Target Tracking
Le 8: RFS tracking

Gustaf Hendeby

Div. Automatic Control
Dept. Electrical Engineering
gustaf.hendeby@liu.se

LINKÖPING UNIVERSITY

LINKÖPING UNIVERSITY

Summary: lecture 7

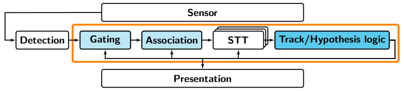

Various methods in MTT

- 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

- 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
- 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
- G. Hendeby and R. Karlsson. Gaussian mixture PHD filtering with variable probability of detection. In *17th International Conference on Information Fusion*, Salamanca, Spain, 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

Classic Multi-Target Tracking: properties

- State of the art: Multi-hypothesis tracker (MHT)
- Combines several more or less independent components:
  - Single target tracking
  - Association
  - Track logic: track creation, maintenance, deletion
  - Outliers

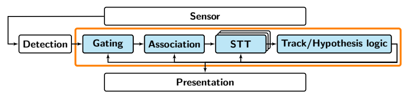

- No common mathematical formulation for all components.
- Separate methods for complexity reduction.
- Methodology heavily dominating practical applications.

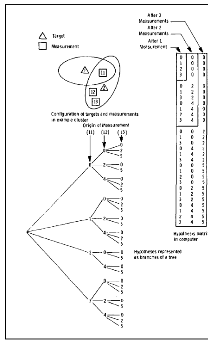

Image: Reid (1979)

Desired Properties

- A unified problem formulation:
  - A mathematical model covering all aspects of target tracking.
  - A single algorithm solving the whole problem.
- Avoid the need for *ad hoc* approximations, *e.g.*, track birth and death.
- Tractability between simplifications of the methods, and approximations of the mathematical model.
- Low complexity.

Tracking Assumptions: revisited

Targets

- The number of targets is unknown and varies over time.
- Independent motion (targets do not influence each other).
- Each target produces at most one measurement in each scan.

Observations

- Unknown origin.
- Independent measurements (independent measurement noise).
- An observation stem from at most one target.
- Measurements can be missed and clutter exist.

Random Finite Sets (RFS)

Mathematical Sets

Definition (Set)

A set is an unordered collection of unique elements. Let a_i, for $i = 1, \dots N$, be unique elements, then $\mathcal{A} = \{a^{(1)}, a^{(2)}, \dots, a^{(n)}\}$ denote a set with these elements.

Properties:

- The same element can only appear once in a set.
- A set can be empty, $\emptyset$.
- The elements in a set are not ordered.
- The carnality, $|\mathcal{A}|$, of a set is the number of elements in the set.

Sets: examples

- Empty set, $|\mathcal{A}| = 0$:
  $\mathcal{A} = \emptyset = \{\}$
- Sets with one element, $|\mathcal{A}| = 1$:
  $\mathcal{A} = \{1\}$, $\mathcal{A} = \{\left(\begin{smallmatrix}1\\3\end{smallmatrix}\right)\}$, $\mathcal{A} = \{1 + 4i\}$
- Sets with two element, $|\mathcal{A}| = 2$:
  $\mathcal{A} = \{2, 1\}$, $\mathcal{A} = \{\left(\begin{smallmatrix}1\\3\end{smallmatrix}\right), \left(\begin{smallmatrix}-1\\3\end{smallmatrix}\right)\}$, $\mathcal{A} = \{\left(\begin{smallmatrix}0.5\\0.5\end{smallmatrix}\right), \left(\begin{smallmatrix}2\\.33\end{smallmatrix}\right)\}$,
- Sets with n elements, $|\mathcal{A}| = n$:
  $\{1, \dots, n\}$, $\{1^3, \dots, n^3\}$

Note

- $\{1, 2\} = \{2, 1\}$ (the order does not matter).
- $\{1, 1\}$ is not a set, a set cannot contain duplicates elements.

Random Finite Set

Definition (Random Finite Set, RFS)

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

Both the cardinality and the elements themselves are stochastic.

- The cardinality is a discrete positive *random variables* (RV), with a *probability mass function* (PMF), $p(|X| = n)$.
- Each element $x^{(i)}$, is a RV, with an underlying PDF, $p(X \mid |X| = n) = p(\{x^{(1)}, \dots, x^{(n)}\}|n)$.

