Target Tracking
Le 8: RFS tracking

Gustaf Hendeby and Rickard Karlsson

Div. Automatic Control
Dept. Electrical Engineering
gustaf.hendeby@liu.se,
rickard.g.karlsson@liu.se

LINKÖPING UNIVERSITY

LINKÖPING UNIVERSITY

Summary: lecture 7

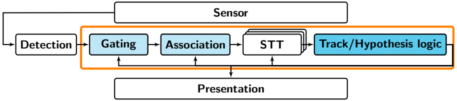

Various methods in MTT

- Performance evaluation using OSPA
- Track to Track fusion
- Track Before Detect
- Extended Target Tracking (ETT)
- Group Tracking

We will now leave the classical MTT for an alternative set representation.

References on Random Finite Set Methods: general

- B.-N. Vo, M. Mallick, Y. Bar-Shalom, S. Coraluppi, R. Osborne, III, R. Mahler, and B.-T. Vo. *Multitarget Tracking.* Wiley Encyclopedia of Electrical and Electronics Engineering, 2015. URL https://www.researchgate.net/publication/283623828_Multitarget_Tracking
- R. P. S. Mahler. Multitarget Bayes filtering via first-order multitarget moments. *IEEE Transactions on Aerospace and Electronic Systems*, 39(4):1152–1178, 2003

References on Random Finite Set Methods: PHD filters

- B.-N. Vo and W.-K. Ma. The Gaussian mixture probability hypothesis density filter. *IEEE Transactions on Signal Processing*, 54(11):4091–4104, 2006. doi: 10.1109/TSP.2006.881190
- D. Fränken, M. Schmidt, and M. Ulmke. "Spooky action at a distance" in the cardinalized probability hypothesis density filter. *IEEE Transactions on Aerospace and Electronic Systems*, 45(4):1657–1664, 2009. ISSN 00189251. doi: 10.1109/TAES.2009.5310327
- G. Hendeby and R. Karlsson. Gaussian mixture PHD filtering with variable probability of detection. In *17th International Conference on Information Fusion (FUSION)*, pages 1–7, 2014

References on Random Finite Set Methods: LMB filters

- S. Reuter, B.-T. Vo, B.-N. Vo, and K. Dietmayer. The labeled multi-Bernoulli filter. *IEEE Transactions on Signal Processing*, 62(12):3246–3260, 2014
- B.-N. Vo, B.-T. Vo, and D. Phung. Labeled random finite sets and the Bayes multi-target tracking filter. *IEEE Transactions on Signal Processing*, 62(24):6554–6567, 2014

References on Random Finite Set Methods: PMBM filters

- J. Williams. Marginal multi-Bernoulli filters: RFS derivation of MHT, JIPDA, and association-based MeMBer. *IEEE Transactions on Aerospace and Electronic Systems*, 51(3):1664–1687, July 2015
- Á. F. García-Fernández, J. L. Williams, K. Granström, and L. Svensson. Poisson multi-Bernoulli mixture filter: Direct derivation and implementation. *IEEE Transactions on Aerospace and Electronic Systems*, 54(4):1883–1901, 2018

Random Finite Set (RFS)

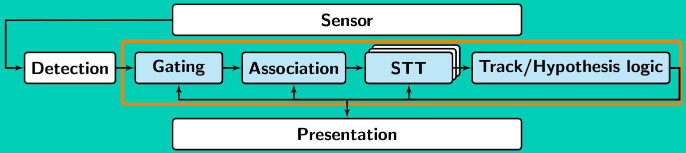

Properties of Classic Tracking Methods

- State of the art: Multi-hypothesis tracker (MHT)
- Combines several more or less independent components:
  - Single target tracking
  - Association
  - Track creation, maintenance, deletion
  - Outliers
- No common mathematical formulation for all components.
- Separate methods for complexity reduction.

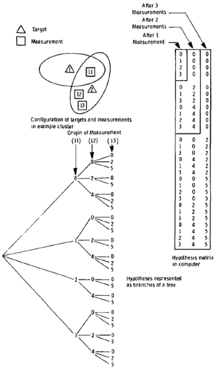

Image: Reid (1979)

Desired Tracking Method Properties

Targets and Measurements

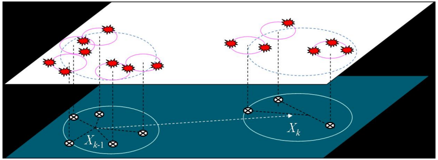

PHDs

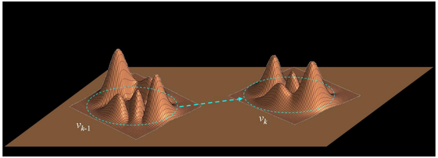

Figures obtained from random set filtering website
`http://randomsets.eps.hw.ac.uk/`

- One unified mathematical formulation.
- Solve the problem with methods that integrate all parts in one.
- Complexity reductions as approximations to mathematical formulation, instead of engineering solutions.

Tracking Assumptions

- Number of targets, unknown and varies over time.
- Origin of observations, unknown.
- Independent measurements (independent measurement noise).
- Independent motion (targets do not influence each other).
- Each target produces at most one observation (point target assumption), and each observation stem from at most one target.
- Measurements can be missed and clutter exist.

