

Recent advances in multi-object estimation

Emmanuel D. Delande
E.D.Delande@hw.ac.uk

Engineering & Physical Sciences
Heriot-Watt University
Edinburgh, UK

UDRC Summer School - July 21, 2015 @ University of Surrey, UK




- 1 Multi-object filtering framework: basics
- 2 Traditional solutions: strengths and weaknesses
 - The track-based approach
 - The RFS-based approach
- 3 Stochastic populations for multi-object filtering
 - Multi-object estimation framework
 - Bayesian filtering
 - Closed-loop sensor management

- 1 Multi-object filtering framework: basics
- 2 Traditional solutions: strengths and weaknesses
- 3 Stochastic populations for multi-object filtering

Bayesian filtering

Principle



- P_{y_t} : “information” known by operator at time t on objects of interest or *targets*
- Z_t : observations produced by the sensor system at time t and collected by the operator

Sensor system for target tracking

What is a sensor, from a tracking perspective?

Sensor system for target tracking

What is a sensor, from a tracking perspective?

- On top of the detection level: produces a *set of observations* Z_t at each time t

Sensor system for target tracking

What is a sensor, from a tracking perspective?

- On top of the detection level: produces a *set of observations* Z_t at each time t
- Known by the operator through a *stochastic description*


Sensor system for target tracking

What is a sensor, from a tracking perspective?

- On top of the detection level: produces a *set of observations* Z_t at each time t
- Known by the operator through a *stochastic description*

Stochastic description

Likelihood $\ell_t(z, x)$: how likely is obs. z to come from a target with state x ?



Sensor system for target tracking


What is a sensor, from a tracking perspective?

- On top of the detection level: produces a *set of observations* Z_t at each time t
- Known by the operator through a *stochastic description*

Stochastic description

Likelihood $\ell_t(z, x)$: how likely is obs. z to come from a target with state x ?

Probability of detection $p_{d,t}(x)$: how likely is a target with state x to be detected?



Sensor system for target tracking

What is a sensor, from a tracking perspective?


- On top of the detection level: produces a *set of observations* Z_t at each time t
- Known by the operator through a *stochastic description*

Stochastic description

Likelihood $\ell_t(z, x)$: how likely is obs. z to come from a target with state x ?

Probability of detection $p_{d,t}(x)$: how likely is a target with state x to be detected?

Probability of false alarm $p_{fa,t}(z)$: how likely is the sensor to produce a false alarm with state z ?




Sensor system for target tracking (cont.)

Sensor system for target tracking (cont.)


 \mathbf{Z}_t

Sensor system for target tracking (cont.)

 \mathbf{Z}_t


- *Discrete* observation space \mathbf{Z}_t

Sensor system for target tracking (cont.)

 \mathbf{Z}_t

- *Discrete* observation space \mathbf{Z}_t
- *Localized* false alarm process: at most one false alarm, per cell and per scan


Sensor system for target tracking (cont.)



\mathbf{Z}_t

- *Discrete* observation space \mathbf{Z}_t
- *Localized* false alarm process: at most one false alarm, per cell and per scan
- In each cell $z \in \mathbf{Z}_t$, false alarm occurs with probability $p_{fa,t}(z)$


Sensor system for target tracking (cont.)



\mathbf{Z}_t


- *Discrete* observation space \mathbf{Z}_t
- *Localized* false alarm process: at most one false alarm, per cell and per scan
- In each cell $z \in \mathbf{Z}_t$, false alarm occurs with probability $p_{fa,t}(z)$

- \mathbf{Z}_t projected onto \mathbf{X} shapes the sensor field of view (FoV)



\mathbf{X}


Sensor system for target tracking (cont.)



\mathbf{Z}_t

- *Discrete* observation space \mathbf{Z}_t
- *Localized* false alarm process: at most one false alarm, per cell and per scan
- In each cell $z \in \mathbf{Z}_t$, false alarm occurs with probability $p_{fa,t}(z)$

- \mathbf{Z}_t projected onto \mathbf{X} shapes the sensor field of view (FoV)
- Outside of the sensor FoV, $p_{d,t}$ is always zero (i.e. no target detection)



\mathbf{X}

Multi-object filtering: common assumptions

Common assumptions (time t)

1. Targets behave independently
2. Observations are produced independently
3. At most one observation per target (if none, target is *miss-detected*)
4. At most one target per observation (if none, obs. is a *false alarm*)

Multi-object filtering: common assumptions

Common assumptions (time t)

1. Targets behave independently
2. Observations are produced independently
3. At most one observation per target (if none, target is *miss-detected*)
4. At most one target per observation (if none, obs. is a *false alarm*)

The assumptions above...

- 1. ... simplify the estimation problem (notably the data association)
- 2. ... will be used in the context of this presentation
- 3. ... are *not* necessary in the general multi-object estimation framework

- 1 Multi-object filtering framework: basics
- 2 Traditional solutions: strengths and weaknesses
 - The track-based approach
 - The RFS-based approach
- 3 Stochastic populations for multi-object filtering

The track-based approach

General principle

“A potential target = one track.”

The track-based approach

General principle

“A potential target = one track.”

Track representation




X

The track-based approach

General principle

“A potential target = one track.”

Track representation



A track y is...


- ... *identified* by its state distribution

The track-based approach

General principle

“A potential target = one track.”

Track representation



X

A track y is...


- ... *identified* by its state distribution (e.g. mean)

The track-based approach

General principle

“A potential target = one track.”

Track representation



X

A track y is...


- ... *identified* by its state distribution (e.g. mean + covariance)

The track-based approach

General principle

“A potential target = one track.”

Track representation



X

A track y is...


- ... *identified* by its state distribution (e.g. mean + covariance)
- ... *described* by its history of past estimates

The track-based approach

General principle

“A potential target = one track.”

Track representation




X

A track y is...

- ... *identified* by its state distribution (e.g. mean + covariance)
- ... *described* by its history of past estimates
- ... *characterised* by its observation path


Track update

Data association


**X**

Track update

Data association





X




X

Track update

Data association




X




X

Track update


Data association




X



X




X




X

Track update


Data association




X




X



X




X




X

Track update


Data association




X




X




X



X




X




X

Track update


Data association




X




X




X



X




X




X

Track update


Data association




X




X




X




X



X




X




X

Track update


Data association




X




X




X




X




X



X




X




X

Track update


Data association




X




X




X




X




X




X



X




X




X

Track update


Data association




X




X




X




X




X




X




X



X




X




X

Track update


Data association




X




X




X




X




X




X




X



X




X




X

Track update


Data association




X




X




X




X




X




X




X



X



X




X


What is propagated to the next step?

Track update


Data association




X




X




X




X




X




X




X



X



X




X

What is propagated to the next step?


All configurations, with associated probabilities → MHT filter

Track update


Data association




X




X




X




X




X




X




X



X



X



X

What is propagated to the next step?

All configurations, with associated probabilities \rightarrow MHT filter

A weighted combination of all configurations \rightarrow JPDA filter

Challenges


Challenges

Since tracks are created upon detections, how to model *yet-to-be-detected* targets?

Challenges

Since tracks are created upon detections, how to model *yet-to-be-detected* targets?


Sensor management problem: explore B_1 or B_2 ?



Challenges


Since tracks are created upon detections, how to model *yet-to-be-detected* targets?

Sensor management problem: explore B_1 or B_2 ?



Challenges

Since tracks are created upon detections, how to model *yet-to-be-detected* targets?

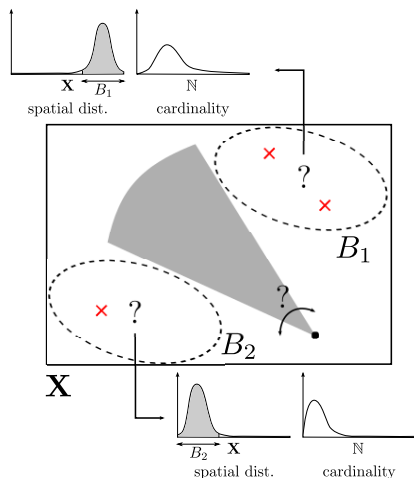


Sensor management problem: explore B_1 or B_2 ?

