Performance metric in closed-loop sensor management for stochastic populations

Emmanuel D. Delande, Jérémie Houssineau, and Daniel E. Clark

{E.D.Delande, jh207, D.E.Clark}@hw.ac.uk

Engineering & Physical Sciences
Heriot-Watt University
Edinburgh, UK

SSPD Conference - September 9, 2014 @ University of Edinburgh, UK

1. Closed-loop sensor management for multi-object filtering

2. Information gain for stochastic populations

3. Further developments

1. Closed-loop sensor management for multi-object filtering
2. Information gain for stochastic populations
3. Further developments

Multi-object Bayesian filtering

Multi-object tracking problem

Multi-object Bayesian filtering

Multi-object tracking problem

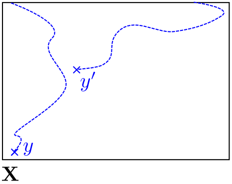

- $\mathbf{X}$: physical space of interest (surveillance area)

Multi-object Bayesian filtering

Multi-object tracking problem

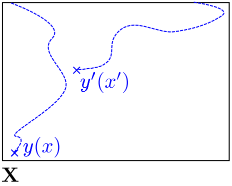

- $\mathbf{X}$: physical space of interest (surveillance area)
- Targets *currently in* $\mathbf{X}$ described by *state* $x \in \mathbf{X}$ (position, velocity, etc.)

Multi-object Bayesian filtering

Multi-object tracking problem

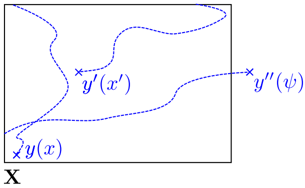

- $\mathbf{X}$: physical space of interest (surveillance area)
- Targets *currently in* $\mathbf{X}$ described by *state* $x \in \mathbf{X}$ (position, velocity, etc.)
- Targets *currently away from* $\mathbf{X}$ assume "empty state" ψ

Multi-object Bayesian filtering

Multi-object tracking problem

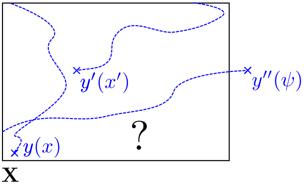

- $\mathbf{X}$: physical space of interest (surveillance area)
- Targets *currently in* $\mathbf{X}$ described by *state* $x \in \mathbf{X}$ (position, velocity, etc.)
- Targets *currently away from* $\mathbf{X}$ assume "empty state" ψ
- Goal: what is the state of y in $\bar{\mathbf{X}} = \{\psi\} \cup \mathbf{X}$?

Multi-object Bayesian filtering

Multi-object tracking problem

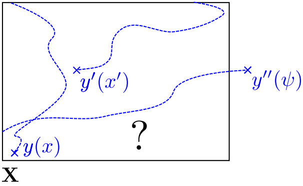

- $\mathbf{X}$: physical space of interest (surveillance area)
- Targets *currently in* $\mathbf{X}$ described by *state* $x \in \mathbf{X}$ (position, velocity, etc.)
- Targets *currently away from* $\mathbf{X}$ assume “empty state” ψ
- Goal: what is the state of y in $\bar{\mathbf{X}} = \{\psi\} \cup \mathbf{X}$?

Bayesian flow

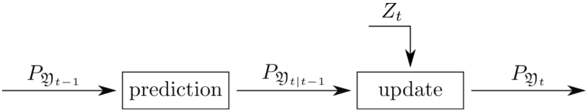

Multi-object Bayesian filtering

Multi-object tracking problem

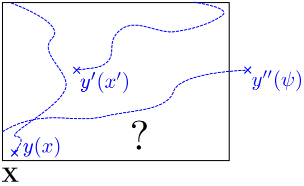

- $\mathbf{X}$: physical space of interest (surveillance area)
- Targets *currently in* $\mathbf{X}$ described by *state* $x \in \mathbf{X}$ (position, velocity, etc.)
- Targets *currently away from* $\mathbf{X}$ assume "empty state" ψ
- Goal: what is the state of y in $\bar{\mathbf{X}} = \{\psi\} \cup \mathbf{X}$?

