RANDOM SETS LLC

Cutting-Edge Information Fusion Innovation

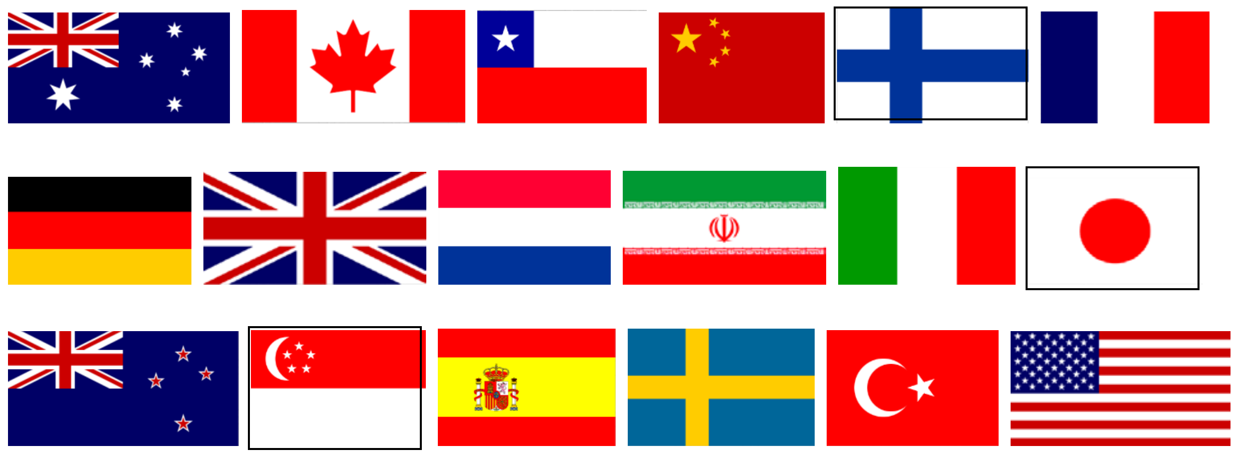

Random Sets, LLC is an independent consultancy established by Ronald Mahler, the founder of the random set (RS) approach to information fusion.

The RS approach has had a major impact on the information fusion field, inspiring a considerable amount of research in at least 19 nations. This research includes a number of novel algorithms that significantly outperform conventional methods. See the position paper that follows.

Random Sets, LLC will address the needs of those seeking expertise in the new RS algorithms, as well as of those interested in RS and other cutting-edge solutions to new information fusion challenges.

Website : http://RandomSets.com
Contact : RandomSets@comcast.net

World Advances in Random Set Information Fusion

Ronald Mahler

1 Introduction

The finite-set statistics (FISST) approach to information fusion was introduced in initial form (random finite sets, belief measures, set derivatives) in the mid-1990s [1]. In its current extended form (probability generating functionals (p.g.fl.'s), functional derivatives), it dates from 2001 [2]. It was first systematically described in 2007 in *Statistical Multisource-Multitarget Information Fusion* [3].

Since 2007, the approach has inspired a considerable amount of research, conducted by many dozens of researchers in at least 19 nations, reported in over a thousand publications. As a result, progress has been rapid and has proceeded in diverse and sometimes unexpected directions, propelled by many clever new ideas. This includes several algorithms that have been shown to significantly outperform traditional methods.

The main purpose of this position paper is to survey the latest and most intriguing aspects of this research, as reported in the 2014 sequel, *Advances in Statistical Multisource-Multitarget Information Fusion* [4], as well as in other publications. It is an update of [36], covering additional advances such as multisensor PHD/CPHD filters, extended-target and cluster-target tracking, and an approach to sensor-network multitarget track-to-track fusion that is fully general, theoretically unified, and immune to unknown double-counting. The paper also refutes several criticisms of FISST.

Besides references [3,4], background information on FISST can be found in the tutorials [5-9], and information on specialized topics in the books [10] (particle filter implementation of RFS filters), and [11] (simultaneous localization and mapping or SLAM). The June 2014 special issue of the *IEEE Robotics & Automation Magazine* is also of interest.

The paper is organized as follows: FISST in a nutshell (Section 2); criticisms of FISST (Section 3); PHD/CPHD filters, including multisensor PHD/CPHD filters (Section 4); the Vo-Vo exact-closed-form multitarget filter (Section 5); unified SLAM (Section 6); unified track-to-track fusion (Section 7), unified multitarget tracking and sensor-bias estimation (Section 8); RFS filters for unknown, dynamic clutter and detection backgrounds (Section 9); track-before-detect (TBD) filters for imaging sensors and superpositional sensors (Section 10); tracking of extended, cluster, and group targets (Section 11); and Bayes-optimal processing of "hard" and "soft" information (Section 12); Conclusions (Section 13).

2 FISST in a Nutshell

The point of finite-set statistics is not that multitarget problems can be formulated in terms of RFSs. The choice of a particular mathematical formalism is of limited practical interest in and of itself. The point is, rather, that FISST techniques provide a carefully constructed practitioner's toolbox of explicit, rigorous, systematic, and general procedures based on systematic statistical multisensor-multitarget modeling and multitarget integro-differential calculus. FISST addresses multisource-multitarget information fusion problems using the following systematic, three-step methodology:

Step 1: Approach information fusion problems in a unified, statistically top-down fashion, by constructing comprehensive statistically accurate models of multitarget-multisensor-multiplatform systems, including top-down, comprehensive statistically accurate models of multitarget sensing and multitarget motion.