Suppose prior information on target population in B_1 and B_2 is available

Challenges

Since tracks are created upon detections, how to model *yet-to-be-detected* targets?



Sensor management problem: explore B_1 or B_2 ?

Suppose prior information on target *population* in B_1 and B_2 is available

How many tracks to create? Where?

Challenges (cont.)

Track creation/deletion

Challenges (cont.)

Track creation/deletion

- When to create tracks?

Challenges (cont.)

Track creation/deletion

- When to create tracks? For every new observation?

Challenges (cont.)

Track creation/deletion

- When to create tracks? For every new observation? For a sequence of “close” unassociated observations?

Challenges (cont.)

Track creation/deletion

- When to create tracks? For every new observation? For a sequence of “close” unassociated observations?
- When to delete tracks?

Challenges (cont.)

Track creation/deletion

- When to create tracks? For every new observation? For a sequence of “close” unassociated observations?
- When to delete tracks? If miss-detected n times during the m last time steps?

Challenges (cont.)

Track creation/deletion

- When to create tracks? For every new observation? For a sequence of “close” unassociated observations?
- When to delete tracks? If miss-detected n times during the m last time steps? If getting away from the surveillance scene?

Challenges (cont.)


Track creation/deletion

- When to create tracks? For every new observation? For a sequence of “close” unassociated observations?
- When to delete tracks? If miss-detected n times during the m last time steps? If getting away from the surveillance scene?
- *Why* deleting a track?

Challenges (cont.)

Track creation/deletion

- When to create tracks? For every new observation? For a sequence of “close” unassociated observations?
- When to delete tracks? If miss-detected n times during the m last time steps? If getting away from the surveillance scene?
- *Why* deleting a track?




X

Because it is “unlikely” to represent a target?


Challenges (cont.)

Track creation/deletion

- When to create tracks? For every new observation? For a sequence of “close” unassociated observations?
- When to delete tracks? If miss-detected n times during the m last time steps? If getting away from the surveillance scene?
- *Why* deleting a track?



X



X

Because it is “unlikely” to represent a target?

Because the target it represents has “probably” left the scene?

The track-based approach: pros and cons

The track-based approach: pros and cons

Can handle spatial information on tracks “optimally”

- Data association allows optimal single-observation/single-track update (e.g. Kalman filter)
- History of past estimates and observation path naturally maintained

The track-based approach: pros and cons

Can handle spatial information on tracks “optimally”

- Data association allows optimal single-observation/single-track update (e.g. Kalman filter)
- History of past estimates and observation path naturally maintained

Lacks probabilistic framework incorporating all system uncertainties

- When to create a track, especially in high clutter environment?
- When and why to delete a track (i.e. what is an “unlikely” track?)

The track-based approach: pros and cons

Can handle spatial information on tracks “optimally”

- Data association allows optimal single-observation/single-track update (e.g. Kalman filter)
- History of past estimates and observation path naturally maintained

Lacks probabilistic framework incorporating all system uncertainties

- When to create a track, especially in high clutter environment?
- When and why to delete a track (i.e. what is an “unlikely” track?)

Lacks natural description on *collective* rather than *individual* level

- Populations of “non-separable” targets (e.g. those yet-to-be-detected) not easily described
- Regional statistics (e.g. mean, variance in target number) unavailable

The RFS-based approach

General principle


“The population of potential targets = one random finite set (RFS).”

The RFS-based approach

General principle

“The population of potential targets = one random finite set (RFS).”

RFS representation



X

The target RFS Ξ ...


- ... is a random object describing *all* the targets in the scene

The RFS-based approach

General principle

“The population of potential targets = one random finite set (RFS).”

RFS representation



X

The target RFS Ξ ...


- ... is a random object describing *all* the targets in the scene
- ... whose realizations $X = \{x_1, \dots, x_n\}$ describe potential multi-target configurations

The RFS-based approach

General principle

“The population of potential targets = one random finite set (RFS).”

RFS representation



X

The target RFS Ξ ...


- ... is a random object describing *all* the targets in the scene
- ... whose realizations $X = \{x_1, \dots, x_n\}$ describe potential multi-target configurations

The RFS-based approach

General principle

“The population of potential targets = one random finite set (RFS).”

RFS representation



$$\mathbf{X} \quad P_{\Xi}(X) = 0.1$$

$$P_{\Xi}(X') = 0.03$$

...


The target RFS Ξ ...

- ... is a random object describing *all* the targets in the scene
- ... whose realizations $X = \{x_1, \dots, x_n\}$ describe potential multi-target configurations
- ... is described by *probability distribution* P_{Ξ}

First-moment density for usual RFS-based filters

First-moment density for usual RFS-based filters


First-moment density μ_{Ξ}



Approximate description of RFS Ξ

First-moment density for usual RFS-based filters

First-moment density μ_{Ξ}




Approximate description of RFS Ξ

Propagated in usual RFS filters (PHD, CPHD)

First-moment density for usual RFS-based filters

First-moment density μ_{Ξ}




Approximate description of RFS Ξ

Propagated in usual RFS filters (PHD, CPHD)

Provides average number of target per volume space, acc. to RFS Ξ

First-moment density for usual RFS-based filters

First-moment density μ_{Ξ}



Approximate description of RFS Ξ


Propagated in usual RFS filters (PHD, CPHD)

Provides average number of target per volume space, acc. to RFS Ξ

Data update

First-moment density for usual RFS-based filters

First-moment density μ_{Ξ}




Approximate description of RFS Ξ

Propagated in usual RFS filters (PHD, CPHD)


Provides average number of target per volume space, acc. to RFS Ξ

Data update



First-moment density for usual RFS-based filters

First-moment density μ_{Ξ}





Approximate description of RFS Ξ

Propagated in usual RFS filters (PHD, CPHD)


Provides average number of target per volume space, acc. to RFS Ξ

Data update



First-moment density for usual RFS-based filters

First-moment density μ_{Ξ}






Approximate description of RFS Ξ

Propagated in usual RFS filters (PHD, CPHD)


Provides average number of target per volume space, acc. to RFS Ξ

Data update



First-moment density for usual RFS-based filters

First-moment density μ_{Ξ}







Approximate description of RFS Ξ

Propagated in usual RFS filters (PHD, CPHD)

Provides average number of target per volume space, acc. to RFS Ξ


Data update



Challenges


Challenges

No equivalent of track history: what are the consequences?




Challenges

No equivalent of track history: what are the consequences?




Challenges

No equivalent of track history: what are the consequences?



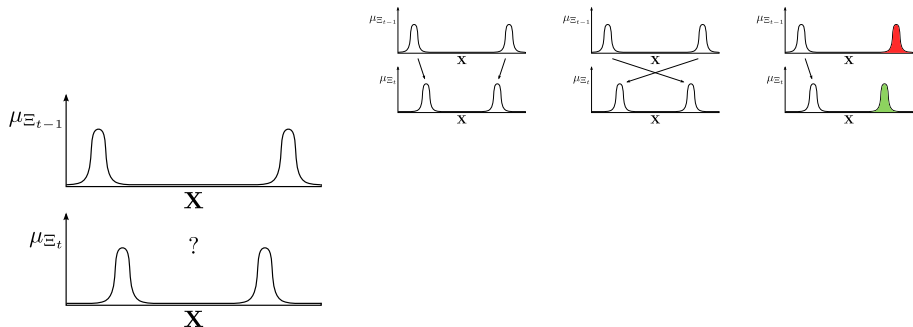
Challenges

No equivalent of track history: what are the consequences?



