UDRC Summer School 2019
Day 2 - Sensing and Tracking

Mengwei Sun, James R. Hopgood
*msun@ed.ac.uk

University of Edinburgh

20^{th} June 2023

Contents

Intro to Data Association

Definition

Data association assigns measurements to existing tracks or existing tracks to measurements (measurement-to-track association vs track-to-measurement association).

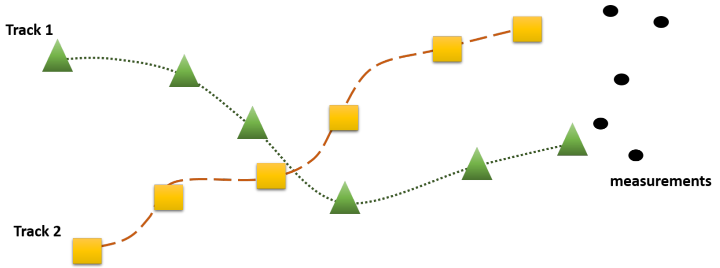

How to determine which measurements to add to which track?

Overall Procedure of Data Association

- **Make observations.** Measurements can be raw data (e.g., processed radar signals) or the output of some target detector (e.g. people detector)
- **Predict the measurements** from the predicted tracks. This yields an area in sensor space where to expect an observation. The area is called the validation gate and is used to narrow the search. Then, check if a measurement lies in the gate.
- Pair a valid measurement candidate to a potential target.

Challenges to Data Association

What makes this a difficult problem

- Multiple targets
- Clutter
- Missing alarm (occlusions, sensor failures, . . .)
- Ambiguities (several measurements in the gate)

Introducing Clutter

Definition

Unwanted measurements do not correspond to the target of interest

Caused by a variety of factors

- Environmental conditions (e.g., rain, fog, snow)
- Natural or man-made objects,
- Other targets that are not of interest.

Methods

- Filtering
- Gating
- Data association

Introducing Clutter

It is usually assumed that clutter measurements typically follow a *Poisson* point process or distribution.

- The mean number of clutter points per time-step is set to be β_{FA}.
- Specifically, at each time, the measurement area is set to be $V = (\bar{x} - \underline{x}) \times (\bar{y} - \underline{y})$, and the number of false alarm is generated by Matlab code '$N_{fa} = \text{poissrnd}(\beta_{\mathrm{FA}} * V)$'.
- The x-axis of each false alarm follows a uniform distribution with region $[\underline{x}, \bar{x}]$. Similarly, the y-axis of each false alarm follows a uniform distribution with region $[\underline{y}, \bar{y}]$.

Data Association in Single Target Tracking (STT)

STT: We need to associate measurement to a track:

- Hard decision
  An explicit decision is made on how to match tracks and measurements, and only those associations are considered further.
  E.g., nearest neighbour (NN) association (or greedy assignment).
- Soft decision
  Several measurements are associated with the track and contribute to the result.
  E.g., probabilistic data association (PDA).

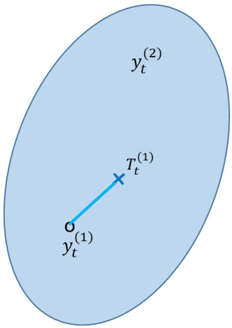

Probabilistic Data Association (PDA)

A soft association method, which uses all the measurements in the gate, is weighted with how well the prediction fits. Measurements in the gate are shown as $Y_t = \{y_t^{\{\text{I}\}}\}_{i=1}^{m_t}$.

- $\mathcal{H}^{(0)} = \{Y_t$ are false alarms, no target-originated measurement.$\}$
- $\mathcal{H}^{(i)} = \{y_t^{\{\text{I}\}}$belongs to the target; all the rest are false alarm.$\}$

Probabilistic Data Association (PDA)

A soft association method, which uses all the measurements in the gate, is weighted with how well the prediction fits. Measurements in the gate are shown as $Y_t = \{y_t^{\{i\}}\}_{i=1}^{m_t}$.

