Loop detection and extended target tracking using laser data

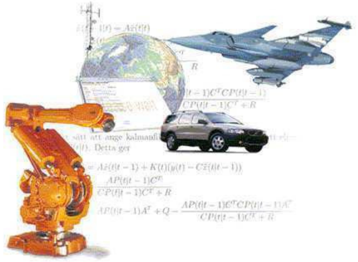

Karl Granström

Division of Automatic Control
Department of Electrical Engineering
Linköping University

Linköping, Sweden

Paper A

"I accept the paper for publication subject to your addressing all the points raised in the review."

Automatic control can be defined as

> making a system behave the way that you want it to.

Example: Automatic Cruise Control (ACC)

Automatic control

Automatic control can be defined as

> making a system behave the way that you want it to.

Example: Automatic Cruise Control (ACC)

Must know the state of the system, e.g. what is the current speed?

Autonomous vehicles / robots

Here – focus on autonomous vehicles / robots.

In the context of robots, knowing the state of the system includes

- knowing where the robot is,
    1. Must be able to recognise previously visited places, i.e. loop detection.

AUTOMATIC CONTROL
REGLERTEKNIK
LINKÖPINGS UNIVERSITET

Here – focus on autonomous vehicles / robots.

In the context of robots, knowing the state of the system includes

- knowing where the robot is,
    1. Must be able to recognise previously visited places, i.e. loop detection.
- knowing where other objects are.
    2. Must be able to follow objects that are moving, i.e. (extended) target tracking.

Autonomous vehicles / robots

Here – focus on autonomous vehicles / robots.

In the context of robots, knowing the state of the system includes

- knowing where the robot is,
  1. Must be able to recognise previously visited places, i.e. loop detection.
- knowing where other objects are.
  2. Must be able to follow objects that are moving, i.e. (extended) target tracking.

A fundamental requirement is that the robot must be able to sense the world – here we have used data from laser range sensors.

AUTOMATIC CONTROL
REGLERTEKNIK
LINKÖPINGS UNIVERSITET

Problem formulation and contributions

In this thesis the following two problems are addressed

1. design of a method that can compare pairs of laser data, such that the method may be used to detect loop closure in SLAM.
   - Feature description of laser data
   - AdaBoost used to learn classifier
   - Evaluated using several experiments

In this thesis the following two problems are addressed

1. design of a method that can compare pairs of laser data, such that the method may be used to detect loop closure in SLAM.
    - Feature description of laser data
    - AdaBoost used to learn classifier
    - Evaluated using several experiments
2. design of a method for target tracking in scenarios where each target possibly gives rise to more than one measurement per time step.
    - Implementation of Gaussian mixture PHD-filter for extended targets
    - Simple method to partition measurements
    - Evaluation using simulations and experiments

AUTOMATIC CONTROL
REGLERTEKNIK
LINKÖPINGS UNIVERSITET

Presentation overview

1. Automatic control
2. Autonomous vehicles / robots
3. Problem formulation and contributions
4. Laser data
5. Loop detection (L.D.)
6. Extended target tracking (E.T.T.)
7. Future work

Laser sensors are common

- Used frequently in robotics for at least a decade.
- Receiving increasing attention from e.g. the automotive industry.

Measures distance to nearest object at certain angles.

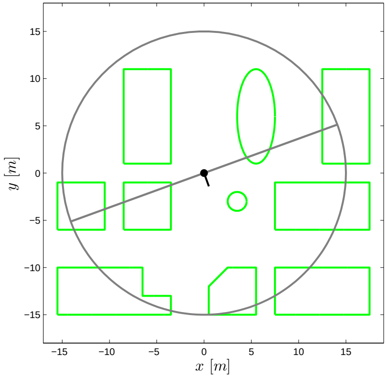

Data is called point clouds.

Measures distance to nearest object at certain angles.

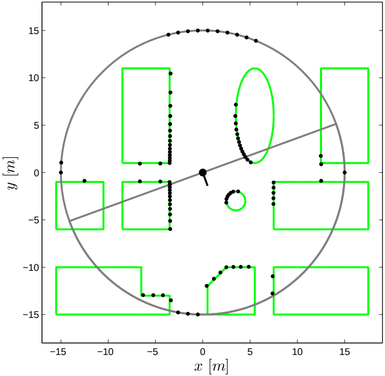

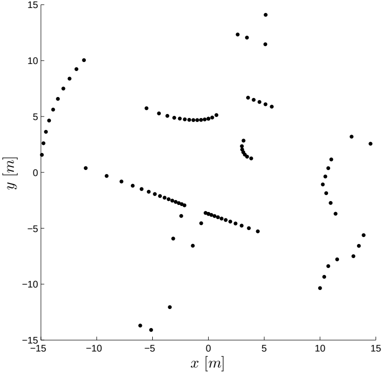

Data is called point clouds.

AUTOMATIC CONTROL
REGLERTEKNIK
LINKÖPINGS UNIVERSITET

2D – laserVPone.avi

3D – laserHannover.avi

SLAM – Map the environment, localise robot in map.