Tracking Problem Seen as a Set Problem

- Should describe the same fundamental problem the classic method, and fulfill the same assumptions.
- The targets in a scene, can be seen as a set of tracks with unknown cardinality (number of elements in the set).
- The observations in a scan is a set of measurements.
- **Sets are unordered.** This has implications for track labeling and the association process.
- The task is to estimate the cardinality, and the state of each track, both assumed stochastic, using Bayesian methods.

Solutions Using the Set Formulation

- Random finite sets (RFSs) describes sets of random elements and random cardinality.
- RFS based approaches propagate the posterior density of the multi-target state recursively in time, hence describing the complete tracking problem.
- Several different approximations of this exist:
  - Probability Hypothesis Density (PHD): propagate first moment
  - **Cardinalized PHD (CPHD):** also propagate the cardinality
  - Multi-Bernoulli filters: propagates the parameters of a multi-Bernoulli distribution that approximate the posterior multi-target density
  - Poisson Multi-Bernoulli Mixture (PMBM): propagates unobserved targets as a PHD and observed targets as a MB mixture.
- Labeling: sometimes disregarded in these methods, but can be added (target id).

RFS Preliminaries

LINKÖPING
UNIVERSITY

Random Finite Sets (RFS)

Examples:

- $\emptyset$, $\{3.14\}$, $\{2.7, 9.82\}$, examples of a RFS of real numbers.
- Let $x_k^i \in \mathbb{R}^{n_x}$ for $i = 1, \ldots, \infty$. Then, some realizations X of the random variable $\mathbf{X}$ can be $\emptyset$, $\{x_k^1\}$, $\{x_k^1, x_k^2\}$, $\{x_k^1, x_k^2, x_k^3\}$, $\ldots$.
  Note: $\{x_k^2, x_k^1\}$ is the same as $\{x_k^1, x_k^2\}$ as sets are unordered.
- Y_t all observations obtained in the scan at time t.
- X_t all targets at time t.

Random Finite Sets (RFS)

Examples:

- $\emptyset$, $\{3.14\}$, $\{2.7, 9.82\}$, examples of a RFS of real numbers.
- Let $x_k^i \in \mathbb{R}^{n_x}$ for $i = 1, \dots, \infty$. Then, some realizations X of the random variable $\mathbf{X}$ can be $\emptyset$, $\{x_k^1\}$, $\{x_k^1, x_k^2\}$, $\{x_k^1, x_k^2, x_k^3\}$, $\dots$.
  Note: $\{x_k^2, x_k^1\}$ is the same as $\{x_k^1, x_k^2\}$ as sets are unordered.
- Y_t all observations obtained in the scan at time t.
- X_t all targets at time t.

Definition: random finite set (RFS)

A random finite set $\mathbf{X}$ is a random variable that has realizations in the form $\mathbf{X} = X \in \mathcal{S}$ where $\mathcal{S}$ is the set of all finite subsets of some underlying space $\mathbb{S}$.

Random Sets: connections to normal random stochastic variables

Random fintite sets

- Every element is a *random stochastic variables* (RSV), with a PDF.
- The cardinality is a discrete positive RSV, with a PDF.

Random finite sets properties

- It is possible to compute the probability of a certain instance, $\Pr(X)$.
- The *belief mass* function has the same purpose as the PDF of normal RSV, but does not sum to 1 but instead the cardinality of the RFS.
- It is possible to integrate over RFS $\int p(X)\,\delta X$, with “minor” modifications to how the integral is computed.

Random Sets: connections to normal random stochastic variables

It is possible to define a PDF for the RFS (with a slight abuse of notation):

$$p(\{x^{(1)}, x^{(2)}, \dots, x^{(n)}\}|n) = n! p(x^{(1)}, x^{(2)}, \dots, x^{(n)})$$

Hence, we also need to define integrals over sets.

Definition (Set Integral)

$$\int p(X)\, \delta X = p(\emptyset) + \int p(x^{(1)})\, dx^{(1)} + \frac{1}{2!} \int p(x^{(1)}, x^{(2)})\, dx^{(1)} dx^{(2)} + \frac{1}{3!} \int p(x^{(1)}, x^{(2)}, x^{(3)})\, dx^{(1)} dx^{(2)} dx^{(3)} + \dots$$

Generic RFS Filter

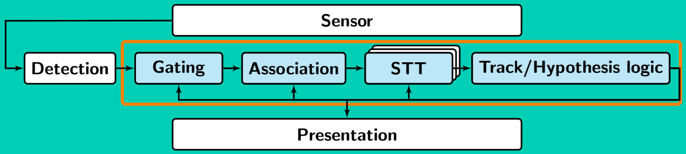

LINKÖPING UNIVERSITY

General Bayesian Filtering Solution

With these tools, we want to approximate the Bayesian filtering solution,

$$p(x_t|\mathbb{Y}_{t-1}) = \int p(x_t|x_{t-1})p(x_{t-1}|\mathbb{Y}_{t-1})\,dx_{t-1} \tag{TU}$$

$$p(x_t|\mathbb{Y}_t) = \frac{p(y_t|x_t)p(x_t|\mathbb{Y}_{t-1})}{p(y_t|\mathbb{Y}_{t-1})}. \tag{MU}$$

