Learning to Close the Loop from 3D Point Clouds

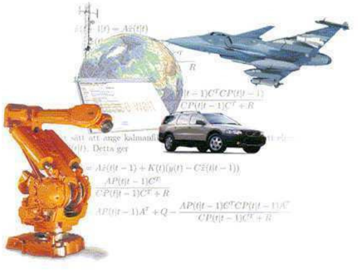

Karl Granström, Thomas B. Schön

Division of Automatic Control
Department of Electrical Engineering
Linköping University, Sweden

Loop closure — Problem definition

Loop closure detection ⇔ place recognition.

Pairwise comparison of data, here point clouds,

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

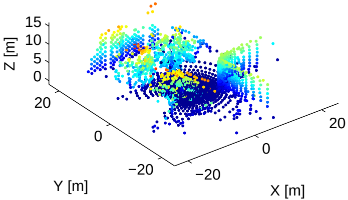

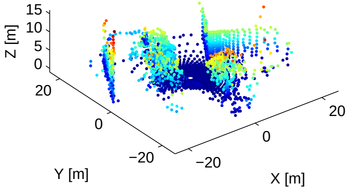

Are $\mathbf{p}_k$ and $\mathbf{p}_l$ from the same location?

Loop closure — Problem definition

Loop closure detection ⇔ place recognition.

Pairwise comparison of data, here point clouds,

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

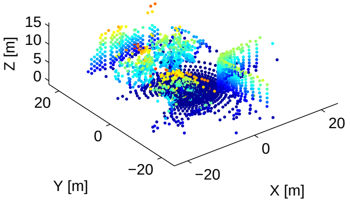

Are $\mathbf{p}_k$ and $\mathbf{p}_l$ from the same location?

Yes

Motivation

Loop closure/place recognition is an important and difficult problem:

- Important in robotics, especially in SLAM.
- Range sensors are common.

Need a method thas is

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

Relation to Related Work

- Raw sensor data instead of landmarks $\Rightarrow$ general representation of environment.
- Same approach, 2D 85% detection at 1% false alarm [Granström et al, 2009].
- NDT-based approach 3D 47% detection at 0% false alarm [Magnusson et al, 2009].

Our results for the same 3D data set:
63% detection at 0% false alarm
99% detection at 1% false alarm

AUTOMATIC CONTROL
REGLERTEKNIK
LINKÖPINGS UNIVERSITET

Approach to loop closing

Learning phase

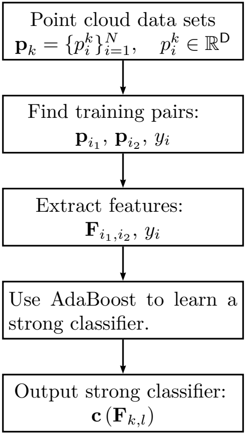

Classification phase (part of SLAM)

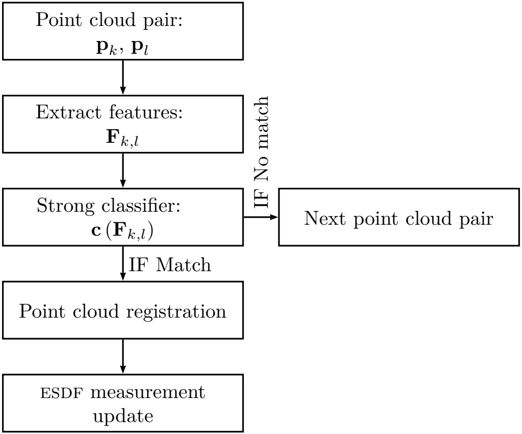

AUTOMATIC CONTROL
REGLERTEKNIK
LINKÖPINGS UNIVERSITET

Point clouds are described with features:

- Meaningful statistics describing shape etc
- Compact description of point cloud, $n_f = 41 \ll N$
- Easy comparison of $\mathbf{p}_k$ and $\mathbf{p}_l$.

Two types of features used, all invariant to rotation.

- Type 1: 32 geometric and statistic properites.
  - f_1 — volume
  - f_3 — average range
  - $\ldots$
- Comparison: $|f_i - f_i|$.
- Same place $\Rightarrow$ similar value $\Rightarrow$ small $|f_i - f_i|$.

