, the expectation $\mathbb{E}[|\Sigma \cap (\mathbb{X} \times \{\ell\})|]$ (i.e., the total attribute PHD mass of ℓ) cannot exceed 1. Hence, the *existence probability* of ℓ and its *attribute (probability) density*, in an LRFS Σ , are given by

$$r_{\Sigma}(\ell) = \int \mathbf{v}_{\Sigma}(x, \ell) dx, \quad p_{\Sigma}(\cdot, \ell) = \mathbf{v}_{\Sigma}(\cdot, \ell) / r_{\Sigma}(\ell). \quad (17)$$

Remark 3. It is imperative to note that $r_{\Sigma}(\ell) \neq w_{\Sigma}(\{\ell\})$ and $p_{\Sigma}(x, \ell) \neq \pi_{\Sigma|\ell}(x)$. Whereas $w_{\Sigma}(\{\ell\})$ is the joint probability that (only) ℓ exists and the other labels do not, $r_{\Sigma}(\ell)$ is the marginal probability that ℓ exists regardless of other labels. Similarly, $\pi_{\Sigma|\ell}(\cdot)$ is the attribute density given that only ℓ exists, whereas $p_{\Sigma}(\cdot, \ell)$ is the attribute density given that ℓ exists regardless of others. The marginal nature of $r_{\Sigma}(\ell)$ and $p_{\Sigma}(x, \ell)$ can be seen from the definition of the PHD $\mathbf{v}_{\Sigma}(x, \ell)$.

The PHD (of LRFS) offers inexpensive sub-optimal JoM and MaM estimates. The sub-optimal JoM estimator seeks the labels with existence probabilities above a prescribed threshold, and estimates their attributes from the corresponding attribute densities (via the modes/means). This estimate depends on the existence threshold similar to the constant c in the JoM estimate. The sub-optimal MaM estimator seeks the most probable cardinality n^* (either from ρ_{Σ} , if available, or from the multi-Bernoulli cardinality approximation using the existence probabilities), and the n^* labels with highest existence probabilities, whose attributes are then estimated

from their attribute densities. The PHD also offers a tractable sub-optimal label-MaM estimate, see Subsection IV-C.

B. Labeled I.I.D. Cluster

Analogous to its unlabeled counterpart, a *labeled i.i.d. cluster* Σ is an LRFS characterized by a cardinality distribution ρ_{Σ} and an attribute probability density f_{Σ} (on \mathbb{X}) [11]. Conditioned on cardinality n , the n (not necessarily distinct) attributes i.i.d. according to f_{Σ} , are marked with distinct labels from $\mathbb{L}(n) \triangleq \{\alpha_i : i = 1:n\}$. A labeled i.i.d. cluster has multi-object density and PGFl [11], [17, pp. 450]

$$\begin{aligned} \pi_{\Sigma}(\mathbf{X}) &= \delta_{\mathbb{L}(|\mathbf{X}|)}[\mathcal{L}(\mathbf{X})] \rho_{\Sigma}(|\mathcal{L}(\mathbf{X})|) (f_{\Sigma} \circ \mathcal{A})^{\mathbf{X}}, \\ \mathbf{G}_{\Sigma}[\mathbf{h}] &= \sum_{n=0}^{\infty} \rho_{\Sigma}(n) \prod_{\ell \in \mathbb{L}(n)} \langle f_{\Sigma}, \mathbf{h}(\cdot, \ell) \rangle, \end{aligned}$$

respectively, for any unitless test function \mathbf{h} on $\mathbb{X} \times \mathbb{L}$, where \circ denotes composition. It is clear from the above descriptors that a labeled i.i.d. cluster is not an i.i.d. cluster. The PHD of a labeled i.i.d. cluster is given by [17, pp. 451]

$$\mathbf{v}_{\Sigma}(x, \ell) = f_{\Sigma}(x) \sum_{n=0}^{\infty} \mathbf{1}_{\mathbb{L}(n)}(\ell) \rho_{\Sigma}(n).$$

The sum over n is the existence probability of ℓ , and f_{Σ} is its attribute density (independent of ℓ). The *labeled Poisson* is the special case with Poisson cardinality distribution. Note that a labeled Poisson RFS is not a Poisson RFS of $\mathbb{X} \times \mathbb{L}$.

C. Labeled Multi-Bernoulli

Similar to a multi-Bernoulli, a *labeled multi-Bernoulli* (LMB) is an LRFS characterized by a collection of independent Bernoulli RFSs (of the attribute space \mathbb{X}) with parameters $\{(r^{(\zeta)}, p^{(\zeta)}) : \zeta \in \Psi\}$, and additionally, a 1-1 (injective) mapping $\sigma : \mathbb{L} \rightarrow \Psi$ that pairs each $\zeta \in \Psi$ with a distinct label $\ell \in \mathbb{L}$ [11]. For each $\zeta \in \Psi$, a *labeled Bernoulli* RFS is constructed by marking the Bernoulli RFS parameterized by $(r^{(\zeta)}, p^{(\zeta)})$ with the associated label $\ell = \sigma^{-1}(\zeta)$. The resulting labeled Bernoulli RFSs are disjoint (due to their distinct labels), and their union is an LMB Σ , whose multi-object density is given by [11]

$$\pi_{\Sigma}(\mathbf{X}) = \Delta(\mathbf{X}) [\mathbf{1}_{\mathcal{D}(\sigma)}]^{\mathcal{L}(\mathbf{X})} [\pi_{\Sigma}(\mathbf{X}; \cdot)]^{\Psi}, \quad (18)$$

where $\mathcal{D}(\sigma)$ is the domain of σ , and

$$\pi_{\Sigma}(\mathbf{X}; \zeta) = \sum_{(x, \ell) \in \mathbf{X}} \delta_{\sigma(\ell)}[\zeta] r^{(\zeta)} p^{(\zeta)}(x) + (1 - \delta_{\sigma(\ell)}[\zeta]) (1 - r^{(\zeta)}).$$

The above sum either takes on: $r^{(\zeta)} p^{(\zeta)}(x)$ if ζ matches the label of $(x, \ell) \in \mathbf{X}$, i.e., $(x, \sigma^{-1}(\zeta)) \in \mathbf{X}$; or $1 - r^{(\zeta)}$ if ζ is not matched with any labels in $\mathcal{L}(\mathbf{X})$, i.e., $\zeta \notin \sigma(\mathcal{L}(\mathbf{X}))$. Hence, it can be written in piece-wise form

$$\pi_{\Sigma}(\mathbf{X}; \zeta) = \begin{cases} r^{(\zeta)} p^{(\zeta)}(x), & \text{if } (x, \sigma^{-1}(\zeta)) \in \mathbf{X} \\ 1 - r^{(\zeta)}, & \text{if } \zeta \notin \sigma(\mathcal{L}(\mathbf{X})) \end{cases}.$$

The cardinality distribution of an LMB is given by (10).

An LMB is parameterized by the existence probability $r_{\Sigma}(\ell) \triangleq \mathbf{1}_{\mathcal{D}(\sigma)}(\ell) r^{(\sigma(\ell))}$, and attribute density $p_{\Sigma}(\cdot, \ell) \triangleq$

$p^{(\sigma(\ell))}$ of each $\ell \in \mathcal{D}(\sigma)$. Indeed, its PGFI and multi-object density can be expressed as follows [17, pp. 456], [11]

$$\mathbf{G}_{\Sigma}[\mathbf{h}] = \prod_{\ell \in \mathcal{D}(\sigma)} (1 - r_{\Sigma}(\ell) + r_{\Sigma}(\ell) \langle p_{\Sigma}(\cdot, \ell), \mathbf{h}(\cdot, \ell) \rangle), \quad (19)$$

$$\pi_{\Sigma}(\mathbf{X}) = \Delta(\mathbf{X}) [1 - r_{\Sigma}]^{\mathcal{D}(\sigma) - \mathcal{L}(\mathbf{X})} r_{\Sigma}^{\mathcal{L}(\mathbf{X})} p_{\Sigma}^{\mathbf{X}}. \quad (20)$$

Note that $\sum_{L \subseteq \mathbb{L}} [1 - r_{\Sigma}]^{\mathcal{D}(\sigma) - L} r_{\Sigma}^L = 1$ [17, pp. 454], and hence π_{Σ} integrates to 1. For brevity, we write the LMB multi-object density π_{Σ} as $\{(r_{\Sigma}(\ell), p_{\Sigma}(\cdot, \ell))\}_{\ell \in \mathcal{D}(\sigma)}$. An LMB permits elements with the same attributes, and unmarking it only yields a multi-Bernoulli if the attributes are distinct.

Interestingly, much like the Poissons and Bernoullis' (both unlabeled and labeled), the PHD of an LMB [17, pp. 457]

$$\mathbf{v}_{\Sigma}(x, \ell) = r_{\Sigma}(\ell) p_{\Sigma}(x, \ell),$$

also provides a complete characterization, since $r_{\Sigma}(\ell)$ and $p_{\Sigma}(\cdot, \ell)$ can be recovered from the PHD via (17). Indeed, assuming $r_{\Sigma}(\ell) \in [0, 1)$ and noting that $[1 - r_{\Sigma}]^{\mathcal{D}(\sigma)} = [1 - r_{\Sigma}]^{\mathbb{L}}$, π_{Σ} can be written as a multi-object exponential

$$\pi_{\Sigma}(\mathbf{X}) = \Delta(\mathbf{X}) [1 - r_{\Sigma}]^{\mathbb{L}} \left[\frac{\mathbf{v}_{\Sigma}}{1 - r_{\Sigma}} \right]^{\mathbf{X}}, \quad (21)$$

which has the same form (neglecting the distinct label indicator) as the Poisson $e^{-(\mathbf{v}_{\Sigma}, 1)} \mathbf{v}_{\Sigma}^{\mathbf{X}}$. Consequently, the LMB shares many analytical properties with the Poisson, and in this sense, is more ‘‘Poisson’’ than the labeled Poisson. However, unlike the Poisson, the LMB cardinality variance (which cannot exceed the mean) can be controlled by the existence probabilities $r_{\Sigma}(\ell), \ell \in \mathcal{D}(\sigma)$, making it more versatile.

Noting that the joint existence probability $w_{\Sigma}(\cdot)$ of an LMB is given by $[1 - r_{\Sigma}]^{\mathcal{D}(\sigma) - L} r_{\Sigma}^L$ in (20), the LMB admits an analytic labeled-MaM estimate, with the mode L^* of $w_{\Sigma}(\cdot)$ as the most probable label set, and the modes of $p_{\Sigma}(\cdot, \ell), \ell \in L^*$ as the most probable attributes (provided all the relevant modes are available). For an arbitrary LRFS, a PHD-based sub-optimal label-MaM estimate can be obtained by applying the label-MaM estimator to an approximate LMB constructed from its PHD.

D. Generalized Labeled Multi-Bernoulli

The labeled i.i.d. cluster and LMB both have tractable multi-object densities of the form $\Delta(\mathbf{X}) w(\mathcal{L}(\mathbf{X})) p^{\mathbf{X}}$. Extending this form to a mixture accommodates a larger class of LRFS that provides trade-offs between tractability and versatility. A *Generalized Labeled Multi-Bernoulli* (GLMB) Σ is an LRFS distributed according to such a mixture [10], [11], i.e.,

$$\pi_{\Sigma}(\mathbf{X}) = \Delta(\mathbf{X}) \sum_{\xi \in \Xi} w^{(\xi)}(\mathcal{L}(\mathbf{X})) \left[p^{(\xi)} \right]^{\mathbf{X}}, \quad (22)$$

where Ξ is a discrete set, each $p^{(\xi)}(\cdot, \ell)$ is a (probability) density on \mathbb{X} , and each $w^{(\xi)}(L)$ is a non-negative weight that satisfies $\sum_{\xi \in \Xi} \sum_{L \subseteq \mathbb{L}} w^{(\xi)}(L) = 1$. Intuitively, $w^{(\xi)}(L)$ can be interpreted as the probability of *hypothesis* (ξ, L) representing the ‘event’ ξ and the joint existence of the labels in L . Conditional on hypothesis (ξ, L) , $p^{(\xi)}(\cdot, \ell)$ is the attribute (probability) density of object $\ell \in L$.

Remark 4. Each term of the GLMB density (22) is rather general, covering a broad class of LRFSs including labeled i.i.d. clusters, LMBs, and their disjoint union (provided their label sets are disjoint). Thus, the unmarked version of a GLMB is a general class of non-simple point processes that includes the Poisson Multi-Bernoulli Mixture [71] as a special case.

The PGFI of a GLMB takes the form [17, pp. 460]

$$\mathbf{G}_{\Sigma}[\mathbf{h}] = \sum_{\xi \in \Xi} \sum_{L \subseteq \mathbb{L}} w^{(\xi)}(L) \prod_{\ell \in L} \langle p^{(\xi)}(\cdot, \ell), \mathbf{h}(\cdot, \ell) \rangle.$$

Note that the closed-form void probability functional in [72] can be obtained by substituting $1 - \mathbf{1}_{\Sigma}$ into the PGFI. Further, the cardinality and PHD are given by [11]

$$\rho_{\Sigma}(n) = \sum_{\xi \in \Xi} \sum_{L \subseteq \mathbb{L}} \delta_n[|L|] w^{(\xi)}(L),$$

$$\mathbf{v}_{\Sigma}(x, \ell) = \sum_{\xi \in \Xi} p^{(\xi)}(x, \ell) \sum_{L \subseteq \mathbb{L}} \mathbf{1}_L(\ell) w^{(\xi)}(L).$$

Hence, the existence probability and attribute density of object ℓ are [26], [73]

$$r_{\Sigma}(\ell) = \sum_{\xi \in \Xi} \sum_{L \subseteq \mathbb{L}} \mathbf{1}_L(\ell) w^{(\xi)}(L), \quad (23)$$

$$p_{\Sigma}(x, \ell) = \frac{1}{r_{\Sigma}(\ell)} \sum_{\xi \in \Xi} p^{(\xi)}(x, \ell) \sum_{L \subseteq \mathbb{L}} \mathbf{1}_L(\ell) w^{(\xi)}(L). \quad (24)$$

For the label-MaM estimate, the most probable label set of a GLMB is the mode L^* of the joint existence probability

$$w_{\Sigma}(L) = \sum_{\xi \in \Xi} w^{(\xi)}(L), \quad (25)$$

(note the distinction between the probability that L exists regardless of other labels, $\sum_{\xi \in \Xi} \sum_{I \supseteq L} w^{(\xi)}(I)$). However, unlike the LMB, finding the mode of the label-conditioned joint attribute density is not tractable in general. A sub-optimal strategy is to find the mode/mean of the attribute density $p_{\Sigma}(\cdot, \ell)$, for each $\ell \in L^*$. Additionally, the structure of the GLMB suggests a tractable and intuitive sub-optimal version of the MaM estimate, especially when each ξ corresponds to an actual probability event. Based on the most probable cardinality n^* , instead of finding the n^* most probable states, the *GLMB estimator* seeks the most probable hypothesis (ξ^*, L^*) with $|L^*| = n^*$, and estimates the attribute of each $\ell \in L^*$ from $p^{(\xi^*)}(\cdot, \ell)$, via the mode (or mean).

In numerical implementations, it is more convenient to write the GLMB density in δ -GLMB form

$$\pi_{\Sigma}(\mathbf{X}) = \Delta(\mathbf{X}) \sum_{(\xi, I) \in \Xi \times \mathcal{F}(\mathbb{L})} w^{(\xi)}(I) \delta_I[\mathcal{L}(\mathbf{X})] \left[p^{(\xi)} \right]^{\mathbf{X}},$$

using the identity $w^{(\xi)}(L) = \sum_{I \subseteq \mathbb{L}} w^{(\xi)}(I) \delta_I[L]$. For brevity, we denote a GLMB density π_{Σ} by the set

$$\{(w^{(\xi)}(I), p^{(\xi)}): (\xi, I) \in \Xi \times \mathcal{F}(\mathbb{L})\} \quad (26)$$

of its basic *components*, corresponding to set-functions that cannot be decomposed as sums of simpler terms. Note that truncating arbitrary basic components still leaves the resulting δ -GLMB a valid set-function, and does not suffer from the *non-symmetric problem* discussed in Subsection II-C.

Remark 5. An important feature of the GLMB family is its closure under the Bayes filtering and posterior recursions [11] (see Subsections V-C and V-F). Moreover, it is furnished with convenient mathematical properties that facilitate principled approximations (see Subsection IV-F).

E. Information Divergence

Analogous to the classical SSM, information divergences measuring statistical similarities/dissimilarities between RFSs are fundamental in multi-object SSMs. Well-known divergences such as f -divergences (or Csiszár-Morimoto, Ali-Silvey) have been extended to RFS [17, pp. 153-160]. This subsection presents a number of divergences that admit tractable analytic expressions for certain LRFS models, including linear complexity divergences for LMBs, which have not been previously published.

1) *Rényi*: The *Rényi divergence* (*Rényi-D*) between the multi-object densities π_1 and π_2 is given by [17, pp. 156]

$$D_R(\pi_1||\pi_2) = \frac{1}{\alpha - 1} \ln \int \pi_1^\alpha(X) \pi_2^{1-\alpha}(X) \delta X. \quad (27)$$

When $\alpha = 0.5$, D_R reduces to the *Bhattacharyya distance*, and is related to the squared *Hellinger distance*.

While the Rényi-D is intractable in general, it admits closed-forms for (unlabeled) Poissons and LMBs. Indeed, for two Poissons with PHDs v_i , $i = 1, 2$, i.e., $\pi_i(X) = K_i v_i^X$, $K_i \triangleq e^{-\langle v_i, 1 \rangle}$, the Rényi-D is given by [17], [74, pp. 158]

$$D_R(\pi_1||\pi_2) = \frac{\langle v_1^\alpha, v_2^{1-\alpha} \rangle + \ln(K_1^\alpha K_2^{1-\alpha})}{\alpha - 1}.$$

For each LMB parameterized by $\{(r_i(\ell), p_i(\cdot, \ell)) : \ell \in \mathcal{D}(\sigma_i)\}$, $i = 1, 2$, let $\tilde{r}_i \triangleq 1 - r_i$, and $\mathbf{f}_i \triangleq r_i p_i / \tilde{r}_i$, so that its multi-object density has the form² $\pi_i(\mathbf{X}) = \Delta(\mathbf{X}) \tilde{r}_i^{\|\mathbf{X}\|} \mathbf{f}_i^{\mathbf{X}}$. It can be shown that (see Supplementary Materials for proof and alternative forms)

$$D_R(\pi_1||\pi_2) = \sum_{\ell \in \mathbb{L}} \frac{\ln[1 + \langle \mathbf{f}_1^\alpha \mathbf{f}_2^{1-\alpha} \rangle(\ell)] + \ln[(\tilde{r}_1^\alpha \tilde{r}_2^{1-\alpha})(\ell)]}{\alpha - 1}. \quad (28)$$

The above expression reduces a series of high-dimensional integrals to a series of integrals on the attribute space \mathbb{X} . The sum over \mathbb{L} is, in fact, only a sum over the union $\mathcal{D}(\sigma_1) \cup \mathcal{D}(\sigma_2) \subset \mathbb{L}$. Outside this union, the existence probabilities $r_1(\ell)$ and $r_2(\ell)$ are zero, which means $(\tilde{r}_1^\alpha \tilde{r}_2^{1-\alpha})(\ell) = 1$, and has no contribution to $D_R(\pi_1||\pi_2)$. Moreover, $\langle \mathbf{f}_1^\alpha \mathbf{f}_2^{1-\alpha} \rangle(\ell)$ only contributes to $D_R(\pi_1||\pi_2)$ on the intersection $\mathcal{D}(\sigma_1) \cap \mathcal{D}(\sigma_2)$, where both $r_1(\ell)$ and $r_2(\ell)$ are non-zero.

2) *Kullback-Leibler*: The limiting case of the Rényi-D when α tends to 1 is the *Kullback-Leibler divergence* (*KL-D*), given by [17, pp. 155] (with $0 \ln 0 = 0$ by convention)

$$D_{KL}(\pi_1||\pi_2) = \int \pi_1(X) \ln \frac{\pi_1(X)}{\pi_2(X)} \delta X. \quad (29)$$

Similar to the Rényi-D, the KL-D is computationally intractable in general, but admits closed-forms for Poissons and LMBs. Indeed, the KL-D for Poissons [17, pp. 157]

$$D_{KL}(\pi_1||\pi_2) = \ln \frac{K_1}{K_2} + \left\langle v_1, \ln \frac{v_1}{v_2} \right\rangle,$$

²We implicitly assume $r_i(\ell) \in [0, 1)$ to illustrate the similarities between LMBs and Poissons. To include $r_i(\ell) = 1$, we need to use the form (20).

resembles that for LMBs [75]

$$D_{KL}(\pi_1||\pi_2) = \sum_{\ell \in \mathbb{L}} \left[\ln \frac{\tilde{r}_1(\ell)}{\tilde{r}_2(\ell)} + \tilde{r}_1(\ell) \left\langle \mathbf{f}_1 \ln \frac{\mathbf{f}_1}{\mathbf{f}_2} \right\rangle(\ell) \right], \quad (30)$$

(see Supplementary Materials for proof and alternative forms). Note that the differential entropy for LMBs derived in [76] follows from the above expression by setting $\pi_2(\mathbf{X})$ to $\Delta(\mathbf{X})U^{-|\mathbf{X}|}$, where U is the unit of hyper-volume.

3) *Chi-squared*: Similar to Rényi-D and KL-D, the *Chi-squared divergence* (χ^2 -D), given by [17, pp. 155]

$$D_{\chi^2}(\pi_1||\pi_2) = \int \frac{\pi_1^2(X)}{\pi_2(X)} \delta X - 1, \quad (31)$$

belongs to the general class of f -divergences that extends to RFS simply by replacing standard density/integration with FISST density/integration. The χ^2 -D is a 2nd-order Taylor's series approximation of the KL-D, and together with the squared Hellinger distance, provides respectively, the upper and lower bounds for the KL-D [77].

For LMBs, the χ^2 -D admits the following closed-form

$$D_{\chi^2}(\pi_1||\pi_2) = \prod_{\ell \in \mathbb{L}} \frac{\tilde{r}_1^2(\ell)}{\tilde{r}_2(\ell)} \left[1 + \left\langle \frac{\mathbf{f}_1^2}{\mathbf{f}_2} \right\rangle(\ell) \right] - 1 \quad (32)$$

(see Supplementary Materials for proof and alternative forms), which incurs similar computational complexity to those of Rényi-D and KL-D. Note that only $\ell \in \mathcal{D}(\sigma_1) \cup \mathcal{D}(\sigma_2)$ contributes to $D_{\chi^2}(\pi_1||\pi_2)$, because outside this union \tilde{r}_1 and \tilde{r}_2 become unity whilst $\mathbf{f}_1^2/\mathbf{f}_2$ vanishes (using the convention $0^2/0 = 0$). Like the Rényi-D and KL-D, the χ^2 -D for LMBs bears some resemblance to that for Poissons [17, pp. 157]

$$D_{\chi^2}(\pi_1||\pi_2) = \frac{K_1^2}{K_2} e^{\langle v_1^2/v_2, 1 \rangle} - 1.$$

4) *Cauchy-Schwarz*: Unlike the f -divergences, the *Cauchy-Schwarz divergence* (*CS-D*) cannot be extended to RFS by replacing standard density/integration with FISST density/integration due to incompatibility with the unit of measurements [78]. Consequently, extension to RFS is accomplished using density/integration w.r.t. the Poisson measure. Nonetheless, using their equivalence with the FISST density/integral, the CS-D can be expressed as [78]

$$D_{CS}(\pi_1, \pi_2) = -\ln \frac{\int U^{|\mathbf{X}|} \pi_1(X) \pi_2(X) \delta X}{\sqrt{\int U^{|\mathbf{X}|} \pi_1^2(X) \delta X \int U^{|\mathbf{X}|} \pi_2^2(X) \delta X}}, \quad (33)$$

where U is the unit of hyper-volume in \mathcal{X} (the factor $U^{|\mathbf{X}|}$ ensures that all constituent integrals of the set integrals are dimensionless). It can also be interpreted as an approximation to the KL-D [79]. Note that the CS-D between the square roots of the multi-object densities is the Bhattacharyya distance (Rényi-D with $\alpha = 0.5$).

Geometrically, the dissimilarity between π_1 and π_2 , according to the CS-D, is based on the angle they subtend in the space of square integrable functions (via the Cauchy-Schwarz inequality). Hence, the CS-D is symmetric, and invariant to the choice of hyper-volume unit U . For Poissons with PHDs v_1 and v_2 , the CS-D simply reduces to [78]

$$D_{CS}(\pi_1, \pi_2) = \frac{U}{2} \|v_1 - v_2\|^2,$$

where $\|f\|^2 \triangleq \langle f, f \rangle$ denotes the squared L_2 -norm. This means the angle subtended by π_1 and π_2 translates to the squared L_2 -distance between their PHDs.