General Bayesian Filtering Solution

With these tools, we want to approximate the Bayesian filtering solution,

$$p(x_t|\mathbb{Y}_{t-1}) = \int p(x_t|x_{t-1})p(x_{t-1}|\mathbb{Y}_{t-1})\,dx_{t-1} \tag{TU}$$

$$p(x_t|\mathbb{Y}_t) = \frac{p(y_t|x_t)p(x_t|\mathbb{Y}_{t-1})}{p(y_t|\mathbb{Y}_{t-1})}. \tag{MU}$$

In this case with RFSs as input and states (with slight abuse of notation):

$$p(X_t|\mathbb{Y}_{t-1}) = \int p(X_t|X_{t-1})p(X_{t-1}|\mathbb{Y}_{t-1})\,\delta X_{t-1} \tag{TU-RFS}$$

$$p(X_t|\mathbb{Y}_t) = \frac{p(Y_t|X_t)p(X_t|\mathbb{Y}_{t-1})}{p(Y_t|\mathbb{Y}_{t-1})}. \tag{MU-RFS}$$

In principal the solutions are very similar, but due to technicalities, the implementation is quite different.

RFS Models: dynamic model

- Denote the set of targets X_t at time t.
- X_t encodes the number of targets, and their positions.
- As $x_t = f(x_{t-1}, w_t)$ describes how targets propagate in time,

$$X_t = F(X_{t-1}) \cup W_t$$

  describes how the RFS X_t evolves over time.
- F yields a RFS, in which each existing target in X_{t-1} has been either been propagated using f (implicitly affected by process noise), or died.
- W_t is a RFS with targets born at time t.

RFS Models: measurement model

- Denote the set of measurements, in the scan at time t, Y_t.
- The equivalent of $y_k = h(x_k) + e_k$ is

$$Y_k = H(X_k) \cup V_k$$

  which is the RFS of all measurements procuded at time t.
- H yields a RFS with the measurements are generated from the targets in X_k using $h(x)$ (implicitly assuming measurement noise). A measurement is generated from at target with probability P_{D}, and all measurements are affected by measurement noise.
- The RFS V_k contains all clutter/false measurements.

RFS Filtering: Bayesian solution

- With these definitions, its possible to compute

$$p(X_t|\mathbb{Y}_{t-1}) = \int p(X_t|X_{t-1})p(X_{t-1}|\mathbb{Y}_{t-1})\,\delta X_{t-1} \qquad \text{(TU-RFS)}$$

$$p(X_t|\mathbb{Y}_t) = \frac{p(Y_t|X_t)p(X_t|\mathbb{Y}_{t-1})}{p(Y_t|\mathbb{Y}_{t-1})}. \qquad \text{(MU-RFS)}$$

where a set integral is needed.

- This filter is computationally prohibitive to implement except few cases.

Important difference

The moments, as heavily used in, *e.g.*, Kalman filter,

$$\cancel{\hat{X}_k = \int X_k p(X_k|Y_{0:k})\delta X_k = \mathsf{E}(X_k|Y_{0:k})}$$

is **not** well-defined! (How are elements with different cardinality combined?)

Random Finite Set Descriptions: Poisson point process

RFS can be represented in many different ways, which in turn leads to different approximations, resulting in different filter algorithms.

Poisson Point Process (PPP) RFS

$$p(X) = e^{-\int \lambda(x)\, dx} \Pi_{x \in X} \lambda(x)$$

Typical usage:

- Targets
- Clutter
- Appearing objects
- Measurements from extended targets

Random Finite Set Descriptions: multi-Bernoulli RFS

RFS can be represented in many different ways, which in turn leads to different filter approximations.

Bernoulli RFS

$$p(X) = \begin{cases} 1 - r, & X = \emptyset \\ rp(x), & X = \{x\} \end{cases}$$

- It can be shown that cardinality probability is $\rho(0) = 1 - r$, $\rho(1) = r$, and $\rho(n > 1) = 0$.

Multi-Bernoulli RFS

$$p(\{x_1, \ldots, x_n\}) = \prod_{j=1}^{M} (1 - r^{(j)}) \sum_{1 \leq i_1 \neq \ldots, \neq i_n \leq M} \prod \frac{r^{(i_j)} p^{(i_j)}(x_j)}{1 - r^{(i_j)}}$$

$$\rho(n) = \prod_{j=1}^{M} (1 - r^{(j)}) \sum_{1 \leq i_1 \neq \ldots, \neq i_n \leq M} \prod \frac{r^{(i_j)}}{1 - r^{(i_j)}}$$

- Union of several independent Bernoulli RFS.
- p_i is the belief function for the Bernoulli RFSs.
- $\rho(n)$ is the cardinality probability.

Relations Between PPP and MB

Note

- A multi-Bernoulli RFS can be approximated by a Poisson point process if the existence probabilities are $r_i < 0.1$.
- Any Poisson point process can be approximated by a multi-Bernoulli RFS, but it may require a large number of Bernoulli components, N.

A Poisson point process is often less computationally expensive than a multi-Bernoulli RFS. However, in a Poisson point process both the mean and the variance of the cardinality is λ, hence the uncertainty grows with the number of components.

