Classifier Grids for Robust Adaptive Object Detection

Sabine Sternig

Institute for Computer Graphics and Vision
Graz University of Technology

Joint work with
Peter M. Roth, Helmut Grabner, and Horst Bischof

Goal

Development of a robust but still **adaptive** real-time object detection system for **static cameras!**

- Why are we interested in object detection from fixed cameras?

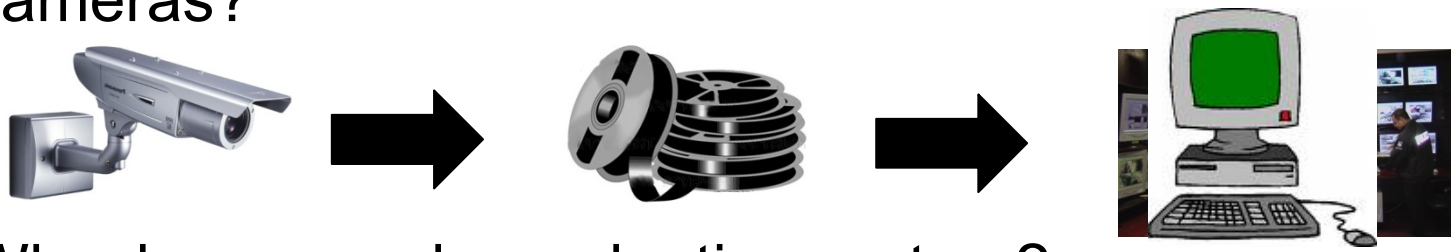

- Why do we need an adaptive system?

Generic object detector

Adaptive object detector

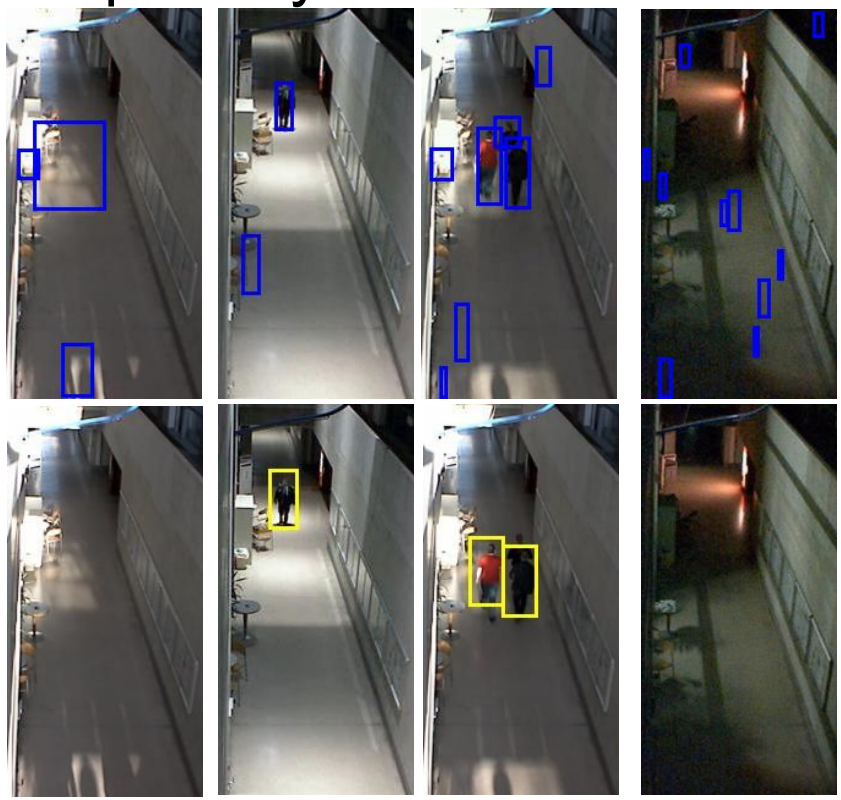

Outline

- Preliminaries
  - Classifier grids
  - Fixed update strategies
  - On-line Boosting for feature selection
- Adaptive grid-based detector
  - Off-line pre-training
  - On-line update strategy
  - On-line feature selection
- Experimental results
- Conclusion

