Finite-Set Statistics and SLAM

Ronald Mahler
Lockheed Martin MS2
Eagan, Minnesota, USA.
ronald.p.mahler@lmco.com

***Abstract*—This tutorial paper describes the elements of a new, practical, Bayes-optimal, and theoretically unified foundation for multisensor-multitarget problems. This foundation, called Finite-Set Statistics (FISST), is the basis for a fundamentally new, Bayes-optimal, and theoretically unified approach to Simultaneous Localization and Mapping (SLAM) and related robotics applications.**

Keywords—finite-set statistics, FISST, multitarget Bayes filter, PHD filter, CPHD filter, random sets, stochastic geometry

I. Introduction

The purpose of this tutorial paper is to describe the elements of a new, practical, Bayes-optimal, and theoretically unified foundation for multisensor-multitarget problems. This foundation, called Finite-Set Statistics (FISST) [11-13], has attracted great worldwide research interest over the last decade because of its applications in multisource-multitarget detection and tracking, sensor management, and information fusion.

In particular, it is the basis for a fundamentally new, Bayes-optimal, and theoretically unified approach to Simultaneous Localization and Mapping (SLAM) and related robotics applications [1, 21-25]. The other contributors to this special issue will describe this work in greater detail. The reader's attention is directed especially to the book [25]. My purpose here is contextual: to provide an overview of FISST and to explain its pertinence for SLAM and similar applications.

As a beginning, consider the SLAM problem at its most general. Multiple moving robots explore an unfamiliar environment without access to GPS or to *a priori* map (terrain, architectural) information. Without human intervention and employing their onboard sensors only, the robots must detect and localize unknown stationary landmarks ("features"), and from these landmarks construct, on-the-fly, a local map of the environment. Then they must situate themselves and each other—as well as any unknown, moving, and possibly noncooperative targets—within this map, again on-the-fly.

Important things to notice are that (1) the landmarks will be unknown and of unknown, varying number; (2) the robots will be unknown and of unknown, varying number; and (3) the sensor measurements—whether generated by robots, landmarks, targets, or clutter—will be varying and of varying number. Additionally, there is no *a priori* way to order the robots, the landmarks, the targets, or the measurements.

Because of this, the mathematical representation of SLAM scenarios using vectors is problematic for a number of reasons. First, vector representation presumes not only that a collection of objects is ordered *a priori*, but that the objects are of *fixed number*. Second, the introduction of *a priori* orderings, into a problem that has none, may unwittingly introduce information extraneous to the problem and—along with it—unrecognized statistical biases.

Third and more crucially, *how does one measure performance—or, more generally, speak of optimal performance?* For example, suppose that we have instrumented a SLAM scenario so that the landmarks are known with some precision. At any given instant a SLAM algorithm will construct a map that is different. Not only the landmarks will be different, but also their number. How can we, in a theoretically defensible manner, measure the degree of deviation between the entity to be estimated (the actual map) and the algorithm's estimate of it? More generally, how can we claim, in a theoretically defensible manner, that the algorithm's estimate is—at least to within the degree of approximation that is assumed—optimal in a Bayesian sense?

The purpose of finite-set statistics is to answer questions such as these in a theoretically rigorous, unified, and yet engineering-practical manner. It consists of the following sequence of conceptual steps [11-13]:

- Formulate multisensor-multitarget problems using the geometrical and mathematically simplest version of point process theory—*random finite set* (RFS) theory.
- Generalize familiar "Statistics 101" concepts—integral, derivative, probability measure, probability density function, to multisensor-multitarget problems. This results in the concepts of set integral and derivative, functional derivative, belief measure, probability generating functional (p.g.fl.), and so on. This "multi-object calculus" includes "turn-the-crank" rules for computing set and functional derivatives: power rules, product rules, chain rules, and so on.
- From formal statistical models of the sensors and targets—including models for sensor noise, probability of detection, field of view (including occlusion), clutter, etc.—create a formal RFS *multisensor-multitarget measurement model*.
- From formal statistical models of target motions—including motion models for individual targets and models for target appearance and disappearance—create a formal RFS *multitarget motion model*.

- From the RFS multisensor-multitarget measurement model, construct a *multisensor-multitarget likelihood function*. This likelihood function is "true" in that it preserves all information in the RFS measurement model, without inadvertently introducing information extraneous to the problem. Its explicit construction requires belief measures and set derivatives.

- From the RFS multitarget motion model, construct a *multitarget Markov transition density*. This Markov density is "true" in that it preserves all information in the RFS motion model without introducing extraneous information. Its construction also requires belief measures and set derivatives.

- From the multisensor-multitarget likelihood function and Markov density, define a *multisensor-multitarget Bayes filter*. Extract information of interest—the number of targets and their states—using *Bayes-optimal multitarget state estimators*.

- Since multisensor-multitarget Bayes filters are computationally intractable in general, derive *approximate multisensor-multitarget filters*. These filters are "principled" in that they are statistical rather than heuristic, and in that they preserve, in approximate form, the information contained in the original RFS measurement and motion models. Construction of these filters requires p.g.fl.'s, functional derivatives, *probability hypothesis densities* (PHDs), and *cardinality distributions*.

