Sensor Management with Regional Statistics

for the PHD Filter

Marian Andrecki, Emmanuel D. Delande, Jérémie Houssineau, and Daniel E. Clark
{ma804, E.D.Delande, J.Houssineau, D.E.Clark}@hw.ac.uk

Engineering & Physical Sciences
Heriot-Watt University
Edinburgh, UK

September 9, 2015 @ Edinburgh, United Kingdom

The research

Multi-Object Tracking and Sensor Management

- Estimate state of targets in surveillance zone
- Control a configurable sensor (mobility, different modes of operation)
- Find suitable policy for the operator to follow (e.g. maximise number of targets in sight)
- Can be used airspace monitoring, border surveillance, etc.

The research

Multi-Object Tracking and Sensor Management

- Estimate state of targets in surveillance zone
- Control a configurable sensor (mobility, different modes of operation)
- Find suitable policy for the operator to follow (e.g. maximise number of targets in sight)
- Can be used airspace monitoring, border surveillance, etc.

Goals

- Explore second order statistics in PHD filter for Sensor Management purposes
- Consider the behaviour of the variance-based sensor manager and its biases
- Prepare for study of advanced filters with Sensor Management

The research

Multi-Object Tracking and Sensor Management

- Estimate state of targets in surveillance zone
- Control a configurable sensor (mobility, different modes of operation)
- Find suitable policy for the operator to follow (e.g. maximise number of targets in sight)
- Can be used airspace monitoring, border surveillance, etc.

Goals

- Explore second order statistics in PHD filter for Sensor Management purposes
- Consider the behaviour of the variance-based sensor manager and its biases
- Prepare for study of advanced filters with Sensor Management

As this is initial analysis of the algorithms all experiments were conducted in simulation

Demonstration

Scenario

- Groups of targets arrive from the West - objects move with near-constant velocity
- Scene is a 2D plane with three regions of interest
- Estimate number of targets in each of the regions of interest at every time step
- Sensor can scan only one of the regions at a time
- Where to sense at each step?
- Range and bearing measurements are distorted by Gaussian noise
- Missed detections and false alarms are possible

Scene

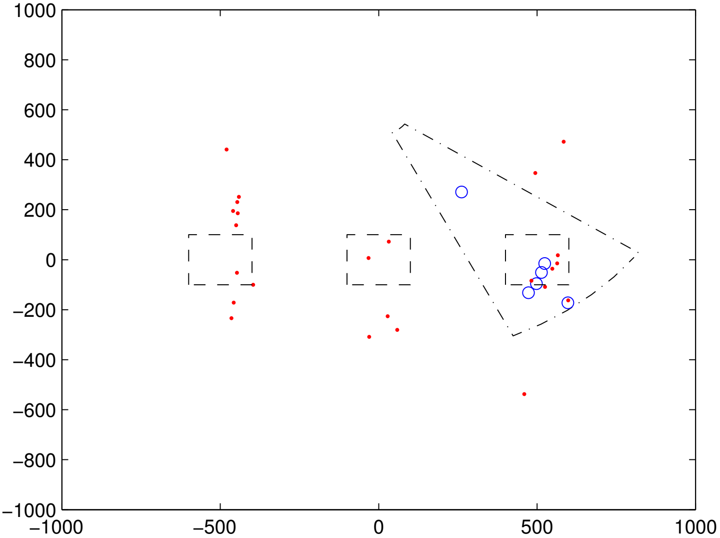

Probability Hypothesis Density filter

- Recursive Bayesian filter that captures statitics of populations
- Propagates mean density of targets in state space

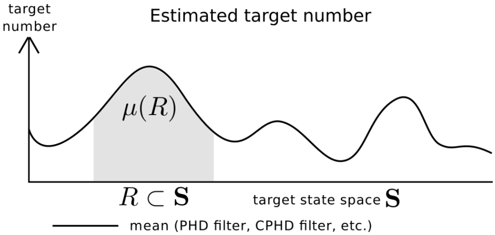

Probability Hypothesis Density filter

- Recursive Bayesian filter that captures statitics of populations
- Propagates mean density of targets in state space

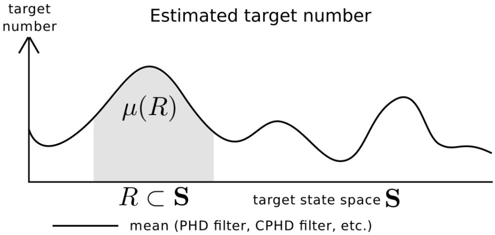

- At some stage filtering assumes density of targets is Poisson distributed $\implies$ propagated variance is equal to the mean
- Accounts for missed detections and false positives
- Gaussian Mixtures implementation was used