- $\mathcal{H}^{(0)} = \{Y_t$ are false alarms, no target-originated measurement.$\}$
- $\mathcal{H}^{(i)} = \{y_t^{\{i\}}$belongs to the target; all the rest are false alarm.$\}$

For $i = 1, \ldots, m_t$, the estimated pdf is calculated using total probability theorem as

$$p(x_t \mid Y_t) = \sum_{i=0}^{m_t} p(x_t \mid \mathcal{H}^{(i)}, Y_t) p(\mathcal{H}^{(i)} \mid Y_t)$$

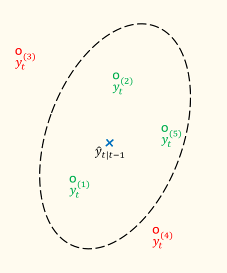

Multi-target Association: Example

- Using STT for each target results in locally optimal solutions, which might be infeasible

  Consider the association hypothesis: $T_t^{(1)} \leftrightarrow y_t^{(1)}$, $T_t^{(2)} \leftrightarrow y_t^{(5)}$, $T_t^{(3)} \leftrightarrow y_t^{(5)}$, which picks the best measurement for each target but violates the assumption that a measurement originated from a single target.

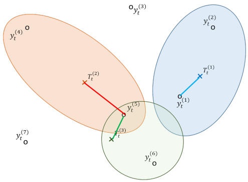

- In Multi-target tracking, the complete association hypothesis is considered only to obtain a **global optimum** and avoid infeasible solutions.

Data Association in Multi-target Tracking

- **Gating**: Get a validation matrix V.
- **Clustering**: Separate tracks that do not share potential measurements.
- Association and updating of confirmed tracks: Associate measurements to confirmed tracks, then update the tracks.
- Association and updating of tentative tracks.
- Update the procedure with the remaining measurements and the tentative Initiate new tentative tracks: Use remaining measurements to start tentative tracks.

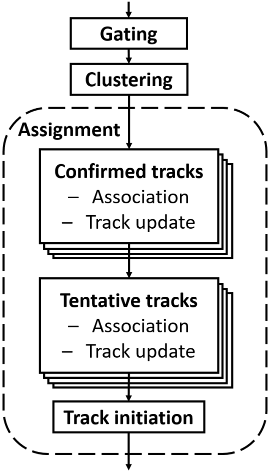

Gating

The purpose of gating is to remove measurements that are very unlikely to originate from a given target, that is:

- Reduce problem complexity by minimising the number of possible measurements for each target.
- A cheap operation:
  - Rectangular gating;
  - Elliptical gating.

	$y_t^{(1)}$	$y_t^{(2)}$	$y_t^{(3)}$	$y_t^{(4)}$	$y_t^{(5)}$	$y_t^{(6)}$	$y_t^{(7)}$
$T_t^{(1)}$	1	1	0	0	0	0	0
$T_t^{(2)}$	0	0	0	1	1	0	0
$T_t^{(3)}$	0	0	0	0	1	0	0

Table: Validation matrix $\mathcal{V}$

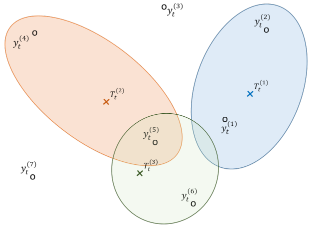

Clustering

- Computational complexity scales exponentially with the number of measurements and targets.
- Tracks that do not share any measurements can be treated separately to reduce the complexity.

An example:

$\mathcal{C}_t^{(1)} = \{T_t^{(2)}, T_t^{(3)}\}, \qquad \mathcal{C}_t^{(2)} = \{T_t^{(1)}\}$

	$y_t^{(4)}$	$y_t^{(5)}$	$y_t^{(6)}$
$T_t^{(2)}$	1	1	0
$T_t^{(3)}$	0	1	0

	$y_t^{(1)}$	$y_t^{(2)}$
$T_t^{(1)}$	1	1

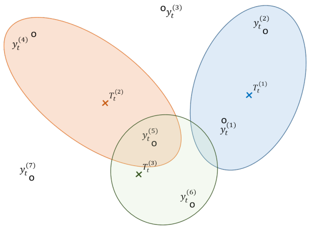

Assignmentt: notation

Measurement origins

- TC–Track Continuation: a measurement will update a track
- FA–False Alarm: a measurement is considered a nuisance
- NT–New Track: a measurement can start a new track

It is reasonable to assume that a measurement can only be used for one of the above.