Step 2: Use these models to construct the optimal solution to the problem at hand—typically, some kind of multitarget recursive Bayes filter. This necessitates the explicit construction of "true" multitarget Markov densities and "true" multitarget likelihood functions, using multitarget calculus.

Step 3: Since the optimal solution will usually be computationally intractable in general, use principled approximation techniques to "trim down" the optimal solution to an approximate one that is tractable and yet preserves, as faithfully as possible, the underlying models and their interrelationships.

A basic aspect of the FISST approach is a systematic methodology for deriving a *scalable family* of increasingly more accurate approximations of the optimal solutions:

Bernoulli Filter: This is the multitarget Bayes filter when the number of targets is known to be either 0 or 1. It is the optimal approach for single-target detection, tracking, and identification in general clutter and detection backgrounds.

Probability Hypothesis Density (PHD) Filter (Poisson approximation): This is the least accurate approximation of the multitarget Bayes filter. It propagates a PHD $D(\mathbf{x}_k|Z_{1:k})$—the the probability that the scene contains a target with state $\mathbf{x}_k$. Its integral is the expected number of targets in the scene. The graph of the PHD $D(\mathbf{x}_k|Z_{1:k})$ provides an *intensity map* of the targets, with the peaks of $D(\mathbf{x}_k|Z_{1:k})$ corresponding to target states.

Cardinalized PHD (CPHD) Filter (i.i.d.c. approximation): This is a generalization of the PHD filter that propagates not only $D(\mathbf{x}_k|Z_{1:k})$ but also the probability distribution $p(n_k|Z_{1:k})$ on the number n_k of targets (the cardinality distribution). The CPHD filter has better performance than the PHD filter, but is more computationally demanding.

Multi-Bernoulli Filter (multi-Bernoulli approximation): Unlike the PHD and CPHD filters, which compress multitarget distributions into summary statistics, multi-Bernoulli filters attempt to accurately model the multitarget distribution. They often have better performance than CPHD filters, especially when implemented using particle methods.

Generalized Labeled Multi-Bernoulli (GLMB) Filter (GLMB approximation): This is the currently most accurate approximation of the multitarget Bayes filter. It is the first-ever provably Bayes-optimal *and* computationally tractable multitarget detection and tracking filter.

The above family of multitarget filters is scalable in the following sense. For combinatorially complex scenarios, such as clusters of currently-unresolved closely-spaced targets, PHD and CPHD filters are probably the computationally most attractive approaches. For scenarios in which high tracking performance is required but combinatorial complexity is not too great, the GLMB filter and its approximations may be most appropriate.

3 Criticisms of FISST

This section addresses various criticisms of FISST that have been raised in recent years. The statement of each criticism will be followed by a scientific refutation consisting of sequences of factually true statements. (See also [4], pp. 10-15.)

Criticism 1: *"Point processes" are a newer, different, and superior approach to multitarget tracking than RFSs*. False. RFSs are a widely accepted formulation of point process theory—see, for example the book [41] by Kingman. More to the point: one corollary of a well-known theorem ([40], p. 138, Prop. 5.4.V) is this: the *instant that a point process is applied to practical multitarget tracking, it becomes "simple"—i.e., it becomes an RFS*. That is: "point process" alternatives to the FISST RFS approach differ from it only in notation and terminology—and so are, in this sense, copies of it.

Criticism 2: *The FISST multitarget tracking approach, based on p.g.fl.'s and functional derivatives, is a mere "corollary" of a 50-year-old paper* [39] *by the pure mathematician Moyal*. False. It is impossible for the FISST multitarget tracking approach to be a "corollary" of a purely measure-theoretic paper that addressed no applications at all (let alone multitarget tracking)—and which appeared at the same time as the Kalman filter and nearly 20 years before Reid's MHT paper [42].

Reverse-engineering is fundamentally different than engineering. It is easy to know the right things to do—and even easier to claim that these things are actually *obvious*—if someone else has previously shown you how to do it all in complete detail.

Criticism 3: *In particular, the FISST calculus concerns functional differentiation of p.g.fl.'s, where functional differentiation has exactly the same meaning as in* [39]. False. [39] addresses abstract multivariate measures defined in terms of Gâteaux derivatives of p.g.fl.'s; and provides no means of deriving concrete formulas for even those measures—let alone for density functions (as is required for practical application). FISST, however, is based on Volterra's *functional derivative* of p.g.fl.'s, which does produce concrete formulas for density functions [3,4]. Neither functional derivatives nor the term "functional derivative" appear anywhere in [39].

Criticism 4: *The FISST probability generating functional version of Bayes' rule* ([3], Eq. (14.280))

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

is"specialized...for multitarget tracking applications" and thus is not "valid for general point processes." False. This central FISST formula was derived as a general theorem of probability—see [3], Eqs. (G.434-G.438). A general theorem of probability does not become less general because it is subsequently applied to multitarget tracking.

Criticism 5: *RFS filters are defined only for Euclidean state and measurement spaces* R^N. False: In FISST, states and measurements can belong to any Hausdorff, locally compact, completely separable topological space. In particular, in application they typically have the form $(\mathbf{x},\ell)$ where $\mathbf{x} \in \mathrm{R}^N$ and $\ell \in L$ where L is a finite set of labels.

Criticism 6: *Because finite sets are order-independent, RFS filters are inherently incapable of constructing time-sequences of labeled tracks—and therefore are not true tracking filters*. False. Target states can have the non-Euclidean form $(\mathbf{x},\ell)$ where ℓ is an identifying label unique to each track. Thus any RFS filter is, at least in principle, capable of maintaining connected tracks—see pp. 505-508 of [3]. The GLMB filter systematizes this fact.