Probability Hypothesis Density (PHD) Filter

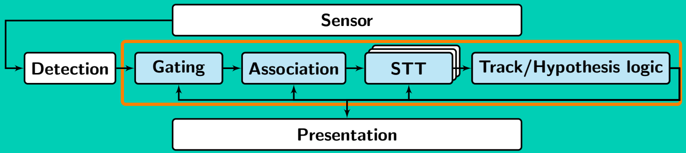

LINKÖPING
UNIVERSITY

Probability Hypothesis Density (PHD)

- Suppose $X_t = \{x_t^{(1)}, x_t^{(2)}, \dots, x_t^{(n)}\}$. Define a scalar valued function from X_t that can be summed,

$$h_{X_t}(x) = \sum_{i=1}^{n} \delta_{x_t^{(i)}}(x).$$

- Then, *the probability hypothesis density* (PHD) is the expectation of $h_{X_t}(x)$ with respect to X_t.

$$\mathsf{PHD}_{t|t}(x) = \mathsf{E}(h_{X_t}(x)|Y_{0:t})$$

Targets and Measurements

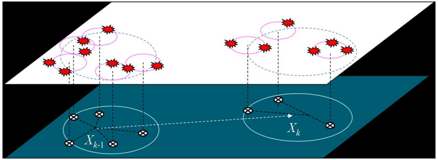

PHDs

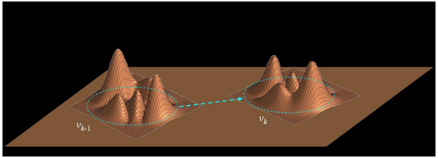

Figures obtained from random set filtering website
`http://randomsets.eps.hw.ac.uk/`

PHD Filter: PHD definition

A usefull definition uses the expected mean of a function (converting sets to vectors) and with $T_{X \cup X'} = T_X + T_{X'}$ if $X \cap X' = \emptyset$.

Commonly used:

$$\mathsf{E}[T] = \int T_X p(X)\, \delta X$$

where the transformation is the Dirac density, so $T_X = \delta_X$, *i.e.*, $\delta_X(x) = 0$ if $X = \emptyset$, otherwise $\delta_X(x) = \sum_{w \in X} \delta_w(x)$.

This definition yields the *probability hypothesis density* (PHD) (intensity function):

$$D(x) = \int \delta_X(x) p(X)\, \delta X.$$

PHD Filter: assumptions

PHD: Model Assumptions

- Motion model PDF: $p_{t+1|t}(x|x')$
- Survival probability for existing targets: $p_{S,t+1|t}(x')$
- Spawning of new targets from existing: $b_{t+1|t}(X|x')$
- Appearance of new targets: $b_{t+1}(X)$
- Probability of detection: p_D
- False alarm model (Poission distribution): probability $c(x)$
- Single target likelihood: $L(x) = p(y|x)$

PHD Filter: filter recursion (1/2)

Actions can be formulated as a RFS, where the possibilities are: survival, spawning of existing target, and spontaneous birth.

$$X_t = \Big[\bigcup_{x \in X_{t-1}} p_{t|t-1}(x'|x) \Big] \bigcup \Big[\bigcup_{x \in X_{t-1}} b_{t|t-1}(x) \Big] \bigcup b_t$$

Given this, its possible to derive the filter recursion with a time and measurement updates.

PHD Filter: filter recursion (2/2)

The PHD filter recursions

$$D_{t+1|t}(x) = \int \Big(\underbrace{p_S(x')p_{t+1|t}(x|x')}_{\text{Survival}} + \underbrace{b_{t+1|t}(x|x')}_{\text{Spawned}}\Big) D_{t|t}(x')\, dx' + \underbrace{b_{t+1}(x)}_{\text{Birth}},$$

$$D_{t|t}(x) = (1 - p_D(x))D_{t|t-1}(x) + \sum_{y \in Y_t} \frac{p_D(x)p(y|x)D_{t|t-1}(x)}{c(y) + \int p_D(x')p(y|x')D_{t|t-1}(x')dx'}.$$

Note:
The first-order moment density (or intensity) is similar to a PDF, but integrates to the number of targets instead of 1!

The GM-PHD Filter

To implement the PHD idea, the inherent exponential complexity must be handled. One way is to use a Gaussian sum filter bank with pruning and merging, yielding the *Gaussian mixture PHD* (GM-PHD).

- Assume p_S and p_D are state-independent (to simplify things).
- Assume that $b_{t+1|t}(x|x')$ and $b_{t+1}(x)$ are Gaussian mixtures.

For details see for instance Vo and Ma (2006) or Hendeby and Karlsson (2014).

The GM-PHD filter uses the following PHD representation

$$v_{t|t}(x) = \sum_{i=1}^{J_{t|t}} w_{t|t}^{(i)} \mathcal{N}(x; m_{t|t}^{(i)}, P_{t|t}^{(i)}), \tag{5}$$

Given that the tracking problem can be modeled using constant $p_{D,t}$, a Gaussian-mixture birth process

$$\gamma_t(x) = \sum_{i=1}^{J_{\gamma,t}} w_{\gamma,t}^{(i)} \mathcal{N}(x; m_{\gamma,t}^{(i)}, P_{\gamma,t}^{(i)}) \tag{6a}$$

and a Gaussian-mixture spawning process

$$\beta_{t|t-1}(x|\zeta) = \sum_{i=1}^{J_{\beta,t}} w_{\beta,t} \mathcal{N}(x; F_{\beta,t-1}^{(i)}\zeta + d_{\beta,t-1}^{(i)}, Q_{\beta,t-1}^{(i)}), \tag{6b}$$

the PHD update are give by the following expressions.