Hypothesis Probabilities: track continuation (TC)

TC

- Detection probability: P_D
- Gate probability: P_G
- Predicted measurent density of j-th target: $p^{(j)}_{t|t-1}(y)$

In the KF case:

$$p^{(j)}_{t|t-1}(y) = \mathcal{N}(y; \hat{y}^{(j)}_{t|t-1}, S^{(j)}_{t|t-1})$$

Hypothesis Probabilities: false alarm (FA)

FA

- Number of false alarms, m_t^{FA}, in V is distributed as:

$$P_{\mathrm{FA}}(m_t^{\mathrm{FA}}) = \frac{(\beta_{\mathrm{FA}} V)^{m_t^{\mathrm{FA}}} e^{-\beta_{\mathrm{FA}} V}}{m_t^{\mathrm{FA}}!}$$

β_{FA}:Poisson distributed with clutter rate.

- False alarm spatial density is $p_{\mathrm{FA}}(y) = 1/V$.

Hypothesis Probabilities: new track (NT)

NT

- Number of false alarms, m_t^{FA}, in V is distributed as:

$$P_{\mathrm{NT}}(m_t^{\mathrm{NT}}) = \frac{(\beta_{\mathrm{NT}} V)^{m_t^{\mathrm{NT}}} e^{-\beta_{\mathrm{NT}} V}}{m_t^{\mathrm{NT}}!}$$

β_{FA}:Poisson distributed with clutter rate.

- New target spatial density is $p_{\mathrm{NT}}(y) = 1/V$.

Hypothesis Probabilities: putting all together

Global logarithmic association probability

$$\log p(\theta_t \mid Y_t) = m_t^{\mathrm{FA}} \log \beta_{\mathrm{FA}} + m_t^{\mathrm{NT}} \log \beta_{\mathrm{NT}} + \sum_{j \in \mathcal{J}} \log \frac{P_D p_{t|t-1}^{(j)}(y_t^{(\theta_t^{-1}(j))})}{1 - P_D P_G}$$

Properties:

- One term per measurement
- The best association is to pick the right contribution from each measurement in a consistent way.

Assignment matrix

The assignment matrix $\mathcal{A}$ consists of all possible measurement contributions to $\log p(\theta_t \mid Y_t)$

	T_2	T_3	FA_4	FA_5	FA_6	NT_4	NT_5	NT_6
$y_t^{(4)}$	l_{42}	$-\infty$	$\log \beta_{\mathrm{FA}}$	$-\infty$	$-\infty$	$\log \beta_{\mathrm{NT}}$	$-\infty$	$-\infty$
$y_t^{(5)}$	l_{52}	l_{53}	$-\infty$	$\log \beta_{\mathrm{FA}}$	$-\infty$	$-\infty$	$\log \beta_{\mathrm{NT}}$	$-\infty$
$y_t^{(6)}$	$-\infty$	l_{63}	$-\infty$	$-\infty$	$\log \beta_{\mathrm{FA}}$	$-\infty$	$-\infty$	$\log \beta_{\mathrm{NT}}$

The gain of assigning measurement $y_t^{(i)}$ to track T_j is

$$l_{ij} = \log \frac{P_D p_{t|t-1}^{(j)}(y_t^{(i)})}{1 - P_D P_G}$$

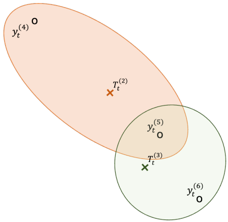

Assignment Problem

Assume a scan with m measurements and n "track hypotheses".