Random Finite Set: properties (1/2)

- It is possible to compute the probability of a certain realization X, $\Pr(X)$.
- 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.
- Independent RFS, X_i, can be combined to joint RFS $X = \bigcup_{i=1}^{N} X^{(i)}$, with the elements of all included RFSs. The resulting distribution can be obtained using convolution:

$$p_X(X) = \sum_{\biguplus_{i=1}^{N} X^{(i)} = X} \prod_{j=1}^{N} p_{X^{(j)}}(X^{(j)})$$

Random Sets: properties (2/2)

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

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

Definition (Set Integral)

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

Random Finite Set Distributions

Common distributions:

- Poisson Point Process (PPP)/Poisson Random Finite Set
- Bernoulli Random Finite Set
- Multi-Bernoulli Random Finite Set

Poisson Point Process

The Poisson Point Process (PPP) is defined in terms of an intensity function $\lambda(x)$ over the domain of possible values in the set.

Poisson Point Process PDF

$$p_X(X) = e^{-\bar{\lambda}} \prod_{x \in X} \lambda(x),$$

where $\bar{\lambda} = -\int \lambda(x)\, dx$.

- $\bar{\lambda}$ is the expected cardinality.
- The cardinality is Poisson distributed
  $$\Pr(|X| = n) = \tfrac{1}{n!} \bar{\lambda}^n e^{-\bar{\lambda}}.$$
- The mean and variance of the cardinality is $\bar{\lambda}$.
- Given cardinality $|X| = n$, all elements x_i in X have the same PDF $\lambda(x_i)/\bar{\lambda}$.

Poisson Point Process: example (1/2)

- Uniform intensity.
- Values in $|x_i| < 5$, $i \in \{1, 2\}$.
- Intensity function $\lambda(x) = 0.013$.
- Accumulated intensity: $\bar{\lambda} = 1.3$.

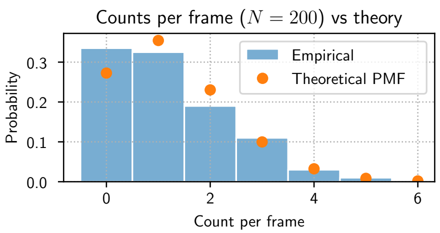

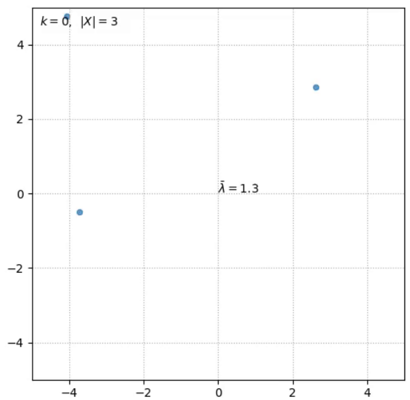

https://www.youtube.com/shorts/HLE_mzGH58s

Poisson Point Process: example (2/2)

- Gaussian sum intensity.
- $$\lambda(x) = 0.4\mathcal{N}\big(x; \left(\begin{smallmatrix}2\\2\end{smallmatrix}\right), \left(\begin{smallmatrix}1&0\\0&1\end{smallmatrix}\right)\big) + 0.9\mathcal{N}\big(x; \left(\begin{smallmatrix}-2\\-2\end{smallmatrix}\right), \left(\begin{smallmatrix}2&0.5\\0.5&1\end{smallmatrix}\right)\big)$$
- $\bar{\lambda} = 0.4 + 0.9 = 1.3$

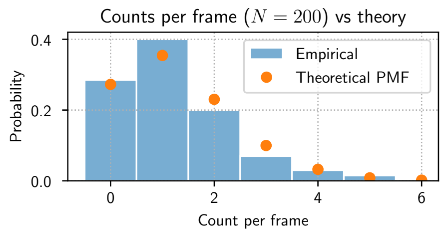

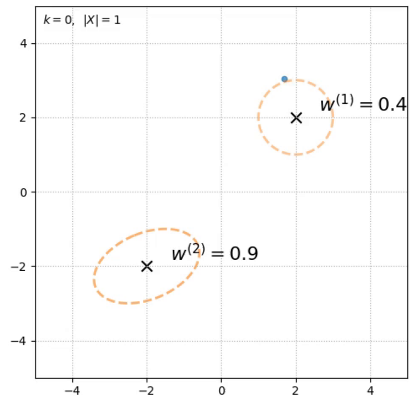

https://youtube.com/shorts/OiPHO4yaubA

Bernoulli Random Finite Set

The Bernoulli random finite set contains either 0 or 1 element.

Bernoulli Random Finite Set PDF

$$p(X) = \begin{cases} 1-r, & \text{if } X = \emptyset \\ r p_x(x), & \text{if } X = \{x\} \\ 0, & \text{if } |X| > 1 \end{cases}$$

- r is the existence probability of the element.
- The cardinality is distributed as $\Pr(|X| = n) = \begin{cases} 1-r, & \text{if } n = 0 \\ r, & \text{if } n = 1 \\ 0, & \text{if } n > 1 \end{cases}$
- The cardinality variance is typically much lower than for a PPP.
- If $r_i < 0.1$, the distribution is well approximated by a PPP.