The versatility of the CS-D in multi-object system lies in the tractable closed-form for the broader GLMB family, which includes LMBs and labeled i.i.d. clusters. Specifically, for $\pi_i = \{(w_i^{(\xi_i)}(I), p_i^{(\xi_i)}) : (\xi_i, I) \in \Xi_i \times \mathcal{F}(\mathbb{L})\}$, $i = 1, 2$, the CS-D is given by [72]

$$D_{CS}(\pi_1, \pi_2) = -\ln \frac{\langle \pi_1, \pi_2 \rangle_U}{\sqrt{\langle \pi_1, \pi_1 \rangle_U} \sqrt{\langle \pi_2, \pi_2 \rangle_U}}, \quad (34)$$

where

$$\langle \pi_i, \pi_j \rangle_U = \sum_{L \subseteq \mathbb{L}} \sum_{\xi_i \in \Xi_i} \sum_{\xi_j \in \Xi_j} w_i^{(\xi_i)}(L) w_j^{(\xi_j)}(L) \langle U p_i^{(\xi_i)} p_j^{(\xi_j)} \rangle^L.$$

The above expression involves summing over all label sets $L \subseteq \mathbb{L}$, though only those with non-zero $w_i^{(\xi_i)}(L)$, $w_j^{(\xi_j)}(L)$ and $\langle U p_i^{(\xi_i)} p_j^{(\xi_j)} \rangle(\ell)$, for all $\ell \in L$, contribute to the sum.

For the LMB special case, the CS-D (34) reduces to a sum of logarithms over $\ell \in \mathbb{L}$, similar to the Rényi-D. Specifically, it can be shown that (see Supplementary Materials)

$$D_{CS}(\pi_1, \pi_2) = -\sum_{\ell \in \mathbb{L}} \ln \frac{1 + \langle U \mathbf{f}_1 \mathbf{f}_2 \rangle(\ell)}{\sqrt{1 + \langle U \mathbf{f}_1^2 \rangle(\ell)} \sqrt{1 + \langle U \mathbf{f}_2^2 \rangle(\ell)}}, \quad (35)$$

though we only need to sum over the union $\mathcal{D}(\sigma_1) \cup \mathcal{D}(\sigma_2)$ of the LMB parameter domains, for the same reason as per the Rényi. Unlike Rényi-D, KL-D, and χ^2 -D, the CS-D for LMBs bears little resemblance to that for Poissons. Note that the CS-D for labeled i.i.d. clusters is more expensive than LMBs, even though both are one-term special cases of the GLMB.

F. Labeled RFS Approximations

While approximations are indispensable for real-world applications, what differentiates heuristics from principled engineering practice is whether the approximation error can be characterized/quantified. The GLMB is an analytic solution to the multi-object Bayes filter that provides trajectory estimation, and is closed under truncation with analytic error characterization [26]. Moreover, the GLMBs also provide principled approximations to other LRFSSs, thus, hitting many birds with one stone. This subsection presents a number of results for LRFSS approximations based on GLMBs.

1) *Truncation of GLMBs*: This is a crucial task in practice, where GLMBs consist of intractably large sums. The key consideration of ensuring validity of the truncated expression as a set-function, discussed in Subsection II-C, is automatically fulfilled (see Subsection IV-D). Moreover, the truncation error can be quantified analytically, similar to that for Fourier series.

For any $\mathbb{H} \subseteq \Xi \times \mathcal{F}(\mathbb{L})$, let $\pi^{(\mathbb{H})}$ denote the unnormalized GLMB $\{(w^{(\xi)}(L), p^{(\xi)}) : (\xi, L) \in \mathbb{H}\}$ (i.e., does not necessarily integrate to 1). Then, the normalizing constant is

$$\int \pi^{(\mathbb{H})}(\mathbf{X}) \delta \mathbf{X} = \|\pi^{(\mathbb{H})}\|_1 = \sum_{(\xi, L) \in \mathbb{H}} w^{(\xi)}(L),$$

where $\|\mathbf{f}\|_1 \triangleq \int |\mathbf{f}(\mathbf{X})| \delta \mathbf{X}$ denotes the L_1 -norm of a function \mathbf{f} on $\mathcal{F}(\mathbb{X} \times \mathbb{L})$. Moreover, the L_1 -error between $\pi^{(\mathbb{H})}$ and its truncated version $\pi^{(\mathbb{T})}$ is given by [26]

$$\|\pi^{(\mathbb{H})} - \pi^{(\mathbb{T})}\|_1 = \sum_{(\xi, L) \in \mathbb{H} - \mathbb{T}} w^{(\xi)}(L), \quad (36)$$

(it is implicit that $\mathbb{T} \subseteq \mathbb{H}$), and for their normalized versions

$$\left\| \frac{\pi^{(\mathbb{H})}}{\|\pi^{(\mathbb{H})}\|_1} - \frac{\pi^{(\mathbb{T})}}{\|\pi^{(\mathbb{T})}\|_1} \right\|_1 \leq 2 \frac{\|\pi^{(\mathbb{H})}\|_1 - \|\pi^{(\mathbb{T})}\|_1}{\|\pi^{(\mathbb{H})}\|_1}. \quad (37)$$

The above result means that the intuitive strategy of discarding basic GLMB components with the smallest weights minimizes the L_1 -error between the actual and approximate multi-object densities. Similar truncations are widely used in MOT algorithms such as MHT, and JPDA, but without mathematical justifications. The LRFS approach characterizes the effect of truncation by the L_1 -error, and provides a mathematical justification for truncating low-weighted components.

2) *Label-Partitioned GLMB Approximation*: A ‘large’ GLMB can be approximated, via label partitioning, as a product of much ‘smaller’ GLMBs that can be processed more efficiently in parallel [27]. The rationale is that such approximations incur minimal information loss when the smaller GLMBs are almost independent, which is usually the case in practice since the objects are not uniformly distributed across the state space, but often in separate groups.

Given a partition \mathcal{L} of the label space \mathbb{L} , each $\mathbf{X} \in \mathcal{F}(\mathbb{X} \times \mathbb{L})$ can be written as $\mathbf{X} = \bigsqcup_{L \in \mathcal{L}} \mathbf{X} \cap (\mathbb{X} \times L)$, and hence $\{\mathcal{F}(\mathbb{X} \times L) : L \in \mathcal{L}\}$ also forms a partition of $\mathcal{F}(\mathbb{X} \times \mathbb{L})$. A labeled multi-object density on $\mathcal{F}(\mathbb{X} \times \mathbb{L})$ is said to be \mathcal{L} -partitioned if it can be written as the product

$$\pi_{\mathcal{L}}(\mathbf{X}) = \prod_{L \in \mathcal{L}} \pi_{\mathcal{L}}^{(L)}(\mathbf{X} \cap (\mathbb{X} \times L)),$$

where each factor $\pi_{\mathcal{L}}^{(L)}$ is a labeled multi-object density on $\mathcal{F}(\mathbb{X} \times L)$ [27]. We denote $\pi_{\mathcal{L}}$, by its factors $\{\pi_{\mathcal{L}}^{(L)}\}_{L \in \mathcal{L}}$, and if each factor $\pi_{\mathcal{L}}^{(L)}$ is a GLMB $\{(w_{\mathcal{L},L}^{(I,\xi)}, p_{\mathcal{L},L}^{(\xi)})\}_{(I,\xi) \in \mathcal{F}(L) \times \Xi(L)}$, then $\pi_{\mathcal{L}}$ is said to be an \mathcal{L} -partitioned GLMB.

An important numerical problem is to approximate a given \mathcal{L} -partitioned GLMB $\pi_{\mathcal{L}} = \{\pi_{\mathcal{L}}^{(L)}\}_{L \in \mathcal{L}}$ using another partition \mathcal{S} of \mathbb{L} . Indeed, the \mathcal{S} -partitioned labeled multi-object density $\pi_{\mathcal{S}} = \{\pi_{\mathcal{S}}^{(S)}\}_{S \in \mathcal{S}}$ that minimizes $D_{KL}(\pi_{\mathcal{L}} \|\pi_{\mathcal{S}})$, is an \mathcal{S} -partitioned GLMB, with GLMB factors [27]

$$\pi_{\mathcal{S}}^{(S)}(\mathbf{X} \cap (\mathbb{X} \times S)) = \prod_{L \in \mathcal{L}} \pi_{\mathcal{L},\mathcal{S}}^{(L,S)}(\mathbf{X} \cap (\mathbb{X} \times S)),$$

where, for each $(L, S) \in \mathcal{L} \times \mathcal{S}$ such that $L \cap S \neq \emptyset$

$$\begin{aligned} \pi_{\mathcal{L},\mathcal{S}}^{(L,S)} &= \{(w_{L,S}^{(H,\xi)}, p_{L,S}^{(\xi)})\}_{(H,\xi) \in \mathcal{F}(L \cap S) \times \Xi(L)}, \\ w_{L,S}^{(H,\xi)} &= \sum_{W \in \mathcal{F}(L-S)} w_{\mathcal{L},L}^{(H \cup W, \xi)}, \end{aligned}$$

$$p_{L,S}^{(\xi)}(x, \ell) = \mathbf{1}_{L \cap S}(\ell) p_{\mathcal{L},L}^{(\xi)}(x, \ell).$$

An efficient algorithm for finding the ‘best’ partitions suitable for large-scale multi-object estimation can be found in [27].

3) *Approximation by GLMB*: Estimation involving non-standard multi-object models invariably results in intractable multi-object densities, see e.g., [70], [80]–[87]. Hence, it is important to make principled approximations by tractable families such as GLMBs. Apart from statistical approximation techniques such as moment matching and information minimization, approximating the cardinality distribution is an important consideration in multi-object estimation [16].

An LRFS, characterized by a multi-object density π , can be approximated by an LMB with matching 1st moment (PHD), by choosing the parameters according to (17). Indeed, this technique was first used in [73] to approximate GLMBs with LMBs by selecting the parameters according to (23), (24). However, LMBs are not versatile enough to match the cardinality or capture the dependence between the attributes of the multi-object state. To this end, approximation by GLMBs can provide trade-offs between tractability and versatility.

A *Marginalized-GLMB (M-GLMB)* is a GLMB with density of the form [70], [88], [89]

$$\bar{\pi}(\mathbf{X}) = \Delta(\mathbf{X}) \sum_{L \subseteq \mathbb{L}} \bar{w}^{(L)} \delta_L[\mathcal{L}(\mathbf{X})] \left[\bar{p}^{(L)} \right]^{\mathbf{X}}. \quad (38)$$

Using a smaller number of components, an M-GLMB can approximate a GLMB $\{(w^{(\xi)}(I), p^{(\xi)}) : (\xi, I) \in \Xi \times \mathcal{F}(\mathbb{L})\}$ with matching PHD and cardinality, by choosing [88]

$$\begin{aligned} \bar{w}^{(L)} &= \sum_{\xi \in \Xi} w^{(\xi)}(L), \\ \bar{p}^{(L)}(\mathbf{x}) &= \frac{\mathbf{1}_L(\mathcal{L}(\mathbf{x}))}{\bar{w}^{(L)}} \sum_{\xi \in \Xi} w^{(\xi)}(L) p^{(\xi)}(\mathbf{x}). \end{aligned}$$

For a labeled multi-object density π , the M-GLMB that matches the PHD and cardinality whilst minimizing the KL-D from π has components given by [70]

$$\begin{aligned} \bar{w}^{(L)} &= \langle \pi \rangle(L), \\ \bar{p}^{(L)}(\mathbf{x}) &= \frac{\mathbf{1}_L(\mathcal{L}(\mathbf{x}))}{\langle \pi \rangle(L)} \langle \pi(\{\mathbf{x}\} \uplus \cdot) \rangle(L - \{\mathcal{L}(\mathbf{x})\}). \end{aligned} \quad (39) \quad (40)$$

Thus, to match the PHD and cardinality of a labeled multi-object density π by an M-GLMB, we set each weight $\bar{w}^{(L)}$ to the joint existence probability $\langle \pi \rangle(L)$, and the $\bar{p}^{\{\ell_{1:n}\}}(\cdot, \ell_i)$'s to the marginals of the label-conditioned joint attribute densities $\pi(\{\cdot, \ell_1\}, \dots, \{\cdot, \ell_n\}) / \langle \pi \rangle(\{\ell_{1:n}\})$. This approximation minimizes the KL-D in a similar way to the approximation of a joint density by the product of its marginals [90].

In certain applications, e.g., [80], [83], [85]–[87], the labeled multi-object density of interest takes on a multi-modal form similar to a GLMB:

$$\pi(\mathbf{X}) = \Delta(\mathbf{X}) \sum_{\xi \in \Xi} w^{(\xi)}(\mathcal{L}(\mathbf{X})) p^{(\xi)}(\mathbf{X}), \quad (41)$$

where each ξ represents a mode, $\sum_{\xi \in \Xi} \sum_{L \subseteq \mathbb{L}} w^{(\xi)}(L) = 1$, and $\langle p^{(\xi)} \rangle(I) = 1$. The M-GLMB approximation cannot capture the modes and the associated information. Nonetheless, the cardinality and PHD matching strategy can be applied to each of the modes [70]. Specifically, a GLMB that matches




Fig. 10: A multi-object trajectory $\mathbf{X}_{0:k}$ on the interval $\{0:k\}$. Individual trajectories are determined by grouping the states of $\mathbf{X}_{0:k}$ according to labels (depicted by the different colors). This representation covers trajectory fragmentation (e.g., at time 4) and crossing (e.g., at time 3).

the PHD and cardinality whilst preserving the modes of (41) is given by $\hat{\pi} = \{(\hat{w}^{(\xi, I)}, \hat{p}^{(\xi, I)}) : (\xi, I) \in \Xi \times \mathcal{F}(\mathbb{L})\}$, where

$$\hat{w}^{(\xi, I)} = w^{(\xi)}(I), \quad (42)$$

$$\hat{p}^{(\xi, I)}(\mathbf{x}) = \mathbf{1}_I(\mathcal{L}(\mathbf{x})) \langle p^{(\xi)}(\{\mathbf{x}\} \uplus \cdot) \rangle(I - \{\mathcal{L}(\mathbf{x})\}). \quad (43)$$

While this approximation requires many more components than the M-GLMB approximation, intuitively, it incurs less information loss, by retaining the information contained in the modes. However, there are no formal results on the KL-D.

G. Spatio-Temporal Modeling and Multi-Scan GLMB

So far, we have only discussed modeling of the multi-object state via LRFS. This subsection extends the discussion to multi-object trajectory modeling. In particular, we present an extension of the GLMB, known as the multi-scan GLMB, as a tractable LRFS model of the multi-object trajectory.

Recall that the multi-object trajectory on an interval $\{j:k\}$ is the sequence $\mathbf{X}_{j:k}$ of labeled multi-object states, and that trajectories in $\mathbf{X}_{j:k}$ are determined by grouping the states according to labels, see Fig. 10. More concisely, the trajectory of each (object with label) $\ell \in \mathcal{L}(\mathbf{X}_{j:k})$ ³ is the time-stamped sequence $\mathbf{x}_{T(\ell)}^{(\ell)} = [(x_i, \ell) \in \mathbf{X}_i]_{i \in T(\ell)}$, where $T(\ell)$ is the set of instants in $\{j:k\}$ such that ℓ exists. This sequence, defines the mapping $\tau_\ell : i \mapsto x_i, i \in T(\ell)$. It is clear that $\mathbf{X}_{j:k}$ can be reconstructed from all the trajectories in $\mathbf{X}_{j:k}$, and hence is equivalently represented as the set of labeled trajectories

$$\mathbf{X}_{j:k} \equiv \left\{ \mathbf{x}_{T(\ell)}^{(\ell)} : \ell \in \mathcal{L}(\mathbf{X}_{j:k}) \right\}. \quad (44)$$

Since the multi-object trajectory on $\{j:k\}$ is represented by a sequence of labeled multi-object states, it is naturally modeled as a sequence $\Sigma_{j:k}$ of LRFSs described by a joint LRFS density $\pi_{\Sigma_{j:k}}$. Beyond multi-object trajectory modeling, statistical characterization of variables/parameters pertaining to the underlying multi-object trajectory ensemble can be computed from the joint LRFS density [91]. Suppose that

³Strictly speaking $\mathcal{L}(\mathbf{X}_{j:k}) = (\mathcal{L}(\mathbf{X}_j), \dots, \mathcal{L}(\mathbf{X}_k))$ and we should write $\ell \in \cup_{i=j}^k \mathcal{L}(\mathbf{X}_i)$. Nonetheless, the notation $\ell \in \mathcal{L}(\mathbf{X}_{j:k})$ is more compact, and also suggestive (that ℓ belongs to any of the sets $\mathcal{L}(\mathbf{X}_j), \dots, \mathcal{L}(\mathbf{X}_k)$).

$f(\mathbf{X}_{j:k})$ is the statistic of a multi-object trajectory $\mathbf{X}_{j:k}$, then the ensemble statistic is the expectation

$$\mathbb{E}_{\Sigma_{j:k}} [f] = \int \dots \int f(\mathbf{X}_{j:k}) \pi_{\Sigma_{j:k}}(\mathbf{X}_{j:k}) \delta \mathbf{X}_j \dots \delta \mathbf{X}_k.$$

Fig. 11 shows some examples of multi-object trajectory statistics. More examples can be found in [91].

1) *Multi-Scan Multi-Object Exponential*: The basic building block for a multi-scan GLMB is a multi-scan version of the multi-object exponential. Noting that $h^{\mathbf{X}}$ is the product of the values of the function h at every state in \mathbf{X} , a natural multi-scan extension to $h^{\mathbf{X}_{j:k}}$ is the product of the values of h at every trajectory in $\mathbf{X}_{j:k}$. More concisely, the *multi-scan multi-object exponential* is defined as [91]

$$h^{\mathbf{X}_{j:k}} \triangleq h^{\{\mathbf{x}_{T(\ell)}^{(\ell)} : \ell \in \mathcal{L}(\mathbf{X}_{j:k})\}} = \prod_{\ell \in \mathcal{L}(\mathbf{X}_{j:k})} h(\mathbf{x}_{T(\ell)}^{(\ell)}), \quad (45)$$

for any real function h on $\bigcup_{I \subseteq \{j:k\}} \mathbb{T}_I$, where

$$\mathbb{T}_{\{i_1, i_2, \dots, i_n\}} \triangleq \prod_{j=1}^n (\mathbb{X} \times \mathbb{L}_{i_j}), \quad \mathbb{I}_{i=j}^k S_i \triangleq S_j \times \dots \times S_k,$$

for $i_1 < i_2 < \dots < i_n \in \{j:k\}$. If $\mathbf{x}_{T(\ell)}^{(\ell)} = [(x_i, \ell)]_{i \in T(\ell)}$, we write $h(\mathbf{x}_{T(\ell)}^{(\ell)})$ as $h(x_{T(\ell)}; \ell)$, and if the trajectory is unfragmented, i.e., $T(\ell) = \{s(\ell) : t(\ell)\}$, where $s(\ell) \triangleq \min T(\ell)$, $t(\ell) \triangleq \max T(\ell)$, we write $h(\mathbf{x}_{T(\ell)}^{(\ell)})$ as $h(x_{s(\ell):t(\ell)}; \ell)$.

The multi-scan multi-object exponential (45) satisfies the exponential-like property:

$$[gh]^{\mathbf{X}_{j:k}} = [g]^{\mathbf{X}_{j:k}} [h]^{\mathbf{X}_{j:k}},$$

where g is another function on $\bigcup_{I \subseteq \{j:k\}} \mathbb{T}_I$, see [91] for additional properties. When $j = k$, (45) reduces to the single-scan multi-object exponential $h^{\mathbf{X}^j}$.

2) *Multi-Scan GLMB*: This model was proposed in [91] for smoothing under the standard multi-object system model that only permits unfragmented trajectories. A *multi-scan GLMB* on the interval $\{j:k\}$ is a sequence $\Sigma_{j:k}$ of LRFSs described by a joint multi-object density on $\prod_{i=j}^k \mathcal{F}(\mathbb{X} \times \mathbb{L}_i)$ of the form

$$\pi_{\Sigma_{j:k}}(\mathbf{X}_{j:k}) = \Delta(\mathbf{X}_{j:k}) \sum_{\xi \in \Xi} w^{(\xi)}(\mathcal{L}(\mathbf{X}_{j:k})) [p^{(\xi)}]^{\mathbf{X}_{j:k}}, \quad (46)$$

where: $\Delta(\mathbf{X}_{j:k}) \triangleq \prod_{i=j}^k \Delta(\mathbf{X}_i)$; $w^{(\xi)}(I_{j:k})$ is non-negative such that $\sum_{\xi, I_{j:k}} w^{(\xi)}(I_{j:k}) = 1$ (it is understood that the sum is taken over $\xi \in \Xi$ and $I_{j:k} \in \prod_{i=j}^k \mathcal{F}(\mathbb{L}_i)$); $\mathcal{L}(\mathbf{X}_{j:k}) = (\mathcal{L}(\mathbf{X}_j), \dots, \mathcal{L}(\mathbf{X}_k))$; and $p^{(\xi)}(x_{s(\ell):t(\ell)}; \ell)$ is a joint density of the attribute sequence $x_{s(\ell):t(\ell)}$, for each $\ell \in I_{j:k}$, with $s(\ell)$, $t(\ell)$ implicitly depend on $(\xi, I_{j:k})$. The joint density (46) indeed integrates to 1 [91]. Similar to the GLMB, the multi-scan GLMB (46) can also be written in delta form

$$\pi_{\Sigma_{j:k}}(\mathbf{X}_{j:k}) = \Delta(\mathbf{X}_{j:k}) \sum_{\xi, I_{j:k}} w^{(\xi, I_{j:k})} \delta_{I_{j:k}}[\mathcal{L}(\mathbf{X}_{j:k})] [p^{(\xi)}]^{\mathbf{X}_{j:k}},$$

and denoted by $\pi_{\Sigma_{j:k}} = \{(w^{(\xi)}(I_{j:k}), p^{(\xi)}(\xi, I_{j:k}))\}$, where it is understood that $\xi \in \Xi$ and $I_{j:k} \in \prod_{i=j}^k \mathcal{F}(\mathbb{L}_i)$.

Conceptually, a multi-scan GLMB can be regarded as a GLMB with the set of labeled states replaced by the set of labeled trajectories (44). Analogous to the GLMB, $w^{(\xi)}(I_{j:k})$ is the probability of hypothesis $(\xi, I_{j:k})$ representing the ‘event’ ξ and the joint existence of the trajectories according to

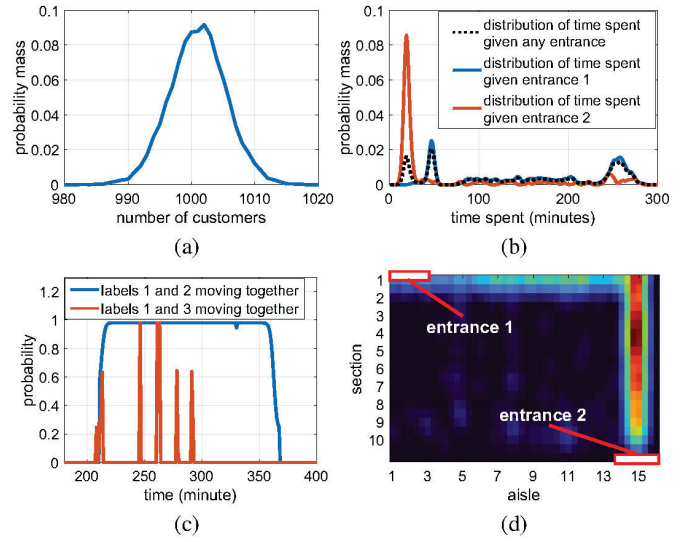


Fig. 11: Multi-object trajectory statistics in a retail store with 15 aisles and 2 entrances. Computed from a hypothetical joint LRFS density of a 1-day scenario: (a) distribution of customer numbers in the store; (b) distribution of the time customers spend in the store given the entrance they use; (c) probability that certain customers move together (within 2m of each other) over time; (d) location intensity of customers who spend less than 60 minutes in the store.

the sequence $I_{j:k}$ of label sets, and conditional on hypothesis $(\xi, I_{j:k})$, $p^{(\xi)}(\cdot, \ell)$ is the attribute density of trajectory $\ell \in I_{j:k}$. The multi-scan GLMB is closed under the Bayes posterior recursion as well as truncation. Indeed, the error expressions (36), (37) also hold for multi-scan GLMBs [91].

Some closed-form multi-object statistics from the multi-scan GLMB are given as follows (see [91] for further details).