- The most widely known and investigated of these approximate filters are the *PHD filter* and its generalization, the *cardinalized PHD* (CPHD) *filter*.

- These filters can then be applied to the SLAM problem, in the manner to be described in Section V-C. In particular, a PHD can itself be interpreted as an "average" SLAM map.

In the remainder of the paper I briefly explain these points. I review single-sensor, single-target Bayes filtering theory in Section II. Multi-object integro-differential calculus is summarized in Section III, and its application to the statistical modeling of multisensor-multitarget systems in Section IV. The multisensor-multitarget Bayes filter is described in Section V, and its PHD/CPHD filter approximations in Section VI. Conclusions are in Section VII.

II. Single-Sensor, Single-Target Theory: Review

In this section I review single-sensor, single-target Bayes filtering theory [12, Chapter 2]: statistical modeling (II-A), the Bayes filter (II-B), and Bayes-optimal state estimation (II-C).

A. Motion and Measurement Models

A *discrete-time nonlinear dynamical system* is described by a target-motion model and a sensor measurement model:

$$\mathbf{X}_{k+1|k} = \varphi_k(\mathbf{x},\mathbf{W}_k) \tag{1}$$

$$\mathbf{Z}_{k+1} = \eta_{k+1}(\mathbf{x},\mathbf{V}_{k+1}) \tag{2}$$

where $\mathbf{V}_{k+1}$ (sensor noise) and $\mathbf{W}_k$ (system or "plant" noise) are random vectors. In the additive-noise model case, simple integral and differential calculus techniques lead to a Markov transition density and a likelihood function of the form

$$f_{k+1|k}(\mathbf{x}\,|\,\mathbf{x}') = f_{\mathbf{W}_k}(\mathbf{x}-\varphi_k(\mathbf{x}')) \tag{3}$$

$$f_{k+1}(\mathbf{z}\,|\,\mathbf{x}) = f_{\mathbf{V}_{k+1}}(\mathbf{z}-\eta_{k+1}(\mathbf{x})) \tag{4}$$

where $\mathbf{z}$ is a measurement in the single-sensor measurement space Z_0 and $\mathbf{x}$ is s state in the single-target state space X_0.

B. Single-Sensor, Single-Target Recursive Bayes Filter

The statistical time-evolution of such a dynamical system is described by a recursive Bayes filter that propagates a probability distribution on the target state $\mathbf{x}$,

$$\ldots \rightarrow f_{k|k}(\mathbf{x}|Z^k) \rightarrow f_{k+1|k}(\mathbf{x}|Z^k) \rightarrow f_{k+1|k+1}(\mathbf{x}|Z^{k+1}) \rightarrow \ldots \tag{5}$$

Here, Z^k: $\mathbf{z}_1,\ldots,\mathbf{z}_k$ is the time-sequence of measurements and the filter is defined by the equations

$$f_{k+1|k}(\mathbf{x}\,|\,Z^k) = \int f_{k+1|k}(\mathbf{x}\,|\,\mathbf{x}')\cdot f_{k|k}(\mathbf{x}'\,|\,Z^k)d\mathbf{x}' \tag{6}$$

$$f_{k+1|k+1}(\mathbf{x}\,|\,Z^{k+1}) = \frac{f_{k+1}(\mathbf{z}_{k+1}\,|\,\mathbf{x})\cdot f_{k+1|k}(\mathbf{x}\,|\,Z^k)}{f_{k+1}(\mathbf{z}_{k+1}\,|\,Z^k)} \tag{7}$$

$$f_{k+1}(\mathbf{z}\,|\,Z^k) = \int f_{k+1}(\mathbf{z}\,|\,\mathbf{x})\cdot f_{k+1|k}(\mathbf{x}\,|\,Z^k)d\mathbf{x}\,. \tag{8}$$

C. Bayes-Optimal State Estimation

Information of interest—target position, velocity, identity—can be extracted from $f_{k|k}(\mathbf{x}|Z^k)$ using a *Bayes-optimal state estimator*. A Bayes-optimal estimator minimizes the posterior expected Bayes risk with respect to a particular cost function ([12], p. 63). Familiar examples are the expected *a posteriori* (EAP) and maximum *a posteriori* (MAP) estimators:

$$\hat{\mathbf{x}}^{EAP}_{k|k} = \int \mathbf{x}\cdot f_{k|k}(\mathbf{x}\,|\,Z^k)d\mathbf{x} \tag{9}$$

$$\hat{\mathbf{x}}^{MAP}_{k|k} = \arg\sup_{\mathbf{x}} f_{k|k}(\mathbf{x}\,|\,Z^k)\,. \tag{10}$$

III. The RFS Multi-Object Calculus

In this section I introduce the elements of the FISST multi-object calculus [12, Chapter 11]: set integrals (III-A), functional, functional derivatives and set derivatives (III-B), random finite sets, belief measures, and multi-object probability distributions (III-C), and probability generating functional and probability hypothesis densities (III-D).