Criticism 7: *Target identifiability is lost in RFS models*. False. An identifying label ℓ can include target-identity information as well as track-label information, thus permitting target identifiability.

Criticism 8: *Because target identifiability is lost in RFS models, RFS filters require that motion models and likelihood functions are the same for all targets*. False. Because single-target likelihood functions are allowed to have the form $f(\mathbf{z}|\mathbf{x},\ell)$, one can specify a different measurement model for each choice of ℓ. Also: because single-target Markov densities are allowed to have the form $f(\mathbf{x},\ell|\mathbf{x}',\ell')$, one can specify a different Markov density $f(\mathbf{x}|\mathbf{x}',\ell,\ell')$, for each choice of ℓ,ℓ'.

Criticism 9: "*The RFS model of the multiple target state is an approximation, because the Bayes posterior RFS is not exact, but is an approximation based on the earlier invocations of the PHD approximation used to close the Bayesian recursion. The Bayes posterior RFS is an approximation even before the PHD approximation is invoked...*" False. Here it is mistakenly believed that, in FISST, the multitarget RFS is always presumed to be Poisson ("PHD approximation"). The Poisson approximation is merely the simplest and least accurate of the various FISST approximations of the random multitarget state. It is not a "model" of it.

Criticism 10: "*…The right model of the multitarget state is that used in the multi-hypothesis tracker (MHT) paradigm, not the RFS paradigm."* False. If the MHT "model" is the right one, then it should be—as is the case with the RFS model—provably Bayes-optimal. But no proof apparently exists that MHT is Bayes-optimal, or even approximately Bayes-optimal. *An algorithm is not Bayes-optimal merely because it employs Bayes' rule in some fashion*. The term "Bayes optimal" has a specific mathematical meaning. In the multitarget case, it requires the minimization of the multitarget Bayes risk corresponding to some multitarget cost function—see Section 5.3 of [4].

Criticism 11: *The RFS approach is questionable because it is inherently computationally intractable, and requires extreme approximations to make it tractable*. This statement appears to reflect the existence of a double standard. "Ideal" MHT is inherently computationally intractable, and can be made practical only by resort to rather extreme approximations that can severely degrade its performance in (for example) heavy clutter.

Criticism 12: *The RFS approach has not "panned out."* False. Provided that we eschew less generous interpretations, this criticism appears to be attributable to unfamiliarity with the FISST literature. Readers are invited to proceed further and draw their own conclusions.

4 PHD and CPHD Filters

The section is organized as follows: CPHD filters for tracking in heavy, known clutter (Section 3.1); and multisensor PHD and CPHD filters (Section 3.2).

4.1 Tracking in Heavy Clutter

This section illustrates the ability of even less accurate RFS-based algorithms to detect and track multiple appearing and disappearing targets in heavy clutter. The baseline computational complexity of the GM-CPHD filter is at most $O(m^3n)$, where m is the current number of measurements and n is the current number of tracks. By way of contrast, the baseline computational complexity of conventional multitarget trackers is roughly $O(m!n!)$.

In the typical scenario pictured in Figure 1, up to 12 targets appear and disappear while moving along straight-line trajectories. The probability of detection is 0.98 and the clutter rate (the average number of clutter measurements per scan) is 100. The target tracks are essentially invisible in any given frame.

Conventional multitarget trackers will tend to exhibit combinatorial breakdown under such situations. A CPHD filter, however, exhibits reasonable tracking performance—see Figure 1.

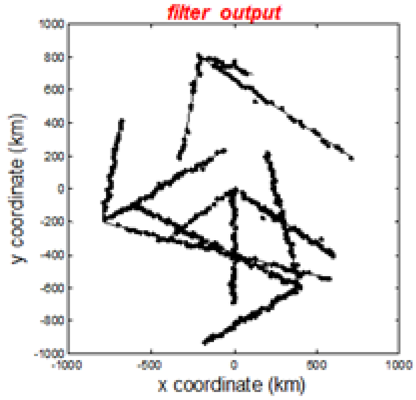

Figure 1*: CPHD filter performance in heavy clutter*

4.2 Multisensor PHD/CPHD Filters

For more details, see Chapter 10 of [4].

The "classical" PHD and CPHD filters are single-sensor filters. Multisensor PHD and CPHD filters exist but are combinatorial. The heuristic "iterated corrector" PHD and CPHD filters are the most commonly-employed approximate multisensor PHD and CPHD filters. They involve repeating the single-sensor measurement-update formulas, once for each sensor in turn. However, this approach produces different answers depending on the order of the sensors. In particular, sensors with larger probabilities of detection should be processed first.

Section 10.6 of [4] describes multisensor PHD and CPHD filters that are principled, computationally tractable, and independent of sensor order.

The basic idea is as follows. Let there be s independent sensors, and let

$$f(Z_k^j \mid X_k)$$

denote the multitarget likelihood function for the jth sensor, where X_k is the multitarget state-set and

$$Z_k^j$$

is the measurement-set generated at time t_k by the jth sensor. Let

$$f(X_k \mid Z_{1:k}^1,\ldots,Z_{1:k}^s)$$

be the multitarget distribution at time t_k, conditioned on all measurements collected from all sensors at that time. Then it can be shown that the total measurement-update is

$$f(X_k \mid Z^1_{1:k},...,Z^s_{1:k}) \propto f(X_k \mid Z^1_k) \cdots f(X_k \mid Z^s_{1:k}) \cdot f(X_k \mid Z^1_{1:k-1},...,Z^s_{1:k-1})^{1-s} \quad (1)$$

where

$$f(X_k \mid Z^1_{1:k-1},...,Z^s_{1:k-1})$$

is the multitarget prior distribution and where

$$f(X_k \mid Z^j_k) = f(X_k \mid Z^1_{1:k-1},...,Z^j_{1:k},...,Z^s_{1:k-1}) \quad (2)$$

is the multitarget distribution "singly-updated" using only the measurement-set collected by the *j*th sensor. In what follows, "i.i.d.c." refers to a particular type of multitarget distribution, one that is used to derive CPHD filters.