Bernoulli Random Finite Set: example

- $r = 0.9$
- $p_x(x) = \mathcal{N}\big(x; \left(\begin{smallmatrix}-2\\-2\end{smallmatrix}\right), \left(\begin{smallmatrix}2 & 0.5\\0.5 & 1\end{smallmatrix}\right)\big)$
- $\bar\lambda = 0.9$

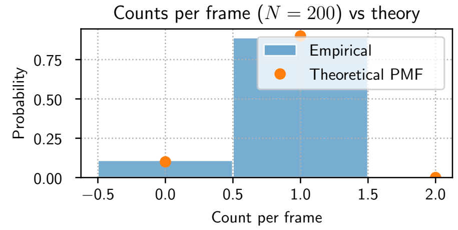

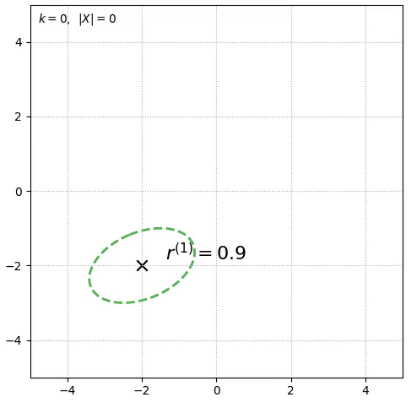

https://youtube.com/shorts/eKwPXKqCCjY

Multi-Bernoulli Random Finite Set

Several independent Bernoulli RFS variables, $X^{(i)}$ $(i = 1, \dots, n)$, can be combined to a multi-Bernoulli random finite set, $X = \bigcup_{i=1}^{N} X^{(i)}$.

Multi-Bernoulli Random Finte Set PDF

$$p_X(X) = \sum_{\biguplus_{i=1}^{N} X^{(i)} = X} \prod_{j=1}^{N} p_{X^{(j)}}(X^{(j)})$$

- The separate components are represented by, r_i and $p_i(x)$.
- $$\Pr(|X| = n) = \prod_{\delta_i \in \{0,1\}, \sum_i \delta_i = n} (r^{(i)})^{\delta_i} (1 - r^{(i)})^{1-\delta_i},$$
  where $\delta_i = 1$ indicates existence of component i.
- $\bar\lambda = \sum_i r^{(i)}$

Multi-Bernoulli Random Process: example

- $r^{(1)} = 0.4$
- $p^{(1)}(x) = \mathcal{N}\big(x; \left(\begin{smallmatrix}2\\2\end{smallmatrix}\right), \left(\begin{smallmatrix}1 & 0\\0 & 1\end{smallmatrix}\right)\big)$
- $r^{(2)} = 0.9$
- $p^{(2)}(x) = \mathcal{N}\big(x; \left(\begin{smallmatrix}-2\\-2\end{smallmatrix}\right), \left(\begin{smallmatrix}2 & 0.5\\0.5 & 1\end{smallmatrix}\right)\big)$
- $\bar\lambda = 1.3$

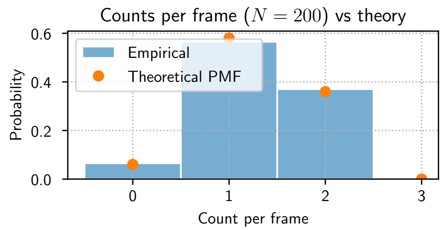

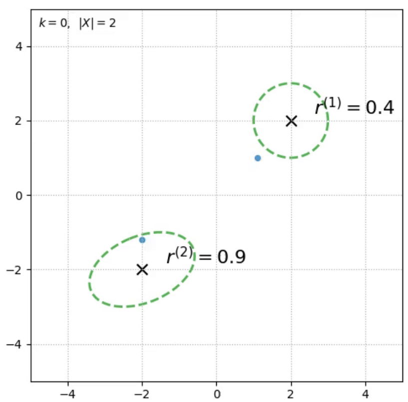

https://youtube.com/shorts/9j2RNycg36Q

Target Tracking using Random Finite Sets

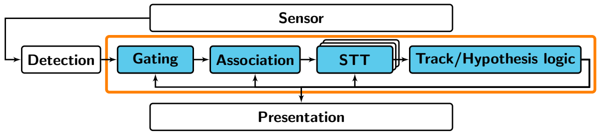

LINKÖPING UNIVERSITY

Target Tracking using Random Finite Sets (1/2)