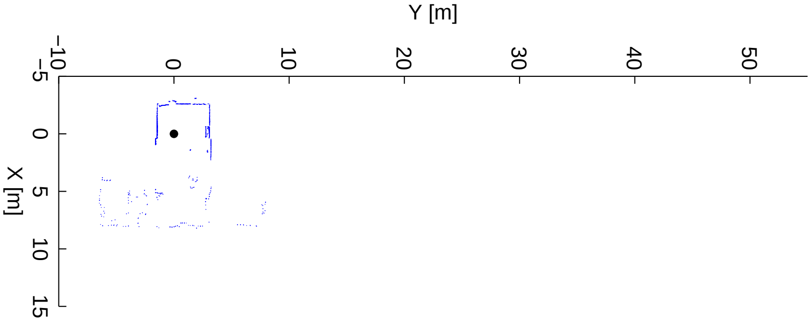

- State vector is history of poses.
- Each pose associated to point cloud describing the location.

SLAM – Map the environment, localise robot in map.

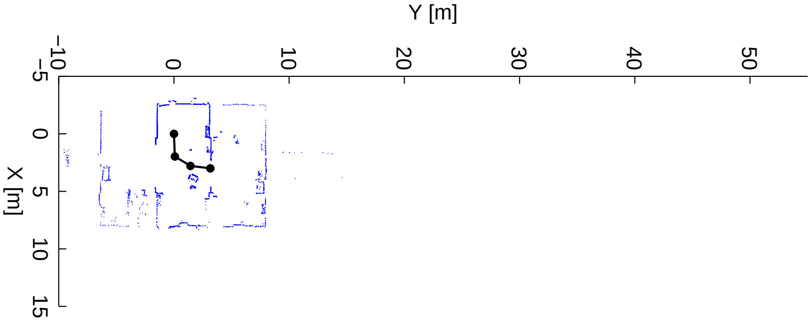

- State vector is history of poses.
- Each pose associated to point cloud describing the location.

SLAM – Map the environment, localise robot in map.

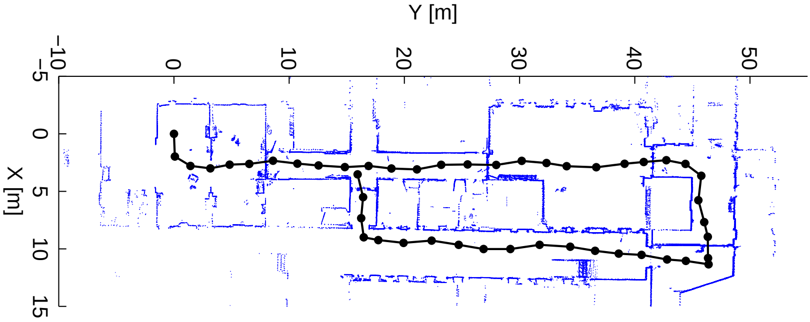

- State vector is history of poses.
- Each pose associated to point cloud describing the location.

Loop closure detection ⇔ place recognition.
Pairwise comparison of data, here point clouds,

$$\mathbf{p}_k = \{p_i^k\}_{i=1}^N, \quad p_i^k \in \mathbb{R}^3$$

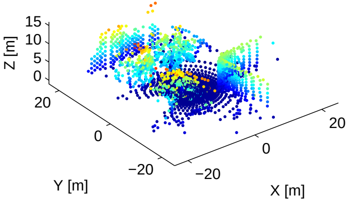

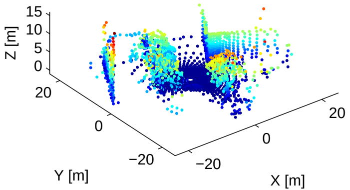

Are $\mathbf{p}_k$ and $\mathbf{p}_l$ from the same location?

Loop closure detection ⇔ place recognition.
Pairwise comparison of data, here point clouds,

$$\mathbf{p}_k = \{p_i^k\}_{i=1}^N, \quad p_i^k \in \mathbb{R}^3$$

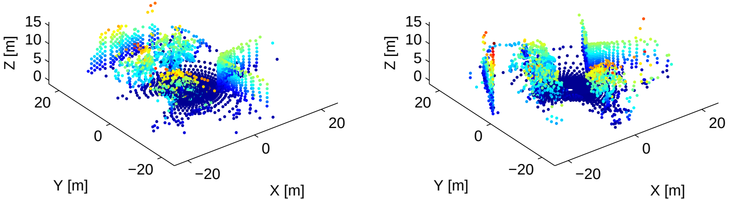

Are $\mathbf{p}_k$ and $\mathbf{p}_l$ from the same location? **Yes**

Loop closure detection ⇔ place recognition.
Pairwise comparison of data, here point clouds,

$$\mathbf{p}_k = \{p_i^k\}_{i=1}^N, \quad p_i^k \in \mathbb{R}^2$$

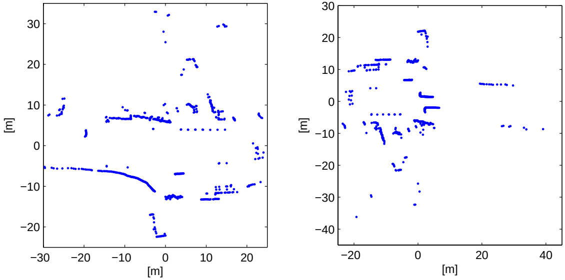

Are $\mathbf{p}_k$ and $\mathbf{p}_l$ from the same location?