A. Multi-Object Density Functions and Set Integrals

Let u denote the units of measurement of the elements $\mathbf{y}$ in the space Y_0 (a single-sensor measurement space or a single-target state space). Let Y denote an arbitrary finite subset of Y_0, the empty set included. Then a nonnegative function $f(Y)$ of Y is a *multi-object density function* if the units of $f(Y)$ are $u^{-|Y|}$. Given such a density function, its *set integral* is

$$\int f(Y)\delta Y = \sum_{n\geq 0}\frac{1}{n!}\int f(\{\mathbf{y}_1,\ldots,\mathbf{y}_n\})d\mathbf{y}_1\cdots d\mathbf{y}_n\,. \tag{11}$$

B. Functionals, Functional Derivatives, and Set Derivatives

A *functional* $F[h]$ is a real-valued function whose argument is a function $h(\mathbf{y})$. Its *directional derivative* at $h'(\mathbf{y})$ is

$$\frac{\partial F}{\partial h'}[h] = \lim_{\varepsilon \to 0} \frac{F[h + \varepsilon \cdot h'] - F[h]}{\varepsilon}. \tag{12}$$

Setting $h' = \delta_{\mathbf{y}}$ we get the *functional derivative* of F at $\mathbf{y}$:[1]

$$\frac{\delta F}{\delta \mathbf{y}}[h] = \lim_{\varepsilon \to 0} \frac{F[h + \varepsilon \cdot \delta_{\mathbf{y}}] - F[h]}{\varepsilon}. \tag{13}$$

If $Y = \{\mathbf{y}_1, \ldots, \mathbf{y}_n\}$ with $|Y| = n$ then the functional derivative is

$$\frac{\delta F}{\delta Y}[h] = F[h] \qquad (Y = \emptyset) \tag{14}$$

$$\frac{\delta F}{\delta Y}[h] = \frac{\delta^n F}{\delta \mathbf{y}_1 \cdots \delta \mathbf{y}_n}[h] \qquad (Y \neq \emptyset) \tag{15}$$

Let $\phi_F(T) = F[\mathbf{1}_T]$ be a set function defined on any closed subset T of Υ_0. Then its *set derivative* is

$$\frac{\delta \phi_F}{\delta Y}(T) = \frac{\delta F}{\delta Y}[\mathbf{1}_T]. \tag{16}$$

C. Random Finite Sets, Belief Measures, and Multi-Object Probability Distributions

Let Υ be the space of all finite subsets of Υ_0, and let Υ be endowed with the Fell-Matheron "hit-and-miss" topology of stochastic geometry [12, p. 711]. A *random finite subset* (RFS) $\Psi \subseteq \Upsilon_0$ is a random variable on Υ. The *belief measure* (a.k.a. belief-mass function) of Ψ is[2]

$$\beta_\Psi(T) = \Pr(\Psi \subseteq T). \tag{17}$$

The *probability distribution* of Ψ is the multi-object probability density function

$$f_\Psi(Y) = \left[\frac{\delta \beta_\Psi}{\delta Y}(T)\right]_{T=\emptyset}. \tag{18}$$

D. Probability Generating Functionals and PHDs

The probability generating functional (p.g.fl.) of Ψ is

$$G_\Psi[h] = \int h^Y \cdot f_\Psi(Y) \delta Y \tag{19}$$

where $0 \le h(\mathbf{y}) \le 1$ is a test function on Υ_0, where the integral is a set integral, and where $h^Y = 1$ if $Y = \emptyset$ and $h^Y = \prod_{\mathbf{y} \in Y} h(\mathbf{y})$ otherwise. Note that

$$\beta_\Psi(T) = G_\Psi[\mathbf{1}_T]. \tag{20}$$

The *cardinality distribution* of an RFS Ψ is

$$p_\Psi(n) = \frac{1}{n!} \int f(\{\mathbf{y}_1, \ldots, \mathbf{y}_n\}) d\mathbf{y}_1 \cdots d\mathbf{y}_n. \tag{21}$$

The *probability hypothesis density* of an RFS Ψ is defined equivalently by the equations

$$D_\Psi(\mathbf{y}) = \left[\frac{\delta G_\Psi}{\delta \mathbf{y}}[h]\right]_{h=1} = \int f_\Psi(\{\mathbf{y}\} \cup Y) \delta Y. \tag{22}$$

The PHD is a multi-object statistical analog of the single-object concept of an expected value. Its integral $N_\Psi = \int D_\Psi(\mathbf{y}) d\mathbf{y}$ is the expected number of objects in the scene. The value $D_\Psi(\mathbf{y})$ is the *density of objects* in Ψ at the point $\mathbf{y}$.

IV. RFS Statistical Modeling of Multisensor-Multitarget Systems

In this section I summarize RFS models: multitarget motion models (IV-A), multisensor-multitarget measurement models (IV-B), multitarget Markov densities (IV-C), and multisensor-multitarget likelihood functions (IV-D).