- Trajectory cardinality distribution over interval $\{j:k\}$:

$$\Pr(|\mathbf{X}_{j:k}| = n) = \sum_{\xi, I_{j:k}} \delta_{|\cup_{i=j}^k I_i|} [n] w^{(\xi)}(I_{j:k}). \quad (47)$$

- Joint existence probability of only the labels in L (and no other labels exist) over interval $\{j:k\}$:

$$\Pr(\text{only } L \text{ exist}) = \sum_{\xi, I_{j:k}} \delta_{\cup_{i=j}^k I_i} [L] w^{(\xi)}(I_{j:k}). \quad (48)$$

- Joint existence probability of the labels in L regardless of other labels, over interval $\{j:k\}$:

$$\Pr(L \text{ exist}) = \sum_{\xi, I_{j:k}} \mathbf{1}_{\mathcal{F}(\cup_{i=j}^k I_i)}(L) w^{(\xi)}(I_{j:k}). \quad (49)$$

- Trajectory length distribution, i.e., the probability that a trajectory has length n , over interval $\{j:k\}$:

$$\Lambda(n) = \sum_{\xi, I_{j:k}} \frac{\sum_{\ell \in \cup_{i=j}^k I_i} \delta_{t(\ell)-s(\ell)+1} [n]}{|\cup_{i=j}^k I_i|} w^{(\xi)}(I_{j:k}). \quad (50)$$

- Length distribution of trajectory ℓ , i.e., the probability that trajectory ℓ has length n , over interval $\{j:k\}$:

$$\Lambda_{\ell}(n) = \sum_{\xi, I_{j:k}} \mathbf{1}_{\cup_{i=j}^k I_i}(\ell) \delta_{t(\ell)-s(\ell)+1} [n] w^{(\xi)}(I_{j:k}). \quad (51)$$

3) *Multi-Scan GLMB Estimator*: Various GLMB estimators can be extended to multi-scan GLMB. The simplest would be to find the most probable hypothesis $(\xi^*, I_{j:k}^*)$ (with highest weight $w^{(\xi^*)}(I_{j:k}^*)$) and compute the most probable or expected trajectory estimate from $p^{(\xi^*)}(\cdot; \ell)$ for each ℓ in $I_{j:k}^*$. Alternatively, instead of the most significant, we can use the most significant among components with the most probable trajectory cardinality n^* determined by maximizing (47).

The label-MaM estimate can also be adapted for multi-scan GLMB, using the most probable label set sequence $I_{j:k}^*$, determined by maximizing $\sum_{\xi} w^{(\xi)}(I_{j:k})$, and computing the trajectory density for each ℓ in $I_{j:k}^*$

$$p(\cdot; \ell) \propto \sum_{\xi} w^{(\xi)}(I_{j:k}^*) p^{(\xi)}(\cdot; \ell),$$

from which the mode or mean trajectory can then be determined. One variation is to choose $I_{j:k}^*$ as the most probable label set sequence with trajectory cardinality n^* .

Another adaptation of the label-MaM estimate is based on the most probable set of labels L^* , determined by maximizing the joint existence probability (48). However, for each $\ell \in L^*$, its trajectory length for different hypotheses $(\xi, I_{j:k})$ (that yield the same L^*) may not be the same, and hence there is no meaningful most probable trajectory (due to different units of measurements in the probability densities of the trajectories). Nonetheless, we can determine the most probable length m^* for each $\ell \in L^*$, by maximizing (51), and estimate the trajectory according to

$$p_{m^*}(\cdot; \ell) \propto \sum_{\xi, I_{j:k}} \mathbf{1}_{\cup_{i=j}^k I_i}(\ell) w^{(\xi)}(I_{j:k}) \delta_{t(\ell)-s(\ell)+1}[m^*] p^{(\xi)}(\cdot; \ell).$$

Variations of this estimator are to use the label set L^* of cardinality n^* with highest joint existence probability, or the set of n^* labels with highest individual existence probabilities.

V. DYNAMIC MULTI-OBJECT ESTIMATION

The apparatus developed in the previous sections enables classical SSM concepts to be applied to multi-object state estimation, as shown in this section. Bayesian estimation for multi-object SSM is formulated in Subsection V-A, while Subsection V-B presents the standard multi-object SSM. An exact solution to the Bayes multi-object filter using GLMB, and an approximate solution using LMB, are presented in Subsections V-C and V-D, while numerical implementations are discussed in Subsection V-E. Extension of the GLMB filter to multi-object smoothing is presented in Subsection V-F, while Subsection V-G discusses its implementation. Extensions to non-standard models, robust and distributed estimation, and control for multi-object system are discussed in Subsections V-H, V-I, V-J, and V-K.

A. Multi-Object State Space Model

Recall that a label $\ell = (s, \iota)$ consists of the starting time s and an index ι , let \mathbb{B}_k denote the space of labels with starting time k . Then, the label space and labeled state space at time k are, respectively, the disjoint union $\mathbb{L}_k = \bigsqcup_{t=0}^k \mathbb{B}_t$ and $\mathbb{X} \times \mathbb{L}_k$.

TABLE IV: Common notations from Section V.

Notation	Description
\mathbb{B}_k	Space of labels (of objects) born at time k
\mathbb{L}_k	$\bigsqcup_{t=0}^k \mathbb{B}_t$, label space at time k
\mathbf{X}_k	Multi-object state at time k
Z_k	Multi-object measurement at time k
$\mathbf{f}_k(\cdot \mathbf{X}_{k-1})$	Multi-object transition density given \mathbf{X}_{k-1}
$\mathbf{g}_k(Z_k \mathbf{X}_k)$	Likelihood of observing Z_k given \mathbf{X}_k
$\pi_{0:k}(\mathbf{X}_{0:k})$	Multi-object posterior density at $\mathbf{X}_{0:k}$
$\pi_k(\mathbf{X}_k)$	Multi-object filtering density at \mathbf{X}_k
$\mathbf{f}_{B,k}$	Density of LRFS of new objects at time k
$P_{B,k}(\ell)$	Birth probability of label ℓ at time k
$p_B(\cdot, \ell)$	Attribute density for newborn label ℓ at time k
$\mathbf{f}_{S,k}(\cdot \mathbf{X}_{k-1})$	Density of LRFS at time k generated by \mathbf{X}_{k-1}
$P_{S,k}(x_{k-1}, \ell)$	Survival probability to time k of state (x_{k-1}, ℓ)
$f_{S,k}(\cdot, \ell)$	Attribute transition density to time k for ℓ
$P_{D,k}(\mathbf{x})$	Detection probability of state \mathbf{x} at time k
κ_k	Clutter intensity function at time k
$\Psi_{Z,k}^{(j)}(\mathbf{x})$	$\mathbf{1}_{\{1: Z \}}(j) \frac{P_{D,k}(\mathbf{x}) g_k(z_j \mathbf{x})}{\kappa_k(z_j)} + \delta_0[j] (1 - P_{D,k}(\mathbf{x}))$
$\Lambda_{S,k}^{(j)}(x \varsigma, \ell)$	$\Psi_{Z,k}^{(j)}(x, \ell) f_{S,k}(x \varsigma, \ell) P_{S,k}(\varsigma, \ell)$
$\Lambda_{B,k}^{(j)}(x, \ell)$	$\Psi_{Z,k}^{(j)}(x, \ell) p_{B,k}(x, \ell) P_{B,k}(\ell)$
θ, θ_k	Association map $\mathbb{L}_k \rightarrow \{0 : Z_k \}$ at time k
$\Theta(I), \Theta_k(I)$	Space of association maps θ_k with domain I
(ξ, I)	Prior GLMB component
(ξ, θ, I)	Updated GLMB component
γ, γ_k	Extended association $I_{k-1} \uplus \mathbb{B}_k \rightarrow \{-1 : Z_k \}$
Γ, Γ_k	Space of extended associations at time k
P, P_k	Number of (hypothesized) labels at time k
M, M_k	Number of measurements at time k

Analogous to traditional SSMs, the multi-object state $\mathbf{X}_k \in \mathcal{F}(\mathbb{X} \times \mathbb{L}_k)$, at time k , evolves from its previous value $\mathbf{X}_{k-1} \in \mathcal{F}(\mathbb{X} \times \mathbb{L}_{k-1})$, and generates an observation $Z_k \in \mathcal{F}(\mathcal{Z})$, according to the multi-object state and observation equations

$$\mathbf{X}_k = \mathbf{S}_k(\mathbf{X}_{k-1}) \cup \mathbf{B}_k, \quad (52)$$

$$Z_k = D_k(\mathbf{X}_k) \cup K_k, \quad (53)$$

where $\mathbf{S}_k(\mathbf{X}_{k-1})$ is the set of states generated from \mathbf{X}_{k-1} , \mathbf{B}_k is the set of newly appearing states, $D_k(\mathbf{X}_k)$ is the set of detections generated from \mathbf{X}_k , and K_k is the set of clutter, (or false alarms).

Under the Bayesian paradigm, the multi-object state and observation are modeled as RFSs. Using the FISST notion of multi-object density, the multi-object state and observation equations (52), (53) can be characterized, respectively, by the *multi-object (Markov) transition density* and *multi-object (observation) likelihood function* [9], [68]

$$\mathbf{f}_k(\mathbf{X}_k | \mathbf{X}_{k-1}), \quad \mathbf{g}_k(Z_k | \mathbf{X}_k).$$

The transition density $\mathbf{f}_k(\cdot | \cdot)$ captures the underlying evolution, appearance and disappearance of the objects. The observation likelihood $\mathbf{g}_k(\cdot | \cdot)$ captures the underlying detections, false negatives/positives and data association uncertainty. Examples of $\mathbf{f}_k(\cdot | \cdot)$ and $\mathbf{g}_k(\cdot | \cdot)$ are given in the standard (commonly used) SSM presented in Subsection V-B.

Given the observation history $Z_{1:k}$, all information on the multi-object trajectory is captured in the *multi-object posterior density* $\pi_{0:k}(\mathbf{X}_{0:k}) \triangleq \pi_{0:k}(\mathbf{X}_{0:k}|Z_{1:k})$ (the dependence on $Z_{1:k}$ is omitted for brevity). Similar to Bayesian estimation for traditional SSMs [34], [37], the (multi-object) posterior density can be propagated forward recursively by [91],

$$\pi_{0:k}(\mathbf{X}_{0:k}) = \frac{\mathbf{g}_k(Z_k|\mathbf{X}_k)\mathbf{f}_k(\mathbf{X}_k|\mathbf{X}_{k-1})\pi_{0:k-1}(\mathbf{X}_{0:k-1})}{\int \mathbf{g}_k(Z_k|\mathbf{Y}_k)\mathbf{f}_k(\mathbf{X}_k|\mathbf{Y}_{k-1})\pi_{0:k-1}(\mathbf{Y}_{0:k-1})\delta\mathbf{Y}_{0:k}}. \quad (54)$$

Since the dimension of $\pi_{0:k}$ increases with k , the computational complexity for each iterate of (54), which is a function of the dimension, grows with time. This growth is far worse than its single-object counterpart because the multi-object state space $\mathcal{F}(\mathbb{X} \times \mathbb{L}_k)$ is far larger than the state space \mathbb{X} .

A cheaper alternative is the *multi-object filtering density*, $\pi_k(\mathbf{X}_k) \triangleq \int \pi_{0:k}(\mathbf{X}_{0:k})\delta\mathbf{X}_{0:k-1}$, which can be propagated by the *multi-object Bayes filter* [17], [68]

$$\pi_k(\mathbf{X}_k) = \frac{\mathbf{g}_k(Z_k|\mathbf{X}_k)\int \mathbf{f}_k(\mathbf{X}_k|\mathbf{Y})\pi_{k-1}(\mathbf{Y})\delta\mathbf{Y}}{\int \mathbf{g}_k(Z_k|\mathbf{X})\int \mathbf{f}_k(\mathbf{X}|\mathbf{Y})\pi_{k-1}(\mathbf{Y})\delta\mathbf{Y}\delta\mathbf{X}}. \quad (55)$$

Similar to the single-object Bayes filter, the computational complexity for each iterate of (55) does not increase with time since the dimension of π_k does not grow with k .

Under the standard multi-object model (to be discussed next), numerical solutions to the posterior and filtering recursions were first developed at around the same time, respectively, in [92], [93] using particle marginal Metropolis-Hasting simulation, and in [10], [11] using GLMBs. The latter is an analytic solution, and was later extended to solve the posterior recursion in [91], [94] (see Subsections V-C, and V-F).

B. Standard Multi-Object Model

1) *Multi-Object State Dynamic*: Let $\mathbf{f}_{B,k}$ denote the density of the LRFS \mathbf{B}_k of new objects, and $\mathbf{f}_{S,k}(\cdot|\mathbf{X}_{k-1})$ the density of the LRFS $\mathbf{S}_k(\mathbf{X}_{k-1})$ of objects generated from the previous multi-object state \mathbf{X}_{k-1} . Assuming $\mathbf{S}_k(\mathbf{X}_{k-1})$ and \mathbf{B}_k are independent, due to the labeling construct, the multi-object transition density is given by [11]

$$\mathbf{f}_k(\mathbf{X}_k|\mathbf{X}_{k-1}) = \mathbf{f}_{B,k}(\mathbf{X}_k \cap (\mathbb{X} \times \mathbb{B}_k)) \mathbf{f}_{S,k}(\mathbf{X}_k \cap (\mathbb{X} \times \mathbb{L}_{k-1})|\mathbf{X}_{k-1}). \quad (56)$$

In the *standard* (commonly used) *multi-object dynamic model*, \mathbf{B}_k is a GLMB on $\mathcal{F}(\mathbb{X} \times \mathbb{B}_k)$. For simplicity (but without loss of generality), we use the one-term GLMB

$$\mathbf{f}_{B,k}(\mathbf{X}) = \Delta(\mathbf{X})w_{B,k}(\mathcal{L}(\mathbf{X}))p_{B,k}^{\mathbf{X}}. \quad (57)$$

Note that $\mathbf{f}_{B,k}(\mathbf{X}) = 0$ if \mathbf{X} contains any element with label outside of \mathbb{B}_k . This birth model is general enough to include labeled i.i.d cluster and LMB, though the latter, given by $\mathbf{f}_{B,k} = \{(P_{B,k}(\ell), p_{B,k}(\cdot, \ell))\}_{\ell \in \mathbb{B}_k}$, is most popular.

Further, for a given \mathbf{X}_{k-1} , each $(x_{k-1}, \ell) \in \mathbf{X}_{k-1}$ either survives with probability $P_{S,k}(x_{k-1}, \ell)$ and evolves to state (x_k, ℓ) at time k , with the same label and attribute transition density $f_{S,k}(x_k|x_{k-1}, \ell)$, or dies with probability $1 - P_{S,k}(x_{k-1}, \ell)$. This means objects keep the same labels for their entire lives. Assuming that conditional on \mathbf{X}_{k-1} each

object survives and evolves independently of one another, $\mathbf{S}_k(\mathbf{X}_{k-1})$ is an LMB on $\mathcal{F}(\mathbb{X} \times \mathbb{L}_{k-1})$, with parameters $\{(P_{S,k}(\zeta), f_{S,k}(\cdot|\zeta))\}_{\zeta \in \mathbf{X}_{k-1}}$, and density given by [11]

$$\mathbf{f}_{S,k}(\mathbf{X}|\mathbf{X}_{k-1}) = \Delta(\mathbf{X})\mathbf{1}_{\mathcal{L}(\mathbf{X}_{k-1})}^{\mathcal{L}(\mathbf{X})} [\mathbf{f}_{S,k}(\mathbf{X}; \cdot)]^{\mathbf{X}_{k-1}}, \quad (58)$$

where $\mathbf{1}_{\mathcal{L}(\mathbf{X}_{k-1})}^{\mathcal{L}(\mathbf{X})} = \prod_{\ell \in \mathcal{L}(\mathbf{X})} \mathbf{1}_{\mathcal{L}(\mathbf{X}_{k-1})}(\ell)$, and

$$\mathbf{f}_{S,k}(\mathbf{X}; y, \ell) = \begin{cases} P_{S,k}(y, \ell) f_{S,k}(x|y, \ell), & \text{if } (x, \ell) \in \mathbf{X} \\ 1 - P_{S,k}(y, \ell), & \text{if } \ell \notin \mathcal{L}(\mathbf{X}) \end{cases}. \quad (59)$$

The standard multi-object transition density only generates unfragmented trajectories, and is completely characterized by the model parameters $w_{B,k}, p_{B,k}, P_{S,k}, f_{S,k}$.

More sophisticated multi-object dynamic models include spawnings [16], [17], [68], [83], division [87], and interactions between objects [85], see also Subsection V-H.

2) *Multi-Object Observation*: At time k , an *association map* $\theta: \mathbb{L}_k \rightarrow \{0: |Z_k|\}$ associates the labels of \mathbf{X}_k with the elements of Z_k , satisfying the *positive 1-1* property that *no two distinct arguments are mapped to the same positive value* [11]. Here, $\theta(\ell) > 0$ means ℓ generates detection $z_{\theta(\ell)} \in Z_k$, and $\theta(\ell) = 0$ means ℓ is misdetected. The positive 1-1 property ensures each detection comes from at most one object. The space of all such association maps is denoted as Θ_k .

In the *standard multi-object observation model*, each $x \in \mathbf{X}_k$, is either detected with probability $P_{D,k}(x)$ and generates a detection z with likelihood $g_k(z|x)$ or missed with probability $1 - P_{D,k}(x)$. Assuming that conditional on \mathbf{X}_k detections are independently generated, the RFS $D_k(\mathbf{X}_k)$ of detections is a multi-Bernoulli with parameters $\{(P_{D,k}(x), g_k(\cdot|x))\}_{x \in \mathbf{X}_k}$. Clutter K_k is modeled as a Poisson RFS with intensity κ_k , and assumed independent of the detections. The standard multi-object likelihood function is given by [11]

$$\mathbf{g}_k(Z_k|\mathbf{X}_k) \propto \sum_{\theta \in \Theta_k(\mathcal{L}(\mathbf{X}_k))} \left[\Psi_{Z_k}^{(\theta \circ \mathcal{L}(\cdot))}(\cdot) \right]^{\mathbf{X}_k}, \quad (60)$$

where $\Theta_k(I) \subseteq \Theta_k$ denotes the collection of association maps with domain I , and

$$\Psi_{k,Z}^{(j)}(\mathbf{x}) = \begin{cases} \frac{P_{D,k}(\mathbf{x})g_k(z_j|\mathbf{x})}{\kappa_k(z_j)}, & j \in \{1:|Z|\} \\ 1 - P_{D,k}(\mathbf{x}), & j = 0 \end{cases}. \quad (61)$$

Note that the likelihood function (60) is characterized by the ‘Signal to Noise Ratio’ (SNR) function $\Psi_{k,Z}^{(\theta \circ \mathcal{L})}$, and does not suffer from the non-symmetry problem discussed in Subsection II-C because *each of its terms is symmetric in the labeled states*. Hence, truncation before or after multiplication by a valid prior, still leaves it a valid set function.

The standard multi-object observation model accommodates non-homogeneous clutter and state-dependent detection probability, covering a broad range of problems. However, it does not address merged measurements [80], occlusions [86], extended measurements, superpositional and image measurements [16], [17], which require more sophisticated models.

3) *Multi-Sensor Multi-Object Observation*: In a multi-sensor setting with V sensors, each sensor registers an observation set $Z_k^{(v)}$, with (standard) multi-object likelihood $\mathbf{g}_k^{(v)}(Z_k^{(v)}|\mathbf{X}_k)$, $v \in \{1:V\}$. Assuming that $Z_k^{(1)}, \dots, Z_k^{(V)}$ are

independent conditional on \mathbf{X}_k , the multi-sensor multi-object likelihood function is given by the product [16], [17]

$$\mathbf{g}_k(Z_k^{(1)}, \dots, Z_k^{(V)} | \mathbf{X}_k) \triangleq \prod_{v=1}^V \mathbf{g}_k^{(v)}(Z_k^{(v)} | \mathbf{X}_k). \quad (62)$$

Let $\Psi_{k, Z^{(v)}}^{(v, j^{(v)})}$ denote the ‘SNR’ function of sensor v , and define the multi-sensor SNR function

$$\Psi_{k, (Z^{(1)}, \dots, Z^{(V)})}^{(j^{(1)}, \dots, j^{(V)})}(\mathbf{x}) \triangleq \prod_{v=1}^V \Psi_{k, Z^{(v)}}^{(v, j^{(v)})}(\mathbf{x}), \quad (63)$$

the multi-sensor observation $Z_k \triangleq (Z_k^{(1)}, \dots, Z_k^{(V)})$, the multi-sensor association map as the V -tuple $\theta \triangleq (\theta^{(1)}, \dots, \theta^{(V)})$ of association maps $\theta^{(v)} \in \Theta_k^{(v)}(I)$ from every sensor, and $\Theta_k(I) \triangleq \prod_{v=1}^V \Theta_k^{(v)}(I)$. Then, the multi-sensor multi-object observation likelihood can be written in the form (60).

4) *Relation with Unlabeled Models:* Historically, only the unlabeled multi-object transition density and observation likelihood were developed (for multi-object localization).

The unlabeled multi-object transition density f_k is given by the convolution (see (7) with $n = 2$)

$$f_k(X_k | X_{k-1}) = \sum_{W \subseteq X_k} f_{B,k}(W) f_{S,k}(X_k - W | X_{k-1}),$$

of the multi-object densities $f_{B,k}$ (of new born objects) and $f_{S,k}(\cdot | X_{k-1})$ (of objects generated from X_{k-1}) [16]. Without labeling, the sum over all subsets of X_k does not reduce to a single term like its labeled counterpart. In the standard multi-object dynamic model, $f_{S,k}(\cdot | X_{k-1})$ is a multi-Bernoulli of the form (9), which is (algebraically) more complex than the LMB (58). Further, the multi-Bernoulli model cannot ensure distinct states (in the transition to time k) necessary for the set representation of the multi-object state (see Remark 2). The LRFS formulation avoids such problems.

The standard unlabeled multi-object observation likelihood g_k takes on a nearly identical form to its labeled counterpart (60), specifically [16, pp. 421]

$$g_k(Z_k | \{x_{1:n}\}) \propto \sum_{\theta \in \Theta_k(\{1:n\})} \prod_{i=1}^n \Psi_{Z_k}^{(\theta(i))}(x_i),$$

where $\Theta_k(\{1:n\})$ denotes the space of association maps with domain $\{1:n\}$. The subtle, but important difference is that $g_k(Z_k | X)$ cannot be written as a sum of multi-object exponentials in X . Further, $g_k(Z_k | \cdot)$ (and hence the *unlabeled multi-object posterior* $\pi_k(\cdot | Z_k) \propto g_k(Z_k | \cdot) \pi_k(\cdot)$) is *not closed under truncation* as illustrated in Example 2.

Example 2. Consider a multi-object state $\{x_1, x_2\}$, with $Z_k = \{z_1\}$, $P_D = 0.5$, and $\kappa = 1$. The possible positive 1-1 mappings from $\{1, 2\}$ to $\{0, 1\}$ are: $[\theta(1), \theta(2)] = [1, 0]$ (x_1 detected and x_2 undetected); $[\theta(1), \theta(2)] = [0, 1]$ (x_1 undetected and x_2 detected); and $[\theta(1), \theta(2)] = [0, 0]$, (x_1 and x_2 undetected), hence $g_k(\{z_1\} | \{x_1, x_2\}) \propto g_k(z_1 | x_1) + g_k(z_1 | x_2) + 1$. Truncating $g_k(z_1 | x_2)$ yields $\hat{g}_k(\{z_1\} | x_1, x_2) = g_k(z_1 | x_1) + 1$, which is not a function of the set $\{x_1, x_2\}$, because $\hat{g}_k(\{z_1\} | x_2, x_1) = g_k(z_1 | x_2) + 1 \neq \hat{g}_k(\{z_1\} | x_1, x_2)$.

Note that the labeled representation avoids this problem because θ maps the label of each state to the detections, making each term invariant to the listing order of the states.

C. GLMB Filter

The crux of multi-object estimation lies in the solutions to the multi-object filtering/posterior recursions. This subsection presents a GLMB-based analytic solution to the multi-object Bayes filter (55), under the standard multi-object SSM model (inclusive of multiple sensors). For convenience, hereon, we omit references to the time index k , and denote $k \pm 1$, with subscripts ‘ \pm ’, e.g., $\mathbb{L}_- \triangleq \mathbb{L}_{k-1}$, $\mathbb{B} \triangleq \mathbb{B}_k$, $\mathbb{L} \triangleq \mathbb{L}_- \cup \mathbb{B}$. Also, when we write $\{P(\xi, I_-, \theta, I) : (\xi, I_-, \theta, I)\}$, it is understood that the variables ξ , I_- , θ , and I , respectively, range over the spaces Ξ , $\mathcal{F}(\mathbb{L}_-)$, Θ , and $\mathcal{F}(\mathbb{L})$, unless otherwise stated.