Loop closure detection ⇔ place recognition.
Pairwise comparison of data, here point clouds,

$$\mathbf{p}_k = \{p_i^k\}_{i=1}^N, \quad p_i^k \in \mathbb{R}^2$$

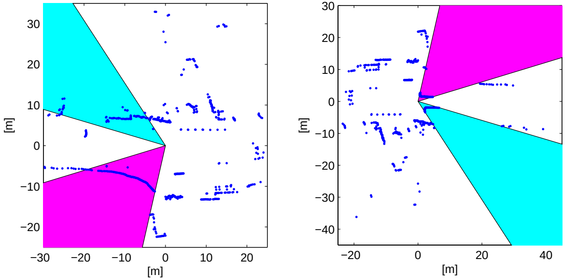

Are $\mathbf{p}_k$ and $\mathbf{p}_l$ from the same location? **Yes**

Loop closure/place recognition is an important and difficult problem, especially in SLAM.

Need a method thas is

- robust against misclassification,
- invariant to rotation and
- computationally inexpensive.

L.D. – our approach

Learning phase | Classification phase (part of SLAM)

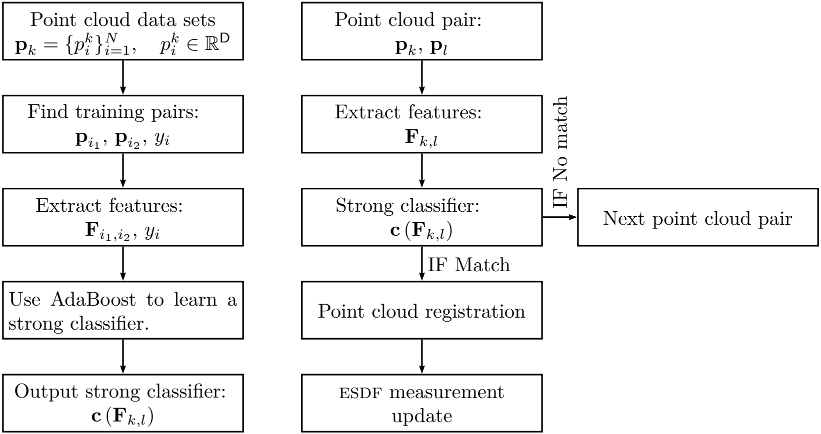

AUTOMATIC CONTROL
REGLERTEKNIK
LINKÖPINGS UNIVERSITET

L.D. – feature motivation

Point clouds are described with features:

- Meaningful statistics describing shape etc
- Compact description of point cloud, $n_f \ll N$
- Easy comparison of $\mathbf{p}_k$ and $\mathbf{p}_l$.

Two types of features used, all invariant to rotation.

AUTOMATIC CONTROL
REGLERTEKNIK
LINKÖPINGS UNIVERSITET

- Type 1: geometric and statistic properites.
  - f_1 — area in 2D, volume in 3D
  - f_3 — average range
  - $\ldots$
- Comparison: $|f_i - f_i|$.
- Same place $\Rightarrow$ similar value $\Rightarrow$ small $|f_i - f_i|$.

L.D. – feature examples

- Type 1: geometric and statistic properites.
  - f_1 — area in 2D, volume in 3D
  - f_3 — average range
  - ...
- Comparison: $|f_i - f_i|$.
- Same place $\Rightarrow$ similar value $\Rightarrow$ small $|f_i - f_i|$.
- Type 2: range histograms.
  - f_{j} — bin size $b_j \in [0.1\mathrm{m}, 3\mathrm{m}]$
- Comparison: Cross correlation of f_{j}:s.
- Same place $\Rightarrow$ similar f_{j} $\Rightarrow$ High cross correlation.

AUTOMATIC CONTROL
REGLERTEKNIK
LINKÖPINGS UNIVERSITET

AdaBoost is used to learn a 2-class classifier.

- Iterative learning. Combination of simple, “weak”, classifiers.

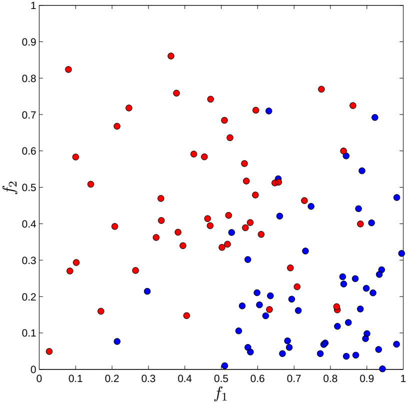

Initialise weights

AdaBoost is used to learn a 2-class classifier.

- Iterative learning. Combination of simple, “weak”, classifiers.

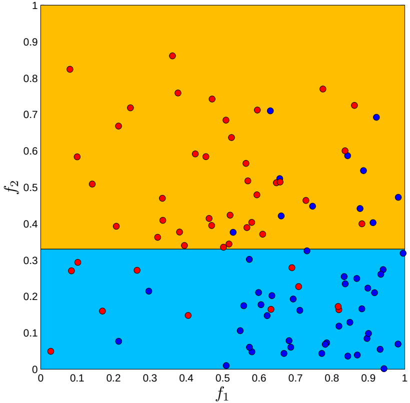

- $c_1 = (f_2 < 0.33)$

Minimise error Increase weights

AdaBoost is used to learn a 2-class classifier.

- Iterative learning. Combination of simple, “weak”, classifiers.