Bayesian flow

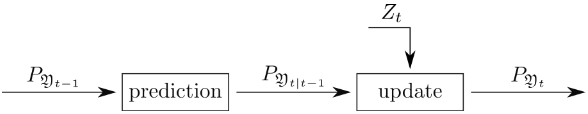

- $P_{\mathfrak{Y}_t}$: "information" known by operator at time t on *all* targets
- Z_t: observations produced and collected at time t by the operator

Sensor modelling

What is a sensor, from a tracking perspective?

Sensor modelling

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 modelling

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 modelling

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?

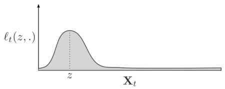

Sensor modelling

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

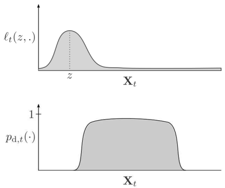

Sensor modelling

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 $p_{\mathrm{d},t}(x)$: how likely is a target with state x to be detected?
- Probability of false alarm $p_{\mathrm{fa},t}(z)$: how likely is the sensor to produce a false alarm with state z?

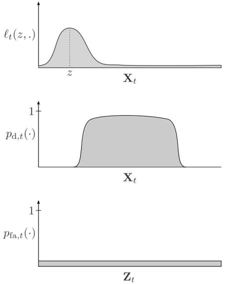

Sensor modelling (cont.)

Sensor modelling (cont.)

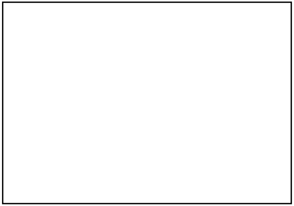

$\mathbf{Z}_t$

Sensor modelling (cont.)

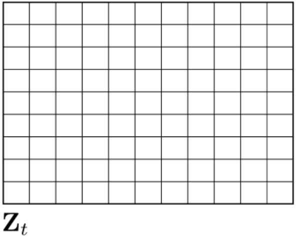

- *Discrete* observation space $\mathbf{Z}_t$

Sensor modelling (cont.)

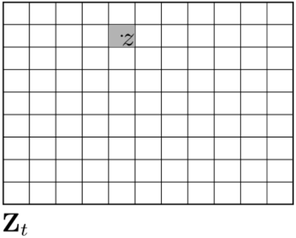

- *Discrete* observation space $\mathbf{Z}_t$
- *Localized* false alarm process: at most one false alarm, per cell and per scan

Sensor modelling (cont.)

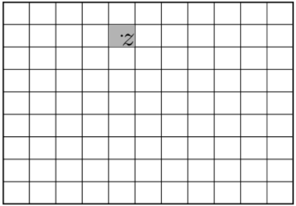

$\mathbf{Z}_t$

- *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 $p_{\mathrm{fa},t}(z)$

Sensor modelling (cont.)

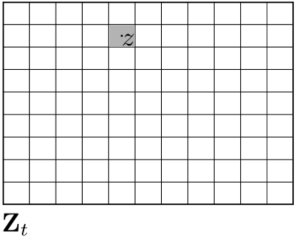

- *Discrete* observation space $\mathbf{Z}_t$
- *Localized* false alarm process: at most one false alarm, per cell and per scan
- In each cell $z \in \mathbf{Z}_t$, false alarm occurs with probability $p_{\mathrm{fa},t}(z)$

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

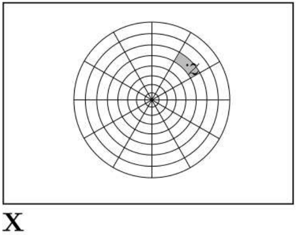

Sensor modelling (cont.)