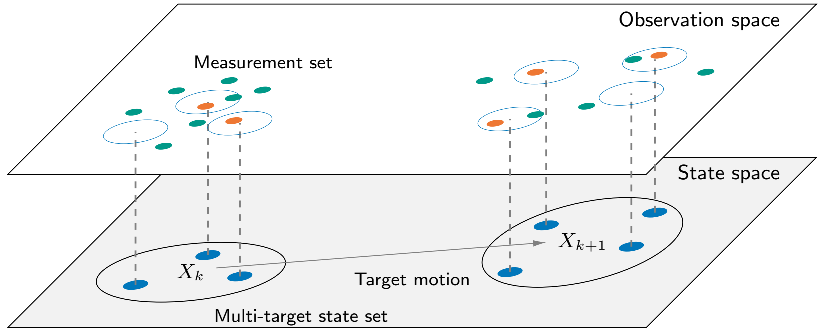

P. Boström-Rost. Sensor Management for Target Tracking Applications. 2021.

Target Tracking using Random Finite Sets (2/2)

Observations

Each scan contains a number of observations without natural ordering. Y_t is well represented by a RFS.

- Each target gives raise to a Bernoulli RFS measurement.
- Clutter is typically a PPP.

Targets

The number of targets is unknown and there is no natural order, represent as a RFS X_t. (Track labels will require some extra care…)

RFS Models: observation model (1/2)

- Y_t represents set of measurements received in the scan at time t.
- The measurement model $p(Y_t|X_t)$ models both measurements from targets and clutter.

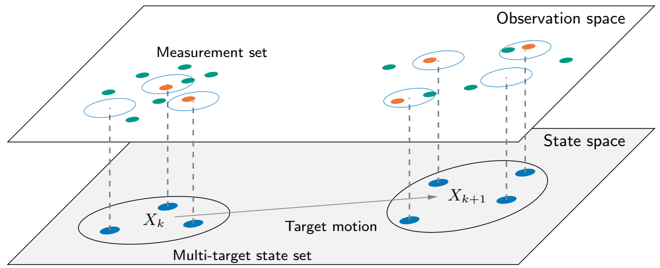

P. Boström-Rost. Sensor Management for Target Tracking Applications. 2021.

RFS Models: observation model (2/2)

- Contributions to Y_t based on X_t:
  - Measurements originating from targets $H(X_t) = \{y_t\}$,

    $$y_t = h(x_t, e_t),$$

    for all $x_t \in X_t$ that are detected, assuming probability of detection P_{D}.
  - Clutter V_t, typically Poisson distributed.
- Y_t comprise both measurements from targets and clutter,

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

RFS Models: target dynamics (1/2)

- The set of targets X_{t+1} encodes **both** number of targets and their states.
- The dynamic model $p(X_{t+1}|X_t)$ encodes **both** changes in the number of targets, and the state of these targets.

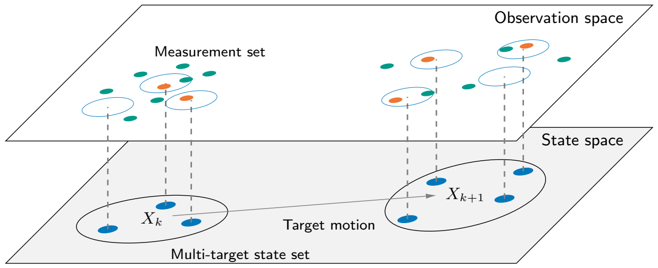

P. Boström-Rost. Sensor Management for Target Tracking Applications. 2021.

RFS Models: target dynamics (2/2)

- Contributions to X_{t+1} given previous X_t:
  - Surviving targets $F(X_t) = \{x_t\}$,

    $$x_t = f(x_{t-1}, w_{t-1}),$$

    for all targets $x_{t-1} \in X_{t-1}$ that survives. Targets survive with probability P_{S}.
  - Appearing targets $W_t(X_{t-1})$. This is a combination of new targets spawning from existing targets, and randomly appearing targets.
- The new set X_{t+1} combines both contributions

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

This describes how the RFS X_t evolves over time.

Bayesian Filtering Solution

The "normal" Bayesian solution holds:

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

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

- $p(Y_t|X_t)$: Describes how observations are created from a set of targets.
- $p(X_t|X_{t-1})$: Describes how targets propagate over time, appear and disappear.
- $p(X_t|\mathbb{Y}_t)$: The posterior tracking solution.

Key Difference: (lack of) expected value of RFS

Limitation

The regular moments, which are heavily used in, *e.g.*, the Kalman filter, are **not** well-defined!

Consider:

$$\hat{X}_t = \int X_t p(X_t|\mathbb{Y}_t)\delta X_t = \mathsf{E}(X_k|\mathbb{Y}_t)$$

How are elements with different cardinality combined?
They cannot, the result would not be well defined!