There are three different multisensor PHD/CPHD filters, base on three sets of approximations:

Parallel Combination Approximate Multisensor (PCAM) *CPHD Filter*: The prior distribution and all singly-updated distributions are i.i.d.c., and all of the sensor clutter processes are i.i.d.c.

PCAM-PHD Filter: The prior distribution is Poisson, all singly-updated distributions are i.i.d.c., and all of the sensor clutter processes are Poisson.

Simplified PCAM-PHD Filter: The prior distribution, all singly-updated distributions, and all sensor clutter processes are Poisson.

Nagappa and Clark conducted performance comparisons of the following multisensor filters: the iterated-corrector PHD filter; the PCAM-CPHD filter; the PCAM-PHD filter; and the multisensor classical PHD filter (the most accurate but also most computationally intensive alternative). They also compared these with a theoretically erroneous "averaged-likelihood" PHD filter (see Section 10.7 of [4]), which causes target localization accuracy to get increasingly worse (rather than better) as the number of sensors increases.

Both sensors had probability of detection 0.95 and clutter rate 10. The performance of the filters was compared using a metric that accounts for the covariances as well as the means of the track distributions. The results were as follows. As expected, the multisensor classical PHD filter performed significantly better than the others. The PCAM-CPHD filter was next while the PCAM-PHD filter and iterated-corrector PHD filter were roughly comparable. As expected, the averaged-likelihood PHD filter performed worse (see Section 10.8 of [4]).

5 The GLMB Filter

For more detail, see [37,38] and Chapter 15 of [4].

It is well-known that the Kalman filter is an exact-closed-form solution of the single-sensor, single-target Bayes filter, given that both the sensor statistics and target statistics are linear-Gaussian. In this case, the linear-Gaussian distributions are the family of distributions that solves the filter.

In like fashion, B.-T. Vo and B.-N. Vo have discovered a *computationally tractable* exact-closed-form solution of the single-sensor, multitarget Bayes filter, given that the sensor and the targets are described by the usual multitarget tracking models. The family of multitarget distributions that solves the filter are called generalized labeled multi-Bernoulli (GLMB) distributions; and the filter is accordingly called the GLMB filter. Although its computational complexity is similar to that of track-oriented MHT, it is the *first-ever provably Bayes-optimal and computationally tractable multitarget detection and tracking filter*. It can be implemented using particle methods or fast Gaussian-sum methods. A fast approximation of the GLMB filter, called the LMB filter, is of increasing interest [14-16].

The GLMB filter and its approximations are being applied in a number of situations. For example, Vo et al. have addressed the problem of simultaneously detecting and tracking very large numbers of moving targets in three dimensions [15]. In this simulation, 1500 appearing and disappearing targets move along straight-line trajectories in three dimensions. The probability of detection is 0.98 and the clutter rate is moderate: 100 clutter measurements per frame. The filter's estimates are essentially perfect, despite the fact that it was run on an ordinary laptop computer in real time.

A second interesting application is to autonomous automobiles. Reuter, Dietmayer et al. have applied the GLMB filter to this problem, using it to track all moving objects (cars, pedestrians, bicycles, etc.) in the fields of view of the sensors [16,17]. A YouTube video of one of the tests can be found at [18]. The autonomous vehicle is an E-class Mercedes equipped with GPS, cameras, laser rangefinders, and radar. It successfully operates in a variety of dynamic real-world conditions along a 6 km test course that includes a roundabout, traffic lights, a rural road, and urban streets with pedestrian crosswalks.

6 Unified SLAM

Adams, Mullane et al. have demonstrated that FISST-based simultaneous localization and mapping (SLAM) algorithms significantly outperform conventional algorithms such as EKF-SLAM and MH-SLAM, in dense-clutter environments [11,19,20]. Typical results are briefly described here.

In one real-data experiment, the robot is a powerboat moving in a region off the southern coast of Singapore, carrying an X-band radar. The boat and landmarks in the area have been ground-truthed using GPS. Clutter is quite heavy. A conventional MH-FastSLAM algorithm was compared to a PHD-SLAM algorithm, with the result that PHD-SLAM did a significantly better job. This was largely due to the fact that MH-FastSLAM generated a large number of false landmarks, whereas PHD-SLAM accurately estimated the number of landmarks.

7 Unified Track-to-Track Fusion

In measurement-to-track fusion, one collects measurement data and then uses it to improve the accuracy of the most recent estimates of the numbers and states of targets. However, new challenges have arisen because of physically dispersed sensors connected by communications networks. It is often not possible to transmit raw measurements in a timely fashion because transmission links are often bandwidth-limited. Emphasis has therefore shifted to the transmission of track data and to track-to-track fusion. However, track-to-track fusion is fundamentally different than measurement-to-track fusion, because the latter is based on two assumptions: (1) measurements are statistically independent from time-step to time-step; and (2) measurements generated by different sensor sources are statistically independent.