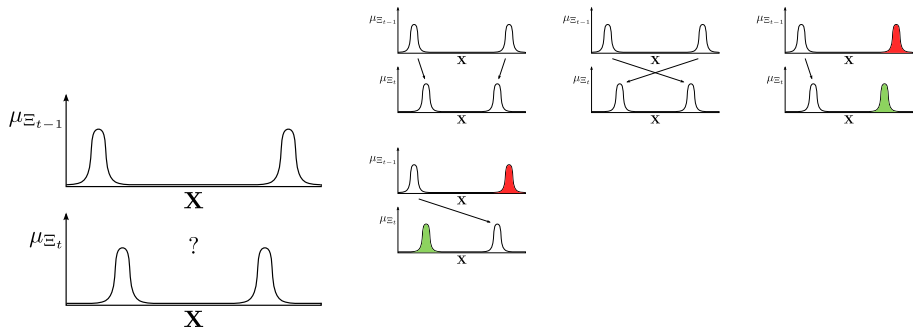
Challenges

No equivalent of track history: what are the consequences?



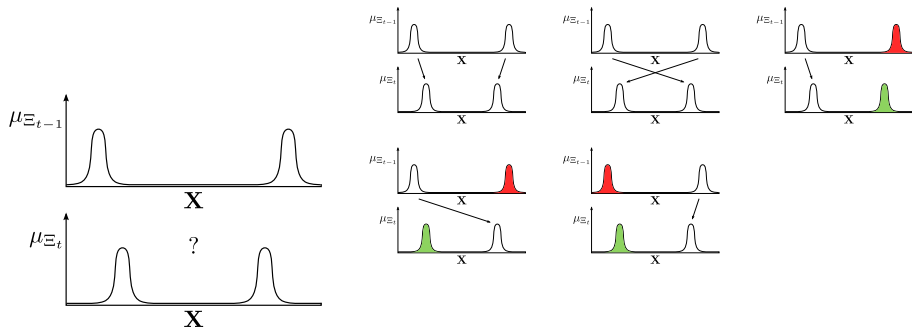
Challenges

No equivalent of track history: what are the consequences?



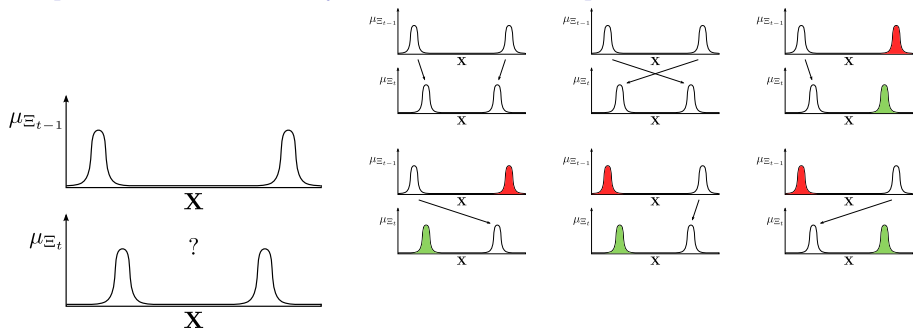
Challenges

No equivalent of track history: what are the consequences?



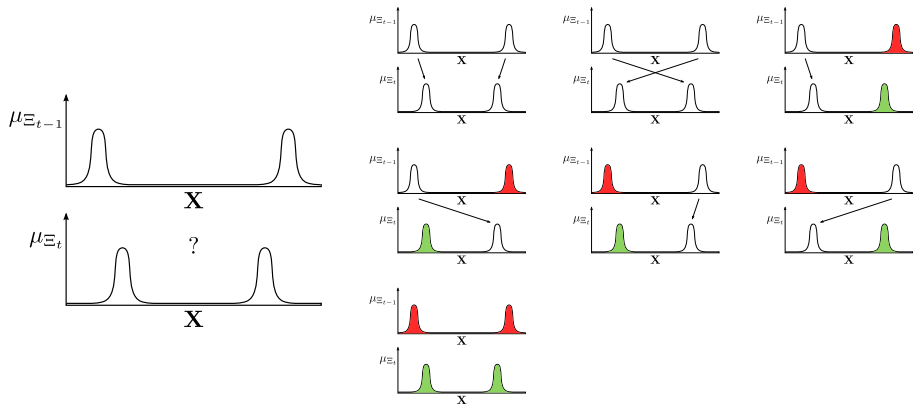
Challenges

No equivalent of track history: what are the consequences?



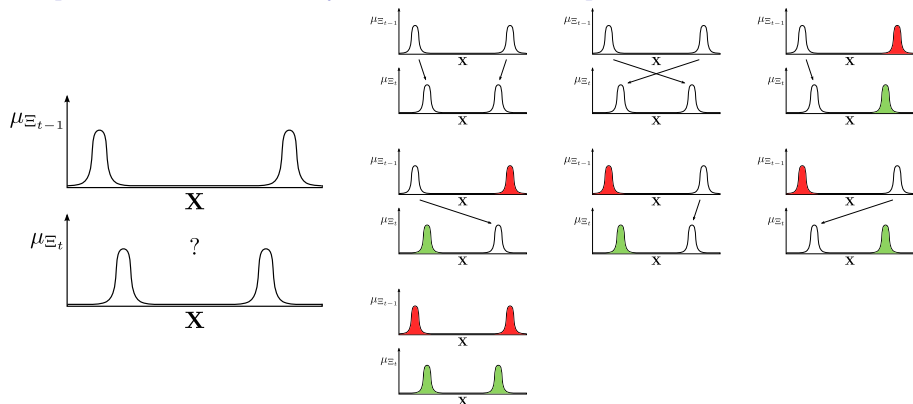
Challenges

No equivalent of track history: what are the consequences?



Challenges

No equivalent of track history: what are the consequences?



No inherent solution to link individuals from successive populations

Introducing labelling on top of RFS framework recently explored (Labeled Multi-Bernoulli filter, Vo et al.)

The RFS-based approach: pros and cons

The RFS-based approach: pros and cons

Provides probability framework incorporating all system uncertainties

- Appearing targets, yet-to-be-detected targets, false alarms, etc. described by RFSs
- No need for track creation/deletion, for data association
- Regional statistics (e.g. mean, variance in target number) naturally available

The RFS-based approach: pros and cons

Provides probability framework incorporating all system uncertainties

- Appearing targets, yet-to-be-detected targets, false alarms, etc. described by RFSs
- No need for track creation/deletion, for data association
- Regional statistics (e.g. mean, variance in target number) naturally available

Lacks natural description on *individual* rather than *collective* level

- Each observation influences the spatial distribution of the whole population (i.e. no “optimal” single-observation/single-target update)
- Track histories or observation paths unavailable (unless through heuristics)

The RFS-based approach: pros and cons

Provides probability framework incorporating all system uncertainties

- Appearing targets, yet-to-be-detected targets, false alarms, etc. described by RFSs
- No need for track creation/deletion, for data association
- Regional statistics (e.g. mean, variance in target number) naturally available

Lacks natural description on *individual* rather than *collective* level

- Each observation influences the spatial distribution of the whole population (i.e. no “optimal” single-observation/single-target update)
- Track histories or observation paths unavailable (unless through heuristics)

Lacks intuitive interpretation (on some points)

- Poisson (resp. i.i.d.) approximation in PHD (resp. CPHD) filter

- 1 Multi-object filtering framework: basics
- 2 Traditional solutions: strengths and weaknesses
- 3 Stochastic populations for multi-object filtering
 - Multi-object estimation framework
 - Bayesian filtering
 - Closed-loop sensor management

Estimation framework for stochastic populations

General principle

“A potential target is represented by a specific amount of information:
not too little, but *not too much either*.”

Estimation framework for stochastic populations


General principle

“A potential target is represented by a specific amount of information: not too little, but *not too much either*.”

Outline


- Well-defined probabilistic framework, developed by J. Houssineau (supervisor: D. Clark)
- Level of description depends whether individual is *distinguishable* or *indistinguishable*
- Ongoing developments beyond tracking (e.g. sensor management, sensor calibration, performance assessment, ...)

Target distinguishability



X


Target distinguishability



Indistinguishable target:

- *unidentified* member of larger population
- no specific information available
- e.g., “*one of the potential individuals that entered 10 time steps ago and has not been detected yet*”

Target distinguishability



Indistinguishable target:

- *unidentified* member of larger population
- no specific information available
- e.g., “*one of the potential individuals that entered 10 time steps ago and has not been detected yet*”

Distinguishable target:

- individual *characterised* by specific information
- e.g., “*the potential individual that entered 10 time steps ago and produced observations z_{21}^3, z_{25}^1* ”

Individual management: where are the targets?