A. Multitarget Motion Models

Suppose that at time-step k the targets have state-set $X' = \{\mathbf{x}'_1, \ldots, \mathbf{x}'_{n'}\}$ with $n' = |X'|$. Then in the transition to time-step $k+1$, some targets will persist while others disappear (target "death"); whereas other targets will appear, either generated by the original targets (target "spawning") or independently of them (target "birth"). Therefore, the random state-set $\Xi_{k+1|k}$ at the next time-step will have the form ([12], Chapter 13)

$$\Xi_{k+1|k} = \Xi_k(\mathbf{x}'_1) \cup \ldots \cup \Xi_k(\mathbf{x}'_{n'}) \cup B_k \tag{23}$$

where $\Xi_k(\mathbf{x}')$ is the RFS of targets originating with a target with state $\mathbf{x}'$; and where B_k is the RFS of targets that appear independently of previous targets. In turn,

$$\Xi_k(\mathbf{x}') = \Pi_k(\mathbf{x}') \cup S_k(\mathbf{x}'). \tag{24}$$

Here $S_k(\mathbf{x}')$ is the RFS of new targets spawned by $\mathbf{x}'$. Also, $\Pi_k(\mathbf{x}')$ is an RFS with $|\Pi_k(\mathbf{x}')| \le 1$—that is, $\Pi_k(\mathbf{x}') = \emptyset$ if $\mathbf{x}'$ disappeared, and $\Pi_k(\mathbf{x}')$ is a singleton set if it persisted.

B. Multisensor-Multitarget Measurement Models

Suppose that at time-step $k+1$ the targets have state-set $X_{k+1} = \{\mathbf{x}, \ldots, \mathbf{x}_n\}$ with $n = |X_{k+1}|$. Then the set of collected measurements will have the form (see [12], Chapter 12)

$$\Sigma_{k+1} = \Sigma_{k+1}(X) \cup C_{k+1} \tag{25}$$

where $\Sigma_{k+1}(X)$ is the set of measurements generated by all of the targets, and where C_k is the RFS of measurements that originate with no target (clutter and/or false alarms). Assume that all target-generated measurements originate with single targets. Then this becomes

$$\Sigma_{k+1} = \Sigma_{k+1}(\mathbf{x}_1) \cup \ldots \cup \Sigma_{k+1}(\mathbf{x}_n) \cup C_{k+1} \tag{26}$$

where $\Sigma_{k+1}(\mathbf{x})$ is the RFS of measurements generated by a target with state $\mathbf{x}$. In the simplest case, $|\Sigma_{k+1}(\mathbf{x})| \le 1$—that is, any target generates either a single measurement or no measurement at all. In this case $p_D(\mathbf{x}) = \Pr(\Sigma_{k+1}(\mathbf{x}) \neq \emptyset)$ is the *probability of detection* of a target with state $\mathbf{x}$.

---

[1] This is a heuristic definition. Functional derivatives can be defined rigorously as a type of Radon-Nikodým derivative.

[2] $\beta_\Psi(T)$ is equivalent to the probability measure $p_\Psi(O) = \Pr(\Psi \in O)$ on the Borel-measurable subsets O of the Fell-Matheron topology [12, p. 711].

C. Multitarget Markov Densities

Suppose that we are given a multitarget motion model $\Xi_{k+1|k}$ as in (23). From (17) its belief measure is

$$\beta_{\Xi_{k+1|K}}(S \mid X') = \Pr(\Xi_{k+1|k} \subseteq S \mid X') \cdot \tag{27}$$

From (18), the multi-object probability distribution that uniquely corresponds to this belief measure is

$$f_{k+1|k}(X \mid X') = \frac{\delta\beta_{\Xi_{k+1|k}}}{\delta X}(\varnothing \mid X') \cdot \tag{28}$$

This is the *multitarget Markov transition density* corresponding to the multitarget motion model (23)—see [12], Chapter 13. It is the probability (density) that targets with state-set X' at time-step k will transition to targets with state-set X at time-step k+1. Using the "turn-the-crank" rules of multi-object calculus—see [12], pp. 386-391, and [2]—it is possible to construct explicit formulas for (28).

D. Multisensor-Multitarget Likelihood Functions

Given a multisensor-multitarget measurement model Σ_{k+1} as in (25), from (17) its belief measure is

$$\beta_{\Sigma_{k+1}}(T \mid X) = \Pr(\Sigma_{k+1} \subseteq T \mid X) \cdot \tag{29}$$

From (18), its multi-object probability distribution is

$$f_{k+1}(Z \mid X) = \frac{\delta\beta_{\Sigma_{k+1}}}{\delta Z}(\varnothing \mid X) \cdot \tag{30}$$

This is the *multisensor-multitarget likelihood function* corresponding to the multisensor-multitarget measurement model (25)—see [12], Chapter 12. It is the probability (density) that a measurement-set Z will be generated at time-step k+1 if targets with state-set X are present. Using the "turn-the-crank" rules of multi-object calculus, it is possible to construct explicit formulas for (30).