1) *Chapman-Kolmogorov Prediction:* Suppose that a multi-object state with density $\pi_- \triangleq \pi_{k-1}$, at time $k-1$, evolves to the current time k according to a multi-object transition density $\mathbf{f} \triangleq \mathbf{f}_k$. Then the predicted multi-object density $\pi \triangleq \pi_k$ is given by the Chapman-Kolmogorov equation

$$\pi(\mathbf{X}) = \int \mathbf{f}(\mathbf{X} | \mathbf{X}_-) \pi_-(\mathbf{X}_-) \delta \mathbf{X}_-, \quad (64)$$

which defines the *prediction operator* $\Pi : \pi_- \mapsto \pi$.

In [11], [26], it was established that the GLMB family is closed under the prediction operator Π with the standard multi-object transition density \mathbf{f} (i.e., model parameters w_B, p_B, P_S, f_S). Specifically, for the GLMB

$$\pi_- = \{(w_-^{(\xi)}(I_-), p_-^{(\xi)}): (\xi, I_-)\}, \quad (65)$$

the prediction density is the GLMB

$$\pi = \Pi(\pi_-) = \{(w^{(\xi)}(I), p^{(\xi)}): (\xi, I)\}, \quad (66)$$

where

$$w^{(\xi)}(I) = w_B(I \cap \mathbb{B}) w_S^{(\xi)}(I \cap \mathbb{L}_-), \quad (67)$$

$$p^{(\xi)}(x, \ell) = \mathbf{1}_{\mathbb{B}}(\ell) p_B(x, \ell) + \mathbf{1}_{\mathbb{L}_-}(\ell) p_S^{(\xi)}(x, \ell), \quad (68)$$

$$w_S^{(\xi)}(L) = [\bar{P}_S^{(\xi)}]^L \sum_{I \supseteq L} [1 - \bar{P}_S^{(\xi)}]^{L-L} w_-^{(\xi)}(I_-), \quad (69)$$

$$p_S^{(\xi)}(x, \ell) = \langle p_-^{(\xi)}(\cdot) f_S(x | \cdot) P_S(\cdot) \rangle(\ell) / \bar{P}_S^{(\xi)}(\ell), \quad (70)$$

$$\bar{P}_S^{(\xi)}(\ell) = \langle P_S p_-^{(\xi)} \rangle(\ell). \quad (71)$$

The prior GLMB components (ξ, I_-) generate the predicted components (ξ, I) . The predicted weight $w^{(\xi)}(I)$ of (ξ, I) is the product of the birth weight for new objects with labels in I , and the survival weight for old objects with labels in I . The predicted attribute density $p^{(\xi)}(\cdot, \ell)$ is either the birth attribute density for a new object with label ℓ , or the predicted attribute density for a surviving object with label ℓ .

2) *Bayes Update:* Suppose that a multi-object state with prior density π generates an observation Z according to a multi-object observation likelihood $\mathbf{g}(Z | \mathbf{X}) \triangleq \mathbf{g}_k(Z | \mathbf{X})$. Then, the multi-object posterior density is given by Bayes rule

$$\pi(\mathbf{X} | Z) = \frac{\mathbf{g}(Z | \mathbf{X}) \pi(\mathbf{X})}{\int \mathbf{g}(Z | \mathbf{Y}) \pi(\mathbf{Y}) \delta \mathbf{Y}}, \quad (72)$$

which defines the *Bayes update operator* $\Upsilon_Z : \pi \mapsto \pi(\cdot | Z)$.

In [11], [26], it was established that the GLMB family is closed under the update operator Υ_Z , i.e., a conjugate prior w.r.t. the standard $g(Z|\mathbf{X})$ (with SNR function $\bar{\Psi}_Z^{(\theta \circ \mathcal{L})}$ and model parameters κ, P_D, g). Specifically, if the prior π is the GLMB (66), then the posterior is the GLMB

$$\pi(\cdot|Z) = \Upsilon_Z(\pi) \propto \{(w_Z^{(\xi, \theta)}(I), p_Z^{(\xi, \theta)}): (\xi, \theta, I)\}, \quad (73)$$

where

$$w_Z^{(\xi, \theta)}(I) = w^{(\xi)}(I) \mathbf{1}_{\Theta(I)}(\theta) \left[\bar{\Psi}_Z^{(\xi, \theta(\cdot))}(\cdot) \right]^I, \quad (74)$$

$$p_Z^{(\xi, \theta)}(x, \ell) = p^{(\xi)}(x, \ell) \Psi_Z^{(\theta(\ell))}(x, \ell) / \bar{\Psi}_Z^{(\xi, \theta)}(\ell), \quad (75)$$

$$\bar{\Psi}_Z^{(\xi, \theta)}(\ell) = \langle p^{(\xi)} \Psi_Z^{(\theta)} \rangle(\ell). \quad (76)$$

Note that each prior GLMB component (ξ, I) generates a series of updated components (ξ, θ, I) . Due to the term $\mathbf{1}_{\Theta(I)}(\theta)$, only components with $\mathcal{D}(\theta) = I$ are needed. The weight $w_Z^{(\xi, \theta)}(I)$ of each (ξ, θ, I) is the product of the validity check for θ , the Bayes evidence from observation Z , and the prior weight $w^{(\xi)}(I)$. The corresponding attribute density $p_Z^{(\xi, \theta)}(\cdot, \ell)$ is simply the Bayes update of the prior attribute with detection $z_{\theta(\ell)} \in Z$ or a misdetection if $\theta(\ell) = 0$.

3) *GLMB Filtering Recursion*: The Bayes filtering recursion (55) is the composition of the prediction and update operators, i.e., $\Upsilon_Z \circ \Pi : \pi_- \mapsto \pi(\cdot|Z)$. Hence, for the standard multi-object SSM, the GLMB family is closed under the Bayes recursion, i.e., starting with an initial GLMB prior, the multi-object filtering density at any time is a GLMB [11], [26]. This result also holds with a GLMB birth model, but for simplicity in the subsequent discussion, we assume an LMB birth model. The direct propagation of the GLMB π_- in (65) to $\pi(\cdot|Z)$ in (73) is given by (see [95])

$$w_Z^{(\xi, \theta)}(I) \propto \sum_{I_- \subseteq \mathbb{L}} w_Z^{(\xi, I_-, \theta, I)} w_-^{(\xi)}(I_-), \quad (77)$$

$$p_Z^{(\xi, \theta)}(x, \ell) \propto \begin{cases} \langle \Lambda_S^{(\theta(\ell))}(x|\cdot) p_-^{(\xi)}(\cdot) \rangle(\ell), & \ell \in \mathbb{L}_- \\ \Lambda_B^{(\theta(\ell))}(x, \ell), & \ell \in \mathbb{B} \end{cases}, \quad (78)$$

where

$$w_Z^{(\xi, I_-, \theta, I)} = \mathbf{1}_{\mathcal{F}(I_- \cup \mathbb{B})}(I) \mathbf{1}_{\Theta(I)}(\theta) \prod_{\ell \in I_- \cup \mathbb{B}} \eta_Z^{(\xi, I, \ell)}(\theta(\ell)), \quad (79)$$

$$\eta_Z^{(\xi, I, \ell)}(j) = \begin{cases} 1 - \langle P_S p_-^{(\xi)} \rangle(\ell), & \ell \in \mathbb{L}_- - I \\ \int \langle \Lambda_S^{(j)}(x|\cdot) p_-^{(\xi)}(\cdot) \rangle(\ell) dx, & \ell \in \mathbb{L}_- \cap I \\ 1 - P_B(\ell), & \ell \in \mathbb{B} - I \\ \int \Lambda_B^{(j)}(x, \ell) dx, & \ell \in \mathbb{B} \cap I \end{cases}, \quad (80)$$

$$\Lambda_S^{(j)}(x|\zeta, \ell) = \Psi_Z^{(j)}(x, \ell) f_S(x|\zeta, \ell) P_S(\zeta, \ell), \quad (81)$$

$$\Lambda_B^{(j)}(x, \ell) = \Psi_Z^{(j)}(x, \ell) p_B(x, \ell) P_B(\ell). \quad (82)$$

A salient feature of the GLMB filter is the provision for smoothed (single-object) trajectory estimates [96], [97]. Suppose that all observations up to the current time follow the standard observation model (LRFS multi-object estimation accommodates updates with different types of observations). Then, in the GLMB filtering density propagation, say from $\{(w(J), p): J \in \mathcal{F}(\mathbb{L}_0)\}$ at time 0, to $\{(w^{(\xi)}(I_-), p^{(\xi)}): (\xi, I_-)\}$ at time $k-1$, we recursively constructs the ‘event’ ξ as the history $\theta_{1:k-1}$ of association maps. Thus, the attribute density

$p^{(\xi, \theta)}(\cdot, \ell) = p^{(\theta_{0:k})}(\cdot, \ell)$ contains the entire history $\theta_{0:k}(\ell)$ of detections associated with trajectory ℓ [11], [26]. This information, encapsulated in the multi-object filtering density, can be used to estimate the entire trajectory (or over a moving window) via smoothing [96], [97].

The GLMB recursion, defined by the above so-called *joint GLMB prediction and update*, is a true Bayesian MOT filter with provably Bayes-optimal track management. Since, each component (ξ, I_-) is propagated forward as a set of children components $\{(\xi, \theta, I): (\theta, I)\}$, the multi-object filtering density accumulates an intractably large number of components with time. This inevitably requires approximation, either by a simpler multi-object density, or by truncation.

D. LMB Filter

This subsection presents the LMB filter, regarded as the PHD filter for trajectories. In the same way that the PHD filter approximates the unlabeled multi-object filtering density by a Poisson with matching PHD, the LMB filter approximates the GLMB filtering density by an LMB with matching PHD. However, unlike the PHD filter, the LMB filter provides trajectory estimates, and does not suffer from high cardinality variance when the number of objects is large.

Using an LMB birth model, the LMB sub-family is also closed under the prediction operator Π . Specifically, given a previous LMB filtering density $\pi_- = \{(r_-(\ell), p_-(\cdot, \ell))\}_{\ell \in \mathbb{L}_-}$, the multi-object prediction density is the LMB [73]

$$\pi = \mathbf{f}_B \cup \{(r_S(\ell), p_S(\cdot, \ell))\}_{\ell \in \mathbb{L}_-}, \quad (83)$$

where \mathbf{f}_B is the birth LMB density/parameter,

$$r_S(\ell) = r_-(\ell) \langle P_S p_- \rangle(\ell), \quad (84)$$

$$p_S(x, \ell) = \frac{\langle f_S(x|\cdot) P_S(\cdot) p_- \rangle(\ell)}{\langle P_S p_- \rangle(\ell)}. \quad (85)$$

While the LMB prediction operation is exact and intuitively appealing, the LMB family is not closed under Bayes rule.

The LMB filter uses the above LMB prediction, and applies the update operator to the predicted LMB π , yielding the GLMB $\Upsilon_Z(\pi) = \{(w_Z^{(\xi)}(I), p_Z^{(\xi)}): (\xi, I)\}$. Further, using (23), (24), the GLMB $\Upsilon_Z(\pi)$ is approximated by the LMB

$$\pi(\cdot|Z) = \{(r_Z(\ell), p_Z(\cdot, \ell))\}_{\ell \in \mathbb{L}}, \quad (86)$$

with matching PHD, by setting

$$r_Z(\ell) = \sum_{\theta \in \Theta} \sum_{I \subseteq \mathbb{L}} \mathbf{1}_I(\ell) w_Z^{(\theta)}(I), \quad (87)$$

$$p_Z(y, \ell) \propto \sum_{\theta \in \Theta} p_Z^{(\theta)}(y, \ell) \sum_{I \subseteq \mathbb{L}} \mathbf{1}_I(\ell) w_Z^{(\theta)}(I). \quad (88)$$

Due to the smaller number of components than the GLMB filter, the LMB filter is faster, albeit with some degradation in tracking performance.

Remark 6. If the weight $w_Z^{(\theta)}(I)$ of every component $(w_Z^{(\theta)}(I), p_Z^{(\theta)})$ of the GLMB $\Upsilon_Z(\pi)$ can be written in the form $K_Z^{(\theta)} [r_Z^{(\theta)} / (1 - r_Z^{(\theta)})]^I$, then $\Delta(\mathbf{X}) w_Z^{(\theta)}(\mathcal{L}(\mathbf{X})) [p_Z^{(\theta)}]^{\mathbf{X}}$

is indeed an (unnormalized) LMB, see (20), and hence $\Upsilon_Z(\boldsymbol{\pi})$ is a mixture of LMBs. However, this is not the case because

$$w_Z^{(\theta)}(I) = \mathbf{1}_{\Theta(I)}(\theta) K_Z^{(\theta)} \left[r_Z^{(\theta)} / (1 - r_Z^{(\theta)}) \right]^I,$$

which cannot be expressed as a product over I due to the term $\mathbf{1}_{\Theta(I)}(\theta)$. Thus, LMB mixtures are not conjugate priors.

E. GLMB Filter Implementation

Implementing the GLMB filter requires truncation of the multi-object filtering density without exhaustive enumeration. Each component (ξ, I_-) of the GLMB filtering density at time $k-1$ propagates a (very large) set $\{(\xi, I_-, \theta, I) : (\theta, I)\}$ of “children” components (before marginalizing out I_-) to the current time. Truncation by selecting “children” with significant weights minimizes the L_1 -approximation error [26], and can be accomplished via solving the rank assignment problem or via Gibbs sampling (GS). Computing $w_Z^{(\xi, \theta)}(I)$ and $p_Z^{(\xi, \theta)}$ can be accomplished via single-object filtering techniques.

Given a fixed component (ξ, I_-) at time $k-1$, the goal is to find a set of $(\theta, I) \in \Theta \times \mathcal{F}(\mathbb{L})$ with significant weight increment $w_Z^{(\xi, I_-, \theta, I)}$. Due to the terms $\mathbf{1}_{\mathcal{F}(I \uplus \mathbb{B})}(I)$ and $\mathbf{1}_{\Theta(I)}(\theta)$ in (79), we only need to consider (θ, I) with $I \subseteq I_- \uplus \mathbb{B}$ and $\mathcal{D}(\theta) = I$. Hence, it is convenient to represent such (θ, I) by an *extended association map* (or simply extended association) $\gamma: I_- \uplus \mathbb{B} \rightarrow \{-1:|Z|\}$, defined by

$$\gamma(\ell) \triangleq \begin{cases} \theta(\ell), & \text{if } \ell \in \mathcal{D}(\theta) \\ -1, & \text{if } \ell \in (I_- \uplus \mathbb{B}) - \mathcal{D}(\theta) \end{cases}. \quad (89)$$

Note that γ inherits the *positive 1-1* property, and the set of all such γ is denoted by Γ . Since (θ, I) is recovered by $I = \{\ell : \gamma(\ell) \geq 0\}$, $\theta(\ell) = \gamma(\ell)$ for each $\ell \in I$, it follows that $\mathbf{1}_{\mathcal{F}(I \uplus \mathbb{B})}(I) = 1$ and $\mathcal{D}(\theta) = I$. Enumerating $I_- \uplus \mathbb{B} = \{\ell_{1:P}\}$, and $Z = \{z_{1:M}\}$, γ can be represented as a P -tuple in $\{-1:M\}^P$. Further, abbreviating

$$\eta^{(i)}(j) \triangleq \begin{cases} 1 - \langle p_-^{(\xi)} P_S \rangle(\ell_i), & \ell_i \in I_-, j < 0 \\ \int \langle \Lambda_S^{(j)}(x|\cdot) p_-^{(\xi)}(\cdot) \rangle(\ell_i) dx, & \ell_i \in I_-, j \geq 0 \\ 1 - P_B(\ell_i), & \ell_i \in \mathbb{B}, j < 0 \\ \int \Lambda_B^{(j)}(x, \ell_i) dx, & \ell_i \in \mathbb{B}, j \geq 0 \end{cases}, \quad (90)$$

(the dependence on ξ, I_- , and Z are suppressed) the weight increment (79) can be expressed as [95]

$$w_Z^{(\xi, I_-, \theta, I)} = \omega(\gamma) \triangleq \mathbf{1}_\Gamma(\gamma) \prod_{i=1}^P \eta^{(i)}(\gamma(\ell_i)). \quad (91)$$

The values of $\eta^{(i)}(j)$ is precomputed as the $P \times (M+2)$ *association score matrix* shown in Fig. 12, and the goal becomes finding a set of γ with significant $\omega(\gamma)$.

1) *Ranked Assignment*: The K best (extended) associations in non-increasing order of $\omega(\gamma)$ can be obtained without exhaustive enumeration by solving a *ranked assignment problem*. Here, each association $\gamma \in \Gamma$ is equivalently represented by a $P \times (M+2P)$ *assignment matrix* S whose entries are either 0 or 1, with every row summing to 1, and every column summing to either 1 or 0. The objective is to find the K assignment matrices with smallest $\text{trace}(S^T C)$, where C is the




Fig. 12: Association score matrix and conditionals (at time k). The i -th unnormalized conditional $\varpi^{(i)}$ is determined by taking the i -th row of the association score matrix, i.e., $\eta^{(i)}(j)$, $j = -1 : M$, and set $\eta^{(i)}(j) = 0$ for each positive j associated with a label other than ℓ_i .

$P \times (M+2P)$ *cost matrix* constructed from the association score matrix in Fig. 12 by [95]

$$C_{i,j} = \begin{cases} -\ln \eta^{(i)}(j), & j \in \{1 : M\} \\ -\ln \eta^{(i)}(0), & j = M+i \\ -\ln \eta^{(i)}(-1), & j = M+P+i \\ \infty, & \text{otherwise} \end{cases}. \quad (92)$$

Since $\exp(-\text{trace}(S^T C)) = \prod_{i=1}^P \eta^{(i)}(\gamma_i) = \omega(\gamma)$, the K best assignment matrices correspond to the K associations with largest weights [95]. A GLMB filter implementation with $\mathcal{O}(K(M+2P)^4)$ complexity was proposed in [98] using Murty’s algorithm [99] to solve the ranked assignment problems. More efficient algorithms [100], [101] can reduce the complexity to $\mathcal{O}(K(M+2P)^3)$.

The initial GLMB filter implementation truncates the predicted and updated multi-object densities at each time [11], [26]. Since truncation of the prediction is separated from the update, information from the observation is not exploited, and a significant portion of the predicted components subsequently generate updated components with negligible weights, which waste computations. The joint prediction and update avoids this problem, thus improving computational speed [95]. For the LMB special case, belief propagation can also be exploited for a fast implementation [102]. A drawback of using ranked assignment is the high computational cost of generating a sequence of components ordered by their weights, whilst such ordering is not needed in the GLMB approximation.

2) *Gibbs Sampling*: An efficient way to generate significant GLMB components is to sample from some probability distribution ϖ on $\{-1:M\}^P$ [95]. To ensure that mostly high-weight positive 1-1 P -tuples are sampled, ϖ is constructed so that only positive 1-1 P -tuples have positive probabilities, and those with high weights are more likely to be chosen than those with low weights. An obvious choice of ϖ is (91).

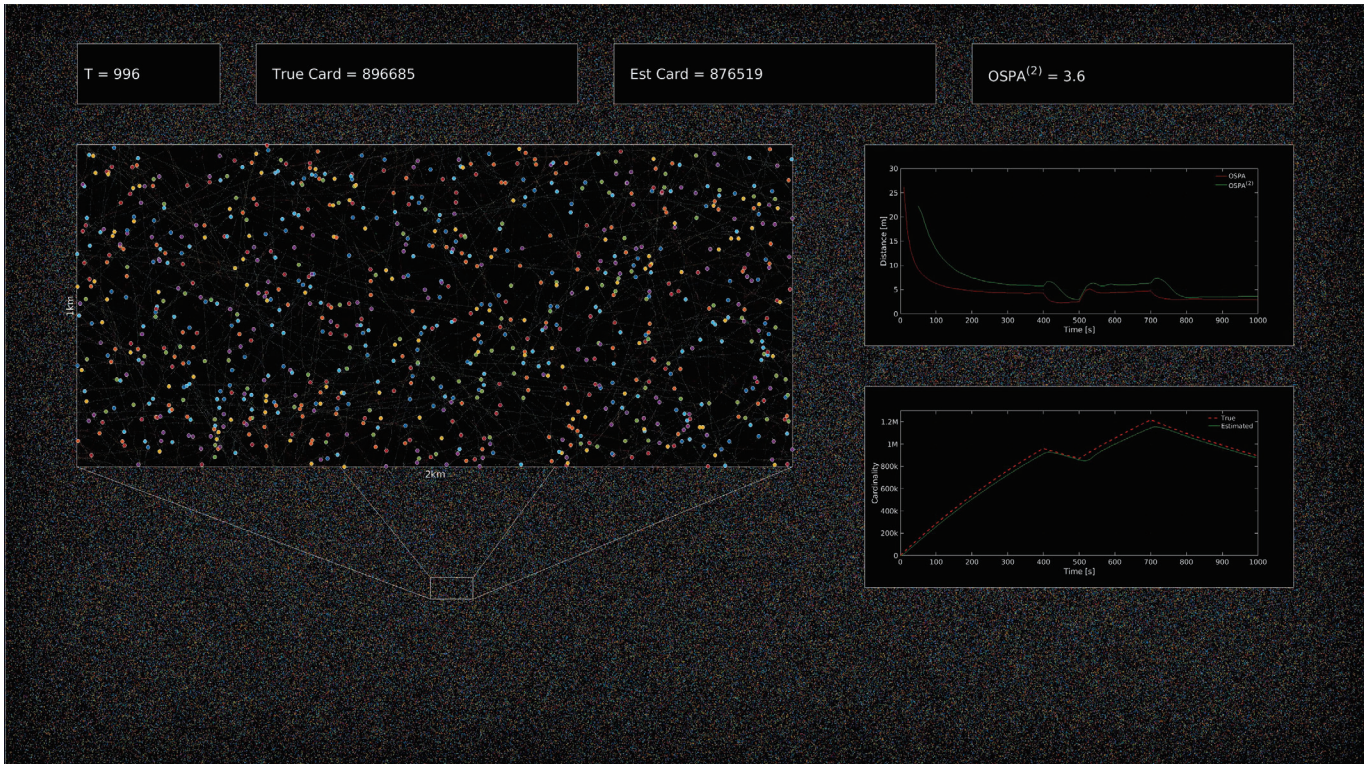


Fig. 13: Large-scale GLMB filter tracking over 1 million objects in a $64\text{km} \times 36\text{km}$ region [27]. The insets show a magnified $2\text{km} \times 1\text{km}$ region, OSPA/OSPA⁽²⁾ errors, and cardinality. Objects can appear anywhere, at a rate of 200-3000 per frame (unknown to filter). The object's position and velocity follow a linear Gaussian motion model with 0.2ms^{-2} process noise standard deviation, and 0.88 detection probability. Position observations are corrupted by Gaussian noise with a 5m standard deviation, and clutter, uniformly distributed on the region at an average of 460800 points. About 1 billion data points accumulate over 1000 frames, each contains about 1 million objects on average, peaking to 1,217,531 at frame 700.

Gibbs Sampling (GS) is an efficient special case of the Metropolis-Hasting algorithm for sampling from an unnormalized distribution ϖ [103], [104]. This algorithm constructs a Markov chain that sequentially generates a new iterate γ' from the current iterate γ . When the chain runs for long enough, i.e., past the burn-in stage, subsequent samples would be distributed according to the target distribution ϖ . Unlike sampling for posterior inference, in GLMB filtering it is not necessary to discard burn-ins because all distinct positive 1-1 P -tuples will reduce the L_1 -approximation error.

The GLMB truncation developed in [95] employs the classical Systematic-scan GS (SGS) strategy, wherein the next iterate γ' is generated from the current iterate γ by sampling each coordinate $\gamma'(\ell_i)$ of γ' , for $i = 1:P$, from the so-called the i -th conditional $\varpi^{(i)}(\cdot | \gamma'(\ell_{1:i-1}), \gamma(\ell_{i+1:P}))$ on $\{-1:M\}$ [103], [104]. It was shown in [95] that to arrive at the target distribution (91), the required unnormalized i -th conditional is given by the i -th row of the association score matrix after zeroing the entry of every positive j that has been paired with a label other than ℓ_i , see Fig. 12 for illustration. More concisely, $\varpi^{(i)}(j | \gamma'(\ell_{1:i-1}), \gamma(\ell_{i+1:P})) \propto \eta^{(i)}(j)$, except at each positive j in the set $\{\gamma'(\ell_{1:i-1}), \gamma(\ell_{i+1:P})\}$ of values associated with labels other than ℓ_i , wherein $\varpi^{(i)}(j | \gamma'(\ell_{1:i-1}), \gamma(\ell_{i+1:P})) = 0$. This remarkably simple result provides an inexpensive way to generate significant positive 1-1 tuples, and was exploited to

implement SGS GLMB truncation in [95] with an $\mathcal{O}(TP^2M)$ complexity, but later reduced to $\mathcal{O}(TPM)$ in [105], where T is the number of iterates of the Gibbs sampler. Note that starting with any positive 1-1 tuple, e.g., all 0's or all -1's, all subsequent iterates are positive 1-1.

