Beyond the PHD filter

Dealing with high variance in the number of detections

Isabel Schlangen

13/04/2017

Outline

Application examples

Video surveillance

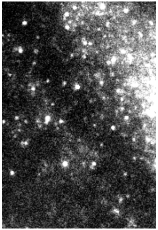

Fluorescence microscopy

Key challenges

- Extraction of detections
  low SNR, ambiguities, out-of-focus blur, ...
- Multi-object filtering
  false positives/negatives, crossings, proximity
- Sensor-level challenges
  drift, unknown noise rate, calibration of multiple inputs

Methodology

Bayesian filtering

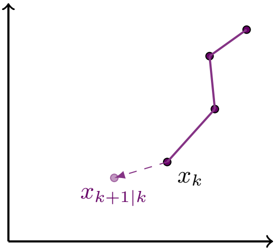

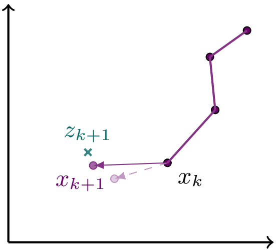

Prediction:

$$p_{k+1|k}(x_{k+1|k}|z_{1:k})$$
$$= \int f_{k+1|k}(x_k|x)p_k(x|z_{1:k})\mathrm{d}x$$

Update:

$$p_{k+1}(x_{k+1}|z_{1:k+1})$$
$$= \frac{g_k(z_{k+1}|x_{k+1|k})p_{k+1|k}(x_{k+1|k}|z_{1:k})}{\int g_k(z_{k+1}|x)p_{k+1|k}(x|z_{1:k})\mathrm{d}x}$$

Challenges of multi-object estimation

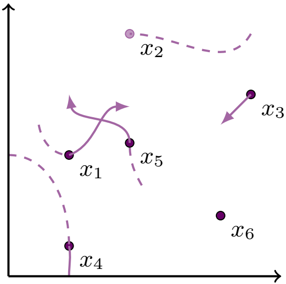

Object state space

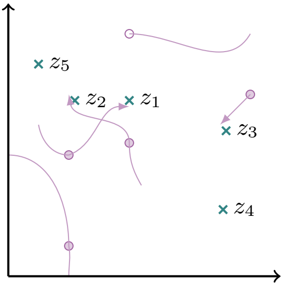

Measurement space

Populations of interest

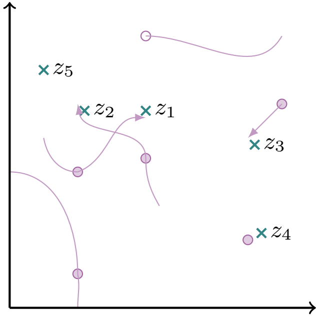

- Target population
  - surviving (prediction)
  - detected (update)
- New-born targets
- Clutter

Point processes

Definition: Point process Φ

- Random variable on some state space $\mathcal{X}$
- Both the target number and locations are random
- Realisation: n-tuple $\varphi = (x_1, \ldots, x_n) \in \mathcal{X}^n$
- Described by probability distribution P_Φ on $(\mathfrak{X}, \mathcal{B}(\mathfrak{X}))$, where $\mathfrak{X} = \prod_{n \geq 0} \mathcal{X}^n$

Three examples

1. Bernoulli point process
2. Poisson point process
3. Panjer point process

1) Bernoulli point process (parameter p, spatial distribution s)

- Binary point process: either one target or no target
- One target with probability p, distributed according to $s(\cdot)$
- No target with probability $(1 - p)$

2) Poisson point process (parameter λ, spatial distribution s)

- Point process whose cardinality distribution is Poisson with parameter λ
- Targets are independently and identically distributed (i.i.d.) according to s
- Target number is characterised by mean μ_{Poisson} and variance $\text{var}_{\text{Poisson}}$:

$$\mu_{\text{Poisson}} = \text{var}_{\text{Poisson}} = \lambda \tag{1}$$

HERIOT WATT UNIVERSITY

3) Panjer point process (parameters α, β, spatial distribution s)