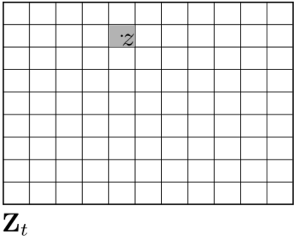

- *Discrete* observation space $\mathbf{Z}_t$
- *Localized* false alarm process: at most one false alarm, per cell and per scan
- In each cell $z \in \mathbf{Z}_t$, false alarm occurs with probability $p_{\mathrm{fa},t}(z)$

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

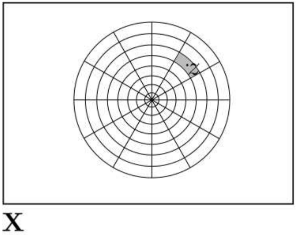

Closed-loop sensor management

Closed-loop sensor management

Closed-loop sensor management

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

Closed-loop sensor management

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

Closed-loop sensor management

- U_t: pool of available *actions* at time t
- $u_i \in U_t$: describes a specific sensor (i.e. ℓ_{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

- U_t: pool of available *actions* at time t
- $u_i \in U_t$: describes a specific sensor (i.e. ℓ_{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?

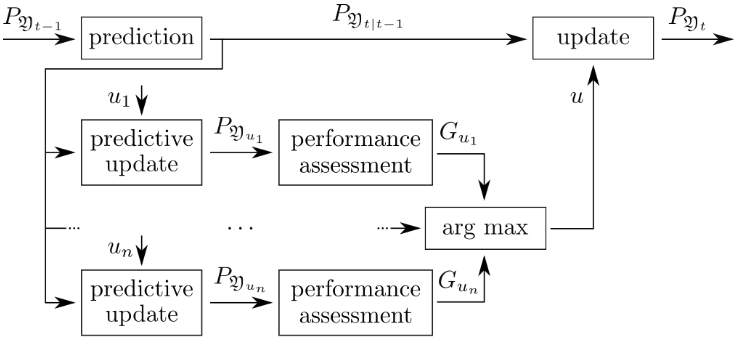

Construction of information gain: principle

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?

Mathematical framework: stochastic populations for Bayesian estimation

- Well-defined probabilistic framework, developed by J. Houssineau (PhD student) and D. Clark (supervisor)
- Tracking algorithm: ISP filter (Delande, Houssineau, Clark)

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

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

The population of appearing targets $\mathfrak{Y}_t^{\mathrm{a}}$

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

The population of appearing targets $\mathfrak{Y}_t^{\mathrm{a}}$

- Principle: no information is available on any specific individual

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

The population of appearing targets $\mathfrak{Y}_t^{\mathrm{a}}$

- Principle: no information is available on any specific individual
- Population size is described by cardinality distribution $c_t^{\mathrm{a}}(\cdot)$

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

The population of appearing targets $\mathfrak{Y}_t^{\mathrm{a}}$

- Principle: no information is available on any specific individual
- Population size is described by cardinality distribution $c_t^{\mathrm{a}}(\cdot)$
- *All* individuals are described by a *unique* probability measure $\hat{p}_t^{\mathrm{a}}(\cdot)$ on $\mathbf{X}$

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

The population of appearing targets $\mathfrak{Y}_t^{\mathrm{a}}$

- Principle: no information is available on any specific individual
- Population size is described by cardinality distribution $c_t^{\mathrm{a}}(\cdot)$
- *All* individuals are described by a *unique* probability measure $\hat{p}_t^{\mathrm{a}}(\cdot)$ on $\mathbf{X}$

Example: the operator expects (eventual) targets entering from the South-West

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

The population of appearing targets $\mathfrak{Y}_t^{\mathrm{a}}$

- Principle: no information is available on any specific individual
- Population size is described by cardinality distribution $c_t^{\mathrm{a}}(\cdot)$
- *All* individuals are described by a *unique* probability measure $\hat{p}_t^{\mathrm{a}}(\cdot)$ on $\mathbf{X}$

Example: the operator expects (eventual) targets entering from the South-West

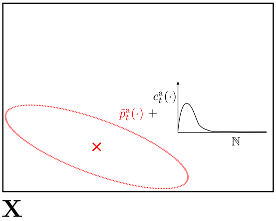

First things first... (cont.)