In [96], SGS was extended to address the NP-hard multi-sensor GLMB truncation problem with linear complexity in the total number of detections across all sensors. This SGS multi-sensor GLMB filter implementation also applies to approximations such as the LMB and marginalized GLMB filters since these filters require full GLMB updates [73], [106]. Such a multi-sensor LMB filter implementation has been extended to address partially overlapping fields of views in [107].

Following the development of SGS GLMB truncation, other sampling-based techniques have been proposed. In [108], the positive 1-1 requirement was neglected to achieve linear complexity in P , albeit at a degradation in performance. A herded GS implementation of the LMB filter was proposed in [109], which turned out to be slower than the SGS implementation. Spatial search and approximation by a product of smaller GLMBs (see Subsection IV-F) that run in parallel were proposed in [27] to reduce computation times. While this approach was demonstrated to track in excess of one million objects from approximately 1.3 million detections at a time, see Fig. 13, the complexity has not been reduced.

Cross-entropy based solutions [110] were developed in [111] and [112], for multi-sensor GLMB filtering and its distributed version, but with higher complexity than the SGS implementation in [96]. The recently developed tempered GS (TGS) technique reduces the complexity to $\mathcal{O}(T(P+M))$ [105].

F. Multi-Scan GLMB Filter/Smother

Similar to the (multi-object) Bayes filtering recursion, the posterior recursion (54) admits an analytic solution in the form of a multi-scan GLMB. This subsection presents the multi-scan GLMB recursion under the standard multi-object SSM model (with LMB birth for simplicity). For convenience, when we write $\{P(\xi, \theta, I_{j:k}) : (\xi, \theta, I_{j:k})\}$, it is understood that the variables ξ , θ , and $I_{j:k}$, respectively, range over the spaces Ξ , Θ , and $\prod_{i=j}^k \mathcal{F}(\mathbb{I}_i)$, unless otherwise stated.

The multi-scan GLMB is closed under the Bayes posterior recursion (54). Indeed, it is closed under: the prediction, defined as $\pi_{0:n}(\mathbf{X}_{0:n}) = \mathbf{f}_n(\mathbf{X}_n | \mathbf{X}_{n-1}) \pi_{0:n-1}(\mathbf{X}_{0:n-1})$; and the Bayes update since it is a GLMB where the argument is a set of labeled trajectories [91]. Thus, starting with an initial multi-scan GLMB prior, the multi-object prediction and posterior at any time are multi-scan GLMBs.

The multi-scan GLMB recursion (or GLMB smoothing recursion) [91, eq. (42)-(46)] propagates the multi-scan GLMB

$$\pi_{0:n-1} = \{(w_{-}^{(\xi)}(I_{0:n-1}), p_{-}^{(\xi)}) : (\xi, I_{0:n-1})\}, \quad (93)$$

at time $n-1$ to the multi-scan GLMB

$$\pi_{0:n}(\cdot | Z) \propto \{(w_{Z}^{(\xi, \theta)}(I_{0:n}), p_{Z}^{(\xi, \theta)}) : (\xi, \theta, I_{0:n})\}, \quad (94)$$

at time n , where

$$w_{Z}^{(\xi, \theta)}(I_{0:n}) = w_{Z}^{(\xi, I_{n-1}, \theta, I_n)} w_{-}^{(\xi)}(I_{0:n-1}), \quad (95)$$

$$p_{Z}^{(\xi, \theta)}(x_{s(\ell):t(\ell)}, \ell) \propto \quad (96)$$

$$\begin{cases} p_{-}^{(\xi)}(x_{s(\ell):t(\ell)}, \ell), & t(\ell) < n-1 \\ (1-P_S(x_{t(\ell)}, \ell)) p_{-}^{(\xi)}(x_{s(\ell):t(\ell)}, \ell), & t(\ell) = n-1 \\ \Lambda_S^{(\theta(\ell))}(x_n | x_{n-1}, \ell) p_{-}^{(\xi)}(x_{s(\ell):n-1}, \ell), & s(\ell) < t(\ell) = n \\ \Lambda_B^{(\theta(\ell))}(x_n, \ell), & s(\ell) = t(\ell) = n \end{cases}$$

$w_{Z}^{(\xi, I_{n-1}, \theta, I_n)}$ is the weight increment given by (79), $s(\ell)$ and $t(\ell)$ are, respectively, the earliest and latest times that ℓ exists on the time window $\{0:n\}$, $\Lambda_S^{(j)}(x_n | x_{n-1}, \ell)$ and $\Lambda_B^{(j)}(x_n, \ell)$ are, respectively, given by (81) and (82).

Observe the similarity between the GLMB filter (77)-(82) and the multi-scan GLMB filter. Indeed, the latter is (algebraically) simpler and more intuitive than the former since no marginalization is needed. The new posterior GLMB weights are simply the old weights scaled by their corresponding weight increments. Further, given the old posterior density $p_{-}^{(\xi)}(\cdot, \ell)$ of trajectory ℓ at time $n-1$: if ℓ died before time $n-1$, then its new posterior is $p_{-}^{(\xi)}(\cdot, \ell)$; if ℓ died at time $n-1$, then its new posterior is $p_{-}^{(\xi)}(\cdot, \ell)$ times the death probability; if ℓ (born before time n) survives to time n , then its new posterior is $p_{-}^{(\xi)}(\cdot, \ell)$ times the survival probability, transition density, and the SNR; if ℓ is born at time n , then its new posterior is the probability of birth times the attribute density of its new state, and the SNR.

Similar to the GLMB filter, the number of components of the GLMB posterior grows super-exponentially with time, and truncation by retaining a prescribed number of components with highest weights minimizes the L_1 -norm approximation error [91]. Unlike the GLMB filter, the multi-scan GLMB truncation problem requires solving large-scale *multi-dimensional ranked assignment problems*, which is NP-hard for more than two dimensions.

G. Multi-Scan GLMB Filter Implementation

An SGS multi-scan GLMB truncation technique has been developed for the multi-scan GLMB filter implementation in [91]. Similar to the single-scan counter-part, we only need to consider (θ, I_k) represented by the (extended) association $\gamma_k : I_{k-1} \uplus \mathbb{B}_k \rightarrow \{-1 : |Z_k|\}$ defined in (89). Note that the domain $\mathcal{D}(\theta)$ is given by $\mathcal{L}(\gamma_k) \triangleq \{\ell \in I_{k-1} \uplus \mathbb{B}_k : \gamma_k(\ell) \geq 0\}$, called the *live labels* of γ_k . Assuming all observations follow the standard model, (θ, I_n) and $(\xi, I_{0:n-1})$ can be equivalently represented by γ_n and $\gamma_{0:n-1}$. Further, for each $k \in \{0:n\}$, enumerating $\mathcal{L}(\gamma_{k-1}) \uplus \mathbb{B}_k = \{\ell_{1:P_k}\}$ and $Z_k = \{z_{1:M_k}\}$, we can represent γ_k as a P_k -tuple in $\{-1:M_k\}^{P_k}$.

Denoting the weight $w_{Z}^{(\xi, \theta)}(I_{0:n})$ by $\omega_{0:n}(\gamma_{0:n})$ to convey the dependence on $\gamma_{0:n}$, and applying the weight propagation (95) iteratively to an initial multi-scan GLMB $\pi_0 = \{(\omega_0(\gamma_0), p_0^{(\gamma_0)}) : \gamma_0 \in \Gamma_0\}$ yields

$$\omega_{0:n}(\gamma_{0:n}) = \prod_{k=1}^n \left[\mathbf{1}_{\Gamma_k}^{(\gamma_{k-1})}(\gamma_k) \prod_{i=1}^{P_k} \eta_k^{(i)}(\gamma_k(\ell_i)) \right] \omega_0(\gamma_0), \quad (97)$$

where $\mathbf{1}_{\Gamma_k}^{(\gamma_{k-1})}(\gamma_k) \triangleq \mathbf{1}_{\Gamma_k}(\gamma_k) \mathbf{1}_{\mathcal{F}(\mathcal{L}(\gamma_{k-1}) \uplus \mathbb{B}_k)}(\mathcal{L}(\gamma_k))$, and $\eta_k^{(i)}(j)$ is defined in (90) with the time k suppressed (keeping in mind its implicit dependence on $\xi = \gamma_{0:k-1}(\ell_i)$). Following the GS strategy, the objective is to generate significant association histories by sampling from some discrete probability distribution ϖ on $\prod_{k=0}^n \{-1:M_k\}^{P_k}$, preferably $\varpi = \omega_{0:n}$, so that only valid association histories have positive probabilities, and those with high weights are more likely to be chosen.

1) *Sequential Factor Sampling*: Decomposing ϖ as

$$\varpi(\gamma_{0:n}) = \prod_{k=1}^n \varpi_k(\gamma_k | \gamma_{0:k-1}) \varpi_0(\gamma_0), \quad (98)$$

and choosing the factors $\varpi_0 = \omega_0$,

$$\varpi_k(\gamma_k | \gamma_{0:k-1}) \propto \mathbf{1}_{\Gamma_k}^{(\gamma_{k-1})}(\gamma_k) \prod_{i=1}^{P_k} \eta_k^{(i)}(\gamma_k(\ell_i)),$$

for $k = 1:n$, yields $\varpi = \omega_{0:n}$. Hence, a simple method to sample from (97) is to sample γ_0 from ϖ_0 , and then for $k = 1:n$, sample $\gamma_k : \mathcal{L}(\gamma_{k-1}) \uplus \mathbb{B}_k \rightarrow \{-1 : |Z_k|\}$ from $\varpi_k(\cdot | \gamma_{0:k-1})$. This sequential generation of $\gamma_{0:n}$ ensures $\mathbf{1}_{\mathcal{F}(\mathcal{L}(\gamma_{k-1}) \uplus \mathbb{B}_k)}(\mathcal{L}(\gamma_k)) = 1$, and hence

$$\varpi_k(\gamma_k | \gamma_{0:k-1}) \propto \mathbf{1}_{\Gamma_k}(\gamma_k) \prod_{i=1}^{P_k} \eta_k^{(i)}(\gamma_k(\ell_i)), \quad (99)$$

for each $k \in \{1:n\}$. Sampling from (99) can be accomplished using the SGS technique described in Subsection V-E2. While any $\gamma_{0:n}$ generated by this method is a valid association history




Fig. 14: Multi-scan association score matrix and conditionals at time k . The i -th unnormalized conditional $\varpi_k^{(i)}$ is determined by taking the i -th row of the k -th association score matrix, i.e., $\eta_{\{k:t(\ell_i)\}}^{(i)}(j)$, $j = -1:M_k$, and set $\eta_{\{k:t(\ell_i)\}}^{(i)}(j) = 0$ for: each positive j associated with a label other than ℓ_i ; and any negative j , if ℓ_i is a live label at time $k + 1$.

[91], to ensure that $\gamma_{0:n}$ is distributed according to (98), it is necessary to run the Gibbs sampler for each k long enough so that γ_k is distributed according to (99). Nonetheless, sequential factor sampling can be used to generate good starting points for the Markov chains in full GS.

2) *SGS*: Sampling from ϖ via SGS involves constructing a Markov chain where a new iterate $\gamma'_{1:n}$ is generated from $\gamma_{1:n}$ by sampling the coordinates $\gamma'_k(\ell_i)$, $k = 1:n$, $i = 1:P_k$ of $\gamma'_{1:n}$, from the conditional distributions $\varpi_k^{(i)}$ defined by

$$\varpi_k^{(i)}(j | \underbrace{\gamma'_{0:k-1}}_{\text{past}}, \underbrace{\gamma'_k(\ell_{1:i-1})}_{\text{current (processed)}}, \underbrace{\gamma_k(\ell_{i+1:P_k})}_{\text{current (unprocessed)}}, \underbrace{\gamma_{k+1:n}}_{\text{future}}) \propto \varpi(\gamma'_{0:k-1}, \gamma'_k(\ell_{1:i-1}), j, \gamma_k(\ell_{i+1:P_k}), \gamma_{k+1:n}). \quad (100)$$

Similar to the single-scan case, the conditional $\varpi_k^{(i)}$ is determined from the association score matrices at times k and $k+1$, see Fig. 14, which can be pre-computed as follows.

Recall from (90) that $\eta_k^{(i)}(j_k)$ depends on ξ , specifically, the indices $j_{s(\ell_i):k-1}$ of the detections associated with ℓ_i up to time $k-1$ (since all observations follow the standard model). To express this dependence explicitly we write $\eta_k^{(i)}(j_k)$ as $\eta_k^{(i)}(j_{s(\ell_i):k})$. The (i, j) -th entry of the multi-scan association score matrix for time k is given by

$$\eta_{\{k:t(\ell_i)\}}^{(i)}(j) \triangleq \prod_{m=k}^{t(\ell_i)} \eta_m^{(i)}(\gamma'_{0:k-1}(\ell_i), j, \gamma_{k+1:m}(\ell_i)),$$

where $t(\ell_i)$ is the latest time ℓ_i exists on $\{0:n\}$, and by convention, $\eta_k^{(i)}(\gamma'_{0:k-1}(\ell_i), j, \gamma_{k+1:k}(\ell_i)) = \eta_k^{(i)}(\gamma'_{0:k-1}(\ell_i), j)$.

It was shown in [91] that to arrive at the target distribution (97), the unnormalized conditional $\varpi_k^{(i)}$ is given by the i -th row of the k -th association score matrix after zeroing the entry of: every positive j that has been paired with a label other than ℓ_i at time k (like the single-scan case); and any negative $j < \gamma_{k+1}(\ell_i)$ (a surviving label at time $k+1$ must be live at time k , i.e., $\gamma'_k(\ell_i) > -1$, because the standard multi-object dynamic model only admits unfragmented trajectories), see Fig. 14 for illustration. This simple result admits a tractable SGS multi-scan GLMB truncation algorithm for the GLMB smoother implementation in [91].

This multi-scan GLMB truncation technique has been extended to multiple sensors in [113], and demonstrated on a multi-object smoothing problem with 100 scans and 4 sensors, which requires solving 400-dimensional ranked assignment problems with approximately 10 variables in each dimension. Moving window-based implementations can achieve a fixed $\mathcal{O}(LTVP^2M)$ complexity per time step using SGS [113], or $\mathcal{O}(LTV(P+M))$ using TGS, where L is the length of the smoothing window, T is the number of iterates of the Gibbs sampler, V is the number of sensors, P is the number of hypothesized trajectories, and M is the maximum number of detections per sensor.

H. Multi-Object Estimation with Non-Standard Models

The standard multi-object dynamic model assumes conditionally independent individual transitions, and is sufficient to cover sophisticated single-object motion, e.g., motion with road constraints [114], or multi-modal motion [115]–[120]. However, it is not adequate to capture inter-object correlations, which requires more general multi-object transition, e.g., interacting objects [85], [121], group targets [122]–[124], and spawning/dividing objects [83], [87]. In the latter, we also seek the ancestries of the objects, which is easily accomplished in the LRFS formulation by augmenting the parents' labels into the object's labels, see Fig. 15. Consequently, the multi-object filtering density is no longer a GLMB, and the exact solution is intractable. M-GLMB approximations (see Subsection IV-F) are used to capture inter-object correlations and ancestries in the multi-object trajectory [83], [87]. An approximate GLMB filter that accommodates objects exiting the state space and reappearing later (along with a non-standard observation model) was developed in [125].

While the standard multi-object observation likelihood is general enough to accommodate a range of noise models, e.g., [126]–[130], it does not cover extended object observations [131], [132], merged/occluded observations [80], [86], and image observations [81], [84], [133], and [134].

Under the extended observation model that allows an object to generate multiple detections, the GLMB filter and smoother




Fig. 15: Labels capture ancestries of object in spawnings. A daughter’s label (ℓ, k, ι) consists of the parent label ℓ , the time of birth k , and an individual identity ι to distinguish it amongst the siblings. At time 4, ℓ_1 spawns 3 daughters with labels $\ell_2 = (\ell_1, 4, 1)$, $\ell_3 = (\ell_1, 4, 2)$, and $\ell_4 = (\ell_1, 4, 3)$. At time 7, ℓ_3 spawns $\ell_5 = (\ell_3, 7, 1)$ and $\ell_6 = (\ell_3, 7, 2)$.

are still exact solutions [131], [135], which can be implemented with Gaussians (and Gaussian mixture) or particle methods [136]. Extended objects are modeled by Gaussian-Gamma-Inverse-Wishart distributions in [131], B-splines in [137], and multiple ellipses in [132].

The standard observation model cannot accommodate multiple objects sharing detections due to occlusions. In this case, the detection probability of an object depends on the entire multi-object state, and requires a non-standard observation model [80], [86], [125]. The resultant multi-object filtering density is not a GLMB, and M-GLMB or LMB approximations have been used in [80], [138]. In [86], [125], efficient GLMB approximations based on the predicted states, were used to update the individual detection probabilities to account for occlusions, while in [139] the dependence of the detection probability on other states are marginalized out.

In applications where the multi-object state generates a single observation, inter-object correlations are introduced through the update into the multi-object filtering density which no longer takes on a GLMB form. The particle implementations for a general observation model proposed in [81], [133] are computationally demanding, and the M-GLMB approximation has been developed in [70]. Efficient approximations exploiting the additivity of superpositional observations are also proposed in [84], [134], [140]. This approach was combined with pseudo-smoothing to address acoustic vector sensor observations in [141]. A parallelizable LMB approximation is developed in [82], while GLMB filtering and smoothing approximations are proposed in [142], [143], [144] (the former also covers extended object observations). An LMB filter for pre-clustered laser range finder image data is proposed in [145]. A salient advantage of LRFS in MOT over legacy approaches such as MHT, is the seamless operation with different observation types such as detections and images. Indeed, a GLMB filter has been applied to an observation model that switches between detections and image observation in [146]. Non-standard models have also been developed for applications with unknown model parameters. These are included in the discussion on robust multi-object estimation.

I. Robust Multi-Object Estimation

Multi-object filtering solutions have also been developed to address unknown multi-object model parameters. In [147], the unknown detection probability is jointly estimated by augmenting it to the attribute state, while the unknown clutter rate is addressed by treating clutter as a different class of objects [148]. This formulation results in a GLMB filter with jump Markov single-object dynamic, which incurs extra computational complexity [147]. A more efficient suboptimal alternative is the bootstrapping approach that uses robust CPHD or multi-Bernoulli filters [149], [150] to estimate the detection probability and clutter rate from the observations, and feed them to the GLMB filter [151], [152], [153]. In [154]–[157], unknown observation noise parameters is addressed by augmenting the object attributes with noise covariance matrices to be jointly estimated by integrating Variational Bayes techniques to the multi-object filter.

Unknown birth model parameters can be addressed by the measurement-driven birth approach [73], where current detections are used to predict the birth parameters at the next time. GLMB filtering with measurement-driven birth is popular for MOT with unknown birth models [158]–[160]. In [161], a measurement-driven birth model for interval observations was developed, while in [162], the Rauch-Tung-Striebel smoother was used to improve the multi-object birth density estimation. An admissible region birth model for space objects was proposed in [163]. In [164], birth parameter estimation was developed for multi-sensor GLMB filtering, where GS is used to select probable detection combinations. A fully robust GLMB filter that combines unknown birth and sensor parameters estimation techniques were developed in [165]. Around the same time an interesting formulation using outer probability measure, which allows GLMB filtering in the absence of model parameters was developed in [166].

J. Distributed Multi-Object Estimation

Encapsulating all information on the multi-object trajectories in the filtering/posterior densities allows density fusion. In [89], the Generalized Covariance Intersection (GCI) fusion rule [167], was applied to fuse the local labeled multi-object densities $\pi^{(s)}$, $s \in \{1:S\}$ into

$$\bar{\pi}_{GCI} = \arg \min_{\pi} \sum_{s \in S} \omega^{(s)} D_{KL}(\pi || \pi^{(s)}),$$

where each $\omega^{(s)}$ is a user-defined fusion weight. It was shown that the M-GLMBs or LMBs, are closed under GCI fusion [89]. For LMB local densities, a technique for determining the weights $\omega^{(s)}$ of each of their component was proposed in [168], using the information gain at the local update step. The fusion weights can also be optimized to fuse GLMBs from sensor nodes with different field of view (FoV) [169]. In [75], the event-trigger strategy was proposed to reduce the communication burden in GCI fusion.

GCI fusion is sensitive to mismatches in the labels due to uncertainty in the object birth times [170]. To avoid this mismatch, GCI was used to fuse multi-Bernoulli approximations of the unlabeled versions of the local labeled multi-object densities, and then augment the labels to the fused multi-object

density [170]. In [171], the label mismatch was quantified by the minimum KL-D between the unlabeled versions of the local labeled densities. Further, it was shown that for local LMB densities, minimizing the label mismatch translates to solving a linear assignment problem. For sensors with different FoVs, it was proposed in [172] to maintain a list of label matches and extend the label space of each node to include labels of objects that cannot be observed by the local sensor. In [173], an additional group variable was introduced to record matching information among sensors, and fusion was performed separately on the surviving and newborn objects.

Exchanging the arguments of the KL-D in GCI fusion results in Minimum Information Loss (MIL) fusion [174]:

$$\bar{\pi}_{MIL} = \arg \min_{\pi} \sum_{s \in S} \omega^{(s)} D_{KL}(\pi^{(s)} || \pi).$$

Unlike GCI fusion, MIL fusion preserves both common and exclusive information of the local densities. MIL fusion formulae for M-GLMBs and LMBs that restrict the fused densities to the same family as the local densities were derived in [174]. The label matching strategy in [171] can be used to ensure all local densities have the same label space. For sensors with different FoVs, the local densities can be decomposed into sub-densities with minimum KL-D from the original ones, and then fused together using MIL [175]. GCI and MIL fusion were combined to exploit their respective advantages in [176], [177], and to enhance resilient to cyber-attack in peer-to-peer sensor network in [178]. GCI and MIL fusion are considered as geometric averaging (of the multi-object densities) approaches. Fusion via arithmetic averaging has also been proposed as a versatile alternative to geometric averaging in [179], [180].

Apart from the KL-D based fusion discussed above, there are also other fusion solutions. In a centralized setting, local densities can be approximated by Poissons with matching first moments, which are then fused together using CS-D based fusion [181]. In [182], the authors proposed to fuse the LMB components (from different nodes) corresponding to a particular label into a single component. In [183], Bayes parallel combination rule was used to compute the global posterior density from synchronous centralized sensors. In [112], instead of fusing the local GLMBs, distributed cross-entropy and average consensus were used to sample multi-sensor assignments with high scores. An approach that directly fuses the local multi-object estimates was proposed in [184], which incurs far lower complexity than density fusion and can be efficiently extended to multiple scans. Further, since only the estimates are fused, this fusion algorithm is not restricted to any tracking algorithms or approaches.

K. Multi-Object Control

In multi-object control, we seek control signals to drive the multi-object state/trajectory by optimizing the cost/reward function, subject to a set of constraints. Cost functions are usually based on information divergence. Due to its closed-form for the GLMB filtering density, the CS-D was used as the cost function for a sensor control problem and void probability constraints in [72], [185]. Moreover, a sensor management framework for the GLMB filter based on the CS-D was

developed in [186]. A passive sensor management solution was also proposed for the GLMB filter in [187]. While the CS-D can be computed analytically for GLMBs, it is still an expensive task. A cheaper alternative was proposed in [188] by using the LMB filter with a cost function based on the CS-D between the Poisson approximations of the LMBs, which can be computed efficiently [78] (see also Subsection IV-E). In [189], Rényi-D was also suggested as the reward function with the GLMB filter. The simple closed-form information divergence for LMBs in Subsection IV-E could be used to accelerate computation of information-based reward functions.

For multi-sensor tasking, a Rényi-D reward was used in [190], while in [191], a task-driven reward function with the LMB filter based on GCI fusion was proposed. In [192], maximizing the task-driven reward was replaced by minimizing the ratio of non-existence to existence probability for a subset of objects of interest in the LMB. A solution to this problem with CS-D as the reward function was given in [193]. Methods to control drones based on Rényi-D and CS-D were proposed in [50]. Inspired by the multi-objective formulation for the multi-object localization problem [194], in [76], a multi-objective reward function for the two conflicting tasks of tracking and discovery with limited FoV sensors was proposed. Using mutual information and differential entropy objectives, it was also established that the multi-objective reward is monotone and sub-modular, which allows the optimal control action to be efficiently computed via greedy algorithms. Metric based reward functions have also been proposed. In [195] the dispersion of the labeled multi-object density about its statistical mean using the OSPA distance was used to maximize the accuracy of the estimates. Similarly, techniques for multi-sensor control/selection that optimize a metric-based objective function were also proposed in [196], [197]. Further, simultaneous localization and mapping can also be performed with the LRFS filters [198], [199].