Spatial distribution and “empty state” ψ

Individual management: where are the targets?

Spatial distribution and “empty state” ψ


$$p_t^y$$

What is the target state?

Individual management: where are the targets?

Spatial distribution and “empty state” ψ

$$p_t^y =$$



What is the target state?

If the target is the scene,
what is its state?

Individual management: where are the targets?

Spatial distribution and “empty state” ψ

$$p_t^y = \underbrace{\frac{p_t^y(\cdot)}{p_t^y(\mathbf{X})}}_{\mathbf{X}} + p_t^y(\{\psi\})$$

What is the target state?

If the target is the scene,
what is its state?

Is the target in the scene?

Individual management: where are the targets?

Spatial distribution and “empty state” ψ

$$p_t^y = \underbrace{\frac{p_t^y(\cdot)}{p_t^y(\mathbf{X})}}_{\mathbf{X}} + p_t^y(\{\psi\})$$

What is the target state?

If the target is the scene,
what is its state?

Is the target in the scene?

Distinguishable target and associated track

Individual management: where are the targets?

Spatial distribution and “empty state” ψ

$$p_t^y = \underbrace{\frac{p_t^y(\cdot)}{p_t^y(\mathbf{X})}}_{\mathbf{X}} + p_t^y(\{\psi\})$$

What is the target state?

If the target is the scene,
what is its state?

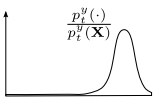
Is the target in the scene?

Distinguishable target and associated track

- *First detection* triggers distinguishability

Individual management: where are the targets?

Spatial distribution and “empty state” ψ

$$p_t^y = \int \frac{p_t^y(\cdot)}{p_t^y(\mathbf{X})} \psi(\cdot) d\mathbf{X} + p_t^y(\{\psi\})$$



What is the target state?

If the target is in the scene,
what is its state?

Is the target in the scene?

Distinguishable target and associated track

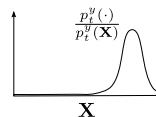
- *First detection* triggers distinguishability
- Associated track y describes observation path and current info on target state



\mathbf{X}

Individual management: where are the targets?

Spatial distribution and “empty state” ψ

$$p_t^y = \int \frac{p_t^y(\cdot)}{p_t^y(\mathbf{X})} \delta(\mathbf{X} - \cdot) + p_t^y(\{\psi\})$$



What is the target state?

If the target is the scene,
what is its state?

Is the target in the scene?

Distinguishable target and associated track

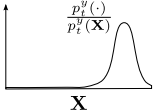
- *First detection* triggers distinguishability
- Associated track y describes observation path and current info on target state
- Usual notion of track, except probability of absence $p_t^y(\{\psi\})$



\mathbf{X}

Individual management: where are the targets?

Spatial distribution and “empty state” ψ

$$p_t^y = \int \frac{p_t^y(\cdot)}{p_t^y(\mathbf{X})} p_t^y(\{\psi\}) d\mathbf{X}$$


What is the target state?


If the target is the scene,
what is its state?

Is the target in the scene?

Distinguishable target and associated track

- *First detection* triggers distinguishability
- Associated track y describes observation path and current info on target state
- Usual notion of track, except probability of absence $p_t^y(\{\psi\})$

Note: $1 - p_t^y(\{\psi\})$ does *not* assess track credibility, but its presence in the scene




\mathbf{X}

Population management: which are the true targets?

Joint probability of existence

Assesses *joint* existence of targets:




X

Population management: which are the true targets?

Joint probability of existence

Assesses *joint* existence of targets:


- $P(\emptyset, 0)$



X

Population management: which are the true targets?

Joint probability of existence




Assesses *joint* existence of targets:

- $P(\emptyset, 0)$
- $P(\{y\}, 1)$

Population management: which are the true targets?

Joint probability of existence




Assesses *joint* existence of targets:

- $P(\emptyset, 0)$
- $P(\{y\}, 1)$
- $P(\{y, y'\}, 2)$

Population management: which are the true targets?

Joint probability of existence




X

Assesses *joint* existence of targets:

- $P(\emptyset, 0)$
- $P(\{y\}, 1)$
- $P(\{y, y'\}, 2)$
- $P(\emptyset, 1)$
- ...

Population management: which are the true targets?

Joint probability of existence



X

Assesses *joint* existence of targets:


- $P(\emptyset, 0)$
- $P(\{y\}, 1)$
- $P(\{y, y'\}, 2)$
- $P(\emptyset, 1)$
- ...

At time t :

- Set of tracks Y_t , m_t populations of indistinguishable targets

Population management: which are the true targets?

Joint probability of existence



X

Assesses *joint* existence of targets:

- $P(\emptyset, 0)$
- $P(\{y\}, 1)$
- $P(\{y, y'\}, 2)$
- $P(\emptyset, 1)$
- ...


At time t :

- Set of tracks Y_t , m_t populations of indistinguishable targets
- $\sum_{Y \subseteq Y_t} \sum_{n_1 \geq 0} \cdots \sum_{n_{m_t} \geq 0} P_t(Y, n_1, \dots, n_{m_t}) = 1$

Individual+population management: what is the true multi-target configuration?


Individual+population management: what is the true multi-target configuration?

Elementary events



Individual+population management: what is the true multi-target configuration?


Elementary events



Probability that:

Individual+population management: what is the true multi-target configuration?

Elementary events




Probability that:

- y exists and lies within B , and

Individual+population management: what is the true multi-target configuration?

Elementary events




Probability that:

- y exists and lies within B , and
- y' exists, and

Individual+population management: what is the true multi-target configuration?

Elementary events




Probability that:

- y exists and lies within B , and
- y' exists, and
- one undetected targets in B , one in B' , none elsewhere.

Individual+population management: what is the true multi-target configuration?


Elementary events



Probability that:

Individual+population management: what is the true multi-target configuration?

Elementary events




Probability that:

- y does not exist, *and*

Individual+population management: what is the true multi-target configuration?

Elementary events




Probability that:

- y does not exist, *and*
- y' exists, but has left the scene, *and*

Individual+population management: what is the true multi-target configuration?

Elementary events




Probability that:

- y does not exist, *and*
- y' exists, but has left the scene, *and*
- two undetected targets in B , none elsewhere.

Individual+population management: what is the true multi-target configuration?

Elementary events




Probability that:

- y does not exist, *and*
- y' exists, but has left the scene, *and*
- two undetected targets in B , none elsewhere.

Compound events

Individual+population management: what is the true multi-target configuration?

Elementary events



Probability that:


- y does not exist, *and*
- y' exists, but has left the scene, *and*
- two undetected targets in B , none elsewhere.

Compound events

- Probability that there are *at least* two undetected targets in B

Individual+population management: what is the true multi-target configuration?

Elementary events



Probability that:


- y does not exist, *and*
- y' exists, but has left the scene, *and*
- two undetected targets in B , none elsewhere.

Compound events

- Probability that there are *at least* two undetected targets in B
- Probability that *either* y exists *or* y' has left the scene, *and* there are no undetected targets in the scene

Individual+population management: what is the true multi-target configuration?

Elementary events



Probability that:


- y does not exist, *and*
- y' exists, but has left the scene, *and*
- two undetected targets in B , none elsewhere.

Compound events

- Probability that there are *at least* two undetected targets in B
- Probability that *either* y exists *or* y' has left the scene, *and* there are no undetected targets in the scene
-

Bayesian filtering with stochastic populations (1/2)


Principle



Maintain and update a dynamical stochastic population reflecting the true multi-target configuration

Bayesian filtering with stochastic populations (1/2)

Principle




Maintain and update a dynamical stochastic population reflecting the true multi-target configuration

Filtering design: a few modelling choices

Bayesian filtering with stochastic populations (1/2)


Principle



Maintain and update a dynamical stochastic population reflecting the true multi-target configuration


Filtering design: a few modelling choices

- How to represent incoming targets at time t ?