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: individual level

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: individual level

- Principle: individuals are *distinguishable* and identified by *tracks*

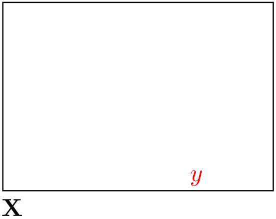

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: individual level

- Principle: individuals are *distinguishable* and identified by *tracks*
- Track y is described by probability measure $p^{y}_{t|t-1}(\cdot)$ on $\bar{\mathbf{X}}$:

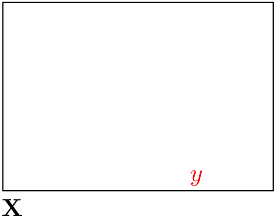

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: individual level

- Principle: individuals are *distinguishable* and identified by *tracks*
- Track y is described by probability measure $p^{y}_{t|t-1}(\cdot)$ on $\bar{\mathbf{X}}$:

$\tilde{p}^{y}_{t|t-1}(\cdot) = \frac{1_{\mathbf{X}}(\cdot) p^{y}_{t|t-1}(\cdot)}{p^{y}_{t|t-1}(1_{\mathbf{X}})}$ is the *spatial distribution* in the scene

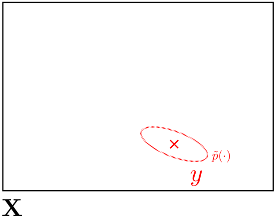

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: individual level

- Principle: individuals are *distinguishable* and identified by *tracks*
- Track y is described by probability measure $p^y_{t|t-1}(\cdot)$ on $\bar{\mathbf{X}}$:

$\tilde{p}^y_{t|t-1}(\cdot) = \frac{1_{\mathbf{X}}(\cdot)p^y_{t|t-1}(\cdot)}{p^y_{t|t-1}(1_{\mathbf{X}})}$ is the *spatial distribution* in the scene

$\bar{p}^y_{t|t-1} = p^y_{t|t-1}(1_{\mathbf{X}})$ is the *probability of presence* in the scene

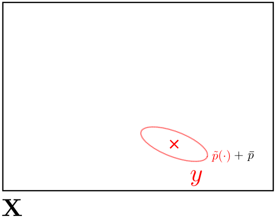

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: individual level

- Principle: individuals are *distinguishable* and identified by *tracks*
- Track y is described by probability measure $p^y_{t|t-1}(\cdot)$ on $\bar{\mathbf{X}}$:

$\tilde{p}^y_{t|t-1}(\cdot) = \frac{1_{\mathbf{X}}(\cdot) p^y_{t|t-1}(\cdot)}{p^y_{t|t-1}(1_{\mathbf{X}})}$ is the *spatial distribution* in the scene

$\bar{p}^y_{t|t-1} = p^y_{t|t-1}(1_{\mathbf{X}})$ is the *probability of presence* in the scene

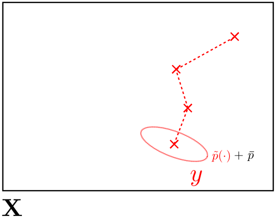

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: individual level

- Principle: individuals are *distinguishable* and identified by *tracks*
- Track y is described by probability measure $p^y_{t|t-1}(\cdot)$ on $\bar{\mathbf{X}}$:

$\tilde{p}^y_{t|t-1}(\cdot) = \frac{1_{\mathbf{X}}(\cdot)p^y_{t|t-1}(\cdot)}{p^y_{t|t-1}(1_{\mathbf{X}})}$ is the *spatial distribution* in the scene

$\bar{p}^y_{t|t-1} = p^y_{t|t-1}(1_{\mathbf{X}})$ is the *probability of presence* in the scene

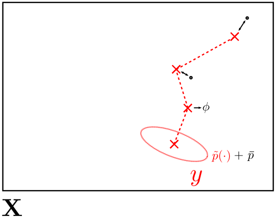

First things first... (cont.)