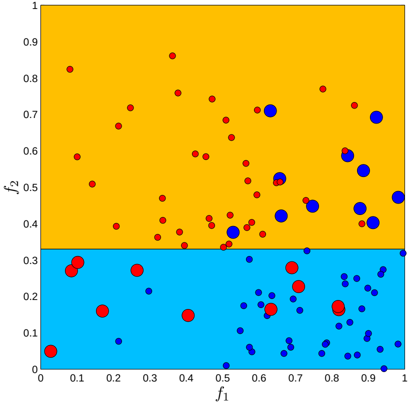

- $c_1 = (f_2 < 0.33)$

Minimise error Increase weights

AdaBoost is used to learn a 2-class classifier.

- Iterative learning. Combination of simple, “weak”, classifiers.

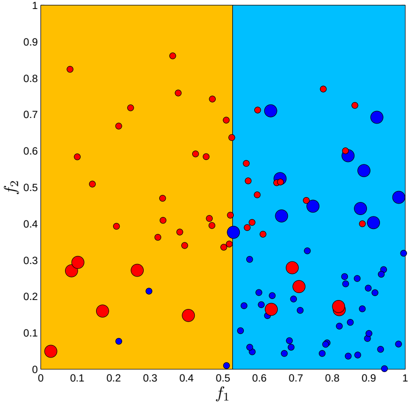

- $c_1 = (f_2 < 0.33)$
- $c_2 = (f_1 > 0.53)$

Minimise error Increase weights

AdaBoost is used to learn a 2-class classifier.

- Iterative learning. Combination of simple, “weak”, classifiers.

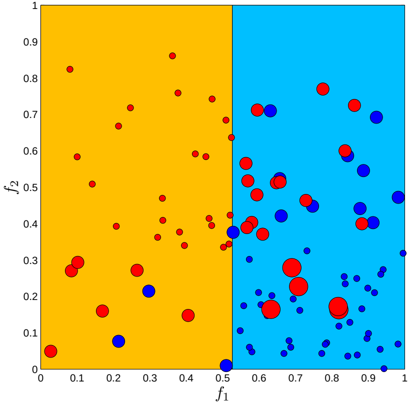

- $c_1 = (f_2 < 0.33)$
- $c_2 = (f_1 > 0.53)$

Minimise error Increase weights

AdaBoost is used to learn a 2-class classifier.

- Iterative learning. Combination of simple, “weak”, classifiers.

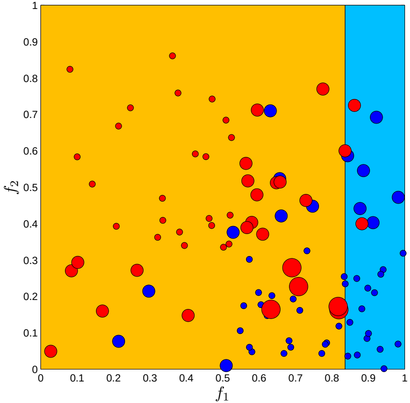

- $c_1 = (f_2 < 0.33)$
- $c_2 = (f_1 > 0.53)$
- $c_3 = (f_1 > 0.84)$

Minimise error Increase weights

AdaBoost is used to learn a 2-class classifier.

- Iterative learning. Combination of simple, “weak”, classifiers.

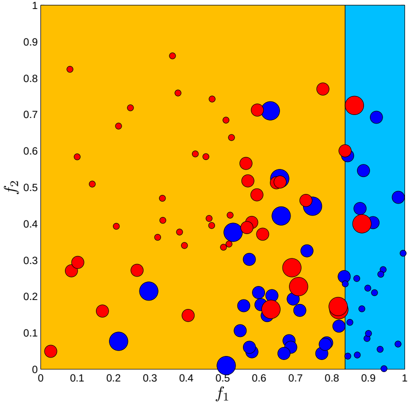

- $c_1 = (f_2 < 0.33)$
- $c_2 = (f_1 > 0.53)$
- $c_3 = (f_1 > 0.84)$

Minimise error Increase weights

AdaBoost is used to learn a 2-class classifier.

- Iterative learning. Combination of simple, “weak”, classifiers.

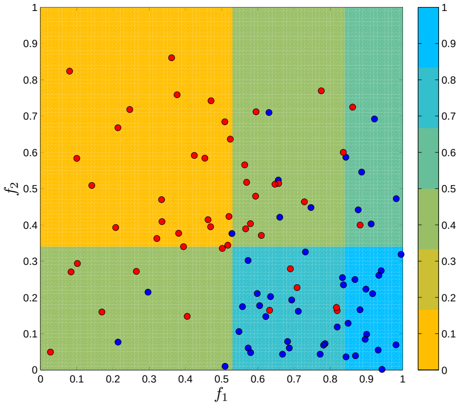

- $c_1 = (f_2 < 0.33)$
- $c_2 = (f_1 > 0.53)$
- $c_3 = (f_1 > 0.84)$
- Likelihood:

$$\mathbf{c} = \frac{\sum_t \alpha_t c_t}{\sum_t \alpha_t}$$

AUTOMATIC CONTROL
REGLERTEKNIK
LINKÖPINGS UNIVERSITET

AdaBoost is used to learn a 2-class classifier.

- Iterative learning. Combination of simple, “weak”, classifiers.