- Point process whose cardinality distribution is Panjer with parameters $\alpha, \beta \in \mathbb{R}_{>0}$ or $\alpha \in \mathbb{Z}_{<0}, \beta \in \mathbb{R}_{<0}$
- Targets are i.i.d. according to s
- Target number is characterised by mean μ_{Panjer} and variance $\text{var}_{\text{Panjer}}$:

$$\mu_{\text{Panjer}} = \frac{\alpha}{\beta}, \quad \text{var}_{\text{Panjer}} = \mu_{\text{Panjer}} \left(1 + \frac{1}{\beta}\right) \tag{2}$$

  - $\mu_{\text{Panjer}} > \text{var}_{\text{Panjer}}$: binomial process
  - $\mu_{\text{Panjer}} < \text{var}_{\text{Panjer}}$: negative binomial process
  - $\mu_{\text{Panjer}} = \text{var}_{\text{Panjer}}$: Poisson process

Poisson, binomial and negative binomial distribution

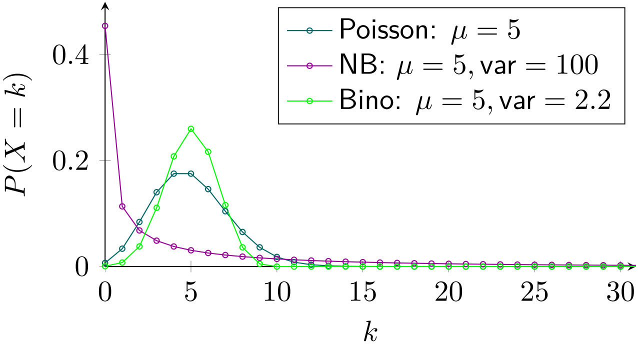

Two multi-object estimation algorithms

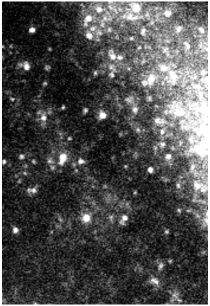

The Probability Hypothesis Density (PHD) filter [Mah03]

$\rightarrow$ propagates only the *first-order moment* or *probability hypothesis density* or *intensity* of the target process.

Filter assumptions

- All targets create observations independently
- Survival and detection processes are Bernoulli processes
- Predicted and false alarm processes are Poisson processes

HERIOT WATT UNIVERSITY

The recursion at time $k-1 \to k$

PHD filter prediction:

$$\mu_{k|k-1}(x) = \int_{\mathcal{X}} p_{\mathrm{s},k}(y) f_{k|k-1}(x|y) \mu_{k-1}(y) \mathrm{d}y + \mu_{\mathrm{b},k}(x) \tag{3}$$

PHD filter update:

$$\mu_k(x) = \Big(1 - p_{\mathrm{d}}(x)\Big)\mu_{k|k-1}(x) + \sum_{z \in Z_k} \frac{p_{\mathrm{d},k}(x) l_k(z|x) \mu_{k|k-1}(x)}{\int_{\mathcal{X}} p_{\mathrm{d},k}(y) l_k(z|y) \mu_{k|k-1}(y) \mathrm{d}y + \mu_{\mathrm{c},k}(z)} \tag{4}$$

Limitation

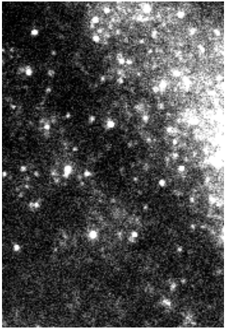

- All targets create observations independently
- Survival and detection processes are Bernoulli processes
- Predicted and false alarm processes are Poisson processes

$\rightarrow$ does not work in the example!

The second-order PHD (SO-PHD) filter [SDHC17]

$\rightarrow$ propagates both the *intensity* and the *variance* of the target process.