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

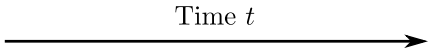

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

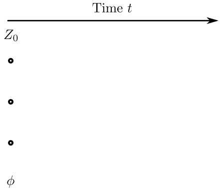

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

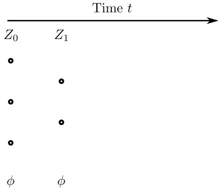

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

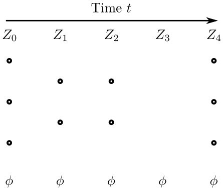

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

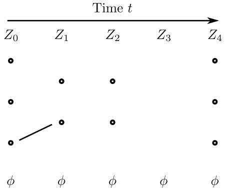

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

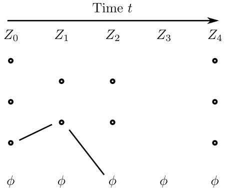

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

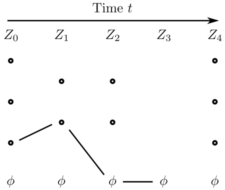

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

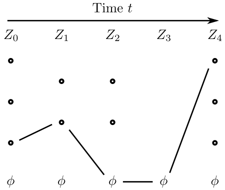

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

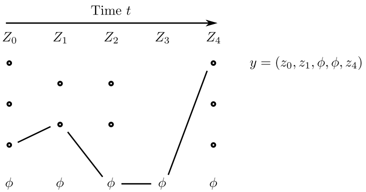

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

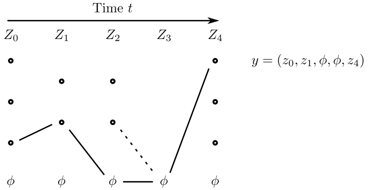

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

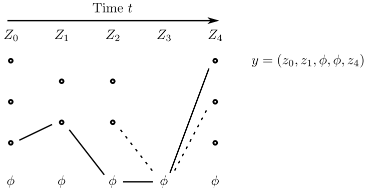

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

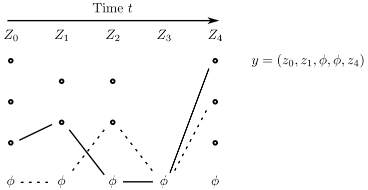

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

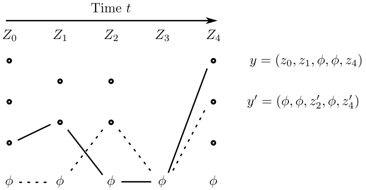

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

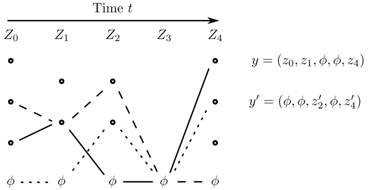

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

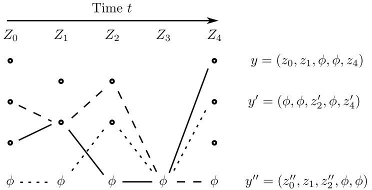

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

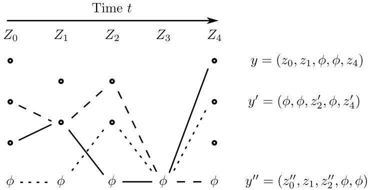

- Hypothesis $h \subseteq Y_{t|t-1}$: any subset of *compatible* tracks

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

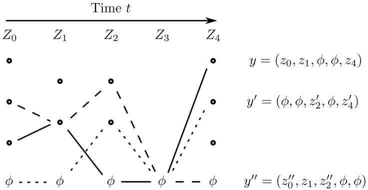

- Hypothesis $h \subseteq Y_{t|t-1}$: any subset of *compatible* tracks
  $\rightarrow \emptyset$, $\{y\}$, $\{y'\}$, $\{y''\}$, $\{y, y'\}$, $\{y', y''\}$ OK,

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

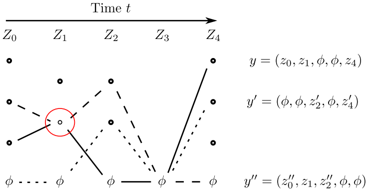

- Hypothesis $h \subseteq Y_{t|t-1}$: any subset of *compatible* tracks
  $\to \emptyset$, $\{y\}$, $\{y'\}$, $\{y''\}$, $\{y, y'\}$, $\{y', y''\}$ OK, $\{y, y''\}$, $\{y, y', y''\}$ not OK