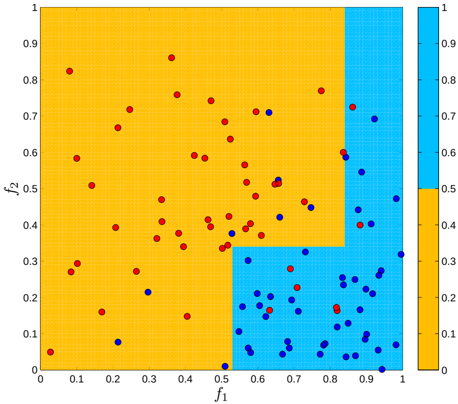

- $c_1 = (f_2 < 0.33)$
- $c_2 = (f_1 > 0.53)$
- $c_3 = (f_1 > 0.84)$
- Likelihood:

$$\mathbf{c} = \frac{\sum_t \alpha_t c_t}{\sum_t \alpha_t}$$

- Decision regions:

$$\mathbf{c} \geq \tau = 0.5$$

L.D. – experiments

1. Execution time
2. Number of weak classifiers needed
3. Most informative features
4. Receiver Operating Characteristic (ROC)
5. Comparison of 2D and 3D performance
6. Dependence to translation
7. Dynamic objects
8. Repetitive structures in the environment
9. Loop closure detection in SLAM

1. Execution time
2. Number of weak classifiers needed
3. Most informative features
4. Receiver Operating Characteristic (ROC)
5. Comparison of 2D and 3D performance
6. Dependence to translation
7. Dynamic objects
8. Repetitive structures in the environment
9. Loop closure detection in SLAM

ROC: detection/false alarm trade-off

Detection at 0% false alarm.

Outdoor 2D:

- 1m: 66% detection
- 2m: 34% detection
- 3m: 18% detection

Outdoor 3D:

- 3m: 63% detection

Indoor 3D:

- 1m: 53% detection

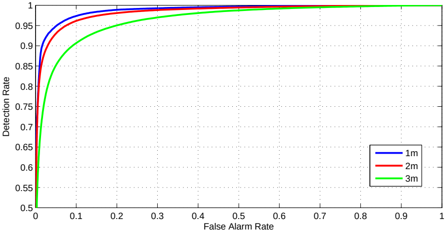

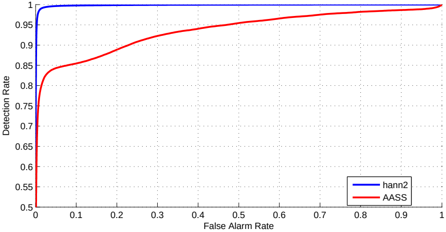

L.D. – 3D SLAM experiment, data

- $\mathbf{c}\left(\mathbf{F}_{k,l}\right)$ trained on outdoor data.
  - $r_{\text{max}} = 30\text{m}$.
  - $N \approx 17'000$.
- SLAM experiment on indoor data.
  - $r_{\text{max}} = 15\text{m}$.
  - $N \approx 110'000$.

Does $\mathbf{c}\left(\mathbf{F}_{k,l}\right)$ work in SLAM experiments?
Does $\mathbf{c}\left(\mathbf{F}_{k,l}\right)$ generalise well between environments?

~50% detection, no false alarms, good environment generalisation.

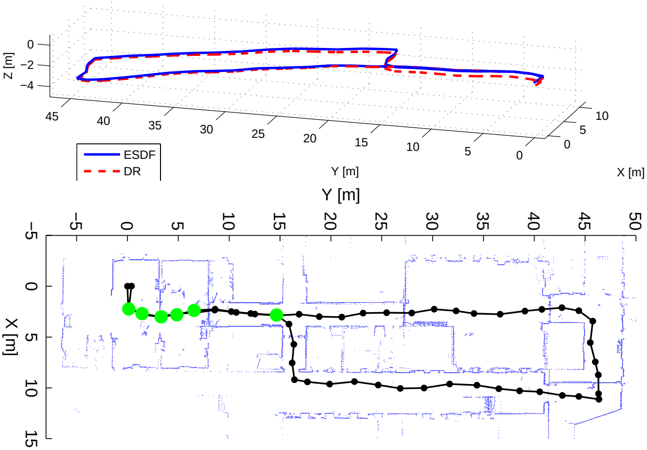

L.D. – comparison

Comparison with loop closure detection methods for point clouds.

- 2D: Higher detection at low false alarm.

Comparison with loop closure detection methods for point clouds.

- 2D: Higher detection at low false alarm.
- 3D outdoor:
  - 2nd highest detection
  - Faster than related work.
  - Rotation invariant without assumptions.

Comparison with loop closure detection methods for point clouds.

- 2D: Higher detection at low false alarm.
- 3D outdoor:
  - 2^{nd} highest detection
  - Faster than related work.
  - Rotation invariant without assumptions.
- 3D indoor:
  - Lower detection
  - Not enough training data from the same location.

AUTOMATIC CONTROL
REGLERTEKNIK
LINKÖPINGS UNIVERSITET

Comparison with loop closure detection methods for point clouds.

- 2D: Higher detection at low false alarm.
- 3D outdoor:
  - 2^{nd} highest detection
  - Faster than related work.
  - Rotation invariant without assumptions.
- 3D indoor:
  - Lower detection
  - Not enough training data from the same location.
- 2D and 3D: More sensitive to translation.

Comparison with loop closure detection methods for point clouds.

- 2D: Higher detection at low false alarm.
- 3D outdoor:
  - 2^{nd} highest detection
  - Faster than related work.
  - Rotation invariant without assumptions.
- 3D indoor:
  - Lower detection
  - Not enough training data from the same location.
- 2D and 3D: More sensitive to translation.

Much work using images – difficult to compare different sensors.

A machine learning approach for the loop closure detection problem using point clouds.