A potential theoretical foundation for unified track-to-track fusion was sketched in [21]. It was based on a version of logarithmic opinion pooling called *exponential-mixture fusion*. Suppose that two sensors are interrogating the same scenario. Suppose that the sensors collect respective measurement-sets Z_1 and Z_2, and that respective multitarget track distributions $f(X|Z_1)$ and $f(X|Z_2)$ are constructed from them. Fuse them as follows:

$$f_\omega(X|Z_1,Z_2)=\frac{f(X|Z_1)^\omega\cdot f(X|Z_2)^{1-\omega}}{\int f(X'|Z_1)^\omega\cdot f(X'|Z_2)^{1-\omega}\,\delta X'}. \qquad (3)$$

Eq. (3) is the multitarget generalization of the well-known covariance-intersection (CI) approach.

In [22], Battistelli et al. adopted this line of investigation to create what appears to be the first-ever fully general, theoretically rigorous and unified approach to track-to-track fusion. It is based on a clever application of the familiar consensus-pooling approach. Battistelli et al. turn the space of multitarget distributions into a linear space, by defining the following concepts of addition and multiplication by a scalar:

$$(f_1\oplus f_2)(X)=\frac{f_1(X)\cdot f_2(X)}{\int f_1(X')\cdot f_2(X')\delta X'} \qquad (4)$$

$$(\omega\odot f)(X)=\frac{f(X)^\omega}{\int f(X')^\omega\delta X'}. \qquad (5)$$

They demonstrated that addition "⊕" and scalar multiplication "⊙" obey the usual rules of vector arithmetic. Suppose that $f_1(X),\ldots, f_s(X)$ are multitarget track distributions supplied by distributed sensor sources. Let $\omega_1,\ldots,\omega_s$ be real-number weights. Then we can fuse the sources with these weightings as:

$$f_{1\ldots s} = (\omega_1\odot f_1)\oplus\ldots\oplus(\omega_s\odot f_s). \qquad (6)$$

Consensus fusion is based on recursively exploiting weighted sums of this form to arrive at a solution that is optimal in the limit. Furthermore, it inherits a very attractive property from exponential-mixture fusion: *it is immune to unknown double-counting of measurements* [22].

In [23], Battistelli et al. devised a "consensus CPHD filter" implementation of the general consensus approach and reported good results.

8 Unified Tracking/Bias Estimation

For more detail, see Chapter 12 of [4].

The reliability of real-world sensors and information sources (such as maps) is often compromised by spatial *biases*. Such misregistrations result in measurements that are not actually located where they appear to be. Conventional bias-estimation approaches are typically limited to a single sensor and require special calibration data sets. However, biases can vary over time; each sensor can have its own bias; and calibration data may not be available.

Recent work using FISST techniques has resulted in the first *statistically unified* multitarget tracking and bias-estimation algorithms. They are based on the fact that "targets of opportunity" can be used to register sensors, even as the biased sensor data can be used to detect and track the targets.

One such algorithm, proposed in [24], is briefly discussed here. The simulation involves four targets moving along curvilinear trajectories, which are observed by three sensors. The first sensor is range-bearing with probabiility of detection 0.9 and clutter rate 50. The second is range-only, with probabiility of detection 0.8 and clutter rate 60. The third is bearing-only with probability of detection 0.7 and clutter rate 40. The joint tracking and registration algorithm not only successfully detected and tracked the targets, but successfully estimated the four unknown translational biases.

9 Unknown, Dynamic Clutter

For more detail, see Chapter 18 of [4].

Essentially all multitarget detection and tracking algorithms, including RFS filters, require an a priori statistical model of the background clutter process. Conventional algorithms, such as MHT, usually (but not always) presume that clutter is spatially uniform and that the probability distribution on the number of measurements (the cardinality distribution) is Poisson. The CPHD filter allows both the clutter spatial distribution and the clutter cardinality distribution to be more general.

In actual application, the statistics of the clutter process will be both unknown and dynamically changing over time. The more that actual clutter statistics differ from modeled clutter statistics, the more that tracker performance will deteriorate.

Recent research has resulted in “clutter-agnostic” CPHD and multi-Bernoulli filters that do not require a priori clutter models [25-28]—and which, moreover, can estimate both the clutter intensity function and the clutter cardinality distribution.

This work has been extended to CPHD and multi-Bernoulli filters that do not require a priori models of the probability of detection. It has also been extended to CPHD filters that can, in addition to clutter estimation, estimate the target-appearance process [29].

In one simulation, a maximum of 10 targets appeared and disappeared in the scene while following curvilinear trajectories. The actual clutter spatial distribution was non-uniform, and the actual clutter cardinality distribution was binomial. The probability of detection was time-varying, ranging from 0.98 to 0.92. Performance results show that the CPHD filter not only successfully tracked the targets, but also successfully estimated the clutter intensity function and clutter cardinality distribution.

10 Track-Before-Detect Filters

Most tracking algorithms process detection data. That is, a detection algorithm is applied to a sensor signature, and point detections are extracted from it. Since this loses information, it is preferable to devise algorithms that can directly process the “raw” signature data. Such algorithms are generically known as “track-before-detect” (TBD) algorithms. Recent research has resulted in RFS multitarget TBD algorithms for two sensor types: imaging (Section 9.1) and superpositional (Section 9.2).

10.1 TBD for Imaging Sensors

For more detail, see Chapter 20 of [4].