Regional Variance in PHD filter

- After update step one can compute variance in the number of targets in a given region
- Provides error bars on the estimated number targets in a given region

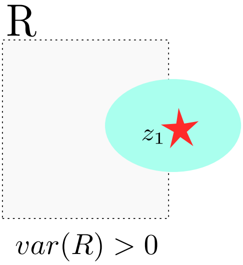

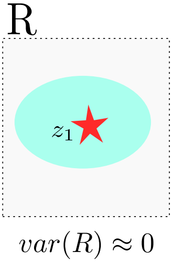

Regional Variance in PHD filter

- After update step one can compute variance in the number of targets in a given region
- Provides error bars on the estimated number targets in a given region

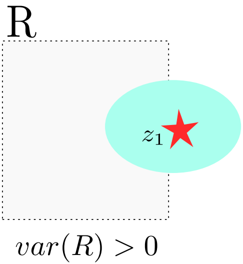

- How certain one is that a measurement does or does not correspond to an object inside of a specific region
- This variance cannot be propagated recursively in PHD filter

Management policy

User's goal

- Estimate number of objects in some regions as well as possible

Manager's goal

- Minimise the total regional variance across the three regions (i.e. keep the error bars on the number of objects as small as possible)

Policy

- Sense in the region with the *highest* variance, (i.e. where uncertainty is the highest)
- Computationally inexpensive

Stimulating exploration

Injecting variance

- We can now implicitly tell the manager to sense more in the leftmost region by injecting high variance into the scene.
- 'You really don't know what's going on there.'
- It is important to propagate this variance so that it accumulates when manager does not scan the region.

Biasing the mean

- In the PHD filter the variance is tied to the mean, so changing one affects the other.
- Solution to this is using a filter which propagates higher order statistics of population (e.g. HISP, DISP)
- Alternatively, point processes which offer variance higher than mean can be used (e.g. Negative Binomial point process)

Experiments

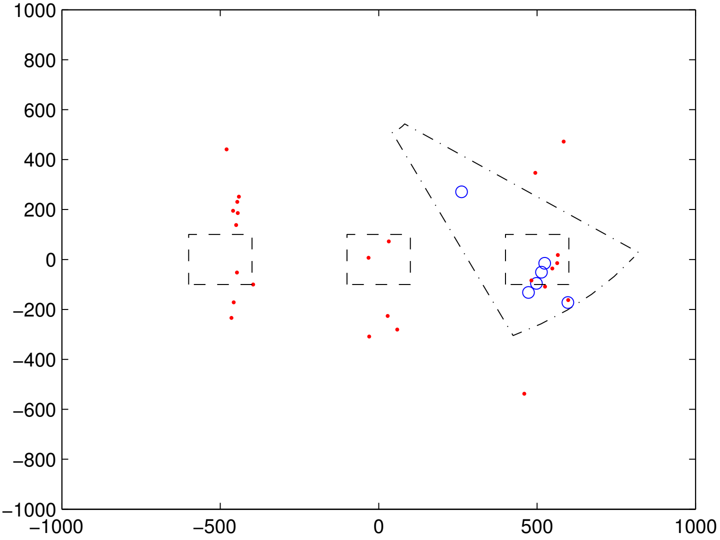

Simulation

Previously shown scenario was used

- Targets arriving from the West
- Three regions of interest
- Manager was implicitly informed about where the objects were coming from

Parameters varied

- Pattern of arrival of the targets (single wave, multiple waves)
- Quality of the sensor - distortion of measurements, probability of detection and false alarm
- Amount of variance injected

Variance minimising manager was compared to naive sequentially swiping manager

Pattern of arrival

Single wave of targets

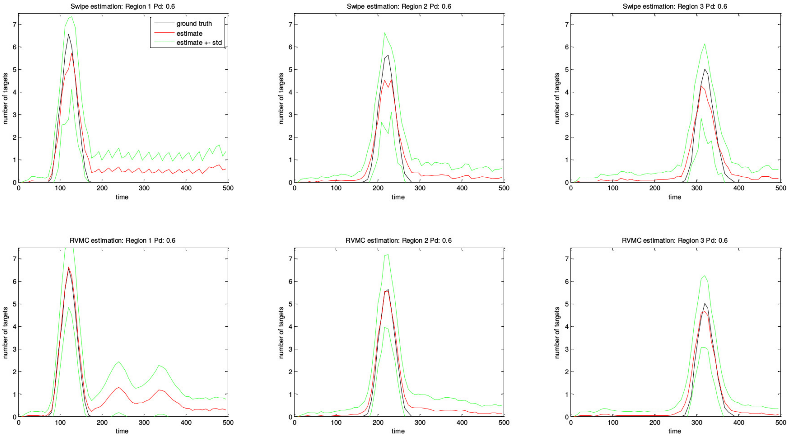

Pattern of arrival

Single wave of targets

- Variance minimising manager outperforms the naive one
- It is clear where to sense
- Naive swiping manager wastes scans by looking into empty regions