- > 40 rotation invariant features.
- Loop closure detected from arbitrary direction.
- Competitive detection for low false alarm (0%).
- Fast to compute.
- Method generalises well between environments.
- SLAM experiments shows the method works in real problems.

E.T.T. – Target tracking?

- Target tracking $\Rightarrow$ Keep track of location of targets
  - Unknown number, not always detected
  - Noise and clutter
  - Difficult data association

E.T.T. – Target tracking?

- Target tracking $\Rightarrow$ Keep track of location of targets
  - Unknown number, not always detected
  - Noise and clutter
  - Difficult data association
- Early RADAR-airplane-tracking
  - Targets behave as points

AUTOMATIC CONTROL
REGLERTEKNIK
LINKÖPINGS UNIVERSITET

E.T.T. – Target tracking?

- Target tracking $\Rightarrow$ Keep track of location of targets
  - Unknown number, not always detected
  - Noise and clutter
  - Difficult data association
- Early RADAR-airplane-tracking
  - Targets behave as points
- Modern sensors – no longer points
  - Multiple measurements

Definition:
Extended targets are targets that potentially give rise to more than one measurement per time step.

AUTOMATIC CONTROL
REGLERTEKNIK
LINKÖPINGS UNIVERSITET

Point target assumption is often not valid, e.g.

- ...laser sensors
- ...automotive radar
- ...camera images

Need method that handles multiple measurements per target.

Approach: Use random finite sets, RFS.

- RFS $\mathbf{Y} = \left\{\mathbf{y}^{(i)}\right\}_{i=1}^{N_y}$
- $\mathbf{y}^{(i)}$ are random vectors, common assumption $\mathbf{y}^{(i)} \in \mathbb{R}^n$
- $N_y < \infty$ is random, common assumption $N_y \in \mathcal{POIS}\left(\beta\right)$
- No order, e.g. $\left\{\mathbf{y}^{(1)}, \mathbf{y}^{(2)}\right\} = \left\{\mathbf{y}^{(2)}, \mathbf{y}^{(1)}\right\}$

- RFS of targets $\mathbf{X}_k = \{\mathbf{x}_k^{(i)}\}_{i=1}^{N_{x,k}}$

$$\mathbf{x}_{k+1}^{(i)} = F_k \mathbf{x}_k^{(i)} + G_k \mathbf{w}_k^{(i)}, \qquad \mathbf{w}_k^{(i)} \in \mathcal{N}\left(\mathbf{0}, Q_k\right).$$

- RFS of measurements $\mathbf{Z}_k = \{\mathbf{z}_k^{(j)}\}_{j=1}^{N_{z,k}}$

$$\mathbf{z}_k^{(j)} = H_k \mathbf{x}_k^{(i)} + \mathbf{e}_k^{(j)}, \qquad \mathbf{e}_k^{(j)} \in \mathcal{N}\left(\mathbf{0}, R_k\right).$$

- Goal:
  Compute target set estimate $\hat{\mathbf{X}}_k$ using measurement sets $\mathbf{Z}_k$.

- Bayes recursion computationally intractable [Mahler (2007)].

- Bayes recursion computationally intractable [Mahler (2007)].
- Approximation: propagate first order moment $D_{k|k}\left(\mathbf{x}|\mathbf{Z}^{(k)}\right)$, called PHD-intensity.

$$\ldots\ D_{k|k}\left(\mathbf{x}|\mathbf{Z}^{(k)}\right) \overset{\text{Prediction}}{\overbrace{\longrightarrow}} D_{k+1|k}\left(\mathbf{x}|\mathbf{Z}^{(k)}\right) \overset{\text{Correction}}{\overbrace{\longrightarrow}} D_{k+1|k+1}\left(\mathbf{x}|\mathbf{Z}^{(k+1)}\right)\ \ldots$$

- Bayes recursion computationally intractable [Mahler (2007)].
- Approximation: propagate first order moment $D_{k|k}\left(\mathbf{x}|\mathbf{Z}^{(k)}\right)$, called PHD-intensity.

$$\ldots\ D_{k|k}\left(\mathbf{x}|\mathbf{Z}^{(k)}\right) \overbrace{\longrightarrow}^{\text{Prediction}} D_{k+1|k}\left(\mathbf{x}|\mathbf{Z}^{(k)}\right) \overbrace{\longrightarrow}^{\text{Correction}} D_{k+1|k+1}\left(\mathbf{x}|\mathbf{Z}^{(k+1)}\right)\ \ldots$$

- Analogous to α-β-filter, which propagates first order moment of random variable (i.e. mean vector).

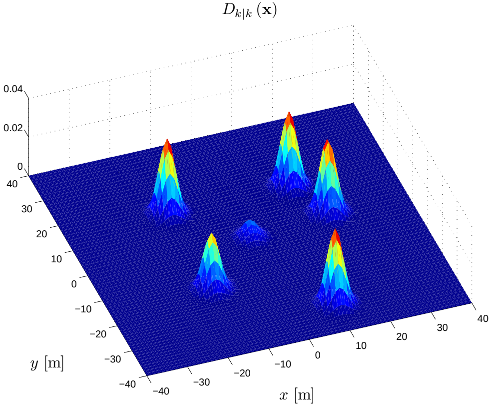

PHD-intensity

PHD-intensity is sum of Gaussians

$$D_{k|k}\left(\mathbf{x}|\mathbf{Z}^{(k)}\right) = \sum_{j=1}^{J_{k|k}} w_{k|k}^{(j)} \mathcal{N}\left(\mathbf{x} \mid m_{k|k}^{(j)}, P_{k|k}^{(j)}\right)$$