Recent multitarget TBD algorithms for pixelized images have been based on Dempster’s expectation-maximization (EM) algorithm. Hoseinezhad et al. have demonstrated that a multi-Bernoulli TBD filter is not only provably Bayes-optimal, but significantly outperformed the previously best EM-based TBD algorithm, the histogram probabilistic multi-hypothesis tracker, H-PMHT [30,31]. Perhaps even more importantly, the computational complexity of this filter does not increase with the number of pixels, whereas that of the H-PMHT algorithm greatly increases.

The simulation scenario described here involves four maneuvering targets of fixed and known number, in a pixelized image with Gaussian background noise. The particle multi-Bernoulli TBD filter presumes that no targets appear or disappear, but is otherwise ignorant of the initial target number. It must determine it, whereas the H-PMHT knows it from the outset.

Since H-PMHT knows the actual target number, it was initially able track the targets better than the RFS filter. Thereafter, the H-PMHT performed considerably worse with respect to all measures. Localization error increased rapidly with time for the H-PMHT, whereas it was essentially constant for the RFS filter. The processing time for H-PMHT was typically greater than that for the RFS filter. It also increased rapidly with the number of pixels in the image, whereas processing time for the RFS filter was essentially flat with respect to pixel number.

The multi-Bernoulli filter has been successfully applied to real data—for example, benchmark streaming videos of soccer and hockey games [31].

10.2 TBD for Superpositional Sensors

For more detail, see Chapter 19 of [4].

A second kind of TBD algorithm is required for *superpositional sensors*, such as acoustic-wave or electromagnetic-wave sensors. In detection-based measurement models, a target can generate at most a single measurement, and any measurement is generated by at most one target. With superpositional sensors, however, the measurement-signature is typically the sum of the signals generated by several targets.

Nannuru et al. have devised a CPHD filter specifically designed for superpositional data [32,33]. They have employed it in two applications: multitarget detection and tracking using a grid of passive-acoustic sensors; and radio-frequency tomography (in which a rectangular array of RF transmitter-receiver pairs is used to interrogate a denied space such as the interior of a building). The superpositional-CPHD filter significantly outperformed conventional Markov chain Monte Carlo (MCMC) techniques while being 32 times (RF tomography) to 80 times (passive-acoustic) faster.

11 Extended Targets

For more detail, see Chapter 21 of [4].

The usual “small-target” model is based on the assumption that targets are neither too near nor too far from the sensor. If a target is near the sensor, however, it may generate multiple detections rather than a single detection. This is usually because multiple measurements are generated by “scatterers” distributed on the target's surface. Such a target is called an extended target.

It is possible to derive filtering equations for a PHD filter designed to track multiple extended targets. In this approach, the measurements generated by a single extended target are modeled as a Poisson process. The measurement-update equation for this filter is combinatorial, since it involves a summation over all partitions of the current measurement-set Z_k. Nevertheless, recent research has shown how certain approximations can render the filter potentially practical.

Here we will briefly describe a "GLO approximation" due to Granström, Lundquist, and Orguner [33,34]. A partition of the measurements will be most "informative" if each of its cells corresponds to the measurement-cluster generated by an actual extended target. In principle, therefore, the non-informative partitions can be neglected. Rather than using a conventional clustering algorithm, Granström et al. proposed what might be called an "*n*-degrees-of-separation" methodology. Define an equivalence relation on the elements of Z_k. as follows. If $\mathbf{z}_1, \mathbf{z}_2 \in Z_k$ then $\mathbf{z}_1 \approx \mathbf{z}_2$ if there is a sequence $\mathbf{w}_1,\ldots,\mathbf{w}_n \in Z_k$ such that $\mathbf{w}_1 = \mathbf{z}_1$ and $\mathbf{w}_1 = \mathbf{z}_2$, and such that, for each $i = 1,\ldots,n-1$, $\mathbf{w}_i$ is "close" to $\mathbf{w}_{i+1}$. The equivalence classes of "$\approx$" are the cells of an informative partition.

This process results in a drastic reduction in the number of partitions that must be considered. For example, suppose that there are four targets, each generating an average of 20 measurements; and let the clutter rate be 50. Then the average number of measurements in a frame is 130, and there are approximately 10^{161} partitions. After applying the GLO approximation, this number is reduced to 27.

Granström et al. have implemented and tested their PHD filter with both simulated and real data. In their baseline simulations, the probability of detection, clutter rate, and target-measurement rate are, respectively, 0.99, 10, and 10. They considered a number of scenarios, including extended targets that cross, travel in parallel, and spawn a new extended target. They have reported good detection and tracking performance.

12 Optimal "Hard + Soft" Fusion

For more detail, see Chapter 22 of [4].

Quantitative evidence ("measurements") $\mathbf{z}$ are mathematically represented as numbers, vectors, and real- or complex-valued functions. Furthermore, $\mathbf{z}$ is mediated by a likelihood function $L_{\mathbf{z}}(\mathbf{x}) = f(\mathbf{z}|\mathbf{x})$. This describes the probability that $\mathbf{z}$ will be observed if an entity with state $\mathbf{x}$ is present. It includes contextual a priori information such as sensor noise statistics and the membership of $\mathbf{x}$ in an ontology of possible target types.

The uncertainty in *qualitative* evidence involves three factors: not just randomness of the observable, but also uncertainty in what is being observed and uncertainty in the modeling of both of these. For example, the simplest qualitative observable is a quantized measurement. The quantitative $\mathbf{z}$ can be observed only as a constraint $\mathbf{z} \in T$ by a quantum T (a subset of the space of observables). General qualitative evidence is similar: $\mathbf{z}$ can be observed only as a qualitative constraint $\mathbf{z} \in \Theta$—that is, as a random constraint (random subset) Θ rather than a fixed constraint T.