Features

- Type 1: 32 geometric and statistic properites.
  - f_1 — volume
  - f_3 — average range
  - ...
- Comparison: $|f_i - f_i|$.
- Same place $\Rightarrow$ similar value $\Rightarrow$ small $|f_i - f_i|$.
- Type 2: 9 range histograms.
  - f_j — bin size $b_j \in [0.1\mathrm{m}, 3\mathrm{m}]$, $j = 33, \ldots, 41$.
- Comparison: Cross correlation of f_j:s.
- Same place $\Rightarrow$ similar f_j $\Rightarrow$ High cross correlation.

AdaBoost used to learn classifier.

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

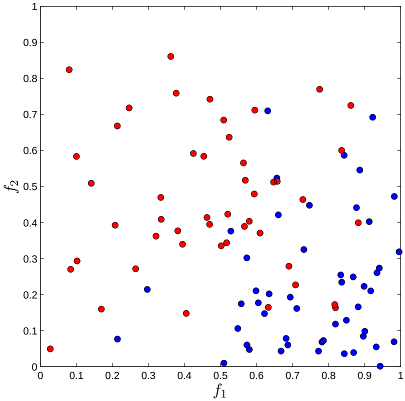

AdaBoost used to learn classifier.

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

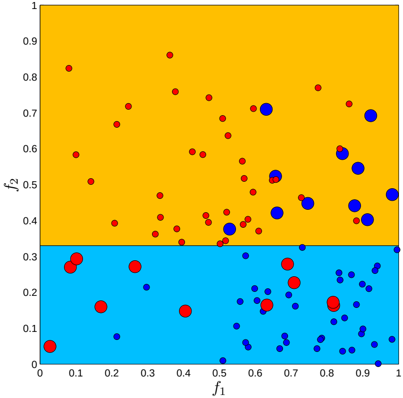

- $c_1 = (f_2 < 0.33)$

AdaBoost used to learn classifier.

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

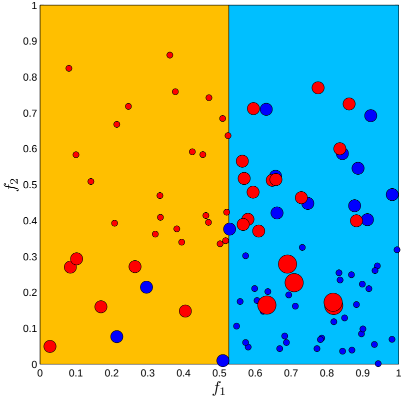

- $c_1 = (f_2 < 0.33)$
- $c_2 = (f_1 > 0.53)$

AdaBoost learned Classifier

AdaBoost used to learn classifier.

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

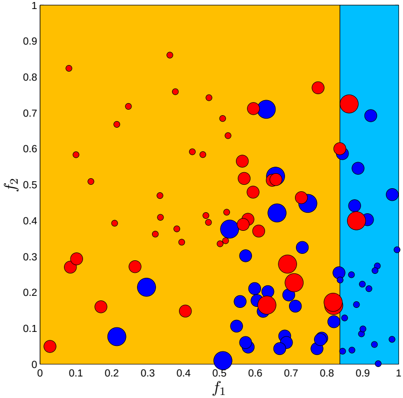

- $c_1 = (f_2 < 0.33)$
- $c_2 = (f_1 > 0.53)$
- $c_3 = (f_1 > 0.84)$

AdaBoost used to learn classifier.

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

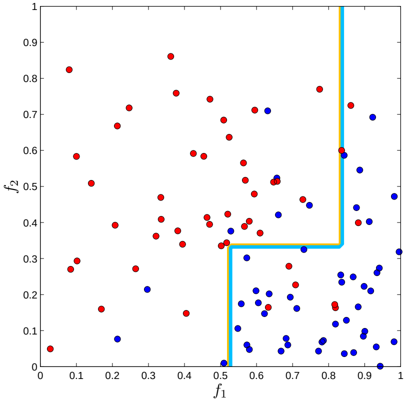

- $c_1 = (f_2 < 0.33)$
- $c_2 = (f_1 > 0.53)$
- $c_3 = (f_1 > 0.84)$
- $\mathbf{c} = \left(\sum_i \alpha_i c_i > \tau \sum_i \alpha_i\right)$

Training and experiment data

Data from two 3D data sets:

- *Hannover 2:* outdoor data set, training data.
  - 924 point clouds,
  - 1.24km trajectory,
  - $3130\mathbf{p}_k$ from same location, $7190\mathbf{p}_k$ from different location.