In the multisensor case, the measurement-set has the form

$$Z = \overset{1}{Z} \cup \ldots \cup \overset{s}{Z} \tag{31}$$

where $\overset{j}{Z}$ denotes the measurement-set collected by the j^{th} sensor. If the sensors are conditionally independent of state,the multisensor-multitarget likelihood function has the form

$$f_{k+1}(Z \mid X) = \overset{1}{f}_{k+1}(\overset{1}{Z} \mid X) \cdots \overset{s}{f}_{k+1}(\overset{s}{Z} \mid X) \tag{32}$$

where $\overset{j}{f}_{k+1}(\overset{j}{Z} \mid X)$ is the multitarget likelihood function for the j^{th} sensor.

V. The Multisensor-Multitarget Bayes Filter

In this section I review the elements of multisensor-multitarget Bayes filtering: the multisensor-multitarget Bayes filter itself (V-A), Bayes-optimal multitarget state estimation (V-B), and the relevance of the multisensor-multitarget Bayes filter to Bayes-optimal SLAM (V-C).

A. The Multisensor-Multitarget Bayes Recursive Filter

This filter propagates a multitarget probability distribution on the multitarget state X ([12], Chapter 14)

$$\ldots \to f_{k|k}(X|Z^{(k)}) \to f_{k+1|k}(X|Z^{(k)}) \to f_{k+1|k+1}(X|Z^{(k+1)}) \to \ldots \tag{33}$$

Here, $Z^{(k)}$: $Z_1,\ldots, Z_k$ is the time-sequence of collected multisensor measurement-sets and X is the multitarget state-set (a finite subset of $\mathfrak{X}_0$). The filter is defined by the following multitarget analogs of (6-8):

$$f_{k+1|k}(X \mid Z^{(k)}) = \int f_{k+1|k}(X \mid X') \cdot f_{k|k}(X' \mid Z^{(k)}) \delta X' \tag{34}$$

$$f_{k+1|k+1}(X \mid Z^{(k+1)}) = \frac{f_{k+1}(Z_{k+1} \mid X) \cdot f_{k+1|k}(X \mid Z^{(k)})}{f_{k+1}(Z_{k+1} \mid Z^{(k)})} \tag{35}$$

$$f_{k+1}(Z \mid Z^{(k)}) = \int f_{k+1}(Z \mid X) \cdot f_{k+1|k}(X \mid Z^{(k)}) \delta X \tag{36}$$

where the indicated integrals are set integrals as in (11).

B. Bayes-Optimal Multitarget State Estimation

As with the single-sensor, single-target Bayes filter of (5-8), one must extract estimates of the quantities of interest. In the multitarget case, it is not only the positions, velocities, identities, etc. of the targets that are of interest, but also their *number*. Bayes-optimal multitarget state estimators determine target state and target number *simultaneously*.

However, there is one difficulty: the multitarget analogs of the EAP and MAP estimators of (9,10) do not exist in general ([12], pp. 494-497). Thus new estimators must be devised. For our purposes, it is enough to describe the closest analog of the MAP estimator, the joint multitarget (JoM) estimator. It is defined by ([12], pp. 498-500):

$$\hat{X}^{JoM}_{k|k} = \arg\sup_X \frac{c^{|X|}}{|X|!} \cdot f_{k|k}(X \mid Z^{(k)}) \tag{37}$$

where $c > 0$ has the same units of measurement as $\mathbf{x}$. Typically, c is chosen to be equal to the accuracy to which the single-target state $\mathbf{x}$ is to be estimated ([12], p. 500).

C. The Multisensor-Multitarget Bayes Filter and SLAM

Consider the simplest case first: a single robot, inserted into an unknown environment which has no moving targets. In this case SLAM is just a multisensor-multitarget detection and tracking problem, in which the coordinate reference frame is not absolute—for example, its origin and orientation is determined by the robot's position and orientation. Using its sensors, the robot detects stationary landmarks and then localizes them. The finite set of landmarks defines the map in the robot's reference frame.

Consequently, the process of recursive map-building is identical to the process of propagating the multitarget probability distribution, as in (33), together with a Bayes-optimal multitarget state estimator as with (37). At time-step k, the multitarget state-estimate $\hat{X}_{k|k}$ is the current estimate—indeed, the *Bayes-optimal*—estimate of the map. The degree to which the estimated map agrees or differs with the actual map can be determined using the optimal sub-pattern assignment (OSPA) metric [27].

Since the multitarget Bayes filter is intractable in general, approximations of it such as the PHD filter—see Section VI—must be employed. This is the approach to SLAM taken in recent publications [1, 21-25]. The PHD $D_{k|k}(\mathbf{x})$ can be regarded as the *average map* at time-step k. The individual landmarks can be estimated from $D_{k|k}(\mathbf{x})$ as follows. Integrate $D_{k|k}(\mathbf{x})$ to get $N_{k|k}$, the expected number of landmarks. Round $N_{k|k}$ off to the nearest integer ν and find the ν largest suprema of $D_{k|k}(\mathbf{x})$. The corresponding states $\mathbf{x}_1,\ldots,\mathbf{x}_\nu$ are an estimate of the landmarks and their number.