Filter assumptions

- All targets create observations independently
- Survival and detection processes are Bernoulli processes
- Predicted and false alarm process are Panjer processes

The recursion at time $k-1 \to k$

Mean and variance prediction (the same intensity as for the PHD filter):

$$\mu_{k|k-1}(x|Z_{1:k}) = p_{\mathrm{s},k} \int_{\mathcal{X}} f_{k|k-1}(x|y)\mu_{k-1}(y)\mathrm{d}y + \mu_{\mathrm{b},k}(x), \tag{5}$$

$$\mathrm{var}_{k|k-1}(\mathcal{X}) = p_{\mathrm{s},k}^2 \mathrm{var}_{k-1}(\mathcal{X}) + p_{\mathrm{s},k}(1-p_{\mathrm{s},k})\mu_{k-1}(\mathcal{X}) + \mathrm{var}_{\mathrm{b},k}(\mathcal{X}). \tag{6}$$

Mean and variance update:

$$\mu_k(x) = \underbrace{\Big(1 - p_{\mathrm{d},k}(x)\Big)\mu_{k|k-1}(x)}_{:=\mu_k^{\phi}(x)} + \sum_{z \in Z_k} \underbrace{p_{\mathrm{d},k}(x) l(z|x) \mu_{k+1|k}(x)}_{:=\mu_k^{z}(x)} \ell(z), \tag{7}$$

$$\begin{aligned} \mathrm{var}_k(B) = {} & \mu_k(B) + \mu_k^{\phi}(B)^2 \left[\ell_2(\phi) - \ell_1(\phi)^2\right] \\ & + 2\mu_k^{\phi}(B) \sum_{z \in Z_k} \mu_k^{z}(B) \left[\ell_2(z) - \ell_1(\phi)\ell_1(z)\right] \\ & + \sum_{z, z' \in Z_k} \mu_k^{z}(B) \hat{\mu}_k^{z'}(B) \left[\ell_2^{\neq}(z, z') - \ell_1(z)\ell_1(z')\right]. \end{aligned} \tag{8}$$

HERIOT WATT UNIVERSITY

Both filters in comparison (intensity update)

PHD filter update:

$$\mu_k(x) = \mu_k^{\phi}(x) + \sum_{z \in Z_k} \frac{\mu_k^z(x)}{\mu_k(\mathcal{X}) + \mu_{\mathrm{c},k}(z)}$$

SO-PHD filter update:

$$\mu_k(x) = \mu_k^{\phi}(x) + \sum_{z \in Z_k} \mu_k^z(x)\ell(z)$$

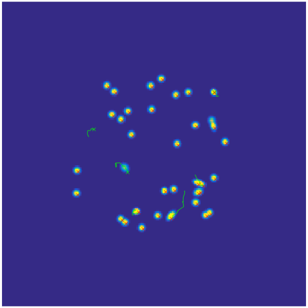

(a) PHD filter with Poisson clutter.

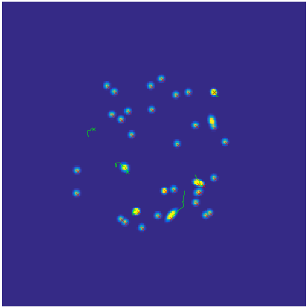

(b) SO-PHD filter with NB clutter.

Experiment 1: Simulated data with high variance in birth

- 100 Monte Carlo runs, 15 time steps
- One target at times 1–14
- Varying number of objects at time 15: $0, 5, 10, ...100$
- Poisson PHD: $\mu_{\mathrm{b}} = 1$
- SO-PHD: $\mu_{\mathrm{b}} = 1$ and $\mathrm{var}_{\mathrm{b}} = 100$

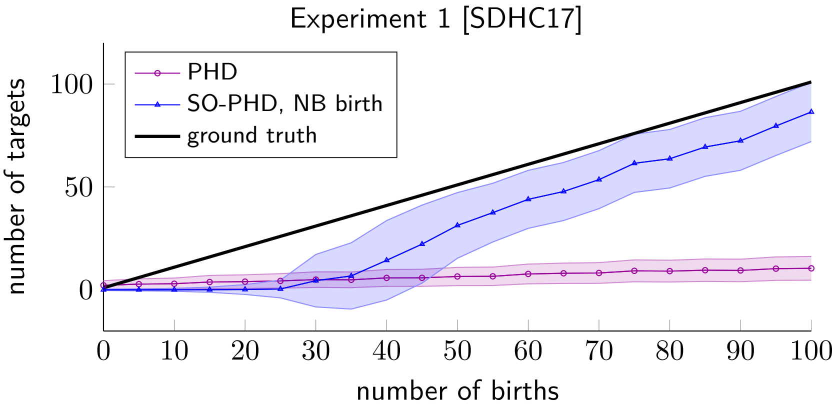
Experiment 1 [SDHC17]
PHD
SO-PHD, NB birth
ground truth
number of targets
100
50
0
0
10
20
30
40
50
60
70
80
90
100
number of births