Training and experiment data

Data from two 3D data sets:

- *Hannover 2:* outdoor data set, training data.
  - 924 point clouds,
  - 1.24km trajectory,
  - $3130\mathbf{p}_k$ from same location, $7190\mathbf{p}_k$ from different location.
- *AASS-loop:* indoor data, training data and SLAM experiment.
  - 60 point clouds,
  - 111m trajectory,
  - $16\mathbf{p}_k$ from same location, $324\mathbf{p}_k$ from different location.

Both publicly available:

`kos.informatik.uni-osnabrueck.de/3Dscans/`

AUTOMATIC CONTROL
REGLERTEKNIK
LINKÖPINGS UNIVERSITET

$\mathbf{c}\left(\mathbf{F}_{k,l}\right)$ learned using the $3130 + 7190$ data pairs, $T \in [1, 250]$.

- $\mathbf{c}\left(\mathbf{F}_{k,l}\right)$ evaluated with 10-fold cross validation.
- Error rates approx. constant after 50 rounds.
- Overfitting not a concern.

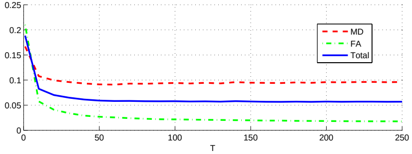

Validation errors

Hannover 2:

- 63% detection at 0% false alarm.
- Area under curve $\approx 99.9\%$.

AASS-loop:

- 53% detection at 0% false alarm.
- Area under curve $\approx 93.6\%$.

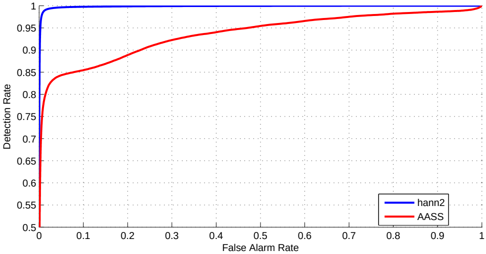

AUTOMATIC CONTROL
REGLERTEKNIK
LINKÖPINGS UNIVERSITET

Time complexity

Matlab implementation on 2.83GHz CPU, 3.48 GB RAM

- Time to compute features (once per point cloud):
    - *Hannover 2:* 19.34ms
    - *AASS-loop:* 225.10ms
- Compare features: 0.845ms
- Compute **c** $(\mathbf{F}_{k,l})$: 0.78ms

- $\mathbf{c}\left(\mathbf{F}_{k,l}\right)$ trained on *Hannover 2* data.
  - Outdoor.
  - $r_{\text{max}} = 30\text{m}$.
  - $N \approx 17'000$.
- SLAM experiment on *AASS-loop* data.
  - Indoor
  - $r_{\text{max}} = 15\text{m}$.
  - $N \approx 110'000$.

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.

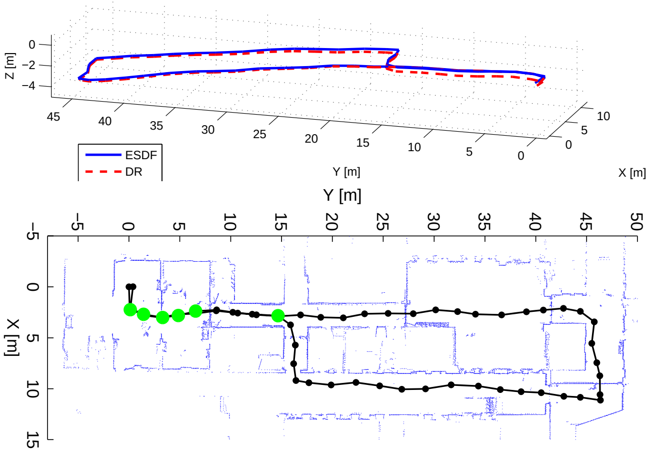

A machine learning approach for the loop closure detection problem using 3D point clouds.

- 41 rotation invariant features.
- Loop closure detected from arbitrary direction.
- Competitive detection 63% for low false alarm 0%.
- Method generalises well between environments.
- SLAM experiment shows the method works in a real problem.

Thank you for listening

Any questions?