Simplifying the Problem

Classifier should be applicable | **Training set** | **Complexity**

Generic detector

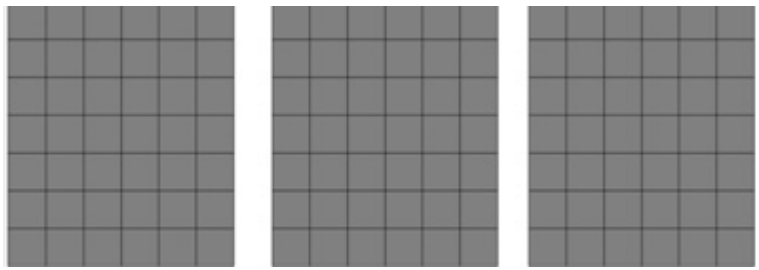

Scene specific

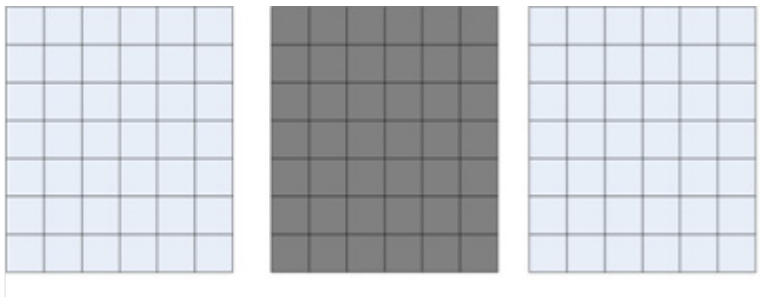

Classifier grid

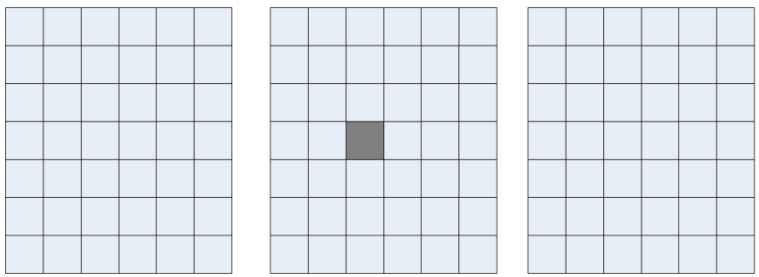

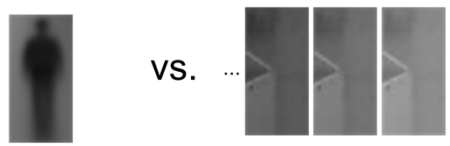

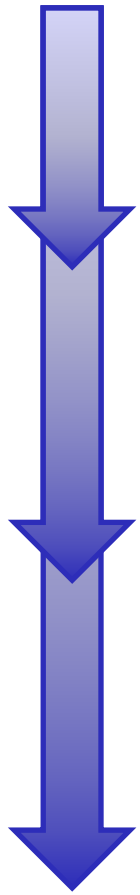

Update Strategies

- Adaptive systems need to be updated all the time
- Typical update schemes for classifier grids
  - Self-training
  - Co-training
- Disadvantage
  - Tend to drift
- Avoid drifting by using fixed update strategy [1]
  - **Positive updates:** hand-labeled set, correct by definition
  - **Negative updates:** low probability $P(x_i = \text{object}) = \frac{\# p_i}{\Delta t}$ of wrong negative update

[1] H. Grabner, P. Roth and H. Bischof. Is pedestrian detection really a hard task? In Proc. IEEE Int. Workshop on PETS'07.

On-line Boosting for Feature Selection [2]