Bayesian filtering with stochastic populations (1/2)

Principle



Maintain and update a dynamical stochastic population reflecting the true multi-target configuration


Filtering design: a few modelling choices



- How to represent incoming targets at time t ? One ind. population?


Bayesian filtering with stochastic populations (1/2)

Principle



Maintain and update a dynamical stochastic population reflecting the true multi-target configuration


Filtering design: a few modelling choices



- How to represent incoming targets at time t ? One ind. population? Two?


Bayesian filtering with stochastic populations (1/2)

Principle



Maintain and update a dynamical stochastic population reflecting the true multi-target configuration

Filtering design: a few modelling choices




X

- How to represent incoming targets at time t ? One ind. population? Two?
- How to mitigate information loss?


Bayesian filtering with stochastic populations (1/2)

Principle



Maintain and update a dynamical stochastic population reflecting the true multi-target configuration


Filtering design: a few modelling choices



- How to represent incoming targets at time t ? One ind. population? Two?
- How to mitigate information loss? Merge “close” ind. populations?


Bayesian filtering with stochastic populations (1/2)

Principle



Maintain and update a dynamical stochastic population reflecting the true multi-target configuration

Filtering design: a few modelling choices




X

- How to represent incoming targets at time t ? One ind. population? Two?
- How to mitigate information loss? Merge “close” ind. populations?


Bayesian filtering with stochastic populations (1/2)

Principle



Maintain and update a dynamical stochastic population reflecting the true multi-target configuration


Filtering design: a few modelling choices



- How to represent incoming targets at time t ? One ind. population? Two?
- How to mitigate information loss?
Merge “close” ind. populations?


Bayesian filtering with stochastic populations (1/2)

Principle



Maintain and update a dynamical stochastic population reflecting the true multi-target configuration


Filtering design: a few modelling choices



- How to represent incoming targets at time t ? One ind. population? Two?
- How to mitigate information loss? Merge “close” ind. populations? Merge “close” tracks?


Bayesian filtering with stochastic populations (1/2)

Principle



Maintain and update a dynamical stochastic population reflecting the true multi-target configuration


Filtering design: a few modelling choices



- How to represent incoming targets at time t ? One ind. population? Two?
- How to mitigate information loss?
Merge “close” ind. populations? Merge “close” tracks?


Bayesian filtering with stochastic populations (1/2)

Principle



Maintain and update a dynamical stochastic population reflecting the true multi-target configuration

Filtering design: a few modelling choices




X

- How to represent incoming targets at time t ? One ind. population? Two?
- How to mitigate information loss?
Merge “close” ind. populations? Merge “close” tracks?


Bayesian filtering with stochastic populations (1/2)

Principle



Maintain and update a dynamical stochastic population reflecting the true multi-target configuration


Filtering design: a few modelling choices



- How to represent incoming targets at time t ? One ind. population? Two?
- How to mitigate information loss? Merge “close” ind. populations? Merge “close” tracks? “Blend” a track *into* a “close” ind. population?


Bayesian filtering with stochastic populations (1/2)

Principle



Maintain and update a dynamical stochastic population reflecting the true multi-target configuration


Filtering design: a few modelling choices



- How to represent incoming targets at time t ? One ind. population? Two?
- How to mitigate information loss? Merge “close” ind. populations? Merge “close” tracks? “Blend” a track *into* a “close” ind. population?


Bayesian filtering with stochastic populations (1/2)

Principle



Maintain and update a dynamical stochastic population reflecting the true multi-target configuration

Filtering design: a few modelling choices




- How to represent incoming targets at time t ? One ind. population? Two?
- How to mitigate information loss? Merge “close” ind. populations? Merge “close” tracks? “Blend” a track *into* a “close” ind. population?

Bayesian filtering with stochastic populations (2/2)

Bayesian filtering with stochastic populations (2/2)


Prediction step



Bayesian filtering with stochastic populations (2/2)

Prediction step


- Time update of spatial distributions



Bayesian filtering with stochastic populations (2/2)

Prediction step


- Time update of spatial distributions
- Single-object time update \rightarrow KF, EKF, UKF, SMC...



Bayesian filtering with stochastic populations (2/2)

Prediction step


- Time update of spatial distributions
- Single-object time update \rightarrow KF, EKF, UKF, SMC...
- No additional observations \rightarrow *no update* of joint probabilities of existence



Bayesian filtering with stochastic populations (2/2)

Prediction step

- Time update of spatial distributions
- Single-object time update \rightarrow KF, EKF, UKF, SMC...
- No additional observations \rightarrow *no update* of joint probabilities of existence




Data update step


Bayesian filtering with stochastic populations (2/2)

Prediction step

- Time update of spatial distributions
- Single-object time update \rightarrow KF, EKF, UKF, SMC...
- No additional observations \rightarrow *no update* of joint probabilities of existence



Data update step




- Data association:

Update of joint probabilities of existence
 Pop. management (indist. \rightarrow tracks, etc.)


Bayesian filtering with stochastic populations (2/2)

Prediction step

- Time update of spatial distributions
- Single-object time update \rightarrow KF, EKF, UKF, SMC...
- No additional observations \rightarrow *no update* of joint probabilities of existence



Data update step



- Data association:


Update of joint probabilities of existence
Pop. management (indist. \rightarrow tracks, etc.)

- Data update of spatial distributions


Bayesian filtering with stochastic populations (2/2)

Prediction step

- Time update of spatial distributions
- Single-object time update \rightarrow KF, EKF, UKF, SMC...
- No additional observations \rightarrow *no update* of joint probabilities of existence



Data update step



- Data association:

Update of joint probabilities of existence
Pop. management (indist. \rightarrow tracks, etc.)

- Data update of spatial distributions
- Single-object data update: \rightarrow KF, EKF, UKF, SMC...

The filter for Distinguishable and Independent Stochastic Populations (DISP)

The filter for Distinguishable and Independent Stochastic Populations (DISP)

Assumptions on the incoming targets, at every time step t :

The filter for Distinguishable and Independent Stochastic Populations (DISP)

Assumptions on the incoming targets, at every time step t :

1. Described by a *single* indistinguishable population

The filter for Distinguishable and Independent Stochastic Populations (DISP)


Assumptions on the incoming targets, at every time step t :

1. Described by a *single* indistinguishable population
2. Detected *at once* (i.e., they become *immediately* distinguishable)

The filter for Distinguishable and Independent Stochastic Populations (DISP)

Assumptions on the incoming targets, at every time step t :


1. Described by a *single* indistinguishable population
2. Detected *at once* (i.e., they become *immediately* distinguishable)



The filter for Distinguishable and Independent Stochastic Populations (DISP)

Assumptions on the incoming targets, at every time step t :

1. Described by a *single* indistinguishable population
2. Detected *at once* (i.e., they become *immediately* distinguishable)




Practical considerations

The filter for Distinguishable and Independent Stochastic Populations (DISP)

Assumptions on the incoming targets, at every time step t :

1. Described by a *single* indistinguishable population
2. Detected *at once* (i.e., they become *immediately* distinguishable)




Practical considerations

- No propagated information on yet-to-be-detected targets

The filter for Distinguishable and Independent Stochastic Populations (DISP)

Assumptions on the incoming targets, at every time step t :

1. Described by a *single* indistinguishable population
2. Detected *at once* (i.e., they become *immediately* distinguishable)




Practical considerations

- No propagated information on yet-to-be-detected targets
- Full data association \rightarrow information on every possible track is maintained

The filter for Distinguishable and Independent Stochastic Populations (DISP)

Assumptions on the incoming targets, at every time step t :

1. Described by a *single* indistinguishable population
2. Detected *at once* (i.e., they become *immediately* distinguishable)




Practical considerations

- No propagated information on yet-to-be-detected targets
- Full data association \rightarrow information on every possible track is maintained
- High computational cost (\approx MHT) \rightarrow parallel implementation ongoing

The filter for Distinguishable and Independent Stochastic Populations (DISP)

Assumptions on the incoming targets, at every time step t :

1. Described by a *single* indistinguishable population
2. Detected *at once* (i.e., they become *immediately* distinguishable)