Bayesian Filtering Solution: methods

- Conjugate prior distributions and approximations are important.
- Several different approximations of this exist:
  - **Probability Hypothesis Density (PHD):**
    Propagate first moment (intensity).
  - **Cardinalized PHD (C-PHD):**
    Also propagate the cardinality of the PHD.
  - **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).

Probability Hypothesis Density (PHD) Filter

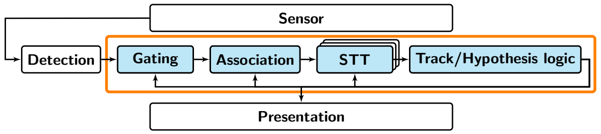

Probability Hypothesis Density (PHD)

Definition (Probability Hypothesis Density)

Let $X_t = \{x_t^{(1)}, x_t^{(2)}, \dots, x_t^{(n)}\}$, and define

$$\boldsymbol{\delta}_X(x) = \sum_{i=1}^{n} \delta_{x^{(i)}}(x),$$

producing a scalar value from an RFS. The *probability hypothesis density* (PHD) D_X is defined as

$$D_X(x) = \mathsf{E}_X(\boldsymbol{\delta}_X(x))$$

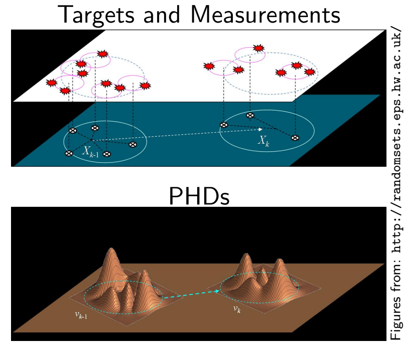

Intuition: $D_X(x)$ is the intensity of targets.

PHD Filter: assumptions

- Motion model PDF: $p(x_{t+1}|x_t)$
- Survival probability for existing targets: $P_\mathrm{S}(x_t)$
- Spawning of new targets from existing: $b^s(x_{t+1}|x_t)$
- Appearance of new targets: $b(x_{t+1})$
- Probability of detection: $P_\mathrm{D}(x_t)$
- False alarm model (Poisson distribution): probability $c(y)$
- 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.

$$X_{t+1} = \Big[\bigcup_{x_t \in X_t} p(x_{t+1}|x_t) \Big] \bigcup \Big[\bigcup_{x_t \in X_t} b^s(x_{t+1}|x_t) \Big] \bigcup b(x_{t+1})$$

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

$$\hat{D}_{t+1|t}(x_{t+1}) = \int \big(\underbrace{p_S(x_t) p(x_{t+1}|x_t)}_{\text{Surviving targets}} + \underbrace{b^s(x_{t+1}|x_t)}_{\text{Spawned targets}} \big) \hat{D}_{t|t}(x_t)\, dx_t + \underbrace{b(x_{t+1})}_{\text{Born targets}} \qquad \text{(TU)}$$

$$\hat{D}_{t|t}(x_t) = \underbrace{(1 - p_D(x_t)) \hat{D}_{t|t-1}(x_t)}_{\text{Undetected}} + \underbrace{\sum_{y_t \in Y_t} \frac{p_D(x_t) p(y_t|x_t) \hat{D}_{t|t-1}(x_t)}{c(y_t) + \int p_D(x_t') p(y|x_t') \hat{D}_{t|t-1}(x_t')\, dx_t'}}_{\text{Detection}} \qquad \text{(MU)}$$

Note

The PHD can be considered the first-order moment density (or intensity) of the RFS PDF. It is similar to a regular 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 PHD approximation, with pruning and merging, yielding the *Gaussian mixture PHD* (GM-PHD) filter.

- 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).

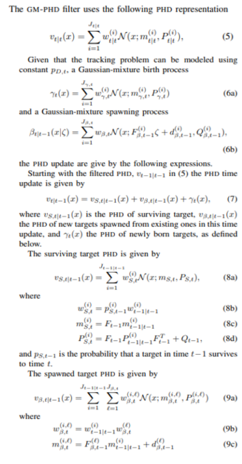

The GM-PHD filter uses the following PHD representation

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

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

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

and a Gaussian-mixture spawning process

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

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

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

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

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

where

$$w_{S,t}^{(i)} = p_{S,t-1}^{(i)} w_{t-1|t-1}^{(i)} \qquad (8b)$$
$$m_{S,t}^{(i)} = F_{t-1} m_{t-1|t-1}^{(i)} \qquad (8c)$$
$$P_{S,t}^{(i)} = F_{t-1} P_{t-1|t-1}^{(i)} F_{t-1}^T + Q_{t-1}, \qquad (8d)$$

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

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

where

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

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.**
- Significant computational gains by realizing that many Kalman gains and covariances are the same.
- The GM-PHD filter requires a non-trivial reduction Gaussian mixture reduction step to alleviate the exponential growth of components.
- The usage of PHD filters peaked in the 2010's, and marked start of the adoption of RFS methods.