- Given the matrix $\mathcal{A} \in \mathbb{R}^{m \times n}$ with $m \leq n$.
- Define the binary value $z_{ij} \in \{0, 1\}$,

Assignment Problem

$$\begin{aligned} \underset{Z}{\text{maximise}} : \quad & \sum_{i,j} z_{i,j} \mathcal{A}_{ij} \\ \text{subject to} : \quad & \sum_{j} z_{i,j} = 1 \ \ \forall i \\ & \sum_{i} z_{i,j} \leq 1 \ \ \forall j \end{aligned}$$

- Each measurement is associated with exactly one track/FA/NT.
- Each track/FA/NT is associated with at most one measurement.

Global Nearest Neighbour (GNN) Tracker

- Select the best association hypothesis, θ_t
  - Munkres algorithm.
  - Auction algorithm (by Bertsekas).
  - JVC algorithm.
- Given θ_t
  - Update all tracks with the associated measurement (usually using an EKF).
  - Update the track logic.

Global Nearest Neighbour (GNN) Tracker

- Makes a hard association decision
- Relative fast and easy to implement
- Works well when targets are well separated
- Should not be used with poorly separated targets
- Heavy clutter and low P_D could cause problems
- Could break down completely with the wrong association

Practical

Jupyter notebook 5 and 6: multiple-target tracking using GNN

Issues

- Examine the output of the previous notebook. What happened?
- Is that your final decision?
  - There's a problem with making a firm decision at each time step
- Running time
  - Surprisingly often some form of brute-force algorithm will be used
  - Gate can be 'tuned.'

Joint Probabilistic Data Association (JPDA)

The JPDA filter is soft decision equivalent of GNN in the way that Probabilistic Data Association (PDA) is a soft version of NN.

- All past is again summarized with a single hypothesis.
- The number of targets is assumed to be fixed in the associations.
- For each previously established target, calculate:

$$p(T_t^{(j)} \leftrightarrow y_t^{(i)}) \text{ and } p(T_t^{(j)} \leftrightarrow \emptyset)$$

  $y_t^{(i)}$ are measurements in the gate. The update is then made with PDA update formulas by using these probabilities instead.

- A separate track initiation logic must run along with JPDAF to detect and initiate new tracks.

JPDA-1

Enumerate all possible measurement hypotheses and compute their respective likelihood. This can be done for each cluster independently.

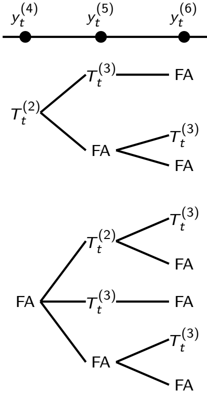

$p_{23\emptyset} = l_{42} l_{53} \log \beta_{\mathrm{FA}}$

$p_{2\emptyset 3} = l_{42} \log \beta_{\mathrm{FA}} l_{63}$

$p_{2\emptyset\emptyset} = l_{42} \log \beta_{\mathrm{FA}} \log \beta_{\mathrm{FA}} \bar{l}_3$

$p_{\emptyset 23} = \log \beta_{\mathrm{FA}} l_{52} l_{63}$

$p_{\emptyset 2\emptyset} = \log \beta_{\mathrm{FA}} l_{52} \log \beta_{\mathrm{FA}} \bar{l}_3$

$p_{\emptyset 3\emptyset} = \log \beta_{\mathrm{FA}} l_{53} \log \beta_{\mathrm{FA}} \bar{l}_2$

$p_{\emptyset\emptyset 3} = \log \beta_{\mathrm{FA}} \beta_{\mathrm{FA}} l_{63} \bar{l}_2$

$p_{\emptyset\emptyset\emptyset} = \log \beta_{\mathrm{FA}} \log \beta_{\mathrm{FA}} \log \beta_{\mathrm{FA}} \bar{l}_3 \bar{l}_3$

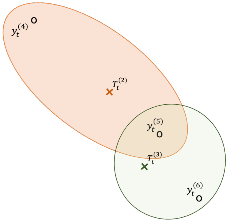

	T_2	T_3	FA_4	FA_5	FA_6
$y_t^{(4)}$	l_{42}	$-\infty$	$\log \beta_{\mathrm{FA}}$	$-\infty$	$-\infty$
$y_t^{(5)}$	l_{52}	l_{53}	$-\infty$	$\log \beta_{\mathrm{FA}}$	$-\infty$
$y_t^{(6)}$	$-\infty$	l_{63}	$-\infty$	$-\infty$	$\log \beta_{\mathrm{FA}}$

Rearrange the hypotheses to be able to compute the probability for each separate track.