VI. CONCLUSIONS

The notion of a State-Space Model (SSM) has been the cornerstone for modern estimation and control theory, in which the celebrated Kalman filter is often recognized as one of the seminal contributions. A multi-object SSM generalizes the state vector to a finite set of vectors for modeling multi-object systems, wherein both the number of objects and their individual states are unknown and dynamically varying. Moreover, to retain the fundamental premise that the system trajectory is given by the history of the system states, it is necessary to use the labeled multi-object representation of the system state. Analogous to how information on a single trajectory is encapsulated in the filtering/posterior density of the state vectors, a Labeled Random Finite Set (LRFS) enables information on the multi-object trajectory to be encapsulated in the multi-object filtering/posterior density.

This article has provided an overview of key developments in the LRFS approach to multi-object SSM. LRFS resolves the theoretical drawbacks of the unlabeled RFS formulation for multi-object SSMs, and offers important advancements. In particular, the LRFS approach:

- Enables principled multi-object trajectory estimation with bounded complexity per time step (without resorting to post-processing heuristics), resolving an important conceptual shortcoming of unlabeled RFS;
- Provides statistical characterization of uncertainty for the underlying multi-object trajectory ensemble, offering new conceptual tools and applications;
- Naturally covers trajectory crossings, fragmentations, and lineages (in spawning objects), as well as seamless operations with different types of observations simultaneously or at different times, thereby resolving the theoretical problem concerning multiple objects occupying the same attribute state and further broadening the application base;
- Admits the concept of joint existence probability for the elements of a multi-object state, and hence, a meaningful notion of the most probable multi-object state(s), fundamental to multi-object estimation;
- Enables principled approximations of multi-object densities (with characterizable approximation errors) and reduces the standard multi-object transition density from a combinatorial sum to a single term, offering new analytical tools as well as facilitating efficient solutions, which are scalable in the numbers of objects, detections, sensors and scans.

Multi-object SSM is not only rich in theory, but also in mathematical and computational tools. While the majority of research in the field has focused on multi-object estimation, the related areas of system identification and control remain largely unexplored. Within the estimation problem there are nonetheless many open theoretical and computational challenges, such as Cramér-Rao like performance bounds for multi-object estimation [200], more sophisticated multi-object models and efficient accompanying solutions.

REFERENCES

- [1] B. D. Anderson and J. B. Moore, *Optimal Filtering*. Courier Corporation, 2012.
- [2] S. Särkkä, *Bayesian Filtering and Smoothing*. Cambridge University Press, 2013.
- [3] A. Jazwinski, *Stochastic Processes and Filtering Theory*. Elsevier Science, 1970.
- [4] N. Gordon, D. J. Salmond, and A. F. Smith, "Novel approach to nonlinear/non-Gaussian Bayesian state estimation," in *IEE F (Radar Signal Process.)*, 140(2):107–113, 1993.
- [5] A. Doucet, S. Godsill, and C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering," *Statist. Comput.*, 10(3):197–208, 2000.
- [6] C. P. Robert et al., *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation*. Springer, 2007.
- [7] S. S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*. Artech House, 1999.
- [8] Y. Bar-Shalom, P. Willett, and X. Tian, *Tracking and Data Fusion: A Handbook of Algorithms*. YBS Publishing, 2011.
- [9] I. Goodman, R. Mahler, and H. Nguyen, *Mathematics of Data Fusion*. Springer Netherlands, 2010.
- [10] B.-T. Vo and B.-N. Vo, "A random finite set conjugate prior and application to multi-target tracking," in *Int. Conf. Intell. Sensors, Sensor Netw. Info.*, 431–436, 2011.
- [11] B.-T. Vo and B.-N. Vo, "Labeled random finite sets and multi-object conjugate priors," *IEEE Trans. Signal Process.*, 61(13):3460–3475, 2013.
- [12] A. Farina and F. A. Studer, *Radar Data Processing Volume 1: Introduction and Tracking*. Research Studies Press, 1985.
- [13] A. Farina and F. A. Studer, *Radar Data Processing Volume 2: Advanced Topics and Applications*. Research Studies Press, 1986.
- [14] S. Challa, M. R. Morelande, D. Mušicki, and R. J. Evans, *Fundamentals of Object Tracking*. Cambridge University Press, 2011.
- [15] W. Koch, *Tracking and Sensor Data Fusion: Methodological Framework and Selected Applications*. Springer Berlin Heidelberg, 2014.
- [16] R. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Artech House, 2007.
- [17] R. Mahler, *Advances in Statistical Multisource-Multitarget Information Fusion*. Artech House, 2014.
- [18] M. Mallick, V. Krishnamurthy, and B.-N. Vo, *Integrated Tracking, Classification, and Sensor Management: Theory and Applications*. John Wiley Sons, 2012.
- [19] L. D. Stone, R. L. Streit, T. L. Corwin, and K. L. Bell, *Bayesian Multiple Target Tracking*. 2nd ed. Artech House, 2013.
- [20] Y. Bar-Shalom and T. E. Fortmann, *Tracking and Data Association*. Academic Press, 1988.
- [21] B.-N. Vo, M. Mallick, Y. Bar-Shalom, S. Coraluppi, R. Osborne, R. Mahler, and B.-T. Vo, "Multitarget tracking," *Wiley Encyclopedia Elect. Electron. Eng.*, 2015.
- [22] I. A. Bol'shakov, V. V. Vatolloy, and V. G. Latysh, "Methods of joint detection and measurement of an unknown number of signals based on the theory of random points," *Radio Eng. Electron. Phys.*, vol. 4, 1964.
- [23] Y. S. Achkasov, "Finding trajectories by a posteriori analysis of flows," *Eng. Cybern.*, vol. 9, no. 5, p. 919, 1971.
- [24] P. A. Bakut, Y. V. Zhulina, and N. A. Ivanchuk, "Detection of moving objects [in Russian]," *Sov. Radio*, 1980.
- [25] R. Washburn, "A random point process approach to multiobject tracking," in *American Control Conf.*, pp. 1846–1852, 1987.
- [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.*, 62(24):6554–6567, 2014.
- [27] M. Beard, B. T. Vo, and B.-N. Vo, "A solution for large-scale multi-object tracking," *IEEE Trans. Signal Process.*, 68:2754–2769, 2020.
- [28] R. Mahler, "Statistics 101" for multisensor, multitarget data fusion," *IEEE Aerosp. Electron. Syst. Mag.*, 19(1):53–64, 2004.
- [29] R. Mahler, "Statistics 102" for multisource-multitarget detection and tracking," *IEEE J. Sel. Top. Signal Process.*, 7(3):376–389, 2013.
- [30] R. Mahler, "A survey of PHD filter and CPHD filter implementations," in *Signal Process. Sensor Fusion Target Recognit.*, 6567:287–298, 2007.
- [31] R. Mahler, "A brief survey of advances in random-set fusion," in *Int. Conf. Control Automat. Inf. Sci.*, 62–67, 2015.
- [32] B. Ristic, *Particle Filters for Random Set Models*. Springer, 2013.
- [33] B. Ristic, M. Beard, and C. Fantacci, "An overview of particle methods for random finite set models," *Inf. Fusion*, 31:110–126, 2016.
- [34] M. Briers, A. Doucet, and S. Maskell, "Smoothing algorithms for state-space models," *Ann. Inst. Statistical Mathematics*, 62(1):61–89, 2010.
- [35] S. Kay, *Fundamentals of Statistical Signal Processing: Estimation theory*. Prentice Hall, 1993.
- [36] J. Meditch, "A survey of data smoothing for linear and nonlinear dynamic systems," *Automatica*, 9(2):151–162, 1973.
- [37] A. Doucet and A. Johansen, "A tutorial on particle filtering and smoothing: Fifteen years later," *Handbook of Nonlinear Filtering*, 12(3):656–704, 2009.
- [38] N. Gordon, B. Ristic, and S. Arulampalam, *Beyond the Kalman filter: Particle Filters for Tracking Applications*. Artech House, 2004.
- [39] G. Kitagawa, "Monte Carlo filter and smoother for nonlinear non Gaussian state models," *J. Comp. Graph. Stat.*, 5:1–25, 1996.
- [40] S. Julier and J. Uhlmann, "Unscented filtering and nonlinear estimation," *IEEE*, 92(3):401–422, 2004.
- [41] H. W. Sorenson and D. L. Alspach, "Recursive Bayesian estimation using Gaussian sums," *Automatica*, 7(4):465–479, 1971.
- [42] I. Arasaratnam and S. Haykin, "Cubature Kalman filters," *IEEE Trans. Autom. Control*, 54(6):1254–1269, 2009.
- [43] B. Jia, M. Xin, and Y. Cheng, "High-degree cubature Kalman filter," *Automatica*, 49(2):510–518, 2013.
- [44] S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian bayesian tracking," *IEEE Trans. Signal Process.*, 50(2):174–188, 2002.
- [45] F. Daum and J. Huang, "How to avoid normalization of particle flow for nonlinear filters, Bayesian decisions, and transport," in *Signal Data Process. Small Targets*, 9092:63–71, 2014.
- [46] Y. Zhang, C. Wang, X. Wang, W. Zeng, and W. Liu, "FairMOT: On the fairness of detection and re-identification in multiple object tracking," *Int. J. Comput. Vision*, 129:3069–3087, 2021.

- [47] A. Makris and C. Prieur, "Bayesian multiple-hypothesis tracking of merging and splitting targets," *IEEE Trans. Geosci. Remote Sens.*, 52(12):7684–7694, 2014.
- [48] A. Andriyenko, S. Roth, and K. Schindler, "An analytical formulation of global occlusion reasoning for multi-target tracking," in *IEEE Int. Conf. Comput. Vision Workshops*, 1839–1846, 2011.
- [49] M. Mallick, S. Rubin, and Y. Zhu, "Track label and classical and quantum probability densities," in *Int. Conf. Control Automat. Inf. Sci.*, 118–123, 2022.
- [50] H. V. Nguyen, H. Rezaatofghi, B.-N. Vo, and D. C. Ranasinghe, "Online UAV path planning for joint detection and tracking of multiple radio-tagged objects," *IEEE Trans. Signal Process.*, 67(20):5365–5379, 2019.
- [51] C. B. Storlie, T. C. Lee, J. Hannig, and D. Nychka, "Tracking of multiple merging and splitting targets: A statistical perspective," *Statistica Sinica*, 1–31, 2009.
- [52] J. F. Henriques, R. Caseiro, P. Martins, and J. Batista, "High-speed tracking with kernelized correlation filters," *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(3):583–596, 2014.
- [53] C. Liang, Z. Zhang, X. Zhou, B. Li, S. Zhu, and W. Hu, "Rethinking the competition between detection and ReID in multiobject tracking," *IEEE Trans. Image Process.*, 31:3182–3196, 2022.
- [54] L. Chen, H. Ai, Z. Zhuang, and C. Shang, "Real-time multiple people tracking with deeply learned candidate selection and person re-identification," in *Int. Conf. Multimedia Expo*, 1–6, 2018.
- [55] B.-N. Vo, B.-T. Vo, N.-T. Pham, and D. Suter, "Joint detection and estimation of multiple objects from image observations," *IEEE Trans. Signal Process.*, 58(10):5129–5141, 2010.
- [56] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, "A consistent metric for performance evaluation of multi-object filters," *IEEE Trans. Signal Process.*, 56(8):3447–3457, 2008.
- [57] T. T. D. Nguyen, H. Rezaatofghi, B.-N. Vo, B.-T. Vo, S. Savarese, and I. Reid, "How trustworthy are performance evaluations for basic vision tasks?," *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(7):8538–8552, 2023.
- [58] B.-N. Vo, N. Dam, D. Phung, Q. Tran, and B.-T. Vo, "Model-based learning for point pattern data," *Pattern Recognit.*, 84:136–151, 2018.
- [59] M. Markou and S. Singh, "Novelty detection: a review-part 1: statistical approaches," *Signal Process.*, 83(12):2481–2497, 2003.
- [60] J. Møller and R. P. Waagepetersen, "Modern statistics for spatial point processes," *Scand. J. Statist.*, 34(4):643–684, 2007.
- [61] S. Chiu, D. Stoyan, W. Kendall, and J. Mecke, *Stochastic Geometry and Its Applications*. John Wiley Sons, 2013.
- [62] D. Daley and D. Vere-Jones, *An Introduction to the Theory of Point Processes*. Springer-Verlag, 1988.
- [63] C. J. Geyer, "Likelihood inference for spatial point processes: Likelihood and computation," in *Stochastic Geometry: Likelihood and Computation*, Chapman and Hall/CRC, 141–172, 1999.
- [64] A. Baddeley, I. Bárány, and R. Schneider, "Spatial point processes and their applications," *Stochastic Geometry: Lectures at the CIME Summer School in Martina Franca, Italy, September, 2004*, 1–75, 2007.
- [65] B.-N. Vo, S. Singh, and A. Doucet, "Sequential Monte Carlo methods for multitarget filtering with random finite sets," *IEEE Trans. Aerosp. Electron. Syst.*, 41(4):1224–1245, 2005.
- [66] D. E. Clark and J. Housineau, "Faa di Bruno's formula for chain differentials," *arXiv:1310.2833*, 2013.
- [67] R. Mahler, "PHD filters of higher order in target number," *IEEE Trans. Aerosp. Electron. Syst.*, 43(4):1523–1543, 2007.
- [68] R. Mahler, "Multitarget Bayes filtering via first-order multitarget moments," *IEEE Trans. Aerosp. Electron. Syst.*, 39(4):1152–1178, 2003.
- [69] M. N. M. Van Lieshout, *Markov Point Processes and Their Applications*. World Scientific, 2000.
- [70] F. Papi, B.-N. Vo, B.-T. Vo, C. Fantacci, and M. Beard, "Generalized labeled multi-Bernoulli approximation of multi-object densities," *IEEE Trans. Signal Process.*, 63(20):5487–5497, 2015.
- [71] A. F. García-Fernández, J. L. Williams, K. Granström, and L. Svensson, "Poisson multi-Bernoulli mixture filter: Direct derivation and implementation," *IEEE Trans. Aerosp. Electron. Syst.*, 54(4):1883–1901, 2018.
- [72] M. Beard, B.-T. Vo, B.-N. Vo, and S. Arulampalam, "Void probabilities and Cauchy-Schwarz divergence for generalized labeled multi-Bernoulli models," *IEEE Trans. Signal Process.*, 65(19):5047–5061, 2017.
- [73] S. Reuter, B.-T. Vo, B.-N. Vo, and K. Dietmayer, "The labeled multi-Bernoulli filter," *IEEE Trans. Signal Process.*, 62(12):3246–3260, 2014.
- [74] 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.*, 47(2):1521–1529, 2011.
- [75] K. Shen, C. Zhang, P. Dong, Z. Jing, and H. Leung, "Consensus-based labeled multi-Bernoulli filter with event-triggered communication," *IEEE Trans. Signal Process.*, 70:1185–1196, 2022.
- [76] H. V. Nguyen, B.-N. Vo, B.-T. Vo, H. Rezaatofghi, and D. C. Ranasinghe, "Multi-objective multi-agent planning for discovering and tracking multiple mobile objects," *IEEE Trans. Signal Process.*, vol. 72, pp. 3669–3685, 2024.
- [77] A. L. Gibbs and F. E. Su, "On choosing and bounding probability metrics," *Int. Statistical Rev.*, 70(3):419–435, 2002.
- [78] H. G. Hoang, B.-N. Vo, B.-T. Vo, and R. Mahler, "The Cauchy-Schwarz divergence for Poisson point processes," *IEEE Trans. Inf. Theory*, 61(8):4475–4485, 2015.
- [79] K. Kampa, E. Hasanbelliu, and J. C. Principe, "Closed-form Cauchy-Schwarz PDF divergence for mixture of Gaussians," in *Int. Joint Conf. Neural Netw.*, 2578–2585, 2011.
- [80] M. Beard, B.-T. Vo, and B.-N. Vo, "Bayesian multi-target tracking with merged measurements using labelled random finite sets," *IEEE Trans. Signal Process.*, 63(6):1433–1447, 2015.
- [81] F. Papi and D. Y. Kim, "A particle multi-target tracker for superpositional measurements using labeled random finite sets," *IEEE Trans. Signal Process.*, 63(16):4348–4358, 2015.
- [82] S. Li, W. Yi, R. Hoseinnezhad, B. Wang, and L. Kong, "Multiobject tracking for generic observation model using labeled random finite sets," *IEEE Trans. Signal Process.*, 66(2):368–383, 2017.
- [83] D. S. Bryant, B.-T. Vo, B.-N. Vo, and B. A. Jones, "A generalized labeled multi-Bernoulli filter with object spawning," *IEEE Trans. Signal Process.*, 66(23):6177–6189, 2018.
- [84] R. Mahler, "A fast labeled multi-Bernoulli filter for superpositional sensors," in *Signal Process. Sensor/Inf. Fusion Target Recognit.*, 10646:113–124, 2018.
- [85] A. K. Gostar, T. Rathnayake, C. Fu, A. Bab-Hadiashar, G. Battistelli, L. Chisci, and R. Hoseinnezhad, "Interactive multiple-target tracking via labeled multi-Bernoulli filters," in *Int. Conf. Control Automat. Inf. Sci.*, 1–6, 2019.
- [86] J. Ong, B.-T. Vo, B.-N. Vo, D. Y. Kim, and S. Nordholm, "A Bayesian filter for multi-view 3D multi-object tracking with occlusion handling," *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(5):2246–2263, 2022.
- [87] T. T. D. Nguyen, B.-N. Vo, B.-T. Vo, D. Y. Kim, and Y. S. Choi, "Tracking cells and their lineages via labeled random finite sets," *IEEE Trans. Signal Process.*, 69:5611–5626, 2021.
- [88] C. Fantacci, B.-T. Vo, F. Papi, and B.-N. Vo, "The marginalized δ -GLMB filter," *arXiv:1501.00926*, 2015.
- [89] C. Fantacci, B.-N. Vo, B.-T. Vo, G. Battistelli, and L. Chisci, "Robust fusion for multisensor multiobject tracking," *IEEE Signal Process. Lett.*, 25(5):640–644, 2018.
- [90] J.-F. Cardoso, "Dependence, correlation and Gaussianity in independent component analysis," *J. Mach. Learning Res.*, 4:1177–1203, 2003.
- [91] B.-N. Vo and B.-T. Vo, "A multi-scan labeled random finite set model for multi-object state estimation," *IEEE Trans. Signal Process.*, 67(19):4948–4963, 2019.
- [92] A.-T. Vu, B.-N. Vo, and R. Evans, "Particle Markov chain Monte Carlo for Bayesian multi-target tracking," in *Int. Conf. Inf. Fusion*, 1–8, 2011.
- [93] T. Vu, B.-N. Vo, and R. Evans, "A particle marginal Metropolis-Hastings multi-target tracker," *IEEE Trans. Signal Process.*, 62(15):3953–3964, 2014.
- [94] B.-T. Vo and B.-N. Vo, "Multi-scan generalized labeled multi-Bernoulli filter," in *Int. Conf. Inf. Fusion*, 195–202, 2018.
- [95] B.-N. Vo, B.-T. Vo, and H. G. Hoang, "An efficient implementation of the generalized labeled multi-Bernoulli filter," *IEEE Trans. Signal Process.*, 65(8):1975–1987, 2017.
- [96] B.-N. Vo, B.-T. Vo, and M. Beard, "Multi-sensor multi-object tracking with the generalized labeled multi-Bernoulli filter," *IEEE Trans. Signal Process.*, 67(23):5952–5967, 2019.
- [97] T. T. D. Nguyen and D. Y. Kim, "GLMB tracker with partial smoothing," *Sensors*, 19(20):4419, 2019.
- [98] H. G. Hoang, B. T. Vo, and B.-N. Vo, "A fast implementation of the generalized labeled multi-Bernoulli filter with joint prediction and update," in *Int. Conf. Inf. Fusion*, 999–1006, 2015.
- [99] K. G. Murty, "An algorithm for ranking all the assignments in order of increasing cost," *Operations Res.*, 16(3):682–687, 1968.
- [100] M. L. Miller, H. S. Stone, and I. J. Cox, "Optimizing Murty's ranked assignment method," *IEEE Trans. Aerosp. Electron. Syst.*, 33(3):851–862, 1997.
- [101] C. R. Pedersen, L. R. Nielsen, and K. A. Andersen, "An algorithm for ranking assignments using reoptimization," *Comput. Operations Res.*, 35(11):3714–3726, 2008.