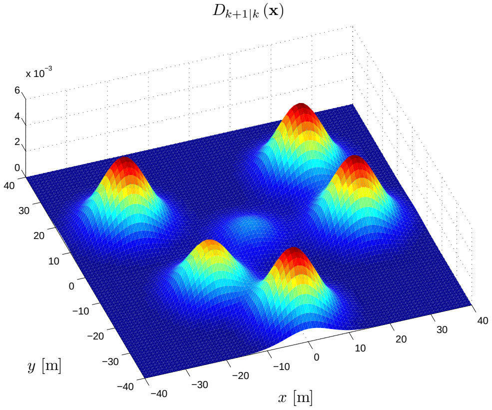

Predicted intensity

PHD-intensity is sum of Gaussians

$$D_{k|k}\left(\mathbf{x}|\mathbf{Z}^{(k)}\right) = \sum_{j=1}^{J_{k|k}} w_{k|k}^{(j)} \mathcal{N}\left(\mathbf{x} \mid m_{k|k}^{(j)}, P_{k|k}^{(j)}\right)$$

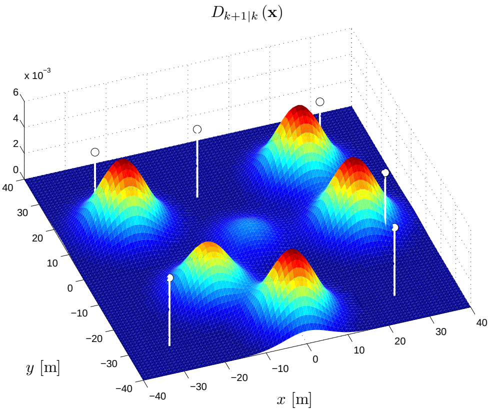

Cluttered set of measurements

PHD-intensity is sum of Gaussians

$$D_{k|k}\left(\mathbf{x}|\mathbf{Z}^{(k)}\right) = \sum_{j=1}^{J_{k|k}} w_{k|k}^{(j)} \mathcal{N}\left(\mathbf{x} \mid m_{k|k}^{(j)}, P_{k|k}^{(j)}\right)$$

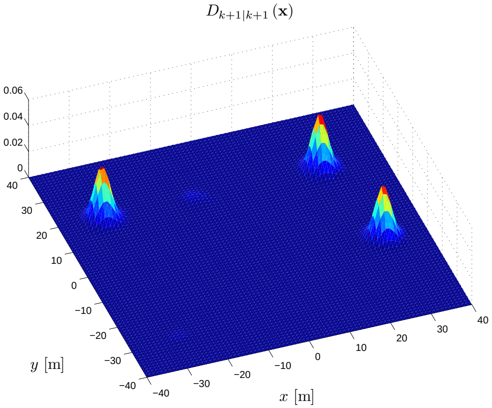

Corrected intensity

PHD-intensity is sum of Gaussians

$$D_{k|k}\left(\mathbf{x}|\mathbf{Z}^{(k)}\right) = \sum_{j=1}^{J_{k|k}} w_{k|k}^{(j)} \mathcal{N}\left(\mathbf{x} \,|\, m_{k|k}^{(j)}, P_{k|k}^{(j)}\right)$$

- Implementationt of prediction shown by [Vo and Ma, 2006].
- $D_{k|k-1}(\mathbf{x}|\mathbf{Z})$ is predicted PHD-intensity. Corrected PHD-intensity

$$D_{k|k}(\mathbf{x}|\mathbf{Z}) = L_{\mathbf{Z}_k}(\mathbf{x}) D_{k|k-1}(\mathbf{x}|\mathbf{Z}),$$

where measurement pseudo-likelihood is given by

$$L_{\mathbf{Z}_k}(\mathbf{x}) = 1 - \left(1 - e^{-\gamma(\mathbf{x})}\right) p_D(\mathbf{x}) + e^{-\gamma(\mathbf{x})} p_D(\mathbf{x}) \sum_{\mathrm{p}\angle\mathbf{Z}_k} \omega_{\mathrm{p}} \sum_{W \in \mathrm{p}} \frac{\gamma(\mathbf{x})^{|W|}}{d_W} \cdot \prod_{\mathbf{z} \in W} \frac{\phi_{\mathbf{z}}(\mathbf{x})}{\lambda_k c_k(\mathbf{z})}.$$

[Mahler, 2009]

E.T.T. – measurement partitioning

- In each time step $\mathbf{Z}_k$ must be partitioned.
- A partition p is a division of $\mathbf{Z}_k$ into cells W.
- Important since more than one measurement can stem from the same target.