Practical considerations

- No propagated information on yet-to-be-detected targets
- Full data association \rightarrow information on every possible track is maintained
- High computational cost (\approx MHT) \rightarrow parallel implementation ongoing
- Approximate DISP: HISP, similar cost to PHD but better tracking performances (early analysis)

Closed-loop sensor management: principle

Closed-loop sensor management: principle

- U_t : pool of available *actions* at time t

Closed-loop sensor management: principle

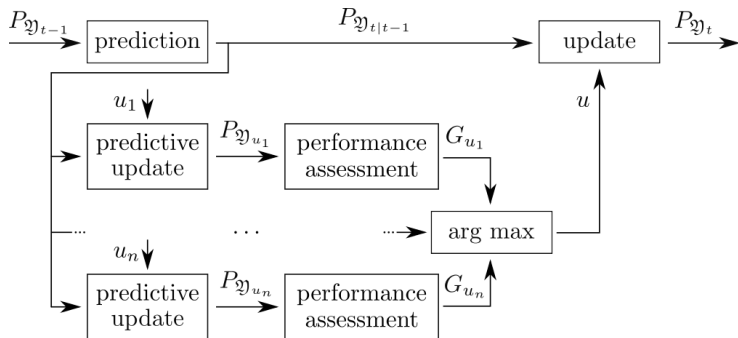
- U_t : pool of available *actions* at time t
- $u_i \in U_t$: describes a specific sensor (i.e. $\ell_{u_i}, p_{d,u_i}, p_{fa,u_i}, \mathbf{Z}_{u_i}$)

Closed-loop sensor management: principle

- U_t : pool of available *actions* at time t
- $u_i \in U_t$: describes a specific sensor (i.e. $\ell_{u_i}, p_{d,u_i}, p_{fa,u_i}, \mathbf{Z}_{u_i}$)
- Question: which action u_i should be selected to collect the “best” observations Z_t ?

Closed-loop sensor management: principle

- U_t : pool of available *actions* at time t
- $u_i \in U_t$: describes a specific sensor (i.e. $\ell_{u_i}, p_{d,u_i}, p_{fa,u_i}, \mathbf{Z}_{u_i}$)
- Question: which action u_i should be selected to collect the “best” observations Z_t ?



Construction of information gain: principle

Construction of information gain: principle

Objective

The operator has access to the predicted information $P_{\mathfrak{Y}_{t|t-1}}$ and considers some action $u \in U_t$ for the next observation. Can we quantify the expected information gain G_u of action u ?

Construction of information gain: principle

Objective

The operator has access to the predicted information $P_{\mathfrak{Y}_{t|t-1}}$ and considers some action $u \in U_t$ for the next observation. Can we quantify the expected information gain G_u of action u ?

Outline

1. Suppose that observations $Z \subseteq \mathbf{Z}_u$ are collected. How does the updated information $P_{\mathfrak{Y}_t}(\cdot|Z)$ look like?

Construction of information gain: principle

Objective

The operator has access to the predicted information $P_{\mathfrak{Y}_{t|t-1}}$ and considers some action $u \in U_t$ for the next observation. Can we quantify the expected information gain G_u of action u ?

Outline

1. Suppose that observations $Z \subseteq \mathbf{Z}_u$ are collected. How does the updated information $P_{\mathfrak{Y}_t}(\cdot|Z)$ look like?
2. How much does the operator learn from $P_{\mathfrak{Y}_{t|t-1}}$ to $P_{\mathfrak{Y}_t}(\cdot|Z)$?

Construction of information gain: principle

Objective

The operator has access to the predicted information $P_{\mathfrak{Y}_{t|t-1}}$ and considers some action $u \in U_t$ for the next observation. Can we quantify the expected information gain G_u of action u ?

Outline

1. Suppose that observations $Z \subseteq \mathbf{Z}_u$ are collected. How does the updated information $P_{\mathfrak{Y}_t}(\cdot|Z)$ look like?
2. How much does the operator learn from $P_{\mathfrak{Y}_{t|t-1}}$ to $P_{\mathfrak{Y}_t}(\cdot|Z)$?
3. How much can he expect to learn from $P_{\mathfrak{Y}_{t|t-1}}$ if he chooses action u ?

Construction of information gain: principle


Objective

The operator has access to the predicted information $P_{\mathfrak{Y}_{t|t-1}}$ and considers some action $u \in U_t$ for the next observation. Can we quantify the expected information gain G_u of action u ?

Outline

1. Suppose that observations $Z \subseteq \mathbf{Z}_u$ are collected. How does the updated information $P_{\mathfrak{Y}_t}(\cdot|Z)$ look like?
2. How much does the operator learn from $P_{\mathfrak{Y}_{t|t-1}}$ to $P_{\mathfrak{Y}_t}(\cdot|Z)$?
3. How much can he expect to learn from $P_{\mathfrak{Y}_{t|t-1}}$ if he chooses action u ?

Filtering method: DISP



First things first... what does the operator know beforehand, i.e. what does $P_{\mathcal{Y}_{t|t-1}}$ look like?

First things first... what does the operator know beforehand, i.e. what does $P_{\mathcal{Y}_{t|t-1}}$ look like?

Incoming targets (indistinguishable)

First things first... what does the operator know beforehand, i.e. what does $P_{\mathcal{Y}_{t|t-1}}$ look like?

Incoming targets (indistinguishable)

- Per assumption, a *single* population describe *all* the incoming targets

First things first... what does the operator know beforehand, i.e. what does $P_{\mathcal{Y}_{t|t-1}}$ look like?

Incoming targets (indistinguishable)

- Per assumption, a *single* population describe *all* the incoming targets
- *All* individuals are described by a *unique* spatial distribution p_t^a

First things first... what does the operator know beforehand, i.e. what does $P_{\mathcal{Y}_{t|t-1}}$ look like?

Incoming targets (indistinguishable)

- Per assumption, a *single* population describe *all* the incoming targets
- *All* individuals are described by a *unique* spatial distribution p_t^a

Previously detected targets (distinguishable)

First things first... what does the operator know beforehand, i.e. what does $P_{\mathcal{Y}_{t|t-1}}$ look like?

Incoming targets (indistinguishable)

- Per assumption, a *single* population describe *all* the incoming targets
- *All* individuals are described by a *unique* spatial distribution p_t^a

Previously detected targets (distinguishable)

- Per construction, all possible tracks are maintained in the set $Y_{t|t-1}$

First things first... what does the operator know beforehand, i.e. what does $P_{\mathcal{Y}_{t|t-1}}$ look like?

Incoming targets (indistinguishable)

- Per assumption, a *single* population describe *all* the incoming targets
- *All* individuals are described by a *unique* spatial distribution p_t^a

Previously detected targets (distinguishable)

- Per construction, all possible tracks are maintained in the set $Y_{t|t-1}$
- Each track $y \in Y_{t|t-1}$ has its *own* spatial distribution $p_{t|t-1}^y$

First things first... what does the operator know beforehand, i.e. what does $P_{\mathcal{Y}_{t|t-1}}$ look like?

Incoming targets (indistinguishable)

- Per assumption, a *single* population describe *all* the incoming targets
- *All* individuals are described by a *unique* spatial distribution p_t^a

Previously detected targets (distinguishable)

- Per construction, all possible tracks are maintained in the set $Y_{t|t-1}$
- Each track $y \in Y_{t|t-1}$ has its *own* spatial distribution $p_{t|t-1}^y$

Joint probability of existence $P_{t|t-1}$

First things first... what does the operator know beforehand, i.e. what does $P_{\mathcal{Y}_{t|t-1}}$ look like?

Incoming targets (indistinguishable)

- Per assumption, a *single* population describe *all* the incoming targets
- *All* individuals are described by a *unique* spatial distribution p_t^a

Previously detected targets (distinguishable)

- Per construction, all possible tracks are maintained in the set $Y_{t|t-1}$
- Each track $y \in Y_{t|t-1}$ has its *own* spatial distribution $p_{t|t-1}^y$

Joint probability of existence $P_{t|t-1}$

- $P_{t|t-1}(Y, n)$: how likely is the configuration with tracks $y \in Y$ and n incoming targets?