When the scene contains moving targets as well as stationary landmarks, matters are only slightly more complex. The moving targets are situated in the relative map specified by the robot and the stationary landmarks.

The greatest degree of complication occurs when there are multiple robots. If they do not communicate with each other, then—from the point of view of any individual robot—the other robots are nothing more than targets. On the other hand, if complete robot-to-robot communication is feasible and allowed, then in principle a single Bayes-optimal map of the environment can be constructed.

VI. Approximate Bayes Multitarget Filters

The multisensor-multitarget Bayes filter of (32-35) is computationally intractable for all but the simplest problems. In this section I summarize principled approximations of this filter: PHD and CPHD filters in the general sense (VI-A), the "classical" PHD and CPHD filters (IV-B), and derivation of these filters using multi-object calculus (IV-C).

A. PHD and CPHD Filters in the General Sense

Approximate multitarget filters are derived by assuming that the distributions $f_{k|k}(X|Z^{(k)})$ and $f_{k+1|k}(X|Z^{(k)})$ have a particular simplified form.

Assume that they are the distributions of *Poisson processes*—i.e., they have the form

$$f(X)=e^{-\mu}\prod_{\mathbf{x}\in X}D(\mathbf{x})=e^{-\mu}\cdot D^X, \qquad \mu=\int D(\mathbf{x})d\mathbf{x} \quad (38)$$

where $D(\mathbf{x})$ is a PHD in the sense of (22). Then the multitarget Bayes filter can be approximated by a filter

$$\ldots\to D_{k|k}(\mathbf{x}|Z^{(k)})\to D_{k+1|k}(\mathbf{x}|Z^{(k)})\to D_{k+1|k+1}(\mathbf{x}|Z^{(k+1)})\to\ldots \quad (39)$$

Even if $f_{k+1|k}(X|Z^{(k)})$ is Poisson this does not mean that $f_{k+1|k+1}(X|Z^{(k+1)})$ is Poisson. So, choose $D_{k+1|k+1}(\mathbf{x}|Z^{(k+1)})$ to be the PHD of $f_{k+1|k+1}(X|Z^{(k+1)})$—that is from (21),

$$D_{k+1|k+1}(\mathbf{x}\mid Z^{(k)})=\int f_{k+1|k+1}(\{\mathbf{x}\}\cup X\mid Z^{(k)})\delta X\,. \quad (40)$$

Similarly, $f_{k+1|k}(X|Z^{(k)})$ is not necessarily Poisson if $f_{k|k}(X|Z^{(k)})$ is Poisson. So choose $D_{k+1|k}(\mathbf{x}|Z^{(k)})$ to be the PHD of $f_{k+1|k}(X|Z^{(k)})$. Any filter (39) arising from this reasoning process is called a *PHD filter in the general sense*. (For PHD filters with general clutter and target-measurement models, see [2]).

The concept of a CPHD filter further generalizes this reasoning. Assume that $f_{k|k}(X|Z^{(k)})$ and $f_{k+1|k}(X|Z^{(k)})$ are the distributions of *independently identically distributed cluster* (i.i.d.c.) *processes*—i.e., they have the form

$$f(X)=|X|!\cdot p(|X|)\cdot\prod_{\mathbf{x}\in X}s(\mathbf{x})=|X|!\cdot p(|X|)\cdot s^X \quad (41)$$

where $p(n)$ is a probability distribution on the number n of targets (the "cardinality distribution") and where $s(\mathbf{x})$ is the *spatial distribution* of the targets. In this case, the multitarget Bayes filter can be approximated by a filter of the form

$$\begin{array}{ccccccc}\ldots\to & s_{k|k}(\mathbf{x}) & \to & s_{k+1|k}(\mathbf{x}) & \to & s_{k+1|k+1}(\mathbf{x}) & \to\ldots\\ & & & & & \uparrow\downarrow & \\ \ldots\to & p_{k|k}(n) & \to & p_{k+1|k}(n) & \to & p_{k+1|k+1}(n) & \to\ldots\end{array} \quad (42)$$

Here, the top row is a filter on the spatial distribution and the bottom row is a filter on the cardinality distribution, as defined in (21). The vertical arrows at the measurement-update step indicate that the two filters are tightly coupled at that transition.