- [102] T. Kropfreiter, F. Meyer, and F. Hlawatsch, "A fast labeled multi-Bernoulli filter using belief propagation," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 56, no. 3, pp. 2478–2488, 2019.
- [103] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Trans. Pattern Anal. Mach. Intell.*, 6:721–741, 1984.
- [104] G. Casella and E. I. George, "Explaining the Gibbs sampler," *Amer. Statistician*, 46(3):167–174, 1992.
- [105] C. Shim, B.-T. Vo, B.-N. Vo, J. Ong, and D. Moratuwage, "Linear complexity Gibbs sampling for generalized labeled multi-Bernoulli filtering," *IEEE Trans. Signal Process.*, 71:1981–1994, 2023.
- [106] C. Fantacci and F. Papi, "Scalable multisensor multitarget tracking using the marginalized δ -GLMB density," *IEEE Signal Process. Lett.*, 23(6):863–867, 2016.
- [107] B. Wang, Y. Xu, S. Li, G. Battistelli, and L. Chisci, "Centralized multi-sensor labeled multi-Bernoulli filter with partially overlapping fields of view," *Signal Process.*, 213:109180, 2023.
- [108] B. Yang, J. Wang, and W. Wang, "An efficient approximate implementation for labeled random finite set filtering," *Signal Process.*, 150:215–227, 2018.
- [109] L. M. Wolf and M. Baum, "Deterministic Gibbs sampling for data association in multi-object tracking," in *IEEE Int. Conf. Multi. Fusion Integ. Intell. Syst.*, 291–296, 2020.
- [110] D. M. Nguyen, H. A. Le Thi, and T. Pham Dinh, "Solving the multidimensional assignment problem by a cross-entropy method," *J. Combinatorial Optim.*, 27(4):808–823, 2014.
- [111] J. Y. Yu, A.-A. Saucan, M. Coates, and M. Rabbat, "Algorithms for the multi-sensor assignment problem in the δ -generalized labeled multi-Bernoulli filter," in *Int. Workshop Comput. Adv. Multi-Sensor Adaptive Process.*, 1–5, 2017.
- [112] A.-A. Saucan and P. K. Varshney, "Distributed cross-entropy δ -GLMB filter for multi-sensor multi-target tracking," in *Int. Conf. Inf. Fusion*, 1559–1566, 2018.
- [113] D. Moratuwage, B.-N. Vo, B.-T. Vo, and C. Shim, "Multi-scan multi-sensor multi-object state estimation," *IEEE Trans. Signal Process.*, 70:5429–5442, 2022.
- [114] C. Yang, X. Cao, and Z. Shi, "Road-map aided Gaussian mixture labeled multi-Bernoulli filter for ground multi-target tracking," *IEEE Trans. Veh. Technol.*, 72(6):7137–7147, 2023.
- [115] S. Reuter, A. Scheel, and K. Dietmayer, "The multiple model labeled multi-Bernoulli filter," in *Int. Conf. Inf. Fusion*, 1574–1580, 2015.
- [116] Y. PUNCHIHewa, B.-N. Vo, and B.-T. Vo, "A generalized labeled multi-Bernoulli filter for maneuvering targets," in *Int. Conf. Inf. Fusion*, 980–986, 2016.
- [117] W. Yi, M. Jiang, and R. Hoseinnezhad, "The multiple model Vo-Vo filter," *IEEE Trans. Aerosp. Electron. Syst.*, 53(2):1045–1054, 2017.
- [118] Z.-X. Liu and B.-J. Huang, "The labeled multi-Bernoulli filter for jump Markov systems under glint noise," *IEEE Access*, 7:92322–92328, 2019.
- [119] S. Wu, X. Dong, J. Zhao, X. Sun, and R. Cai, "A fast implementation of interactive-model generalized labeled multi-Bernoulli filter for interval measurements," *Signal Process.*, 164:345–353, 2019.
- [120] C. Cao and Y. Zhao, "A multiple-model generalized labeled multi-Bernoulli filter based on blocked Gibbs sampling for tracking maneuvering targets," *Signal Process.*, 186:108119, 2021.
- [121] N. Ishtiaq, A. K. Gostar, A. Bab-Hadiashar, and R. Hoseinnezhad, "Interaction-aware labeled multi-Bernoulli filter," *IEEE Trans. Intell. Transp. Syst.*, 24(11):11668–11681, 2023.
- [122] W. Liu and Y. Chi, "Resolvable group state estimation with maneuver based on labeled RFS and graph theory," *Sensors*, 19(6):1307, 2019.
- [123] X. Xue, S. Huang, J. Xie, J. Ma, and N. Li, "Resolvable cluster target tracking based on the DBSCAN clustering algorithm and labeled RFS," *IEEE Access*, 9:43364–43377, 2021.
- [124] G. Li, G. Li, and Y. He, "Labeled multi-Bernoulli filter based multiple resolvable group targets tracking with leader-follower model," *IEEE Trans. Aerosp. Electron. Syst.*, 59(5):6683–6694, 2023.
- [125] L. Van Ma, T. T. D. Nguyen, B.-N. Vo, H. Jang, and M. Jeon, "Track initialization and re-identification for 3D multi-view multi-object tracking," *Inf. Fusion*, p. 102496, 2024.
- [126] P. Dong, Z. Jing, H. Leung, K. Shen, and M. Li, "The labeled multi-Bernoulli filter for multitarget tracking with glint noise," *IEEE Trans. Aerosp. Electron. Syst.*, 55(5):2253–2268, 2018.
- [127] P. Dong, Z. Jing, H. Leung, K. Shen, and J. Wang, "Student-t mixture labeled multi-Bernoulli filter for multi-target tracking with heavy-tailed noise," *Signal Process.*, 152:331–339, 2018.
- [128] W. Zhang, F. Yang, Y. Liang, and Z. Liu, "A heavy-tailed noise tolerant labeled multi-Bernoulli filter," in *Int. Conf. Inf. Fusion*, 1–7, 2018.
- [129] J. Ong, B. T. Vo, and S. Nordholm, "Blind separation for multiple moving sources with labeled random finite sets," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, 29:2137–2151, 2021.
- [130] J. Ong, B. T. Vo, S. Nordholm, B.-N. Vo, D. Moratuwage, and C. Shim, "Audio-visual based online multi-source separation," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, 30:1219–1234, 2022.
- [131] M. Beard, S. Reuter, K. Granström, B.-T. Vo, B.-N. Vo, and A. Scheel, "Multiple extended target tracking with labeled random finite sets," *IEEE Trans. Signal Process.*, 64(7):1638–1653, 2015.
- [132] Z. Liang, F. Liu, L. Li, and J. Gao, "Improved generalized labeled multi-Bernoulli filter for non-ellipsoidal extended targets or group targets tracking based on random sub-matrices," *Digit. Signal Process.*, 99:102669, 2020.
- [133] A. F. García-Fernández, "Track-before-detect labeled multi-Bernoulli particle filter with label switching," *IEEE Trans. Aerosp. Electron. Syst.*, 52(5):2123–2138, 2016.
- [134] A.-A. Saucan, Y. Li, and M. Coates, "Particle flow superpositional GLMB filter," in *Signal Process. Sensor/Inf. Fusion Target Recognit.*, 10200:116–127, 2017.
- [135] K. Granström, M. Baum, and S. Reuter, "Extended object tracking: Introduction, overview, and applications," *J. Adv. Inf. Fusion*, 12(2):139–174, 2017.
- [136] X. Cheng, H. Ji, and Y. Zhang, "Multiple extended target tracking based on gamma box particle and labeled random finite sets," *Digit. Signal Process.*, 134:103902, 2023.
- [137] A. Daniyan, S. Lambotcharan, A. Deligiannis, Y. Gong, and W.-H. Chen, "Bayesian multiple extended target tracking using labeled random finite sets and splines," *IEEE Trans. Signal Process.*, 66(22):6076–6091, 2018.
- [138] A.-A. Saucan and M. Z. Win, "On the labeled multi-Bernoulli filter with merged measurements," in *Int. Conf. Commun.*, 1–5, 2020.
- [139] K. Dai, Y. Wang, J.-S. Hu, K. Nam, and C. Yin, "Intertarget occlusion handling in multiextended target tracking based on labeled multi-Bernoulli filter using laser range finder," *IEEE/ASME Trans. Mechatronics*, 25(4):1719–1728, 2020.
- [140] X. Dong, J. Zhao, M. Sun, X. Zhang, and Y. Wang, "A modified δ -generalized labeled multi-Bernoulli filtering for multi-source DOA tracking with coprime array," *IEEE Trans. Wirel. Commun.*, 22(12):9424–9437, 2023.
- [141] J. Zhang, M. Bao, J. Yang, Z. Chen, and H. Hou, "DOA tracking algorithm based on AVS pseudo-smoothing for coherent acoustic targets," *IEEE Trans. Aerosp. Electron. Syst.*, 59(6):8175–8193, 2023.
- [142] C. Cao and Y. Zhao, "A generalized labeled multi-Bernoulli filter based on track-before-detect measurement model for multiple-weak-target state estimate using belief propagation," *Remote Sens.*, 14(17):4209, 2022.
- [143] C. Cao and Y. Zhao, "A multi-frame GLMB smoothing based on the image-observation sensor for tracking multiple weak targets using belief propagation," *Remote Sens.*, 14(22):5666, 2022.
- [144] B. Yu and E. Ye, "Track-before-detect labeled multi-Bernoulli smoothing for multiple extended objects," in *Int. Conf. Inf. Fusion*, 1–8, 2020.
- [145] K. Dai, Y. Wang, Q. Ji, H. Du, and C. Yin, "Multiple vehicle tracking based on labeled multiple Bernoulli filter using pre-clustered laser range finder data," *IEEE Trans. Veh. Technol.*, 68(11):10382–10393, 2019.
- [146] D. Y. Kim, B.-N. Vo, B.-T. Vo, and M. Jeon, "A labeled random finite set online multi-object tracker for video data," *Pattern Recognit.*, 90:377–389, 2019.
- [147] Y. G. PUNCHIHewa, B.-T. Vo, B.-N. Vo, and D. Y. Kim, "Multiple object tracking in unknown backgrounds with labeled random finite sets," *IEEE Trans. Signal Process.*, 66(11):3040–3055, 2018.
- [148] R. Mahler, "A clutter-agnostic generalized labeled multi-Bernoulli filter," in *Signal Process. Sensor/Inf. Fusion Target Recognit.*, 10646:101–112, 2018.
- [149] R. Mahler, B.-T. Vo, and B.-N. Vo, "CPHD filtering with unknown clutter rate and detection profile," *IEEE Trans. Signal Process.*, 59(8):3497–3513, 2011.
- [150] B.-T. Vo, B.-N. Vo, R. Hoseinnezhad, and R. Mahler, "Robust multi-Bernoulli filtering," *IEEE J. Sel. Top. Signal Process.*, 7(3):399–409, 2013.
- [151] C.-T. Do, T. T. D. Nguyen, D. Moratuwage, C. Shim, and Y. D. Chung, "Multi-object tracking with an adaptive generalized labeled multi-Bernoulli filter," *Signal Process.*, 196:108532, 2022.
- [152] C.-T. Do, T. T. D. Nguyen, and H. V. Nguyen, "Robust multi-sensor generalized labeled multi-Bernoulli filter," *Signal Process.*, 192:108368, 2022.

- [153] D. Y. Kim, "Visual multiple-object tracking for unknown clutter rate," *IET Comput. Vision*, 12(5):728–734, 2018.
- [154] W. Zhang, Y. Liang, F. Yang, and L. Xu, "A robust Student's t-based labeled multi-Bernoulli filter," in *Int. Conf. Inf. Fusion*, 1–6, 2019.
- [155] L. Hou, F. Lian, S. Tan, C. Xu, and G. T. F. de Abreu, "Robust generalized labeled multi-Bernoulli filter for multitarget tracking with unknown non-stationary heavy-tailed measurement noise," *IEEE Access*, 9:94438–94453, 2021.
- [156] G. Li, P. Wei, Y. Li, L. Gao, and H. Zhang, "A robust fast LMB filter for superpositional sensors," *Signal Process.*, 174:107606, 2020.
- [157] A. Scheel and K. Dietmayer, "Tracking multiple vehicles using a variational radar model," *IEEE Trans. Intell. Transp. Syst.*, 20(10):3721–3736, 2019.
- [158] K. A. LeGrand and K. J. DeMars, "The data-driven delta-generalized labeled multi-Bernoulli tracker for automatic birth initialization," in *Signal Process. Sensor/Inf. Fusion Target Recognit.*, 10646:36–55, 2018.
- [159] S. Lin, "Robust pitch estimation and tracking for speakers based on subband encoding and the generalized labeled multi-Bernoulli filter," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, 27(4):827–841, 2019.
- [160] H. Cai, S. Gehly, Y. Yang, and K. Zhang, "Modeling birth for the labeled multi-Bernoulli filter using a boundary-value approach," *J. Guid. Control Dyn.*, 43(1):162–169, 2020.
- [161] S. Zhu, B. Yang, and S. Wu, "Measurement-driven multi-target tracking filter under the framework of labeled random finite set," *Digit. Signal Process.*, 112:103000, 2021.
- [162] P. Hoher, T. Baur, S. Wirtensohn, and J. Reuter, "A detection driven adaptive birth density for the labeled multi-Bernoulli filter," in *Int. Conf. Inf. Fusion*, 1–8, 2020.
- [163] J. A. Gaebler, P. Axelrad, and P. W. Schumacher Jr, "Cubesat cluster deployment track initiation via a radar admissible region birth model," *J. Guid. Control Dyn.*, 43(10):1927–1934, 2020.
- [164] A. Trezza, D. J. Bucci, and P. K. Varshney, "Multi-sensor joint adaptive birth sampler for labeled random finite set tracking," *IEEE Trans. Signal Process.*, 70:1010–1025, 2022.
- [165] T. T. D. Nguyen, C.-T. Do, and H. V. Nguyen, "An adaptive multi-sensor generalised labelled multi-Bernoulli filter for linear Gaussian models," in *Int. Conf. Control Automat. Inf. Sci.*, 84–89, 2022.
- [166] H. Cai, J. Houssineau, B. A. Jones, M. Jah, and J. Zhang, "Possibility generalized labeled multi-Bernoulli filter for multi-target tracking under epistemic uncertainty," *IEEE Trans. Aerosp. Electron. Syst.*, 59:1312–1326, 2022.
- [167] R. Mahler, "Optimal/robust distributed data fusion: a unified approach," in *Signal Process. Sensor Fusion Target Recognit.*, 4052:128–138, 2000.
- [168] X. Wang, A. K. Gostar, T. Rathnayake, B. Xu, A. Bab-Hadiashar, and R. Hoseinnezhad, "Centralized multiple-view sensor fusion using labeled multi-Bernoulli filters," *Signal Process.*, 150:75–84, 2018.
- [169] Y. Jin and J. Li, "A discrepant information selection mechanism for cooperative sensors," *Signal Process.*, 203:108772, 2023.
- [170] S. Li, W. Yi, R. Hoseinnezhad, G. Battistelli, B. Wang, and L. Kong, "Robust distributed fusion with labeled random finite sets," *IEEE Trans. Signal Process.*, 66(2):278–293, 2017.
- [171] S. Li, G. Battistelli, L. Chisci, W. Yi, B. Wang, and L. Kong, "Computationally efficient multi-agent multi-object tracking with labeled random finite sets," *IEEE Trans. Signal Process.*, 67(1):260–275, 2018.
- [172] K. Shen, P. Dong, Z. Jing, and H. Leung, "Consensus-based labeled multi-Bernoulli filter for multitarget tracking in distributed sensor network," *IEEE Trans. Cybern.*, 52(12):12722–12733, 2021.
- [173] H. Liu, J. Yang, Y. Xu, and L. Yang, "Optimizing distributed multi-sensor multi-target tracking algorithm based on labeled multi-Bernoulli filter," in *Int. Conf. Acoust. Speech Signal Process.*, IEEE, 1–5, 2023.
- [174] L. Gao, G. Battistelli, and L. Chisci, "Fusion of labeled RFS densities with minimum information loss," *IEEE Trans. Signal Process.*, 68:5855–5868, 2020.
- [175] L. Gao, G. Battistelli, and L. Chisci, "Fusion of labeled RFS densities with different fields of view," *IEEE Trans. Aerosp. Electron. Syst.*, 58(6):5908–5924, 2022.
- [176] L. Gao, G. Battistelli, L. Chisci, and A. Farina, "Fusion-based multidetection multitarget tracking with random finite sets," *IEEE Trans. Aerosp. Electron. Syst.*, 57(4):2438–2458, 2021.
- [177] G. Li, G. Battistelli, L. Chisci, L. Gao, and P. Wei, "Distributed joint detection, tracking, and classification via labeled multi-Bernoulli filtering," *IEEE Trans. Cybern.*, 2022.
- [178] L. Gao, G. Battistelli, and L. Chisci, "Resilient labeled multi-Bernoulli fusion with peer-to-peer sensor network," *Inf. Fusion*, 100:101965, 2023.
- [179] T. Li, "Arithmetic average density fusion - part II: Unified derivation for (labeled) RFS fusion," *IEEE Trans. Aerosp. Electron. Syst.*, 2024.
- [180] T. Li, R. Yan, K. Da, and H. Fan, "Arithmetic average density fusion-part III: Heterogeneous unlabeled and labeled RFS filter fusion," *IEEE Trans. Aerosp. Electron. Syst.*, 2023.
- [181] A. K. Gostar, T. Rathnayake, R. Tennakoon, A. Bab-Hadiashar, G. Battistelli, L. Chisci, and R. Hoseinnezhad, "Cooperative sensor fusion in centralized sensor networks using Cauchy-Schwarz divergence," *Signal Process.*, 167:107278, 2020.
- [182] A. K. Gostar, T. Rathnayake, R. Tennakoon, A. Bab-Hadiashar, G. Battistelli, L. Chisci, and R. Hoseinnezhad, "Centralized cooperative sensor fusion for dynamic sensor network with limited field-of-view via labeled multi-Bernoulli filter," *IEEE Trans. Signal Process.*, 69:878–891, 2020.
- [183] M. Herrmann, C. Hermann, and M. Buchholz, "Distributed implementation of the centralized generalized labeled multi-Bernoulli filter," *IEEE Trans. Signal Process.*, 69:5159–5174, 2021.
- [184] H. V. Nguyen, H. Rezatofghi, B.-N. Vo, and D. C. Ranasinghe, "Distributed multi-object tracking under limited field of view sensors," *IEEE Trans. Signal Process.*, 69:5329–5344, 2021.
- [185] M. Beard, B.-T. Vo, B.-N. Vo, and S. Arulampalam, "Sensor control for multi-target tracking using Cauchy-Schwarz divergence," in *Int. Conf. Inf. Fusion*, 937–944, 2015.
- [186] R. Mahler, "A GLMB filter for unified multitarget multisensor management," in *Signal Process. Sensor/Inf. Fusion Target Recognit.*, 11018:123–134, 2019.
- [187] Y. Zhu, S. Liang, M. Gong, and J. Yan, "Decomposed POMDP optimization-based sensor management for multi-target tracking in passive multi-sensor systems," *IEEE Sensors J.*, 22(4):3565–3578, 2022.
- [188] A. K. Gostar, R. Hoseinnezhad, T. Rathnayake, X. Wang, and A. Bab-Hadiashar, "Constrained sensor control for labeled multi-Bernoulli filter using Cauchy-Schwarz divergence," *IEEE Signal Process. Lett.*, 24(9):1313–1317, 2017.
- [189] G. Zhang, F. Lian, C. Han, H. Chen, and N. Fu, "Two novel sensor control schemes for multi-target tracking via delta generalised labelled multi-Bernoulli filtering," *IET Signal Process.*, 12(9):1131–1139, 2018.
- [190] H. Cai, S. Gehly, Y. Yang, R. Hoseinnezhad, R. Norman, and K. Zhang, "Multisensor tasking using analytical Rényi divergence in labeled multi-Bernoulli filtering," *J. Guid. Control Dyn.*, 42(9):2078–2085, 2019.
- [191] X. Wang, R. Hoseinnezhad, A. K. Gostar, T. Rathnayake, B. Xu, and A. Bab-Hadiashar, "Multi-sensor control for multi-object Bayes filters," *Signal Process.*, 142:260–270, 2018.
- [192] S. Panicker, A. K. Gostar, A. Bab-Hadiashar, and R. Hoseinnezhad, "Tracking of targets of interest using labeled multi-Bernoulli filter with multi-sensor control," *Signal Process.*, 171:107451, 2020.
- [193] N. Ishtiaq, S. Panicker, A. K. Gostar, A. Bab-Hadiashar, and R. Hoseinnezhad, "Selective sensor control via Cauchy Schwarz divergence," in *Smart Trends Comput. Commun.*, 113–124, 2021.
- [194] Y. Zhu, J. Wang, and S. Liang, "Multi-objective optimization based multi-Bernoulli sensor selection for multi-target tracking," *Sensors*, 19(4):980, 2019.
- [195] A. K. Gostar, R. Hoseinnezhad, A. Bab-Hadiashar, and W. Liu, "Sensor-management for multitarget filters via minimization of posterior dispersion," *IEEE Trans. Aerosp. Electron. Syst.*, 53(6):2877–2884, 2017.
- [196] F. Lian, L. Hou, J. Liu, and C. Han, "Constrained multi-sensor control using a multi-target MSE bound and a δ -GLMB filter," *Sensors*, 18(7):2308, 2018.
- [197] F. Lian, L. Hou, B. Wei, and C. Han, "Sensor selection for decentralized large-scale multi-target tracking network," *Sensors*, 18(12):4115, 2018.
- [198] H. Deusch, S. Reuter, and K. Dietmayer, "The labeled multi-Bernoulli SLAM filter," *IEEE Signal Process. Lett.*, 22(10):1561–1565, 2015.
- [199] D. Moratuwage, M. Adams, and F. Inostroza, " δ -generalized labeled multi-Bernoulli simultaneous localization and mapping with an optimal kernel-based particle filtering approach," *Sensors*, 19(10):2290, 2019.
- [200] D. E. Clark, "A Cramér Rao bound for point processes," *IEEE Trans. Inf. Theory*, vol. 68, no. 4, pp. 2147–2155, 2022.

VII. SUPPLEMENTARY MATERIALS

This supplementary section presents the proofs and additional discussions of the divergence results for Labeled Multi-Bernoullis (LMBs) in Subsection IV-E. We recall the notation $\langle \mathbf{f} \rangle(\ell) \triangleq \int \mathbf{f}(x, \ell) dx$, and for each LMB parameterized by $\{(r_i(\ell), p_i(\cdot, \ell)) : \ell \in \mathcal{D}(\sigma_i)\}$, with $r_i < 1$, $i = 1, 2$ and $\mathcal{D}(\sigma)$ denoting the domain of σ . Moreover, we use the multi-object density of the form $\pi_i(\mathbf{X}) = \Delta(\mathbf{X}) K_i \mathbf{f}_i^{\mathbf{X}}$, where $K_i = \tilde{r}_i^{\mathbb{L}}$, $\tilde{r}_i \triangleq 1 - r_i$, and $\mathbf{f}_i \triangleq r_i p_i / \tilde{r}_i$. Note that $\langle \mathbf{f}_i \rangle(\ell) = r_i(\ell) / \tilde{r}_i(\ell)$, $\tilde{r}_i(\ell) = (1 + \langle \mathbf{f}_i \rangle(\ell))^{-1}$, and $\tilde{r}_i^{\mathbb{L}} = \tilde{r}_i^{\mathcal{D}(\sigma_i)}$, because $r_i(\ell) = 0$, i.e., $\tilde{r}_i(\ell) = 1$ for $\ell \notin \mathcal{D}(\sigma_i)$.

A. Preliminary Lemmas

The following lemmas facilitate the arguments in the proofs.

Lemma 1. For functions g, h defined on a finite set S ,

$$\sum_{L \subseteq S} g^L h^{S-L} = (h + g)^S.$$

Proof. Let $S = \{\ell_1, \dots, \ell_n\}$, $f = g/h$, and recall Vieta's formula for elementary symmetric functions $e_i(f(\ell_1), \dots, f(\ell_n))$:

$$\sum_{i=0}^n e_i(f(\ell_1), \dots, f(\ell_n)) (-1)^i z^{n-i} = \prod_{j=1}^n (z - f(\ell_j)),$$

Substituting $z = -1$ gives

$$\begin{aligned} (-1)^n \sum_{i=0}^n e_i(f(\ell_1), \dots, f(\ell_n)) &= \prod_{j=1}^n (-1 - f(\ell_j)) \\ \sum_{i=0}^n e_i(f(\ell_1), \dots, f(\ell_n)) &= \prod_{j=1}^n (1 + f(\ell_j)). \end{aligned}$$

Noting that

$$e_i(f(\ell_1), \dots, f(\ell_n)) \triangleq \sum_{L \subseteq S, |L|=i} f^L,$$

we have

$$\begin{aligned} \sum_{L \subseteq S} f^L &= \sum_{i=0}^n e_i(f(\ell_1), \dots, f(\ell_n)) \\ &= \prod_{j=1}^n (1 + f(\ell_j)) = (1 + f)^S. \end{aligned}$$

Substituting $f = g/h$, and multiplying both sides by h^S gives

$$\sum_{L \subseteq S} \left(\frac{g}{h}\right)^L h^S = \left(1 + \frac{g}{h}\right)^S h^S.$$

Hence, $\sum_{L \subseteq S} g^L h^{S-L} = (h + g)^S$. \square

Lemma 2. For $\mathbf{f} : \mathbb{X} \times \mathbb{L} \rightarrow \mathbb{R}$ integrable on \mathbb{X} ,

$$\int \Delta(\mathbf{X}) \mathbf{f}^{\mathbf{X}} \delta \mathbf{X} = \sum_{L \subseteq \mathbb{L}} \langle \mathbf{f} \rangle^L = (1 + \langle \mathbf{f} \rangle)^{\mathbb{L}}.$$

Proof. The result follows by applying Lemma 3 in [11] to the set integral, and then Lemma 1. \square

Lemma 3. For $f, g : \mathbb{X} \times \mathbb{L} \rightarrow \mathbb{R}$ integrable on \mathbb{X} , with g unitless

$$\int \Delta(\mathbf{X}) \mathbf{f}^{\mathbf{X}} \ln \mathbf{g}^{\mathbf{X}} \delta \mathbf{X} = \sum_{L \subseteq \mathbb{L}} \sum_{\ell \in L} \langle \mathbf{f} \rangle^{L-\{\ell\}} \langle \mathbf{f} \ln \mathbf{g} \rangle(\ell)$$

The proof of this lemma is given in the Appendix of [76].

Lemma 4. For $f : \mathcal{F}(\mathbb{L}) \rightarrow \mathbb{R}$, $g : \mathbb{L} \rightarrow \mathbb{R}$, and $S \subseteq \mathbb{L}$,

$$\sum_{L \subseteq S} f(L) \sum_{\ell \in L} g(\ell) = \sum_{\ell \in \mathbb{L}} g(\ell) \sum_{L \subseteq S - \{\ell\}} f(L \cup \{\ell\}).$$

The proof of this lemma is given in the Appendix of [76].