The random set Θ is a *generalized measurement*. It can represent both quantitative and qualitative a posteriori information, including attributes, features, natural-language statements, and inference rules. This is because expert-system approaches, such as fuzzy logic and Dempster-Shafer theory, can be used to construct Θ. For example, the natural-language statement "Gustav is near the tower" can be encoded as a fuzzy membership function $g(x,y)$ on the plane. Its RS representation is $\Theta = \{(x,y) : A \leq g(x,y)\}$ where A is a uniformly distributed random number on the unit interval [0,1].

Qualitative measurements are processed in the same manner as quantitative measurements. Generalized measurements Θ are mediated by *generalized likelihood functions* (GLFs) $\rho(\Theta|\mathbf{x})$. Moreover, they can be processed in a *provably Bayes optimal* fashion [4, pp. 779-782] using the generalized Bayes' rule $f(\mathbf{x}|\Theta) \propto \rho(\Theta|\mathbf{x}) \cdot f_0(\mathbf{x})$. Furthermore, $\rho(\Theta|\mathbf{x})$ can be constructed even if the target ontology is qualitative—e.g., if entity types are only imprecisely known or even if an observed entity has never been previously observed. The GLF $\rho(\Theta|\mathbf{x})$ is the fundamental knowledge representation of a posteriori knowledge Θ, as well as of certain aspects of contextual information. The general formula for $\rho(\Theta|\mathbf{x})$ is:

$$\rho(\Theta \mid \mathbf{x}) = \int \mu_\Theta(\mathbf{z}) \cdot f(\mathbf{z} \mid \mathbf{x}) d\mathbf{z} \tag{7}$$

where $\mu_\Theta(\mathbf{z}) = \Pr(\mathbf{z} \in \Theta)$. For example, for the natural-language statement "Gustav is near the tower," $\mu_\Theta(\mathbf{z}) = g(\mathbf{z})$. Or, $\mu_\Theta(\mathbf{z}) = (g \wedge g')(\mathbf{z}) + ½(1-g(\mathbf{z}))$ for a rule $g \Rightarrow g'$ with fuzzy antecedent g and fuzzy consequent g' (where "$\wedge$" denotes fuzzy conjunction).

The effectiveness of the approach has been verified by Bishop and Ristic [35]. The positions of five observers are known, as are the locations of various landmarks such as a wall, a tower, a tree, etc. The observers provide the following:

1. *Observer 1*: "The target is in the field."
2. *Observer 2*: "If the sun is shining then the target is near the pool or the garage."
3. *Observer 3*: "I do not see the target."
4. *Observer 4*: "The target is in front of the tower."
5. *Observer 5*: "The target is at one o'clock."

Note that the second statement is an inference rule.

These statements were processed using a particle-filter implementation of the FISST generalization of Bayes' rule. The unknown target was localized fairly accurately despite the ambiguity of the evidence.

13 Conclusion

This position paper has surveyed some of the most intriguing advances in random set information fusion since 2007. In particular, it has identified several random set algorithms that have been proven to significantly outperform conventional techniques. It also provided refutations of the most common criticisms of FISST.

References

[1] R. Mahler, "Global optimal sensor allocation," *Proc. 9th National Symp. on Sensor Fusion*, Vol. I (Unclassified), pp. 347-366, Naval Postgraduate School, Monterey CA, Mar. 12-14 1996.

[2] R. Mahler, "Multitarget moments and their application to multitarget tracking," *Proc. Workshop on Estimation, Tracking, and Fusion: A Tribute to Y. Bar-Shalom*, pp. 134-166, Naval Postgraduate School, Monterey CA, May 17, 2001. http://www.dtic.mil/ cgibin/GetTRDoc?AD= ADA414365.

[3] R. Mahler, *Statistical Multisource-Multitarget Information Fusion*, Artech House Publishers, Norwood, MA, 2007.

[4] R. Mahler, *Advances in Statistical Multisource-Multitarget Information Fusion*, Artech House Publishers, Norwood, MA, 2014.

[5] R. Mahler, "'Statistics 101' for Multisensor, Multitarget Data Fusion," *IEEE Aerospace & Electronics Sys. Mag*., Part 2: Tutorials, 19(1): 53-64, 2004.

[6] R. Mahler, "'Statistics 102' for multisensor-multitarget tracking," *IEEE J. Selected Topics in Sign. Proc*., 7(3): 376-389, 2013.

[7] B.-N. Vo, B.-T. Vo, and D. Clark, "Bayesian multiple target filtering using random finite sets," Chapter 3 of M. Mallick, V. Krishnamurthy, and B.-N. Vo (eds.), *Integrated Tracking, Classification, and Sensor Management*, Wiley, 2013.

[8] 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.

[9] B. Ristic, B.-T. Vo, B.-N. Vo, and A. Farina, "A tutorial on Bernoulli filters: Theory, implementation, and applications," *IEEE Trans. Sign. Proc*., 61(13): 3406 - 3430, 2012.

[10] B. Ristic, *Particle Filters for Random Set Models*, Springer, New York, 2013.

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

[12] R. Mahler, "Multitarget filtering via first-order multitarget moments," *IEEE Trans. Aerospace and Electronics Systems*, 39(4): 1152-1178, 2003.

[13] R. Mahler, "PHD filters of higher order in target number," *IEEE Trans. Aerospace and Electronic Systems*, 43(4): 1523-1543, 2007.