E.T.T. – measurement partitioning example

Partition the measurement set $\mathbf{Z}_k = \left\{ \mathbf{z}_k^{(1)}, \mathbf{z}_k^{(2)}, \mathbf{z}_k^{(3)} \right\}$

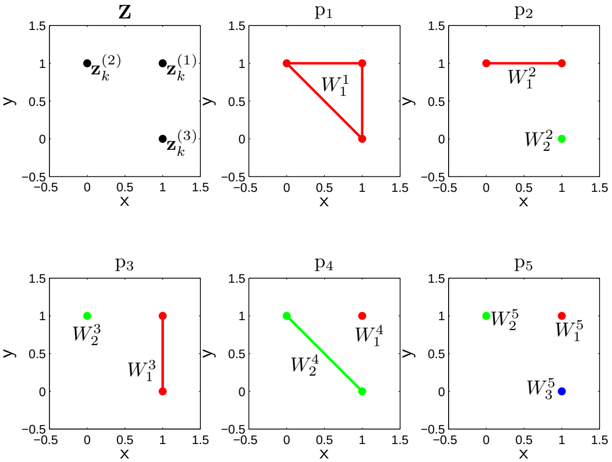

AUTOMATIC CONTROL
REGLERTEKNIK
LINKÖPINGS UNIVERSITET

- Measurements belong to same cell W if distance is “small”.
- Partions p_i where cells contain measurements $< d_i \mathrm{m}$ apart.

- Measurements belong to same cell W if distance is “small”.
- Partions p_i where cells contain measurements $< d_i\mathrm{m}$ apart.
- Let $\{d_i^m\}_{i=1}^{N_{z,k}}$ be set of measurement to measurement distances.
- Good partitions for d_i corresponding to

$$d_{\text{min}} \leq d_i^m < d_{\text{max}}$$

- Measurements belong to same cell W if distance is “small”.
- Partions p_i where cells contain measurements $< d_i \mathrm{m}$ apart.
- Let $\{d_i^m\}_{i=1}^{N_{z,k}}$ be set of measurement to measurement distances.
- Good partitions for d_i corresponding to

$$d_{\text{min}} \leq d_i^m < d_{\text{max}}$$

- Use knowledge about scenario to determine d_{min} and d_{max}

E.T.T. – measurement partitioning method example

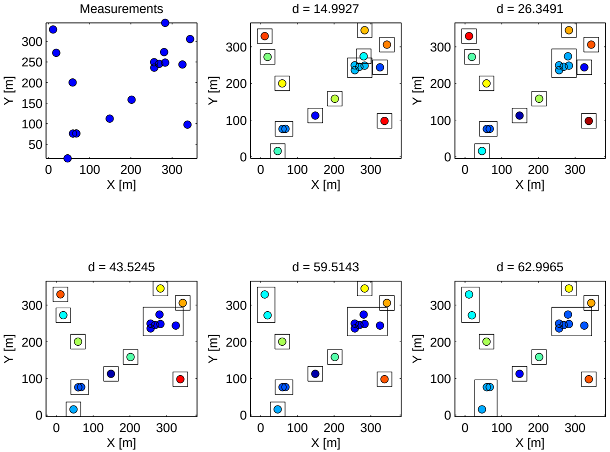

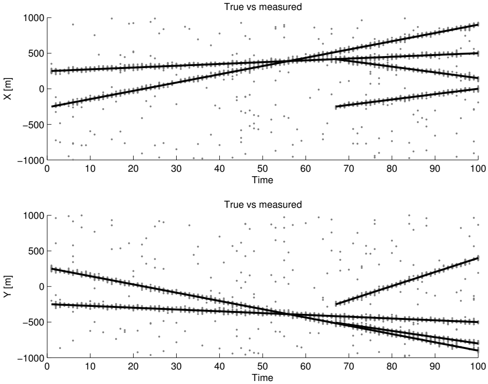

- True target track crossing at time $k = 56$.
- New target birth and target spawned at time $k = 66$.
- Evaluation against standard GM-PHD using OSPA metric.

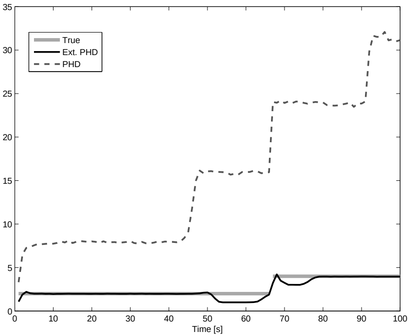

Cardinality

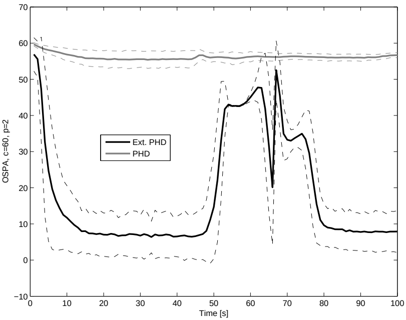

OSPA

- 2D laser range scanner
- Two human targets
- One target occluded by the other
- Variable probability of detection to handle occlusion

targetTracking.avi

– Implementation of GM-PHD-filter for extended targets.
– Simple method for measurement set partitioning.
– Variable p_{D} to handle occlusion.

The suggested implementation handles...

- ...unknown number of targets.
- ...noisy and cluttered measurement sets.

Future work

- Loop detection:
  - Decrease sensitivity to translation

Future work

- Loop detection:
    - Decrease sensitivity to translation

- Extended target tracking:
    - Estimate target shape and size
    - Cardinalised PHD-filter – robust estimate of number of targets
    - Compare to other extended target tracking solutions

Future work

- Loop detection:
  - Decrease sensitivity to translation
- Extended target tracking:
  - Estimate target shape and size
  - Cardinalised PHD-filter – robust estimate of number of targets
  - Compare to other extended target tracking solutions
- Further extensions
  - Environment labeling of data
  - Separation of background environment and dynamic targets
  - Semi-supervised learning instead of supervised learning

Thank you for listening!