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

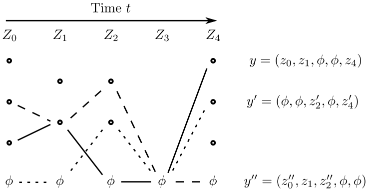

- Hypothesis $h \subseteq Y_{t|t-1}$: any subset of *compatible* tracks
  $\rightarrow \emptyset$, $\{y\}$, $\{y'\}$, $\{y''\}$, $\{y, y'\}$, $\{y', y''\}$ OK, $\{y, y''\}$, $\{y, y', y''\}$ not OK
- $c_{t|t-1}(h)$: *probability* of hypothesis h ($\sum_{h \in H_{t|t-1}} c_{t|t-1}(h) = 1$)

First things first... (cont.)

The population of previously detected targets $\mathfrak{Y}^{\mathrm{d}}_{t|t-1}$: hypothesis level

- $Y_{t|t-1}$: all possible tracks, up to time t

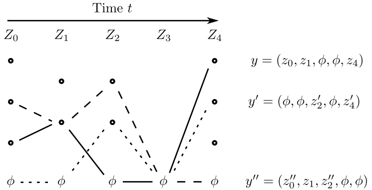

- Hypothesis $h \subseteq Y_{t|t-1}$: any subset of *compatible* tracks
  $\rightarrow$ $\emptyset$, $\{y\}$, $\{y'\}$, $\{y''\}$, $\{y, y'\}$, $\{y', y''\}$ OK, $\{y, y''\}$, $\{y, y', y''\}$ not OK
- $c_{t|t-1}(h)$: *probability* of hypothesis h ($\sum_{h \in H_{t|t-1}} c_{t|t-1}(h) = 1$)
  $\rightarrow$ i.e. how likely is h to represent the true target configuration?

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 (ISP filter, unpublished)

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 (ISP filter, unpublished)

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?

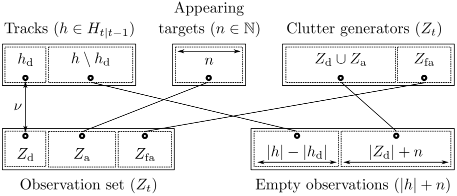

If some $Z \subseteq \mathbf{Z}_u$ is collected, how does $P_{\mathfrak{Y}_t}(\cdot|Z)$ look like?

Data association (ISP filter, unpublished)

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?

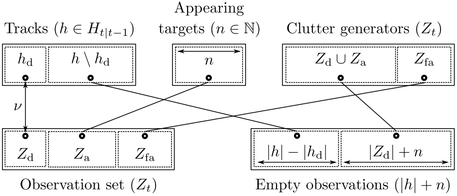

Each association $\mathbf{a} = (h, n, \mathbf{h} \in \mathrm{Adm}_{Z_t}(h, n))$ leads to a unique hyp. $\hat{h} \in H_t$:

- Assessed by prob. $P_u^{\mathbf{a}}$ (i.e. how likely is the association producing $\hat{h}$?)
- Composed of tracks $\hat{h} = \bigcup_{y \in h_\mathrm{d}} \{y : \nu(y)\} \cup \bigcup_{y \in h \setminus h_\mathrm{d}} \{y : \phi\} \cup \bigcup_{z \in Z_\mathrm{a}} \{\mathrm{a} : z\}$
- Update from $p_{t|t-1}^{y}$ to $p_u^{y:z}$: usual single-measurement/single-target update (e.g. Kalman)

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 $p^{y}_{t|t-1}$ to $p^{y:z}_{u}$?