[14] S. Reuter, B.-T. Vo, B.-N. Vo, and K. Dietmayer, "The labelled multi-Bernoulli filter, *IEEE Trans. Sign. Proc.*, 62(12): 3246-3260, 2014.

[15] B.-T. Vo, B.-N. Vo, S. Reuter, Q. Lam, and K. Dietmayer, "Towards large scale multi-target tracking," *Proc. SPIE*. 9085, 2014.

[16] S. Reuter, "Multi-Object Tracking Using Random Finite Sets," Ph.D. Dissertation, University of Ulm, 2014. Available at http://vts.uni-ulm.de/doc.asp?id=9214.

[17] U. of Ulm autonomous automobile website, https://www.uni-ulm.de/index.php?id=53474.

[18] U. of Ulm autonomous automobile Youtube video, http://youtu.be/R44p69-UHDg.

[19] J. Mullane, B.-N. Vo, M. Adams, and B.-T. Vo, "A random-finite-set approach to Bayesian SLAM," *IEEE Trans. Robotics and Automation*, 27(2): 268-282, 2011.

[20] M. Adams, B.-N. Vo, R. Mahler, and J. Mullane, "SLAM gets a PHD: New concepts in map estimation," *IEEE Robotics & Automation Magazine*, 21(2): 26-37, 2014.

[21] R. Mahler, "Toward a theoretical foundation for distributed fusion," Chapter 8 in D. Hall, M. Liggins II, C.-Y. Chong, and J. Llinas (eds.), *Distributed Data Fusion for Network-Centric Operations*, CRC Press, Boca Raton, 2012.

[22] G. Battistelli, L. Chisci, C. Fantacci, A. Farina, A. Graziano, and R. Mahler, "Distributed fusion of multitarget densities and consensus PHD filters," *Proc. SPIE* 9474, 2015.

[23] G. Battistelli, L. Chisci, C. Fantacci, A. Farina, and A. Graziano, "Consensus CPHD filter for distributed multitarget tracking," *IEEE J. Selected Topics in Sign. Proc*., 7(3): 508-520, 2013.

[24] F. Lian, C. Han, W. Liu, and H. Chen, H, "Joint spatial registration and multi-target tracking using an extended probability hypothesis density filter," *IET Radar, Sonar, and Navigation*, 5(4): 441-448, 2011.

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

[26] R. Mahler, B.-T. Vo, and B.-N. Vo, "CPHD filtering with unknown clutter rate and detection profile," *IEEE Trans. Sign. Proc*., 59(6): 3497-3513, 2011.

[27] R. Mahler and B.-T. Vo, "An improved CPHD filter for unknown clutter backgrounds," *Proc. SPIE* 9091, 2014.

[28] R. Mahler, "CPHD filters for unknown clutter and target-birth processes," *Proc. SPIE* 9091, 2014.

[29] A. El-Fallah, A. Zatezalo, R. Mahler, R. Mehra, and W. Pereira, "Background agnostic CPHD tracking of dim targets in heavy clutter," *Proc. SPIE* 8745, 2013.

[30] R. Hoseinnezhad, B.-N. Vo, and B.-T. Vo, "Visual tracking in background subtracted image sequence via multi-Bernoulli filtering," *IEEE Trans. Sign. Proc*., 61(2): 392-397, 2012.

[31] R. Hoseinnezhad, B.-N.Vo, B.-T. Vo, and D. Suter, "Visual tracking of numerous targets via multi-Bernoulli filtering of image data," *Pattern Recognition*, 45(10): 3625-3635, 2012.

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

[33] S. Nannuru, M. Coates, and R. Mahler, "Computationally-tractable approximate PHD and CPHD filters for superpositional sensors," *IEEE J. Selected Topics in Sign. Proc*., 7(3): 410-420, 2013.

[33] K. Granström, C. Lundquist, and O. Orguner, "Extended target tracking using a Gaussian-mixture PHD filter," *IEEE Trans. Aerospace & Electr. Sys*., 48(4): 3268-3286, 2012.

[34] K. Granström and U. Orguner, "A PHD filter for tracking multiple extended targets using random matrices," *IEEE Trans. Sign. Proc*., 60(11): 5675-5671, 2012.

[35] A. Bishop and B. Ristic, "Fusion of spatially referring natural language statements with random set theoretic likelihoods," *IEEE Trans. Aerospace and Electr. Sys*., 49(2): 932-944, 2013.

[36] R. Mahler, "Recent world advances in random set information fusion," *Proc. 2014 Nat'l Symp. on Sensor and Data Fusion*, Vol. I (Unclassified), Springfield, MA, Oct. 28-31, 2014.

[37] B.-T. Vo and B.-N. Vo, "Labelled random finite setts and multi-object conjugate priors," *IEEE Trans. Sign. Proc.* 61(13): 3460-3475, 2013.

[38] B.-N. Vo, B.-T-. Vo, and D. Phung, "The Bayes multi-target tracking filter," *IEEE Trans. Sign. Proc.*, 62(24): 6554-6567, 2014.

[39] J. Moyal, "The general theory of stochastic population processes," *Acta Mathematica*, pp. 108: 1-31, 1962.

[40] D. Daley and D. Vere-Jones, *An Introduction to the Theory of Point Processes*, First Edition, Springer-Verlag, New York, 1988.

[41] J. Kingman, *Poisson Processes*, Oxford University Press, London, 1992.

[42] D. B. Reid, "An Algorithm for Tracking Multiple Targets," *IEEE Trans. Auto. Contr*., 24(6): 843-854, 1979.