Even if $f_{k+1|k}(X|Z^{(k)})$ is i.i.d.c., $f_{k+1|k+1}(X|Z^{(k+1)})$ is not necessarily i.i.d.c. So, choose $s_{k+1|k+1}(\mathbf{x})$ to be the normalized PHD of $f_{k+1|k+1}(X|Z^{(k+1)})$ and choose $p_{k+1|k+1}(n)$ to be the cardinality distribution of $f_{k+1|k+1}(X|Z^{(k+1)})$. Similarly, $f_{k+1|k}(X|Z^{(k)})$ is not necessarily i.i.d.c. if $f_{k|k}(X|Z^{(k)})$ is i.i.d.c. So choose $s_{k+1|k}(\mathbf{x})$ to be the normalized PHD of $f_{k+1|k}(X|Z^{(k)})$ and $p_{k+1|k}(n)$ to be the cardinality distribution of $f_{k+1|k}(X|Z^{(k)})$. Any filter (41) that results from this line of reasoning is called a *CPHD filter in the general sense*.

B. The "Classical" PHD Filter

The "classical" PHD filter ([6] and Chapter 16 of [12]) is a PHD filter with the following simplifying assumptions: (1) a single sensor; (2) all target motions are independent; (3) measurements are conditionally independent of the target states; (4) the clutter process is Poisson; and (5) target-generated measurements are Bernoulli. By the latter is meant a measurement process that is governed by the following single-target likelihood function:

$$f_{k+1}(Z\mid\mathbf{x})=\begin{cases}1-p_D(\mathbf{x}) & if \quad Z=\emptyset\\ p_D(\mathbf{x})\cdot f_{k+1}(\mathbf{z}\mid\mathbf{x}) & if \quad Z=\{\mathbf{z}\}\\ 0 & if \quad |Z|>1\end{cases} \quad (43)$$

where $p_D(\mathbf{x})$ is the probability of detection and $f_{k+1}(\mathbf{z}|\mathbf{x})$ is the conventional sensor likelihood function.

C. Derivation of PHD and CPHD Filters Using Multitarget Calculus

The explicit measurement-update equations for PHD and CPHD filters ([2,6,8] and Chapter 16 of [12]) are derived using the following procedure (similar reasoning is used for the time-update equations):

Step 1: Derive a closed-form formula for the p.g.fl.

$$F_{k+1}[g,h]=\int h^X\cdot G_{k+1}[g\mid X]\cdot f_{k+1|k}(X\mid Z^{(k)})\delta X \quad (44)$$

where

$$G_{k+1}[g \mid X] = \int g^Z \cdot f_{k+1}(Z \mid X)\delta Z \tag{45}$$

is the p.g.fl. of the multitarget likelihood function $f_{k+1}(Z|X)$. Also, the power notation g^Z was defined in (19). The p.g.fl. (44) typically ends up having the form (see [12], pp. 651-652)

$$F_{k+1}[g,h] = \kappa_{k+1}[g] \cdot G_{k+1|k}[T_h[g]], \tag{46}$$

where $\kappa_{k+1}[g]$ is the p.g.fl. of the clutter process; $G_{k+1|k}[h]$ is the predicted-target p.g.fl.; and $T_h[g]$ is a functional transformation derived from the single-target measurement-generation model (26).

Step 2: Derive formulas for the functional derivatives

$$\frac{\delta F_{k+1}}{\delta Z_{k+1}}[g,h] \tag{47}$$

of $F_{k+1}[g,h]$ with respect to g, where Z_{k+1} is the current measurement-set. Then the posterior p.g.fl. is ([12], p. 651)

$$G_{k+1|k+1}[h] = \frac{\frac{\delta F_{k+1}}{\delta Z_{k+1}}[0,h]}{\frac{\delta F_{k+1}}{\delta Z_{k+1}}[0,1]} \tag{48}$$

and the posterior PHD is then derived from (48) using (22).

Evaluation of (48) usually requires the product rule for functional derivatives (see [12], p. 386) because of the product in (46); and a chain rule for functional derivatives because of the nested p.g.fl. $G_{k+1|k}[T_h[g]]$ in (46).

VII. CONCLUSIONS

Finite-Set Statistics (FISST) is the basis for a fundamentally new, Bayes-optimal, and theoretically unified approach to Simultaneous Localization and Mapping (SLAM) and related robotics applications [1,21-25]. In this tutorial paper I summarized the elements of FISST and briefly described its relevance to Bayes-optimal SLAM (Section V-C).

Further information about FISST can be found in the tutorial [13], the book [12], the book chapter [11], and in the following papers describing specific topics: multisensor PHD and CPHD filters [2,3,5], CPHD filters for superpositional sensors [4,15], jump-Markov CPHD filters [7], PHD filters for extended and unresolved targets [9,10], sensor and platform management [14], CPHD filters for unknown clutter and detection backgrounds [16,17,20], PHD filters for joint sensor-bias estimation and multitarget tracking [19], and multitarget filters for nontraditional information [12, Chapters 3-6], [18].

REFERENCES

[1] D. Clark, "SLAM with single cluster PHD filters," *Proc. 2012 IEEE Int'l Conf. on Robotics and Automation* (ICRA2012), St. Paul, MN, May 14-18, 2012

[2] D. Clark and R. Mahler, "Generalized PHD filters via a general chain rule," *Proc. 15th Int'l Conf. on Information Fusion*, Singapore, July 9-12, 2012.