First things first... what does the operator know beforehand, i.e. what does $P_{\mathcal{Y}_{t|t-1}}$ look like?

Incoming targets (indistinguishable)

- Per assumption, a *single* population describe *all* the incoming targets
- *All* individuals are described by a *unique* spatial distribution p_t^a

Previously detected targets (distinguishable)

- Per construction, all possible tracks are maintained in the set $Y_{t|t-1}$
- Each track $y \in Y_{t|t-1}$ has its *own* spatial distribution $p_{t|t-1}^y$

Joint probability of existence $P_{t|t-1}$


- $P_{t|t-1}(Y, n)$: how likely is the configuration with tracks $y \in Y$ and n incoming targets?
- $\sum_{Y \subseteq Y_{t|t-1}} \sum_{n \geq 0} P_{t|t-1}(Y, n) = 1$

First things first... (cont.)

Are all the subsets of tracks $Y \subseteq Y_{t|t-1}$ worth considering?


First things first... (cont.)

Are all the subsets of tracks $Y \subseteq Y_{t|t-1}$ worth considering?



First things first... (cont.)


Are all the subsets of tracks $Y \subseteq Y_{t|t-1}$ worth considering?



$\rightarrow \emptyset, \{y\}, \{y'\}, \{y''\}, \{y, y'\}, \{y', y''\}$ OK

First things first... (cont.)

Are all the subsets of tracks $Y \subseteq Y_{t|t-1}$ worth considering?




$\rightarrow \emptyset, \{y\}, \{y'\}, \{y''\}, \{y, y'\}, \{y', y''\}$ OK

$\rightarrow \{y, y''\}, \{y, y', y''\}$ not OK

First things first... (cont.)

Are all the subsets of tracks $Y \subseteq Y_{t|t-1}$ worth considering?




$\rightarrow \emptyset, \{y\}, \{y'\}, \{y''\}, \{y, y'\}, \{y', y''\}$ OK

$\rightarrow \{y, y''\}, \{y, y', y''\}$ not OK

Set of hypotheses $H_{t|t-1}$

First things first... (cont.)

Are all the subsets of tracks $Y \subseteq Y_{t|t-1}$ worth considering?



$\rightarrow \emptyset, \{y\}, \{y'\}, \{y''\}, \{y, y'\}, \{y', y''\}$ OK


$\rightarrow \{y, y''\}, \{y, y', y''\}$ not OK

Set of hypotheses $H_{t|t-1}$

- The hypotheses $h \subseteq Y_{t|t-1}$ are the subsets of *compatible* tracks

First things first... (cont.)

Are all the subsets of tracks $Y \subseteq Y_{t|t-1}$ worth considering?



$\rightarrow \emptyset, \{y\}, \{y'\}, \{y''\}, \{y, y'\}, \{y', y''\}$ OK

$\rightarrow \{y, y''\}, \{y, y', y''\}$ not OK

Set of hypotheses $H_{t|t-1}$

- The hypotheses $h \subseteq Y_{t|t-1}$ are the subsets of *compatible* tracks
- $\sum_{Y \subseteq Y_{t|t-1}} \sum_{n \geq 0} P_{t|t-1}(Y, n) = 1 \rightarrow \sum_{h \in H_{t|t-1}} \sum_{n \geq 0} P_{t|t-1}(h, n) = 1$

If some $Z \subseteq \mathbf{Z}_u$ is collected, how does $P_{\mathfrak{y}_t}(\cdot|Z)$ look like?

If some $Z \subseteq \mathbf{Z}_u$ is collected, how does $P_{\mathfrak{Y}_t}(\cdot|Z)$ look like?


Data association

Given a possible configuration $(h \in H_{t|t-1}, n \in \mathbb{N})$ of the target population, what are the possible associations with the collected observations Z ?

If some $Z \subseteq \mathbf{Z}_u$ is collected, how does $P_{\mathfrak{Y}_t}(\cdot|Z)$ look like?

Data association


Given a possible configuration ($h \in H_{t|t-1}, n \in \mathbb{N}$) of the target population, what are the possible associations with the collected observations Z ?



If some $Z \subseteq \mathbf{Z}_u$ is collected, how does $P_{\mathfrak{Y}_t}(\cdot|Z)$ look like?

Data association

Given a possible configuration $(h \in H_{t|t-1}, n \in \mathbb{N})$ of the target population, what are the possible associations with the collected observations Z ?



Each association $\mathbf{a} = (h, n, \mathbf{h} \in \text{Adm}_{Z_t}(h, n))$ leads to a unique hyp. $\hat{h} \in H_t$:

- Assessed by prob. $P_u^{\mathbf{a}}$ (i.e. how likely is the association producing \hat{h} ?)
- Composed of tracks $\hat{h} = \bigcup_{y \in h_d} \{y: \nu(y)\} \cup \bigcup_{y \in h \setminus h_d} \{y: \phi\} \cup \bigcup_{z \in Z_a} \{\mathbf{a}: z\}$

Then, what is the information gain for track y ?

Then, what is the information gain for track y ?

Rényi divergence

Then, what is the information gain for track y ?

Rényi divergence

- Suppose track y has been updated with observation $z \in Z \cup \{\phi\}$

Then, what is the information gain for track y ?

Rényi divergence

- Suppose track y has been updated with observation $z \in Z \cup \{\phi\}$
- What have we learnt from $p_{t|t-1}^y$ to $p_u^{y:z}$?

Then, what is the information gain for track y ?

Rényi divergence

- Suppose track y has been updated with observation $z \in Z \cup \{\phi\}$
- What have we learnt from $p_{t|t-1}^y$ to $p_u^{y:z}$?
- We define $G_u^{y:z} = \frac{1}{\alpha-1} \log \left[\int [p_{t|t-1}^y(x)]^\alpha [p_u^{y:z}(x)]^{1-\alpha} \mu(dx) \right]$

Then, what is the information gain for track y ?

Rényi divergence

- Suppose track y has been updated with observation $z \in Z \cup \{\phi\}$
- What have we learnt from $p_{t|t-1}^y$ to $p_u^{y:z}$?
- We define $G_u^{y:z} = \frac{1}{\alpha-1} \log \left[\int [p_{t|t-1}^y(x)]^\alpha [p_u^{y:z}(x)]^{1-\alpha} \mu(dx) \right]$
- The track gain $G_u^{y:z}$ is non-negative, and equals zero iff:

Then, what is the information gain for track y ?

Rényi divergence

- Suppose track y has been updated with observation $z \in Z \cup \{\phi\}$
- What have we learnt from $p_{t|t-1}^y$ to $p_u^{y:z}$?
- We define $G_u^{y:z} = \frac{1}{\alpha-1} \log \left[\int [p_{t|t-1}^y(x)]^\alpha [p_u^{y:z}(x)]^{1-\alpha} \mu(dx) \right]$
- The track gain $G_u^{y:z}$ is non-negative, and equals zero iff:
 $p_{t|t-1}^y = p_u^{y:z}$ on \mathbf{X} (i.e. nothing learnt on target localization)

Then, what is the information gain for track y ?

Rényi divergence

- Suppose track y has been updated with observation $z \in Z \cup \{\phi\}$
- What have we learnt from $p_{t|t-1}^y$ to $p_u^{y:z}$?
- We define $G_u^{y:z} = \frac{1}{\alpha-1} \log \left[\int [p_{t|t-1}^y(x)]^\alpha [p_u^{y:z}(x)]^{1-\alpha} \mu(dx) \right]$
- The track gain $G_u^{y:z}$ is non-negative, and equals zero iff:
 $p_{t|t-1}^y = p_u^{y:z}$ on \mathbf{X} (i.e. nothing learnt on target localization), *and*

Then, what is the information gain for track y ?