Starting with the filtered PHD, $v_{t-1|t-1}$ in (5) the PHD time update is given by

$$v_{t|t-1}(x) = v_{S,t|t-1}(x) + v_{\beta,t|t-1}(x) + \gamma_t(x), \tag{7}$$

where $v_{S,t|t-1}(x)$ is the PHD of surviving target, $v_{\beta,t|t-1}(x)$ the PHD of new targets spawned from existing ones in this time update, and $\gamma_t(x)$ the PHD of newly born targets, as defined below.

The surviving target PHD is given by

$$v_{S,t|t-1}(x) = \sum_{i=1}^{J_{t-1|t-1}} w_{S,t}^{(i)} \mathcal{N}(x; m_{S,t}, P_{S,t}), \tag{8a}$$

where

$$w_{S,t}^{(i)} = p_{S,t-1}^{(i)} w_{t-1|t-1}^{(i)} \tag{8b}$$

$$m_{S,t}^{(i)} = F_{t-1} m_{t-1|t-1}^{(i)} \tag{8c}$$

$$P_{S,t}^{(i)} = F_{t-1} P_{t-1|t-1}^{(i)} F_{t-1}^T + Q_{t-1}, \tag{8d}$$

and $p_{S,t-1}$ is the probability that a target in time $t-1$ survives to time t.

The spawned target PHD is given by

$$v_{\beta,t|t-1}(x) = \sum_{i=1}^{J_{t-1|t-1}} \sum_{\ell=1}^{J_{\beta,t}} w_{\beta,t}^{(i,\ell)} \mathcal{N}(x; m_{\beta,t}^{(i,\ell)}, P_{\beta,t}^{(i,\ell)}) \tag{9a}$$

where

$$w_{\beta,t}^{(i,\ell)} = w_{t-1|t-1}^{(i)} w_{\beta,t}^{(\ell)} \tag{9b}$$

$$m_{\beta,t}^{(i,\ell)} = F_{\beta,t-1}^{(\ell)} m_{t-1|t-1}^{(i)} + d_{\beta,t-1}^{(\ell)} \tag{9c}$$

The GM-PHD Filter: example

Illustration of PHD at time: k

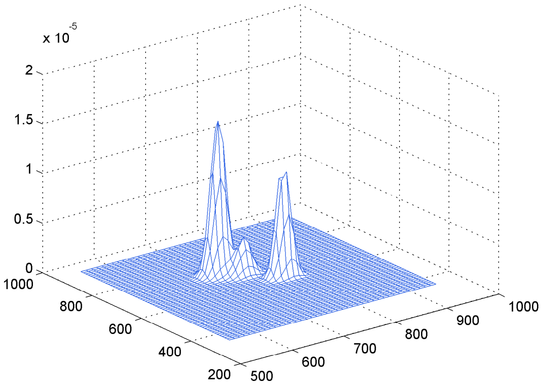

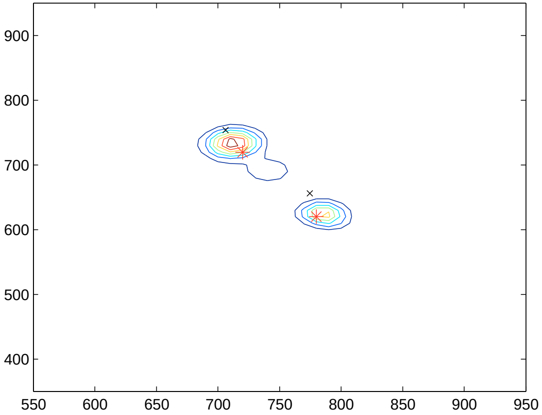

Note: The expected mean over the intensity function sums to the number of targets.

The GM-PHD Filter: example

Illustration of PHD at time: $k + 1$

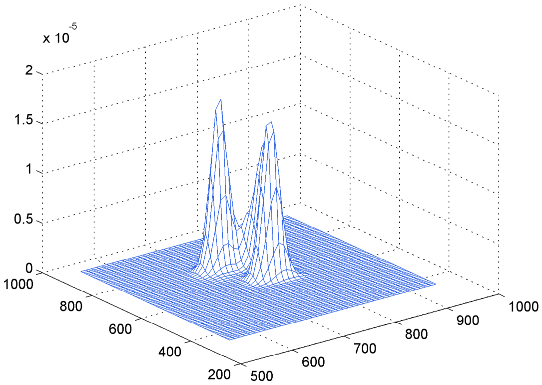

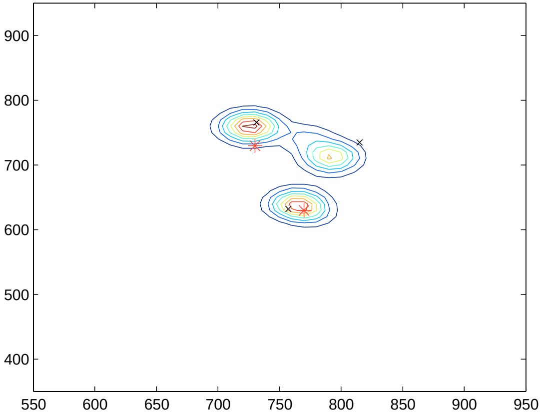

Note: The expected mean over the intensity function sums to the number of targets.