Experiment 2: Real data (surveillance)

Exploiting second-order information

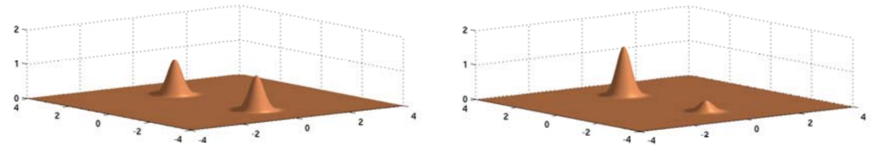

The "spooky effect" in the CPHD filter: mass is transferred from one object to the other due to miss-detection [VV12].

Covariance, variance and correlation [SDHC17]

Covariance:

$$\mathrm{cov}(B, B') = \mu^{(2)}(B \times B') - \mu(B)\mu(B'). \tag{9}$$

Variance:

$$\mathrm{var}(B) = \mu^{(2)}(B \times B) - [\mu(B)]^2 . \tag{10}$$

Correlation:

$$\mathrm{corr}(B, B') = \frac{\mathrm{cov}(B, B')}{\sqrt{\mathrm{var}(B)}\sqrt{\mathrm{var}(B')}}. \tag{11}$$

Experiment 3: Spooky effect [SDHC17]

- 100 Monte Carlo runs, 100 time steps
- Two regions A (red) and B (green), 10 targets each
- Every 10 time steps: all objects in B miss-detected
- Clutter rate in each region: $\mu_{\mathrm{c}}(A) = \mu_{\mathrm{c}}(B) = 20$
- SO-PHD and CPHD: $\mu_{\mathrm{b}} = 1$ and $\mathrm{var}_{\mathrm{b}} = 100$

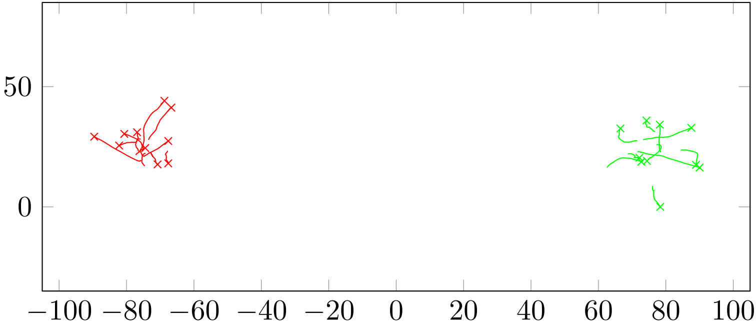

ground truth

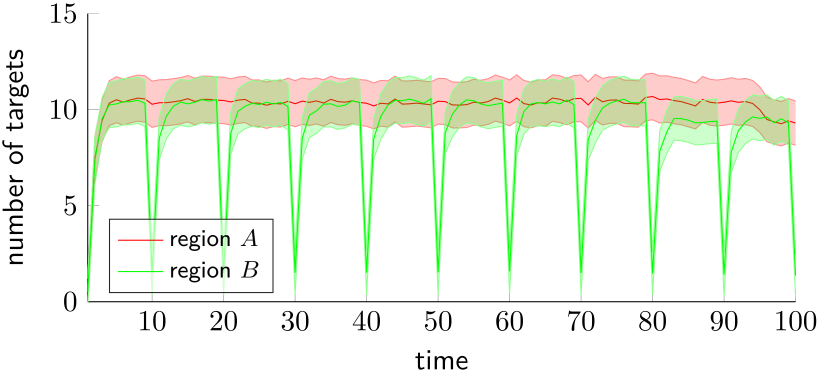
15
10
5
0
number of targets
10
20
30
40
50
60
70
80
90
100
time
region A
region B