[3] R. Mahler, "Approximate multisensor CPHD and PHD filters," *Proc. 13th Int'l Conf. on Information Fusion*, Edinburgh, Scotland, July 26-29, 2010.

[4] R. Mahler, "CPHD filters for superpositional sensors," *SPIE Proc.*, vol. 7445, 2009.

[5] R. Mahler, "The multisensor PHD filter, I: General solution via multitarget calculus," *SPIE Proc.* Vol. 7336, 2009.

[6] R. Mahler, "Multitarget filtering via first-order multitarget moments," *IEEE Trans. AES*, vol. 39, no. 4, pp. 1152-1178, 2003.

[7] R. Mahler, "On multitarget jump-Markov filters," *Proc. 15th Int'l Conf. on Information Fusion*, Singapore, July 9-12, 2012.

[8] R. Mahler, "PHD filters of higher order in target number," *IEEE Trans. AES*, vol. 43, no. 4, pp. 1523-1543, 2007.

[9] R. Mahler, "PHD filters for nonstandard targets, I: Extended targets," *Proc. 12th Int'l Conf. on Information Fusion*, Seattle, pp. 915-921, WA, July 6-9, 2009.

[10] R. Mahler, "PHD filters for nonstandard targets, II: Unresolved targets," *Proc. 12th Int'l Conf. on Information Fusion*, pp. 922-929, Seattle, WA, July 6-9, 2009.

[11] R. Mahler, "Random Set Theory for Target Tracking and Identification," Chapter 16 of D.L. Hall and J. Llinas (eds.), *Handbook of Multisensor Data Fusion: Theory and Practice*, Second Edition, CRC Press, Boca Raton FL, 2008.

[12] R. Mahler, *Statistical Multisource-Multitarget Information Fusion*, Boston, Artech House Publishers, 2007.

[13] R. Mahler, "'Statistics 101' for Multisensor, Multitarget Data Fusion," *IEEE Aerospace & Electronics Sys. Mag., Part 2: Tutorials*, vol. 19, no. 1, pp. 53-64, 2004

[14] R. Mahler, "A unified approach to sensor and platform management," *Proc. 2011 Nat'l Symp. on Sensor and Data Fusion*, Washington D.C., October 24-26, 2011.

[15] R. Mahler and A. El-Fallah, "An approximate CPHD filter for superpositional sensors," *SPIE Proc.*, vol. 8392, 2012.

[16] R. Mahler and A. El-Fallah, "CPHD and PHD filters for unknown backgrounds, III: Tractable multitarget filtering in dynamic clutter," *SPIE Proc.*, vol. 7698, 2010.

[17] R. Mahler and A. El-Fallah, "CPHD filtering with unknown probability of detection," *SPIE Proc.*, vol. 7697, 2010.

[18] R. Mahler and A. El-Fallah, "The random set approach to nontraditional measurements is rigorously Bayesian," *SPIE Proc.*, vol. 8392, Baltimore, MD, April 23-27, 2012.

[19] R. Mahler and A. El-Fallah, "Unified Bayesian registration and tracking," *SPIE Proc.*, vol. 8050, Orlando, FL, April 26-28, 2011.

[20] R. Mahler, B.-T. Vo, and B.-N. Vo, "CPHD filtering with unknown clutter rate and detection profile," *Proc. 14th Int'l Conf. on Information Fusion*, Chicago, July 5-8, 2011.

[21] J. Mullane, B.-N. Vo, and M. Adams, "Rao-Blackwellised PHD SLAM," *Proc. 2010 Int'l Conf. on Robotics & Automation* (ICRA2010), Anchorage, Alaska, May 3-8, 2010

[22] J. Mullane, B.-N. Vo, M. Adams, and M. Wijesoma, "A random set approach to SLAM," *Proc. 2009 IEEE Int'l Conf. on Robotics and Automation* (ICRA09), Kobe, Japan, May 12-17, 2009

[23] J. Mullane, B.-T. Vo, M. Adams. and M. Wijesoma, "A random set formulation for Bayesian SLAM," *Proc. 2008 IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems* (IROS2008), Nice, France, pp. 1043 - 1049, Sept. 22-26, 2008.

[24] J. Mullane, B.-N. Vo, M. Adams, and B.-T. Vo, "A random-finite-set approach to Bayesian SLAM," *IEEE Trans. Robotics*, vol. 27, no. 2, pp. 268-282, 2011.

[25] J. Mullane, B.-N. Vo, M. Adams, and B.-T. Vo, *Random Finite Sets in Robotic Map Building and SLAM*, Springer, 2011

[26] S. Nagappa, D. Clark, and R. Mahler, "Incorporating track uncertainty into the OSPA metric," *Proc. 14th Int'l Conf. on Information Fusion*, Chicago, July 5-8, 2011.

[27] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, "A consistent metric for performance evaluation of multi-object filters," *IEEE Trans. SP*, vol. 86 No. 8, pp. 3447-3457, 2008.