Remarks and Extensions

There exist several approximation/implementation ideas as well as other important issues for multi-target tracking using PHD:

- Birth at a fixed source (can be relaxed).
- No association, but pruning and mixing.
- There exist other PHD implementations, for instance one based on the particle filter.
- The cardinalized PHD (CPHD):
  We want second order approximation, but instead of a full implementation try to improve upon existing PHD. The CPHD-filter will not fully model the second order moment but instead it propagates both the intensity $D_{t|t}(x)$ and the cardinal distribution $p_{t|t}(n)$. Hence, it propagates the entire probability density for the number of targets.
- PHD with track labeling.

Efficient Gaussian Sum Implementations

Given a Gaussian sum assumption of the PHD, this boils down to more or less a tracking filter where all associations are attempted and then merged in each time step. (Without considering target identity.)

PHD Filter Example: Tracking divers

- Example of tracking divers using sonar.
- Modified GM-PHD to handle varying p_{D},

$$p_{\mathrm{D}}(x) = 0.9 - 4{\cdot}10^{-7}R^2 - 1.6{\cdot}10^{-9}R^3,$$

- Fairly high clutter.
- For details: Hendeby and Karlsson (2014)

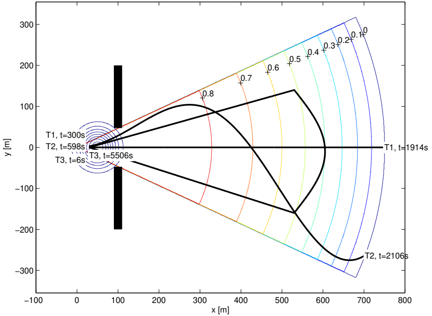

PHD Filter Example

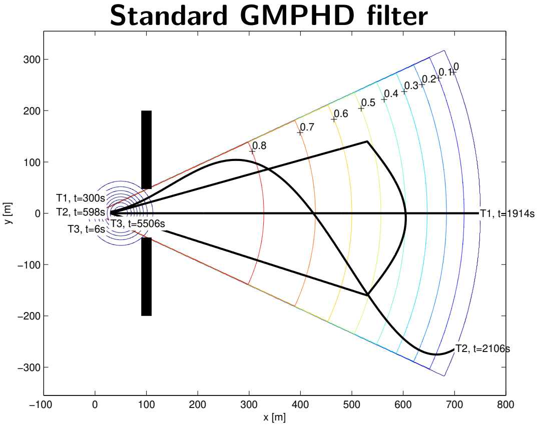

https://youtu.be/PJimgDB3X88

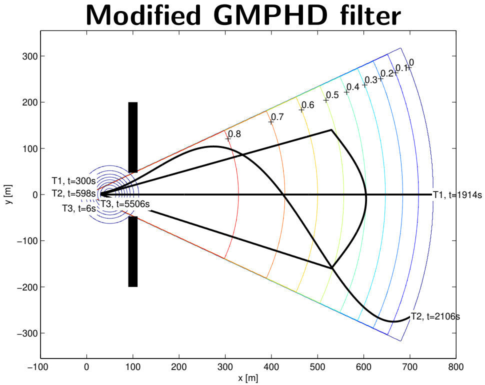

https://youtu.be/cAL3ynfVZAk

- Probability of detection dies off as a 3rd-degree polynomial, inspired by real data.

PHD Filter Example

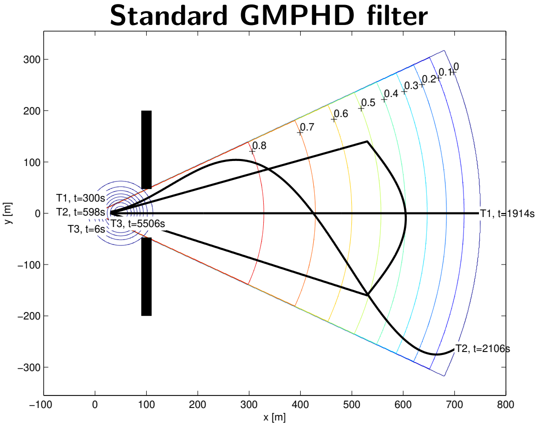

https://youtu.be/PJimgDB3X88

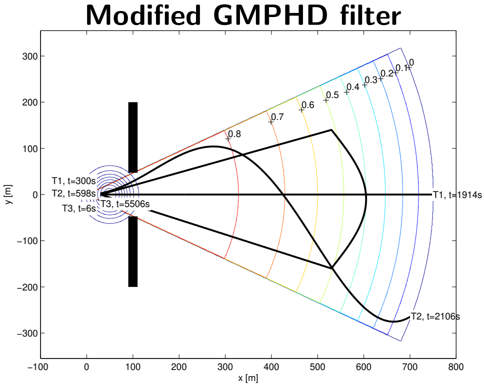

https://youtu.be/cAL3ynfVZAk

- Probability of detection dies off as a 3rd-degree polynomial, inspired by real data.