B. Proof of Rényi Divergence (28)

Proof. The Rényi Divergence (Rényi-D) between the LMBs π_1 and π_2 is given by

$$\begin{aligned} D_R(\pi_1 || \pi_2) &= \frac{1}{\alpha - 1} \ln \int \Delta(\mathbf{X}) \left(K_1 \mathbf{f}_1^{\mathbf{X}}\right)^\alpha \left(K_2 \mathbf{f}_2^{\mathbf{X}}\right)^{1-\alpha} \delta \mathbf{X} \\ &= \frac{1}{\alpha - 1} \ln \left(K_1^\alpha K_2^{1-\alpha} \int \Delta(\mathbf{X}) (\mathbf{f}_1^\alpha)^{\mathbf{X}} (\mathbf{f}_2^{1-\alpha})^{\mathbf{X}} \delta \mathbf{X} \right) \\ &= \frac{1}{\alpha - 1} \ln \left((\tilde{r}_1^{\mathbb{L}})^\alpha (\tilde{r}_2^{\mathbb{L}})^{1-\alpha} \int \Delta(\mathbf{X}) (\mathbf{f}_1^\alpha \mathbf{f}_2^{1-\alpha})^{\mathbf{X}} \delta \mathbf{X} \right) \\ &= \frac{1}{\alpha - 1} \ln \left[(\tilde{r}_1^\alpha \tilde{r}_2^{1-\alpha})^{\mathbb{L}} (1 + \langle \mathbf{f}_1^\alpha \mathbf{f}_2^{1-\alpha} \rangle)^{\mathbb{L}} \right] \\ &= \sum_{\ell \in \mathbb{L}} \frac{\ln [(\tilde{r}_1^\alpha \tilde{r}_2^{1-\alpha})(\ell) (1 + \langle \mathbf{f}_1^\alpha \mathbf{f}_2^{1-\alpha} \rangle(\ell))]}{\alpha - 1} \\ &= \sum_{\ell \in \mathbb{L}} \frac{\ln [(\tilde{r}_1^\alpha \tilde{r}_2^{1-\alpha})(\ell)] + \ln [1 + \langle \mathbf{f}_1^\alpha \mathbf{f}_2^{1-\alpha} \rangle(\ell)]}{\alpha - 1}, \end{aligned}$$

where the 3rd last line follows from Lemma 2. Note that $\tilde{r}_1^\alpha \tilde{r}_2^{1-\alpha} > 0$, $\langle \mathbf{f}_1^\alpha \mathbf{f}_2^{1-\alpha} \rangle(\ell) \geq 0$ and hence each term of the sum is well-defined. \square

Remark: Since $\tilde{r}_i(\ell) = (1 + \langle \mathbf{f}_i \rangle(\ell))^{-1}$, the above Rényi-D can be written completely in terms of \mathbf{f}_1 and \mathbf{f}_2 as

$$\begin{aligned} D_R(\pi_1 || \pi_2) &= \sum_{\ell \in \mathbb{L}} \frac{\ln [1 + \langle \mathbf{f}_1^\alpha \mathbf{f}_2^{1-\alpha} \rangle(\ell)]}{\alpha - 1} \\ &\quad - \sum_{\ell \in \mathbb{L}} \frac{\alpha \ln [1 + \langle \mathbf{f}_1 \rangle(\ell)] + (1 - \alpha) \ln [1 + \langle \mathbf{f}_2 \rangle(\ell)]}{\alpha - 1}. \end{aligned}$$

Noting that $\mathbf{f}_i = r_i p_i / \tilde{r}_i$, and hence

$$\langle \mathbf{f}_1^\alpha \mathbf{f}_2^{1-\alpha} \rangle(\ell) = \frac{r_1^\alpha(\ell) r_2^{1-\alpha}(\ell)}{\tilde{r}_1^\alpha(\ell) \tilde{r}_2^{1-\alpha}(\ell)} \langle p_1^\alpha p_2^{1-\alpha} \rangle(\ell),$$

we can write the Rényi-D in terms of the LMB parameters

$$D_R(\pi_1 || \pi_2) = \sum_{\ell \in \mathbb{L}} \frac{\ln [\tilde{r}_1^\alpha \tilde{r}_2^{1-\alpha}(\ell) + r_1^\alpha r_2^{1-\alpha}(\ell) \langle p_1^\alpha p_2^{1-\alpha} \rangle(\ell)]}{\alpha - 1}.$$

It also can be written in terms of the Probability Hypothesis

Density (PHD) $\mathbf{v}_i = r_i p_i = \tilde{r}_i \mathbf{f}_i$,

$$D_R(\boldsymbol{\pi}_1 || \boldsymbol{\pi}_2) = \sum_{\ell \in \mathbb{L}} \frac{\ln \left[1 + \frac{\langle \mathbf{v}_1^\alpha \mathbf{v}_2^{1-\alpha} \rangle(\ell)}{(1 - \langle \mathbf{v}_1 \rangle(\ell))^\alpha (1 - \langle \mathbf{v}_2 \rangle(\ell))^{1-\alpha}} \right]}{\alpha - 1} \\ + \sum_{\ell \in \mathbb{L}} \frac{\alpha \ln [1 - \langle \mathbf{v}_1 \rangle(\ell)] + (1 - \alpha) \ln [1 - \langle \mathbf{v}_2 \rangle(\ell)]}{\alpha - 1},$$

since $\tilde{r}_i(\ell) = 1 - \langle \mathbf{v}_i \rangle(\ell)$, and $\mathbf{f}_i = \mathbf{v}_i / (1 - \langle \mathbf{v}_i \rangle)$.

C. Proof of Kullback-Leibler Divergence (30)

Proof. The Kullback-Leibler Divergence (KL-D) between the LMBs $\boldsymbol{\pi}_1$ and $\boldsymbol{\pi}_2$ is given by

$$D_{KL}(\boldsymbol{\pi}_1 || \boldsymbol{\pi}_2) = \int \Delta(\mathbf{X}) K_1 \mathbf{f}_1^{\mathbf{X}} \ln \frac{K_1 \mathbf{f}_1^{\mathbf{X}}}{K_2 \mathbf{f}_2^{\mathbf{X}}} \delta \mathbf{X} \\ = \int \Delta(\mathbf{X}) K_1 \mathbf{f}_1^{\mathbf{X}} \left(\ln \frac{K_1}{K_2} + \ln \frac{\mathbf{f}_1^{\mathbf{X}}}{\mathbf{f}_2^{\mathbf{X}}} \right) \delta \mathbf{X} \\ = \ln \frac{K_1}{K_2} \int \Delta(\mathbf{X}) K_1 \mathbf{f}_1^{\mathbf{X}} \delta \mathbf{X} \\ + K_1 \int \Delta(\mathbf{X}) \mathbf{f}_1^{\mathbf{X}} \ln \left(\frac{\mathbf{f}_1}{\mathbf{f}_2} \right)^{\mathbf{X}} \delta \mathbf{X} \\ = \ln \frac{K_1}{K_2} + K_1 \int \Delta(\mathbf{X}) \mathbf{f}_1^{\mathbf{X}} \ln \left(\frac{\mathbf{f}_1}{\mathbf{f}_2} \right)^{\mathbf{X}} \delta \mathbf{X} \\ = \ln \frac{K_1}{K_2} + K_1 \sum_{L \subseteq \mathbb{L}} \sum_{\ell \in L} \frac{r_1^{L-\{\ell\}}}{\tilde{r}_1^{L-\{\ell\}}} \left\langle \mathbf{f}_1 \ln \frac{\mathbf{f}_1}{\mathbf{f}_2} \right\rangle(\ell),$$

where the last line follows from Lemma 3 and $\langle \mathbf{f}_1 \rangle = r_1 / \tilde{r}_1$. Since $\mathbf{f}_1(\cdot, \ell) = r_1(\ell) = 0$ for $\ell \notin \mathcal{D}(\sigma_1)$, and $0 \ln 0 = 0$ by convention, we can exchange \mathbb{L} and $\mathcal{D}(\sigma_1)$ in the sum. Further, applying Lemma 4 gives

$$D_{KL}(\boldsymbol{\pi}_1 || \boldsymbol{\pi}_2) \\ = \ln \frac{\tilde{r}_1^{\mathbb{L}}}{\tilde{r}_2^{\mathbb{L}}} + \tilde{r}_1^{\mathbb{L}} \sum_{L \subseteq \mathcal{D}(\sigma_1)} \frac{r_1^L}{\tilde{r}_1^L} \sum_{\ell \in L} \frac{\langle \mathbf{f}_1 \ln \frac{\mathbf{f}_1}{\mathbf{f}_2} \rangle(\ell)}{r_1(\ell) / \tilde{r}_1(\ell)} \\ = \ln \frac{\tilde{r}_1^{\mathbb{L}}}{\tilde{r}_2^{\mathbb{L}}} + \tilde{r}_1^{\mathbb{L}} \sum_{\ell \in \mathcal{D}(\sigma_1)} \frac{\langle \mathbf{f}_1 \ln \frac{\mathbf{f}_1}{\mathbf{f}_2} \rangle(\ell)}{r_1(\ell) / \tilde{r}_1(\ell)} \sum_{L \subseteq \mathcal{D}(\sigma_1) - \{\ell\}} \frac{r_1^{L \cup \{\ell\}}}{\tilde{r}_1^{L \cup \{\ell\}}} \\ = \ln \frac{\tilde{r}_1^{\mathbb{L}}}{\tilde{r}_2^{\mathbb{L}}} + \tilde{r}_1^{\mathbb{L}} \sum_{\ell \in \mathcal{D}(\sigma_1)} \left\langle \mathbf{f}_1 \ln \frac{\mathbf{f}_1}{\mathbf{f}_2} \right\rangle(\ell) \sum_{L \subseteq \mathcal{D}(\sigma_1) - \{\ell\}} \frac{r_1^L}{\tilde{r}_1^L} \\ = \ln \frac{\tilde{r}_1^{\mathbb{L}}}{\tilde{r}_2^{\mathbb{L}}} + \sum_{\ell \in \mathbb{L}} \left\langle \mathbf{f}_1 \ln \frac{\mathbf{f}_1}{\mathbf{f}_2} \right\rangle(\ell) \tilde{r}_1(\ell) \sum_{L \subseteq \mathbb{L} - \{\ell\}} r_1^L \tilde{r}_1^{L-\{\ell\}} \\ = \ln \frac{\tilde{r}_1^{\mathbb{L}}}{\tilde{r}_2^{\mathbb{L}}} + \sum_{\ell \in \mathbb{L}} \tilde{r}_1(\ell) \left\langle \mathbf{f}_1 \ln \frac{\mathbf{f}_1}{\mathbf{f}_2} \right\rangle(\ell) \\ = \sum_{\ell \in \mathbb{L}} \left[\ln \frac{\tilde{r}_1(\ell)}{\tilde{r}_2(\ell)} + \tilde{r}_1(\ell) \left\langle \mathbf{f}_1 \ln \frac{\mathbf{f}_1}{\mathbf{f}_2} \right\rangle(\ell) \right],$$

where the 2nd last equation follows from Lemma 1 and $r_1 + \tilde{r}_1 = 1$. \square

Remark: Using $\tilde{r}_i(\ell) = (1 + \langle \mathbf{f}_i \rangle(\ell))^{-1}$, the above KL-D can be written completely in terms of \mathbf{f}_1 and \mathbf{f}_2 as

$$D_{KL}(\boldsymbol{\pi}_1 || \boldsymbol{\pi}_2) = \sum_{\ell \in \mathbb{L}} \left[\ln \frac{1 + \langle \mathbf{f}_2 \rangle(\ell)}{1 + \langle \mathbf{f}_1 \rangle(\ell)} + \frac{\langle \mathbf{f}_1 \ln \frac{\mathbf{f}_1}{\mathbf{f}_2} \rangle(\ell)}{1 + \langle \mathbf{f}_1 \rangle(\ell)} \right].$$

Further, using $\mathbf{f}_i = r_i p_i / \tilde{r}_i$, we can write

$$\left\langle \mathbf{f}_1 \ln \frac{\mathbf{f}_1}{\mathbf{f}_2} \right\rangle(\ell) \\ = \int \frac{r_1(\ell)}{\tilde{r}_1(\ell)} p_1(x, \ell) \ln \left(\frac{r_1(\ell) \tilde{r}_2(\ell) p_1(x, \ell)}{\tilde{r}_1(\ell) r_2(\ell) p_2(x, \ell)} \right) dx \\ = \frac{r_1(\ell)}{\tilde{r}_1(\ell)} \int p_1(x, \ell) \left[\ln \frac{r_1(\ell) \tilde{r}_2(\ell)}{\tilde{r}_1(\ell) r_2(\ell)} + \ln \frac{p_1(x, \ell)}{p_2(x, \ell)} \right] dx \\ = \frac{r_1(\ell)}{\tilde{r}_1(\ell)} \left[\ln \frac{r_1(\ell) \tilde{r}_2(\ell)}{\tilde{r}_1(\ell) r_2(\ell)} + D_{KL}(p_1(\cdot, \ell) || p_2(\cdot, \ell)) \right].$$

Hence, the KL-D can be written completely in terms of the LMB parameters as follows

$$D_{KL}(\boldsymbol{\pi}_1 || \boldsymbol{\pi}_2) = \sum_{\ell \in \mathbb{L}} \ln \frac{\tilde{r}_1(\ell)}{\tilde{r}_2(\ell)} + r_1(\ell) \left[\ln \frac{r_1(\ell) \tilde{r}_2(\ell)}{\tilde{r}_1(\ell) r_2(\ell)} + D_{KL}(p_1(\cdot, \ell) || p_2(\cdot, \ell)) \right].$$

Alternatively, using $\tilde{r}_i(\ell) = 1 - \langle \mathbf{v}_i \rangle(\ell)$ and $\mathbf{f}_i = \mathbf{v}_i / (1 - \langle \mathbf{v}_i \rangle)$,

$$\tilde{r}_1(\ell) \left\langle \mathbf{f}_1 \ln \frac{\mathbf{f}_1}{\mathbf{f}_2} \right\rangle(\ell) = \left\langle \mathbf{v}_1 \ln \frac{\mathbf{v}_1 (1 - \langle \mathbf{v}_2 \rangle)}{\mathbf{v}_2 (1 - \langle \mathbf{v}_1 \rangle)} \right\rangle(\ell) \\ = \left\langle \mathbf{v}_1 \ln \frac{\mathbf{v}_1}{\mathbf{v}_2} \right\rangle(\ell) - \langle \mathbf{v}_1 \rangle(\ell) \ln \frac{1 - \langle \mathbf{v}_1 \rangle(\ell)}{1 - \langle \mathbf{v}_2 \rangle(\ell)},$$

hence, the KL-D can be written in terms of the PHD as

$$D_{KL}(\boldsymbol{\pi}_1 || \boldsymbol{\pi}_2) \\ = \sum_{\ell \in \mathbb{L}} \ln \frac{1 - \langle \mathbf{v}_1 \rangle(\ell)}{1 - \langle \mathbf{v}_2 \rangle(\ell)} - \sum_{\ell \in \mathbb{L}} \langle \mathbf{v}_1 \rangle(\ell) \ln \frac{1 - \langle \mathbf{v}_1 \rangle(\ell)}{1 - \langle \mathbf{v}_2 \rangle(\ell)} \\ + \sum_{\ell \in \mathbb{L}} \left\langle \mathbf{v}_1 \ln \frac{\mathbf{v}_1}{\mathbf{v}_2} \right\rangle(\ell) \\ = \sum_{\ell \in \mathbb{L}} \left[\ln \frac{1 - \langle \mathbf{v}_1 \rangle(\ell)}{1 - \langle \mathbf{v}_2 \rangle(\ell)} - \langle \mathbf{v}_1 \rangle(\ell) \ln \frac{1 - \langle \mathbf{v}_1 \rangle(\ell)}{1 - \langle \mathbf{v}_2 \rangle(\ell)} \right] \\ + \sum_{\ell \in \mathbb{L}} \left\langle \mathbf{v}_1 \ln \frac{\mathbf{v}_1}{\mathbf{v}_2} \right\rangle(\ell) \\ = \sum_{\ell \in \mathbb{L}} \left[(1 - \langle \mathbf{v}_1 \rangle(\ell)) \ln \frac{1 - \langle \mathbf{v}_1 \rangle(\ell)}{1 - \langle \mathbf{v}_2 \rangle(\ell)} + \left\langle \mathbf{v}_1 \ln \frac{\mathbf{v}_1}{\mathbf{v}_2} \right\rangle(\ell) \right].$$

D. Proof of χ^2 Divergence (32)

Proof. The χ^2 Divergence (χ^2 -D) between the LMBs $\boldsymbol{\pi}_1$ and $\boldsymbol{\pi}_2$ is given by

$$D_{\chi^2}(\boldsymbol{\pi}_1 || \boldsymbol{\pi}_2) = \int \Delta(\mathbf{X}) \frac{(K_1 \mathbf{f}_1^{\mathbf{X}})^2}{K_2 \mathbf{f}_2^{\mathbf{X}}} \delta \mathbf{X} - 1 \\ = \frac{K_1^2}{K_2} \int \Delta(\mathbf{X}) \frac{(\mathbf{f}_1^{\mathbf{X}})^{\mathbf{X}}}{\mathbf{f}_2^{\mathbf{X}}} \delta \mathbf{X} - 1 \\ = \frac{K_1^2}{K_2} \left(1 + \left\langle \frac{\mathbf{f}_1^2}{\mathbf{f}_2} \right\rangle \right)^{\mathbb{L}} - 1$$

$$\begin{aligned}
&= \left[\frac{\tilde{r}_1^2}{\tilde{r}_2} \left(1 + \left\langle \frac{\mathbf{f}_1^2}{\mathbf{f}_2} \right\rangle \right) \right]^{\mathbb{L}} - 1 \\
&= \prod_{\ell \in \mathbb{L}} \frac{\tilde{r}_1^2(\ell)}{\tilde{r}_2(\ell)} \left[1 + \left\langle \frac{\mathbf{f}_1^2}{\mathbf{f}_2} \right\rangle(\ell) \right] - 1,
\end{aligned}$$

where the 3rd last equation follows from Lemma 2. \square

Remark: The above χ^2 -D can be written completely in terms of \mathbf{f}_1 , \mathbf{f}_2 , or the LMB parameters, or the PHDs, as follows

$$\begin{aligned}
D_{\chi^2}(\boldsymbol{\pi}_1 || \boldsymbol{\pi}_2) &= \left[\frac{1 + \langle \mathbf{f}_2 \rangle}{(1 + \langle \mathbf{f}_1 \rangle)^2} \left(1 + \left\langle \frac{\mathbf{f}_1^2}{\mathbf{f}_2} \right\rangle \right) \right]^{\mathbb{L}} - 1 \\
&= \left[\frac{\tilde{r}_1^2}{\tilde{r}_2} \left(1 + \frac{r_1^2 \tilde{r}_2}{\tilde{r}_1^2 r_2} \left\langle \frac{p_1^2}{p_2} \right\rangle \right) \right]^{\mathbb{L}} - 1 \\
&= \left[\frac{(1 - \langle \mathbf{v}_1 \rangle)^2}{1 - \langle \mathbf{v}_2 \rangle} \left(1 + \left\langle \frac{\mathbf{v}_1^2 (1 - \langle \mathbf{v}_2 \rangle)}{\mathbf{v}_2 (1 - \langle \mathbf{v}_1 \rangle)^2} \right\rangle \right) \right]^{\mathbb{L}} - 1.
\end{aligned}$$

E. Proof of Cauchy-Schwarz Divergence (35)

Proof. For the multi-object exponentials $\mathbf{f}_1^{(\cdot)}$, $\mathbf{f}_2^{(\cdot)}$, let

$$\left\langle \Delta \mathbf{f}_1^{(\cdot)}, \Delta \mathbf{f}_2^{(\cdot)} \right\rangle_U = \int \Delta(\mathbf{X}) U^{|\mathbf{X}|} \mathbf{f}_1^{\mathbf{X}} \mathbf{f}_2^{\mathbf{X}} \delta \mathbf{X},$$

where U is the unit of hyper-volume. Applying Lemma 2 gives

$$\begin{aligned}
\left\langle \Delta \mathbf{f}_1^{(\cdot)}, \Delta \mathbf{f}_2^{(\cdot)} \right\rangle_U &= \int \Delta(\mathbf{X}) (U \mathbf{f}_1 \mathbf{f}_2)^{\mathbf{X}} \delta \mathbf{X} \\
&= (1 + \langle U \mathbf{f}_1 \mathbf{f}_2 \rangle)^{\mathbb{L}}.
\end{aligned}$$

Using the above result, the Cauchy-Schwarz Divergence (CS-D) between two LMBs $\boldsymbol{\pi}_1$ and $\boldsymbol{\pi}_2$ can be written as

$$\begin{aligned}
D_{CS}(\boldsymbol{\pi}_1, \boldsymbol{\pi}_2) &= -\ln \frac{\left\langle K_1 \Delta \mathbf{f}_1^{(\cdot)}, K_2 \Delta \mathbf{f}_2^{(\cdot)} \right\rangle_U}{\left\langle K_1 \Delta \mathbf{f}_1^{(\cdot)}, K_1 \Delta \mathbf{f}_1^{(\cdot)} \right\rangle_U^{1/2} \left\langle K_2 \Delta \mathbf{f}_2^{(\cdot)}, K_2 \Delta \mathbf{f}_2^{(\cdot)} \right\rangle_U^{1/2}} \\
&= -\ln \frac{\left\langle \Delta \mathbf{f}_1^{(\cdot)}, \mathbf{f}_2^{(\cdot)} \right\rangle_U}{\left\langle \Delta \mathbf{f}_1^{(\cdot)}, \mathbf{f}_1^{(\cdot)} \right\rangle_U^{1/2} \left\langle \Delta \mathbf{f}_2^{(\cdot)}, \mathbf{f}_2^{(\cdot)} \right\rangle_U^{1/2}} \\
&= -\ln \frac{(1 + \langle U \mathbf{f}_1 \mathbf{f}_2 \rangle)^{\mathbb{L}}}{[(1 + \langle U \mathbf{f}_1 \mathbf{f}_1 \rangle)^{\mathbb{L}}]^{1/2} [(1 + \langle U \mathbf{f}_2 \mathbf{f}_2 \rangle)^{\mathbb{L}}]^{1/2}} \\
&= -\ln \frac{(1 + \langle U \mathbf{f}_1 \mathbf{f}_2 \rangle)^{\mathbb{L}}}{\left(\sqrt{1 + \langle U \mathbf{f}_1 \mathbf{f}_1 \rangle} \right)^{\mathbb{L}} \left(\sqrt{1 + \langle U \mathbf{f}_2 \mathbf{f}_2 \rangle} \right)^{\mathbb{L}}} \\
&= -\ln \left(\frac{(1 + \langle U \mathbf{f}_1 \mathbf{f}_2 \rangle)}{\sqrt{1 + \langle U \mathbf{f}_1 \mathbf{f}_1 \rangle} \sqrt{1 + \langle U \mathbf{f}_2 \mathbf{f}_2 \rangle}} \right)^{\mathbb{L}} \\
&= -\sum_{\ell \in \mathbb{L}} \ln \left(\frac{(1 + \langle U \mathbf{f}_1 \mathbf{f}_2 \rangle)(\ell)}{\sqrt{1 + \langle U \mathbf{f}_1 \mathbf{f}_1 \rangle(\ell)} \sqrt{1 + \langle U \mathbf{f}_2 \mathbf{f}_2 \rangle(\ell)}} \right) \\
&= -\sum_{\ell \in \mathbb{L}} \ln \frac{1 + \langle U \mathbf{f}_1 \mathbf{f}_2 \rangle(\ell)}{\sqrt{1 + \langle U \mathbf{f}_1^2 \rangle(\ell)} \sqrt{1 + \langle U \mathbf{f}_2^2 \rangle(\ell)}}.
\end{aligned}$$

Note that the logarithms in the above sum are well-defined for all $\ell \in \mathbb{L}$. \square

Remark: While the CS-D depends on the unit U , the Bhattacharyya distance $d_B(\boldsymbol{\pi}_1, \boldsymbol{\pi}_2) = D_{CS}(\sqrt{\boldsymbol{\pi}_1}, \sqrt{\boldsymbol{\pi}_2})$ is

invariant to it, because the unit of $\sqrt{\boldsymbol{\pi}_1(\mathbf{X})\boldsymbol{\pi}_2(\mathbf{X})}$ cancels out the unit of $\delta \mathbf{X}$. Noting that $\langle \sqrt{\boldsymbol{\pi}_i}, \sqrt{\boldsymbol{\pi}_i} \rangle = \int \boldsymbol{\pi}_i(\mathbf{X}) \delta \mathbf{X} = 1$, the Bhattacharyya distance between two LMBs is

$$\begin{aligned}
d_B(\boldsymbol{\pi}_1, \boldsymbol{\pi}_2) &= -\ln \langle \sqrt{\boldsymbol{\pi}_1}, \sqrt{\boldsymbol{\pi}_2} \rangle \\
&= -\ln \left\langle \sqrt{K_1 \Delta \mathbf{f}_1^{(\cdot)}}, \sqrt{K_2 \Delta \mathbf{f}_2^{(\cdot)}} \right\rangle \\
&= -\ln \left(\sqrt{K_1 K_2} \left\langle \sqrt{\Delta \mathbf{f}_1^{(\cdot)}}, \sqrt{\Delta \mathbf{f}_2^{(\cdot)}} \right\rangle \right) \\
&= -\ln \left(\sqrt{K_1 K_2} \int \Delta(\mathbf{X}) \sqrt{\mathbf{f}_1^{\mathbf{X}} \mathbf{f}_2^{\mathbf{X}}} \delta \mathbf{X} \right) \\
&= -\ln \left((\sqrt{\tilde{r}_1 \tilde{r}_2})^{\mathbb{L}} \int \Delta(\mathbf{X}) (\sqrt{\mathbf{f}_1 \mathbf{f}_2})^{\mathbf{X}} \delta \mathbf{X} \right) \\
&= -\ln \left((\sqrt{\tilde{r}_1 \tilde{r}_2})^{\mathbb{L}} (1 + \langle \sqrt{\mathbf{f}_1 \mathbf{f}_2} \rangle)^{\mathbb{L}} \right) \\
&= -\ln \left[\sqrt{\tilde{r}_1 \tilde{r}_2} (1 + \langle \sqrt{\mathbf{f}_1 \mathbf{f}_2} \rangle) \right]^{\mathbb{L}} \\
&= -\sum_{\ell \in \mathbb{L}} \ln \sqrt{\tilde{r}_1(\ell) \tilde{r}_2(\ell)} \left[1 + \langle \sqrt{\mathbf{f}_1 \mathbf{f}_2} \rangle(\ell) \right],
\end{aligned}$$

where 3rd last line follows from Lemma 2. This is the Rényi-D with $\alpha = 0.5$.