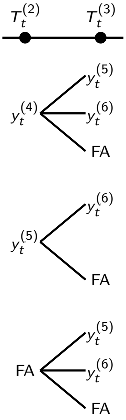

$$p_{23\emptyset} = l_{42} l_{53} \log \beta_{\mathrm{FA}}$$

$$p_{2\emptyset 3} = l_{42} \log \beta_{\mathrm{FA}} l_{63}$$

$$p_{2\emptyset\emptyset} = l_{42} \log \beta_{\mathrm{FA}} \log \beta_{\mathrm{FA}} \bar{l}_3$$

$$p_{\emptyset 23} = \log \beta_{\mathrm{FA}} l_{52} l_{63}$$

$$p_{\emptyset 2\emptyset} = \log \beta_{\mathrm{FA}} l_{52} \log \beta_{\mathrm{FA}} \bar{l}_3$$

$$p_{\emptyset 3\emptyset} = \log \beta_{\mathrm{FA}} l_{53} \log \beta_{\mathrm{FA}} \bar{l}_2$$

$$p_{\emptyset\emptyset 3} = \log \beta_{\mathrm{FA}} \beta_{\mathrm{FA}} l_{63} \bar{l}_2$$

$$p_{\emptyset\emptyset\emptyset} = \log \beta_{\mathrm{FA}} \log \beta_{\mathrm{FA}} \log \beta_{\mathrm{FA}} \bar{l}_3 \bar{l}_3$$

$$\Pr(\theta^{-1}(2) = 4) = \frac{1}{C}(p_{23\emptyset} + p_{2\emptyset 3} + p_{2\emptyset\emptyset})$$

$$\Pr(\theta^{-1}(2) = 5) = \frac{1}{C}(p_{\emptyset 23} + p_{\emptyset 2\emptyset})$$

$$\Pr(\theta^{-1}(2) = \emptyset) = \frac{1}{C}(p_{\emptyset 3\emptyset} + p_{\emptyset\emptyset 3} + p_{\emptyset\emptyset\emptyset})$$

$$\Pr(\theta^{-1}(3) = 5) = \frac{1}{C}(p_{23\emptyset} + p_{\emptyset 3\emptyset})$$

$$\Pr(\theta^{-1}(3) = 6) = \frac{1}{C}(p_{2\emptyset 3} + p_{\emptyset 23} + p_{\emptyset\emptyset 3})$$

$$\Pr(\theta^{-1}(3) = \emptyset) = \frac{1}{C}(p_{2\emptyset\emptyset} + p_{\emptyset 2\emptyset} + p_{\emptyset\emptyset\emptyset})$$

JPDA-Properties

- Makes no hard association decision:
- More robust in heavily cluttered environments with low Pd.
- Sub-optimal compared to using the correct associations.
- Works well when targets are well separated!
- Closely separated targets suffer from coalescence, i.e., neighbouring tracks become identical.
- More complicated and more computationally complex than GNN.
- Consideration required when implementing the tracking logic

Practical

Jupyter notebook 7, 8: joint probabilistic data association - compare with GNN

Summary

- Methods for gating, clustering, and association were presented, yielding the validation and association matrix.
- SHT: One measurement association hypothesis is used

  GNN: A hard decision; choose the most likely association hypothesis. The association problem can be solved with many off-the-shelf algorithms, e.g., auction, after constructing the association (cost) matrix.

  JPDA: A soft decision; marginalize all possible associations. How to combine the possible measurements depends on the association matrix