Cardinalized Probability Hypothesis Density (CPHD)

- In the PHD filter the number of target are captured by integrating the intensity function.
- The PHD can be seen as a first order moment filter.
- The CPHD is the equivalent to second order moment.
- CPHD: the cardinality is explicitly estimated in combination with the intensity.
- Motivations: why use CPHD instead of PHD?
  - The assumption of Poisson target cardinality makes the PHD sensitive to clutter.
  - "Spooky action at a distance" (Fränken et al., 2009):
    Missed measurements shifts the PHD from unrelated areas to detected parts.

The last years other methods are more important, for instance Multi-Bernoulli filters, and Poisson multi-Bernoulli mixture filters.

Multi-Bernoulli Filters

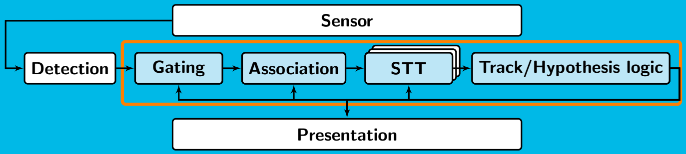

Labeled Multi-Bernoulli Representation

- Bernoulli RFS:

$$p(X) = \begin{cases} 1 - r & X = \emptyset \\ rp(x) & X = \{x\} \end{cases}$$

- Multi-Bernoulli RFS

$$p(\{x_1, \ldots, x_n\}) = \prod_{j=1}^{M} (1 - r^{(j)}) \sum_{1 \leq i_1 \neq \ldots, \neq i_n \leq M} \prod \frac{r^{(i_j)} p^{(i_j)}(x_j)}{1 - r^{(i_j)}}$$

$$\rho(n) = \prod_{j=1}^{M} (1 - r^{(j)}) \sum_{1 \leq i_1 \neq \ldots, \neq i_n \leq M} \prod \frac{r^{(i_j)}}{1 - r^{(i_j)}}$$

- Labeled Multi-Bernoulli RFS
  Add a unique label ℓ to the state, and the RFS becomes $\{(r^{(\ell)}, x^{(\ell)})\}_{\ell \in \mathcal{L}}$.

Labeled Multi-Bernoulli: algorithm (1/2)

Time update

$$p_+(X_+) = \int f(X_+)p(X)\delta X$$

$$p_+ = \{(r^{(\ell)}_{+,S}, p^{(\ell)}_{+,S}\}_{\ell\in\mathcal{L}} \bigcup \{(r^{(\ell)}_B, p^{(\ell)}_B)\}_{\ell\in\mathcal{B}}$$

where

$$r_{+,S} = \eta_S(\ell) r^{(\ell)}$$

$$p^{(\ell)}_{+,S} = \frac{\int p_S(\xi,\ell) f(x|\xi,\ell) p(\xi|\ell)\, d\xi}{\eta_S(\ell)}$$

$$\eta_S(\ell) = \int p_S(\xi,\ell) p(\xi|\ell)\, d\xi$$

$\{(r^{(\ell)}_B, p^{(\ell)}_B)\}_{\ell\in\mathcal{B}}$ is a birth model.

Labeled Multi-Bernoulli: algorithm (2/2)

Measurement update

$$
\begin{aligned}
p(\cdot|Z) &\approx \bigcup_{i=1}^{N} \{(r^{(\ell,i)}, p^{(\ell,i)})\}_{\ell\in\mathcal{L}_+^{(i)}} \\
r^{(\ell,i)} &= \sum_{(I_+,\theta)\in\mathcal{F}(\mathcal{L}_+^{(i)})\times\Theta_{I_+}} \omega^{(I_+,\theta)}(Z^{(i)})\mathbf{1}_{I_+}(\ell) \\
p^{(\ell,i)}(x) &= \frac{1}{r^{(\ell,i)}} \sum_{(I_+,\theta)\in\mathcal{F}(\mathcal{L}_+^{(i)})\times\Theta_{I_+}} \omega^{(I_+,\theta)}(Z^{(i)})\mathbf{1}_{I_+}(\ell)p^{(\theta)}(x,\ell|Z^{(i)}) \\
\omega^{(I_+,\theta)}(Z^{(i)}) &\propto \omega_{+,i}^{(I_+)}[\eta_{Z^{(i)}}^{(\theta)}]^{I_+} \\
&= \prod_{\ell\in\mathcal{L}_+^{(i)}-I_+} (1-r_+^{\ell}) \prod_{\ell'\in I_+^a} r_+^{(\ell')}\eta_{Z^{(i)}}^{(\theta)}(\ell') \prod_{\ell''\in I_+^n} r_+^{(\ell'')}\eta_{Z^{(i)}}^{(\theta)}(\ell'')
\end{aligned}
$$

Observation

The labeled Multi-Bernoulli filter is similar to the MHT

- All potential targets can be handled separately, extra book keeping for existence probabilities etc.
- The δ*-generalized labeled multi-Bernoulli* ($\delta - GLMB$) filter, which essentially is a TO-MHT.

Poisson Multi-Bernoulli Mixture (PMBM) Filter

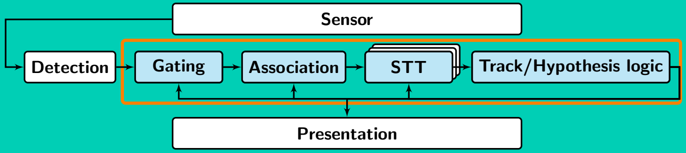

Poisson Multi-Bernoulli Mixture (PMBM) Filter

- Relatively new, Williams (2015).
- Can be considered the new state-of-the-art MTT method. (This can be debated.)
- Has one description of unknown and unobserved targets.
- Has one description of observed targets.
- The PMBM is a conjugate prior, hence suitable as a recursive filter.