Pattern of arrival

Single wave of targets

- Variance minimising manager outperforms the naive one
- It is clear where to sense
- Naive swiping manager wastes scans by looking into empty regions

Multiple waves of targets

- Variance minimising manager only slightly outperforms the naive one
- Important to balance exploration and exploitation

Quality of the sensor

High false alarm rate

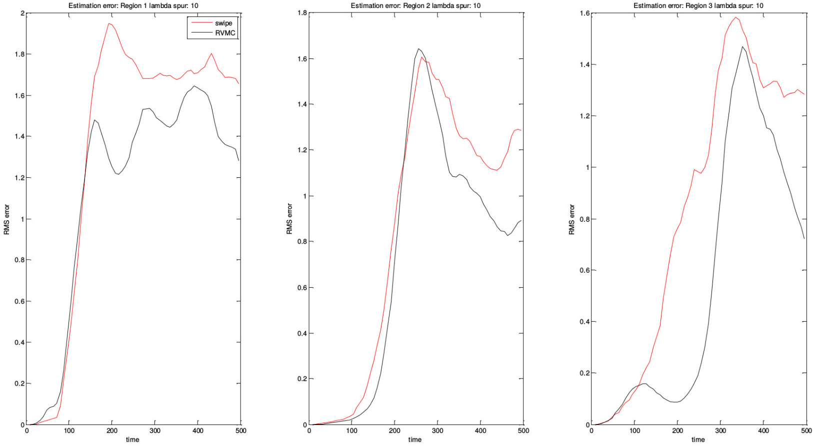

Quality of the sensor

Resilience to overestimation

- False positives contribute to higher variance in a region
- Manager chooses to scan the regions with heightened variance and corrects its belief

Quality of the sensor

Resilience to overestimation

- False positives contribute to higher variance in a region
- Manager chooses to scan the regions with heightened variance and corrects its belief

Overall performance

- Improvement in estimate for low quality sensor (e.g. $P_d = 0.6$) even for multiple waves of targets
- When information is scarce it is more important to choose regions for scanning wisely

Amount of variance injected

Managers with different levels of injected variance

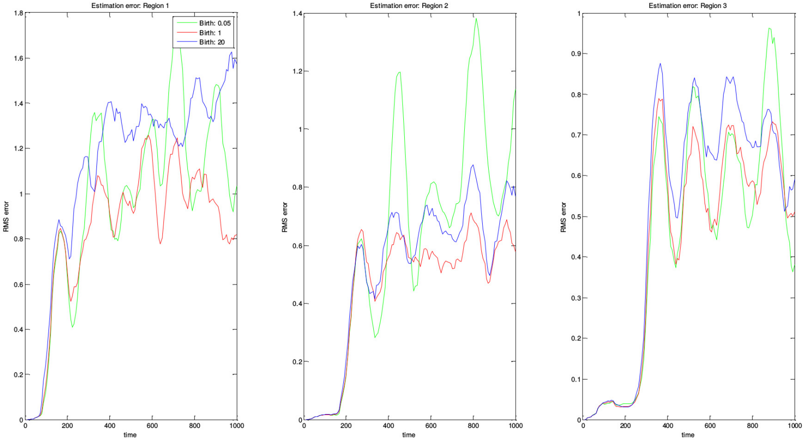

Amount of variance injected

- Too low variance level - manager does not explore enough and misses waves
- Too high variance level - manager focuses on exploration too much and has no time to converge on the targets

Conclusions

- Injecting variance leads to biasing the estimate in case of the PHD filter
- Tuning the value of injected variance is required to balance exploration and exploitation
- Higher order statitics driven sensor manager can lead to performance improvements in the PHD filter

Further work

- Implement the manager for a filter that propagates higher order statistics of the population (e.g. HISP)
- Explore new statistical models for populations with variance higher than the mean (e.g. Negative Binomial point process)