GM-PHD Filter: example

Illustration of PHD at time: $k+1$

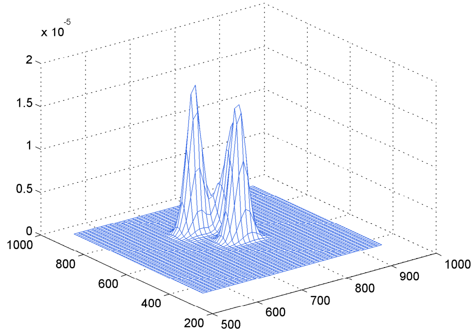

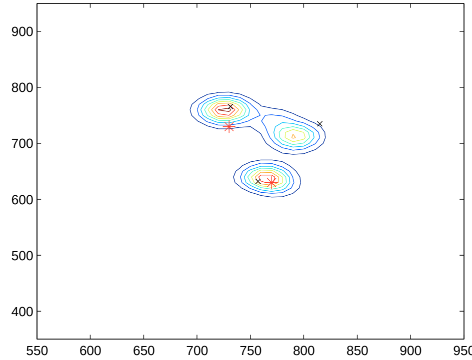

Note: The expected mean over the intensity function sums to the number of targets.

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)

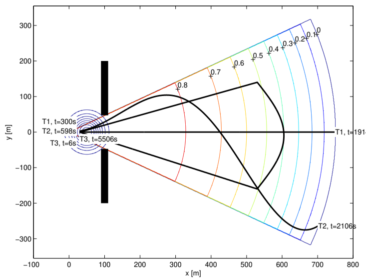

PHD Filter Example

Standard GMPHD filter

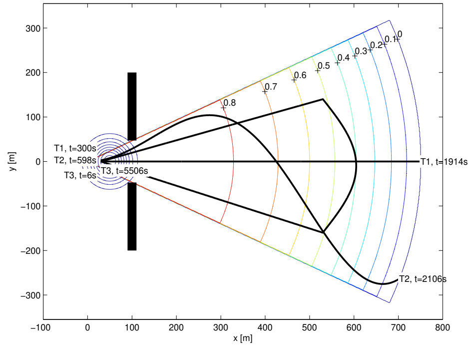

https://youtu.be/PJimgDB3X88

Modified GMPHD filter

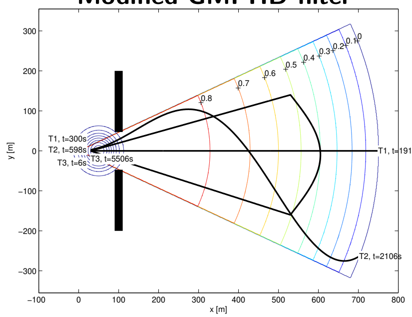

https://youtu.be/cAL3ynfVZAk

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

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 based on the particle filter.
- The cardinalized PHD (CPHD).
- PHD with track labeling.

Cardinalized Probability Hypothesis Density (CPHD)

- In the PHD filter, the number of target is obtained by integrating $\hat{D}(x)$. This estimate is known to have a **high variance**.
- The PHD filter can be seen as a first order moment filter.
- The CPHD filter is the equivalent to second order moment.
- The CPHD explicitly estimates the cardinality jointly with the intensity.
- Motivation: 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.

Multi-Bernoulli Filters

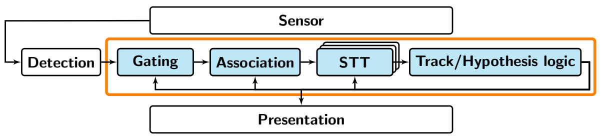

(Labeled) Multi-Bernoulli Representation

- Only estimating the PHD is limiting.
- The PPP RFS is fairly uninformative, alternatives are interesting.
- Use a multi-Bernoulli RFS to represent the MTT state:

$$p_X(X) = \sum_{\uplus_{i=1}^{N} X^{(i)}=X} \prod_{j=1}^{N} p_{X^{(j)}}(X^{(j)})$$

  represented by: $\{r^{(i)}, p_x^{(i)}(x)\}_i$ for $i = 1, \dots, n$.
- Multi-Bernoulli filters represent both targets and the birth process as multi-Bernoulli RFSs.
- **Labeled Multi-Bernoulli** RFS: Add a unique label $\ell \in \mathcal{L}$ to the state, and the RFS becomes $\{(r^{(\ell)}, x^{(\ell)})\}_{\ell \in \mathcal{L}}$.