Rényi divergence

- Suppose track y has been updated with observation $z \in Z \cup \{\phi\}$
- What have we learnt from $p_{t|t-1}^y$ to $p_u^{y:z}$?
- We define $G_u^{y:z} = \frac{1}{\alpha-1} \log \left[\int [p_{t|t-1}^y(x)]^\alpha [p_u^{y:z}(x)]^{1-\alpha} \mu(dx) \right]$
- The track gain $G_u^{y:z}$ is non-negative, and equals zero iff:
 - $p_{t|t-1}^y = p_u^{y:z}$ on \mathbf{X} (i.e. nothing learnt on target localization), *and*
 - $p_{t|t-1}^y(\psi) = p_u^{y:z}(\psi)$ (i.e. nothing learnt on target presence)

Finally, what is the expected gain from action u ?

Finally, what is the expected gain from action u ?

What is the expected gain from action u , given Z ?

Finally, what is the expected gain from action u ?

What is the expected gain from action u , given Z ?

1. Gain from y to $y:z$: $G_u^{y:z} = \frac{1}{\alpha-1} \log \left[\int [p_{t|t-1}^y(x)]^\alpha [p_u^{y:z}(x)]^{1-\alpha} \mu(dx) \right]$

Finally, what is the expected gain from action u ?

What is the expected gain from action u , given Z ?

1. Gain from y to $y:z$: $G_u^{y:z} = \frac{1}{\alpha-1} \log \left[\int [p_{t|t-1}^y(x)]^\alpha [p_u^{y:z}(x)]^{1-\alpha} \mu(dx) \right]$
2. Gain from (h, n) to \hat{h} : $G_u^{\mathbf{a}} = \sum_{y \in h_d} G_u^{y:\nu(y)} + \sum_{y \in h \setminus h_d} G_u^{y:\phi} + \sum_{z \in Z_a} G_u^{\mathbf{a}:z}$

Finally, what is the expected gain from action u ?

What is the expected gain from action u , given Z ?

1. Gain from y to $y:z$: $G_u^{y:z} = \frac{1}{\alpha-1} \log \left[\int [p_{t|t-1}^y(x)]^\alpha [p_u^{y:z}(x)]^{1-\alpha} \mu(dx) \right]$
2. Gain from (h, n) to \hat{h} : $G_u^{\mathbf{a}} = \sum_{y \in h_d} G_u^{y:\nu(y)} + \sum_{y \in h \setminus h_d} G_u^{y:\phi} + \sum_{z \in Z_a} G_u^{\mathbf{a}:z}$
3. Expected gain from (h, n) , given Z : $G_u^{h,n}(\cdot|Z) = \sum_{\mathbf{h} \in \text{Adm}_Z(h,n)} P_u^{\mathbf{a}} G_u^{\mathbf{a}}$

Finally, what is the expected gain from action u ?

What is the expected gain from action u , given Z ?

1. Gain from y to $y:z$: $G_u^{y:z} = \frac{1}{\alpha-1} \log \left[\int [p_{t|t-1}^y(x)]^\alpha [p_u^{y:z}(x)]^{1-\alpha} \mu(dx) \right]$
2. Gain from (h, n) to \hat{h} : $G_u^{\mathbf{a}} = \sum_{y \in h_d} G_u^{y:\nu(y)} + \sum_{y \in h \setminus h_d} G_u^{y:\phi} + \sum_{z \in Z_a} G_u^{\mathbf{a}:z}$
3. Expected gain from (h, n) , given Z : $G_u^{h,n}(\cdot|Z) = \sum_{\mathbf{h} \in \text{Adm}_Z(h,n)} P_u^{\mathbf{a}} G_u^{\mathbf{a}}$
4. Expected gain from $P_{\mathfrak{Y}_{t|t-1}}$, given Z :
 $G_u(\cdot|Z) = \sum_{h \in H_{t|t-1}} \sum_{n \geq 0} P_{t|t-1}(h, n) G_u^{h,n}(\cdot|Z)$

Finally, what is the expected gain from action u ?

What is the expected gain from action u , given Z ?

1. Gain from y to $y:z$: $G_u^{y:z} = \frac{1}{\alpha-1} \log \left[\int [p_{t|t-1}^y(x)]^\alpha [p_u^{y:z}(x)]^{1-\alpha} \mu(dx) \right]$
2. Gain from (h, n) to \hat{h} : $G_u^{\mathbf{a}} = \sum_{y \in h_d} G_u^{y:\nu(y)} + \sum_{y \in h \setminus h_d} G_u^{y:\phi} + \sum_{z \in Z_a} G_u^{\mathbf{a}:z}$
3. Expected gain from (h, n) , given Z : $G_u^{h,n}(\cdot|Z) = \sum_{\mathbf{h} \in \text{Adm}_Z(h,n)} P_u^{\mathbf{a}} G_u^{\mathbf{a}}$
4. Expected gain from $P_{\mathfrak{Y}_{t|t-1}}$, given Z :
 $G_u(\cdot|Z) = \sum_{h \in H_{t|t-1}} \sum_{n \geq 0} P_{t|t-1}(h, n) G_u^{h,n}(\cdot|Z)$

What is the expected gain from action u ?

Finally, what is the expected gain from action u ?

What is the expected gain from action u , given Z ?

1. Gain from y to $y:z$: $G_u^{y:z} = \frac{1}{\alpha-1} \log \left[\int [p_{t|t-1}^y(x)]^\alpha [p_u^{y:z}(x)]^{1-\alpha} \mu(dx) \right]$
2. Gain from (h, n) to \hat{h} : $G_u^{\mathbf{a}} = \sum_{y \in h_d} G_u^{y:\nu(y)} + \sum_{y \in h \setminus h_d} G_u^{y:\phi} + \sum_{z \in Z_a} G_u^{\mathbf{a}:z}$
3. Expected gain from (h, n) , given Z : $G_u^{h,n}(\cdot|Z) = \sum_{\mathbf{h} \in \text{Adm}_Z(h,n)} P_u^{\mathbf{a}} G_u^{\mathbf{a}}$
4. Expected gain from $P_{\mathfrak{y}_{t|t-1}}$, given Z :

$$G_u(\cdot|Z) = \sum_{h \in H_{t|t-1}} \sum_{n \geq 0} P_{t|t-1}(h, n) G_u^{h,n}(\cdot|Z)$$

What is the expected gain from action u ?

1. Expected gain from $P_{\mathfrak{y}_{t|t-1}}$: $G_u = \sum_{Z \subseteq \mathbf{Z}_u} G_u(\cdot|Z)$

Region-specific and/or track-specific information gain

Region-specific and/or track-specific information gain

Information gain G_u global by nature, but core element is *track*-based Rényi divergence

$$G_u^{y:z} = \frac{1}{\alpha - 1} \log \left[\int [p_{t|t-1}^y(x)]^\alpha [p_u^{y:z}(x)]^{1-\alpha} \mu(dx) \right]$$

Region-specific and/or track-specific information gain

Information gain G_u global by nature, but core element is *track*-based Rényi divergence

$$G_u^{y:z} = \frac{1}{\alpha - 1} \log \left[\int [p_{t|t-1}^y(x)]^\alpha [p_u^{y:z}(x)]^{1-\alpha} \mu(dx) \right]$$


Elementary changes in the divergence operator allow emphasis on specific regions of the target state space and/or specific tracks, e.g.

Region-specific and/or track-specific information gain

Information gain G_u global by nature, but core element is *track*-based Rényi divergence

$$G_u^{y:z} = \frac{1}{\alpha - 1} \log \left[\int [p_{t|t-1}^y(x)]^\alpha [p_u^{y:z}(x)]^{1-\alpha} \mu(dx) \right]$$

Elementary changes in the divergence operator allow emphasis on specific regions of the target state space and/or specific tracks, e.g.




Exclusion of *regions* from decision policy

Region-specific and/or track-specific information gain

Information gain G_u global by nature, but core element is *track*-based Rényi divergence


$$G_u^{y:z} = \frac{1}{\alpha - 1} \log \left[\int [p_{t|t-1}^y(x)]^\alpha [p_u^{y:z}(x)]^{1-\alpha} \mu(dx) \right]$$

Elementary changes in the divergence operator allow emphasis on specific regions of the target state space and/or specific tracks, e.g.



Exclusion of *regions* from decision policy

Exclusion of *tracks* from decision policy



Thank you for your attention!