PMBM: unobserved targets

- Maintain a representation of unobserved targets.
- "P" in PMBM is a Poisson point process for unobserved targets.
- Often handled as a Gaussian mixture PHD, much similar to a GM-PHD filter but removing detected targets.
- Alternatives to GM-PHD exists, *e.g.*, Boström-Rost et al. (2021).

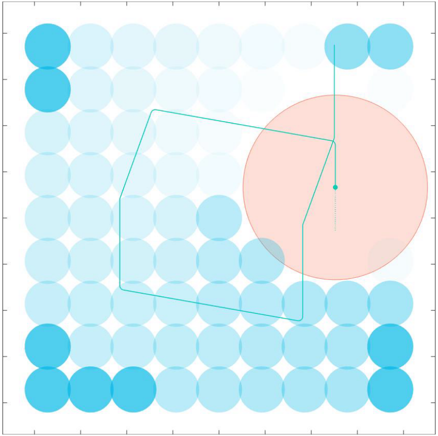

https://youtu.be/e0ebz08vgTo

PMBM: unobserved targets

- Maintain a representation of unobserved targets.
- "P" in PMBM is a Poisson point process for unobserved targets.
- Often handled as a Gaussian mixture PHD, much similar to a GM-PHD filter but removing detected targets.
- Alternatives to GM-PHD exists, *e.g.*, Boström-Rost et al. (2021).

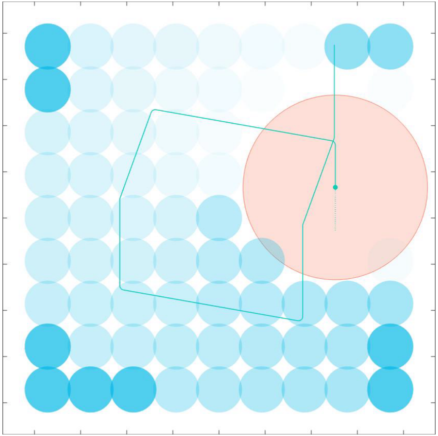

https://youtu.be/e0ebz08vgTo

PMBM: observed targets

- Observed targets are extracted from the unobserved PPP and handled separately.
- "MBM" in PMBM is a multi-Bernoulli mixture for observed targets.
- Basically a δ-*generalized LMB* (δ-GLMB)filter, which is a type of LMB with preferable properties.
- Can be efficiently implemented using a TO-MHT.

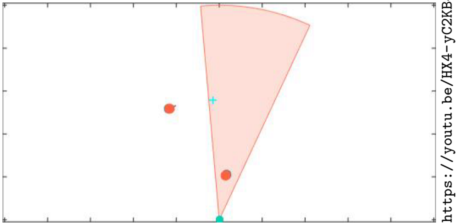

PMBM: observed targets

- Observed targets are extracted from the unobserved PPP and handled separately.
- "MBM" in PMBM is a multi-Bernoulli mixture for observed targets.
- Basically a δ-*generalized LMB* (δ-GLMB)filter, which is a type of LMB with preferable properties.
- Can be efficiently implemented using a TO-MHT.

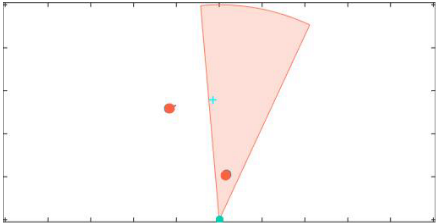

https://youtu.be/HX4-yC2KB4

PMBM Algorithm

- Conceptually:
  - Run a PHD filter too keep track of unobserved targets.
  - Run a TO-MHT too keep track of observed targets.
  - When an unobserved target is observed, remove it from PHD filter, and insert into the MHT again.
  - Weights are update to make this correct.
- For details see, *e.g.*, Williams (2015); García-Fernández et al. (2018)

Summary

- *Random finite set* (RFS) represents a "new" view of target tracking, but are mostly very similar to classic methods.
- Many different approximations, resulting in different filters:
  - *Probability hypothesis distribution* (PHD) filter
  - Multi-Bernoulli distribution filters
  - *Poisson multi-Bernoulli mixture* (PMBM) filter
- Utilize a random set formulation to include the full MTT problem in a single mathematical formulation.

Course Summary

- Focus on classic multi-target tracking problems, with outlooks to common extensions and the RFS formulation.
- **Examination** (how many intend to get credits?):
  - Exam, 2 ETCS credits:
    Take home exam, 1 h exam. Tests the understanding of the principles discussed in the course.

    Due: End of January (contact us to schedule the exam).
  - Exercises, 4 ETCS credits:
    Show hands on experience of important discussed methods.

    Due: December 19, 2021
  - Project, 3 ETCS credits:
    More advanced utilization of MTT techniques, preferably related to your research.

    Discuss details with us.

Gustaf Hendeby and Rickard Karlsson

www.liu.se

LINKÖPING
UNIVERSITY