Labeled Multi-Bernoulli: time update

Time update

$$p(X_{t+1}|\mathbb{Y}_t) = \int p(X_{t+1}|X_t)p(X_t|\mathbb{Y}_t)\,\delta X_t$$

$$X_{t+1|t} = \big\{(r^{(\ell)}_{t+1|t,\mathrm{S}}, p^{(\ell)}_{t+1|t,\mathrm{S}}(x_{t+1}))\big\}_{\ell\in\mathcal{L}} \bigcup \big\{(r^{(\ell)}_{\mathrm{B}}, p^{(\ell)}_{\mathrm{B}}(x_{t+1}))\big\}_{\ell\in\mathcal{B}}$$

where:

- $\{(r^{(\ell)}_{\mathrm{B}}, p^{(\ell)}_{\mathrm{B}}(x_{t+1}))\}_{\ell\in\mathcal{B}}$ is the birth process.
- $r_{t+1|t,\mathrm{S}} = \eta_{t,S}(\ell) r^{(\ell)}_{t|t}$
- $p^{(\ell)}_{t+1|t,\mathrm{S}}(x_{t+1}) = \frac{\int p_S(x_t,\ \ell)p(x_{t+1}|x_t,\ell)p(x_t|\ell)\,dx_t}{\eta_{t,\mathrm{S}}(\ell)}$
- $\eta_{t,\mathrm{S}}(\ell) = \int P_{\mathrm{S}}(x_t,\ell)p(x_t|\ell)\,dx_t$

Labeled Multi-Bernoulli: measurement update

Measurement update

$$p(X_t|\mathbb{Y}_1) \approx \bigcup_{i=1}^{N} \{(r^{(\ell,i)}_{t|t}, p^{(\ell,i)}_{t|t}(x_t))\}_{\ell\in\mathcal{L}^{(i)}_{t|t-1}}$$

$$r^{(\ell,i)}_{t|t} = \sum_{(I_{t|t-1},\theta)\in\mathcal{F}(\mathcal{L}^{(i)}_{t|t-1})\times\Theta_{I_{t|t-1}}} \omega^{(I_{t|t-1},\theta)}(y^{(i)}_t)\mathbf{1}_{I_{t|t-1}}(\ell)$$

$$p^{(\ell,i)}_{t|t}(x_t) = \frac{1}{r^{(\ell,i)}} \sum_{(I_{t|t-1},\theta)\in\mathcal{F}(\mathcal{L}^{(i)}_{t|t-1})\times\Theta_{I_{t|t-1}}} \omega^{(I_{t|t-1},\theta)}(y^{(i)}_t)\mathbf{1}_{I_{t|t-1}}(\ell)p^{(\theta)}(x_t,\ell|y^{(i)}_t)$$

$$\omega^{(I_{t|t-1},\theta)}(Y^{(i)}) \propto \omega^{(I_{t|t-1})}_{t|t-1,i}[\eta^{(\theta)}_{y^{(i)}_t}]^{I_{t|t-1}}$$

$$= \prod_{\ell\in\mathcal{L}^{(i)}_{t|t-1}\setminus I_{t|t-1}} (1 - r^{\ell}_{t|t-1}) \prod_{\ell'\in I^a_{t|t-1}} r^{(\ell')}_{t|t-1}\eta^{(\theta)}_{Y^{(i)}}(\ell') \prod_{\ell''\in I^n_{t|t-1}} r^{(\ell'')}_{t|t-1}\eta^{(\theta)}_{y^{(i)}_t}(\ell'')$$

Warning!!!

Assume there are typos in this, look up the equations elsewhere! :)

Observations

- The notations quickly becomes horrific, we would need some lectures to sort that out properly!
- 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* (δ-GLMB) filter (a LMB variation), is essentially a TO-MHT.

Example: ice tracking, Kongsfjorden, Svalbard (1/2)

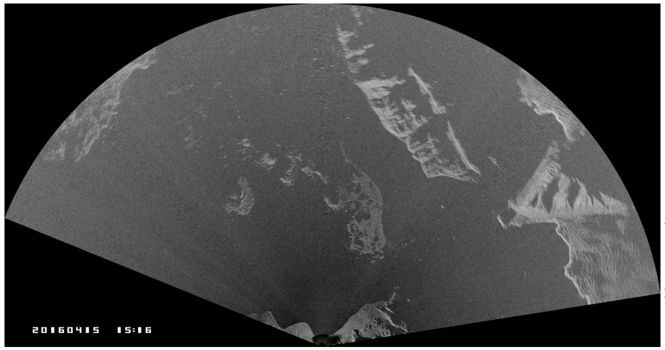

https://youtu.be/GWD9Dw6cHkY

Illustration of measurements

- Radar facing out over the fjord.
- Plenty of clutter.
- Fairly low P_{D}.
- Objects splitting and disappearing.
- LMB for tracking.