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 $p^y_{t|t-1}$ to $p^{y:z}_u$?
- We define $G^{y:z}_u = \frac{1}{\alpha-1} \log \left[\int \left[p^y_{t|t-1}(x) \right]^{\alpha} \left[p^{y:z}_u(x) \right]^{1-\alpha} \mu(\mathrm{d}x) \right]$

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 $p^{y}_{t|t-1}$ to $p^{y:z}_{u}$?
- We define $G^{y:z}_{u} = \frac{1}{\alpha-1} \log \left[\int \left[p^{y}_{t|t-1}(x) \right]^{\alpha} \left[p^{y:z}_{u}(x) \right]^{1-\alpha} \mu(\mathrm{d}x) \right]$
- The track gain $G^{y:z}_{u}$ 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 $p^{y}_{t|t-1}$ to $p^{y:z}_{u}$?
- We define $G^{y:z}_{u} = \frac{1}{\alpha-1} \log \left[\int \left[p^{y}_{t|t-1}(x) \right]^{\alpha} \left[p^{y:z}_{u}(x) \right]^{1-\alpha} \mu(\mathrm{d}x) \right]$
- The track gain $G^{y:z}_{u}$ is non-negative, and equals zero iff:

  $p^{y}_{t|t-1} = p^{y:z}_{u}$ 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 $p^{y}_{t|t-1}$ to $p^{y:z}_{u}$?
- We define $G^{y:z}_{u} = \frac{1}{\alpha-1} \log \left[\int \left[p^{y}_{t|t-1}(x)\right]^{\alpha} \left[p^{y:z}_{u}(x)\right]^{1-\alpha} \mu(\mathrm{d}x) \right]$
- The track gain $G^{y:z}_{u}$ is non-negative, and equals zero iff:

  $p^{y}_{t|t-1} = p^{y:z}_{u}$ 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 $p^{y}_{t|t-1}$ to $p^{y:z}_{u}$?
- We define $G^{y:z}_{u} = \frac{1}{\alpha-1} \log \left[\int \left[p^{y}_{t|t-1}(x)\right]^{\alpha} \left[p^{y:z}_{u}(x)\right]^{1-\alpha} \mu(\mathrm{d}x) \right]$
- The track gain $G^{y:z}_{u}$ is non-negative, and equals zero iff:

  $p^{y}_{t|t-1} = p^{y:z}_{u}$ on $\mathbf{X}$ (i.e. nothing learnt on target localization), *and*

  $p^{y}_{t|t-1}(\psi) = p^{y:z}_{u}(\psi)$ (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 $y\!:\!z$: $G_u^{y:z} = \frac{1}{\alpha-1} \log \left[\int \left[p_{t|t-1}^{y}(x)\right]^{\alpha} \left[p_u^{y:z}(x)\right]^{1-\alpha} \mu(\mathrm{d}x) \right]$

Finally, what is the expected gain from action u?

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

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

Finally, what is the expected gain from action u?

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

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

Finally, what is the expected gain from action u?

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

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

Finally, what is the expected gain from action u?

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

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

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

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

$$G_u^{y:z} = \frac{1}{\alpha - 1} \log \Big[\int \big[p_{t|t-1}^{y}(x)\big]^{\alpha} \big[p_u^{y:z}(x)\big]^{1-\alpha} \mu(\mathrm{d}x) \Big]$$

Region-specific and/or track-specific information gain

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

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

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

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

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

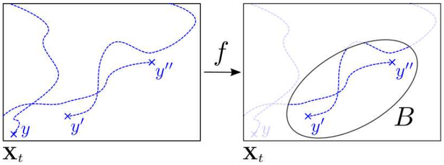

Exclusion of *regions* from decision policy

Region-specific and/or track-specific information gain

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

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

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

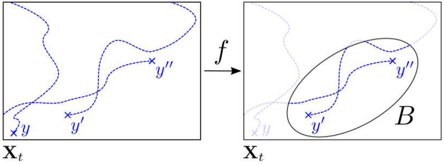

Exclusion of *regions* from decision policy

Exclusion of *tracks* from decision policy

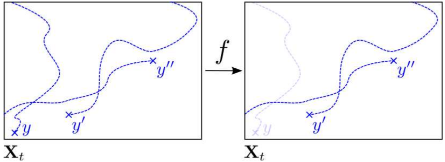

Thank you for your attention!