PHD

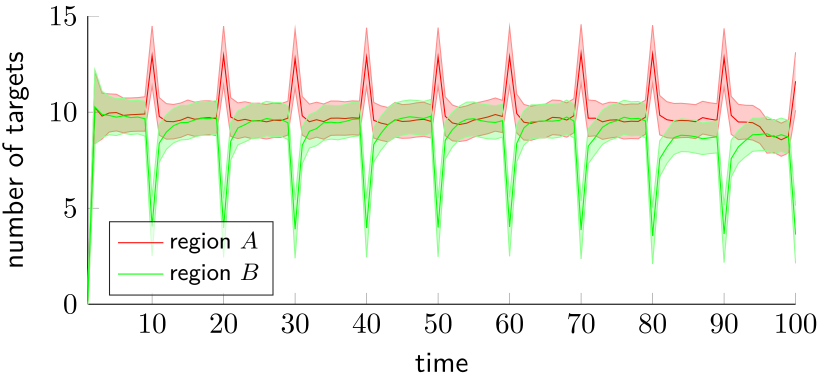
15
10
5
0
number of targets
10
20
30
40
50
60
70
80
90
100
time
region A
region B

CPHD

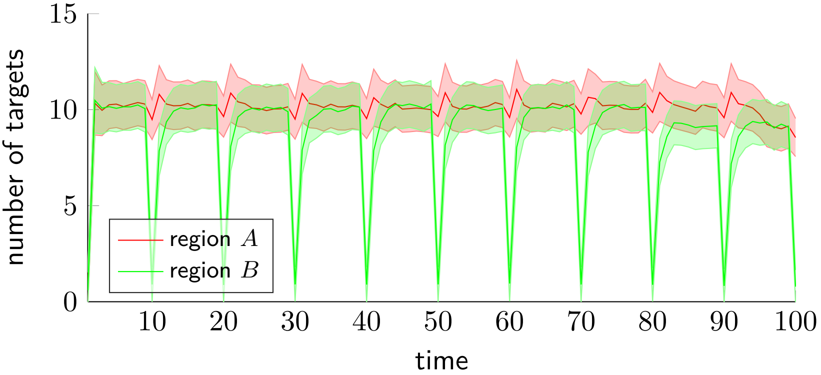
15
10
5
0
number of targets
10
20
30
40
50
60
70
80
90
100
time
region A
region B

SO-PHD

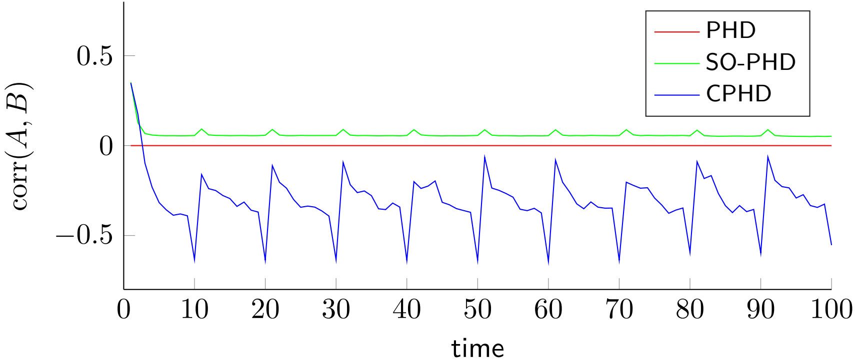
PHD
SO-PHD
CPHD
0.5
0
−0.5
corr(A, B)
0
10
20
30
40
50
60
70
80
90
100
time

Conclusion

- SO-PHD filter propagating second-order information
  $\rightarrow$ more flexible modelling of challenging scenarios

- Second-order information as a means of assessing the "spooky effect"
  $\rightarrow$ Regional information yields valuable insight in the performance of the filters

References

[Mah03] R. P. S. Mahler. Multitarget Bayes filtering via first-order multitarget moments. *Aerospace and Electronic Systems, IEEE Transactions on*, 39(4):1152–1178, 2003.

[SDHC17] I. Schlangen, E. D. Delande, J. Houssineau, and D. E. Clark. A second-order PHD filter with mean and variance in target number, 2017. arXiv:1704.02084.

[VV12] Ba Tuong Vo and Ba Ngu Vo. The para-normal bayes multi-target filter and the spooky effect. In *Information Fusion (FUSION), 2012 15th International Conference on*, pages 173–180. IEEE, 2012.