Example: ice tracking, Kongsfjorden, Svalbard (2/2)

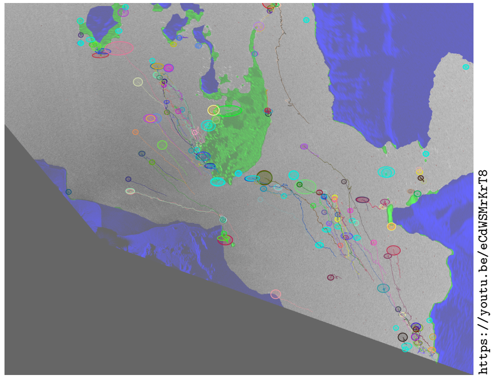

LMB results after 7 hours

- Radar facing out over the fjord.
- Plenty of clutter.
- Fairly low P_{D}.
- Objects splitting and disappearing.
- LMB for tracking.

Poisson Multi-Bernoulli Mixture (PMBM) Filter

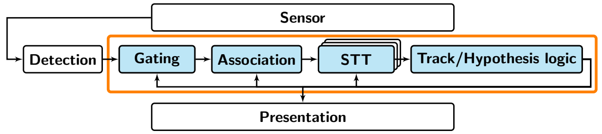

Poisson Multi-Bernoulli Mixture (PMBM) Filter

- Relatively new, Williams (2015).
- Can be considered the new state-of-the-art MTT method. (This can be debated, there are still new MHT development going on.)
- 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, though grows over time and has to be reduced.

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 detections are handled in the MBM part, and removed from the PHD.
- Alternatives to GM-PHD exists, *e.g.*, Boström-Rost et al. (2021).

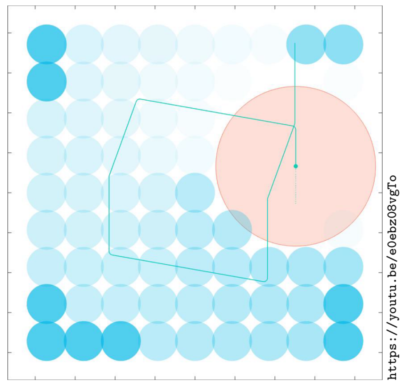

PMBM: observed targets

- 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.

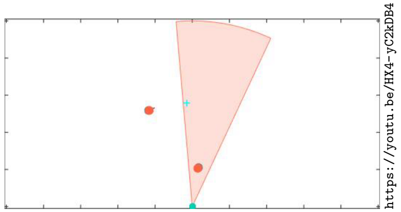

PMBM Algorithm

- Conceptually:
  - Predict a PHD filter to 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.
  - Weights are update to make this correct.
  - Reduce the MBM part by pruning, merging, and recycling (reintroduce pruned tracks in the PHD representation).
- For details see, *e.g.*, Williams (2015); García-Fernández et al. (2018)

Summary

Summary

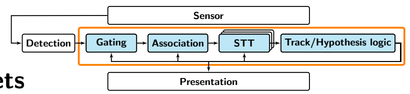

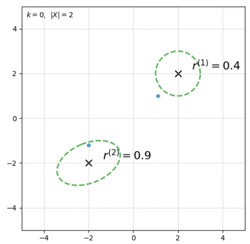

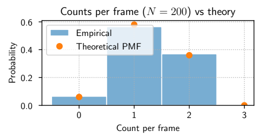

Random Finite Sets

- Can model the full multi-target tracking problem.
- Random finite sets (RFSs) are sets of random variables of random cardinality.
- RFS examples:
  - Poisson point process (PPP)
  - Bernoulli random finite set
  - Multi-Bernoulli random finite set
- Methods (will be discussed next time):
  PHD filter, C-PHD filter,
  Multi-Bernoulli filters, PMBM filter, ...

Examination

- **Examination** (how many intend to get credits?):
  - Ethics session, required for all credits:
    Seminar/Workshop at Lecture 9, January 26, 2026, 13–15. Expect more information and instructions early January.
  - Exam, 2 ETCS credits:
    Take home exam, 2 h exam. Tests the understanding of the principles discussed in the course.

    Due: End of January, 2026, (contact me to schedule the exam).
  - Exercises, 4 ETCS credits:
    Show hands on experience of important discussed methods.

    Due: January 15, 2026
  - Project, 3 ETCS credits:

    More advanced utilization of MTT techniques, preferably related to your research. Discuss details with me.