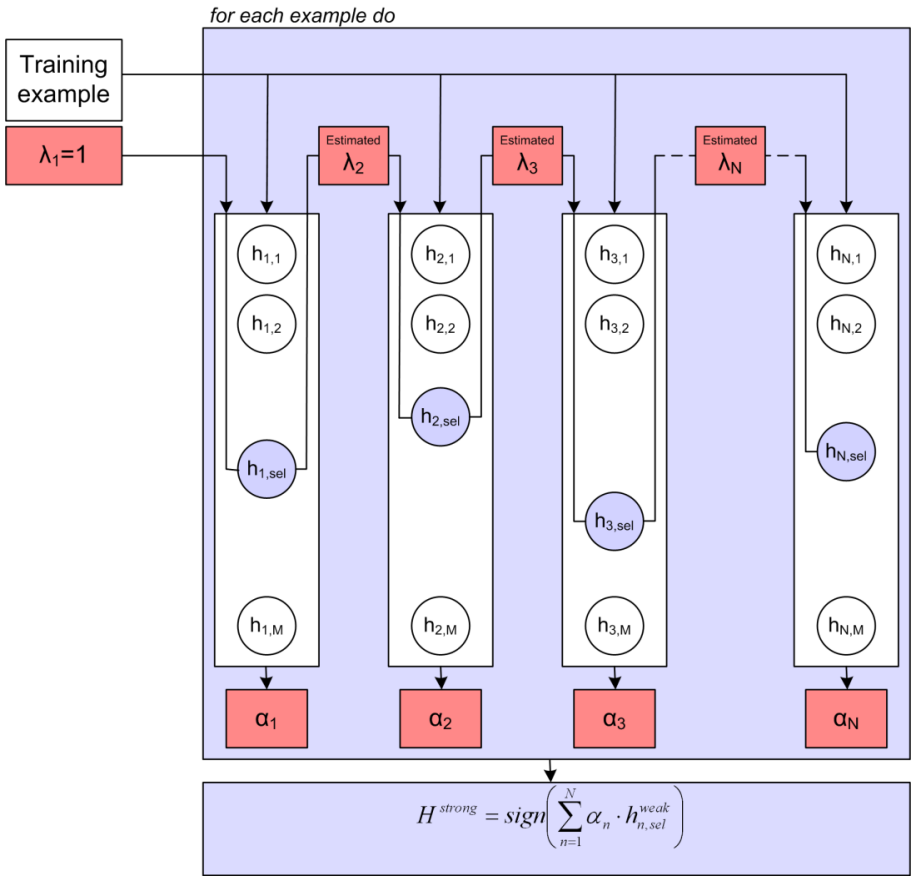

[2] H. Grabner, H. Bischof. On-line boosting and vision. CVPR'06

- Weak classifier $h_{i,j}(\mathbf{x})$
  - corresponds to one feature
  - are based on two distributions $D_{i,j}^{+}$ and $D_{i,j}^{-}$

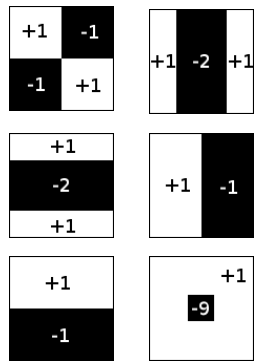

Haar-like features

Robust Adaptive Grid-based Classification

Fixed Generative Model for Object

Adaptive Generative Model for Background

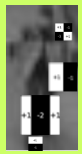

Discriminative Combination at Feature Level

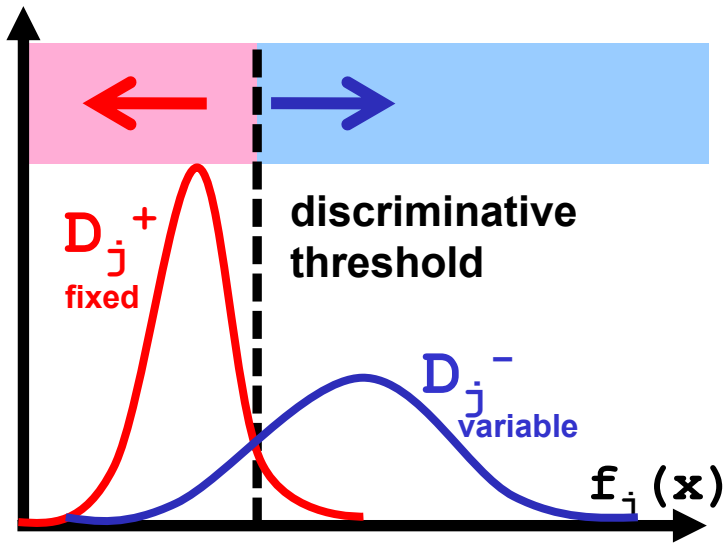

Robust Adaptive Grid-based Classification

- Three steps for estimating an adaptive discriminative model by combining two generative models:
  - Off-line pre-training
  - On-line weak classifier update
  - On-line feature selection

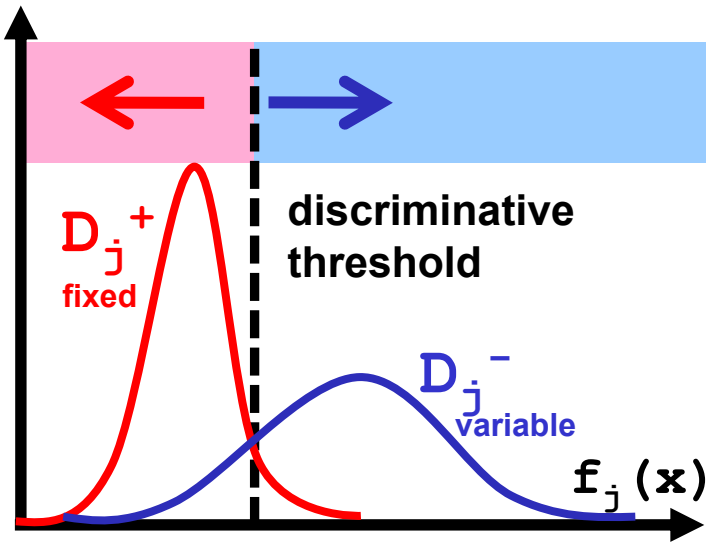

Off-line Pre-training for Feature Selection

- Performance of on-line classifier depends on randomly initialized features
- Include prior knowledge
- Keep statistics D_j^+ and error $e_+^{off-line}$ from off-line training

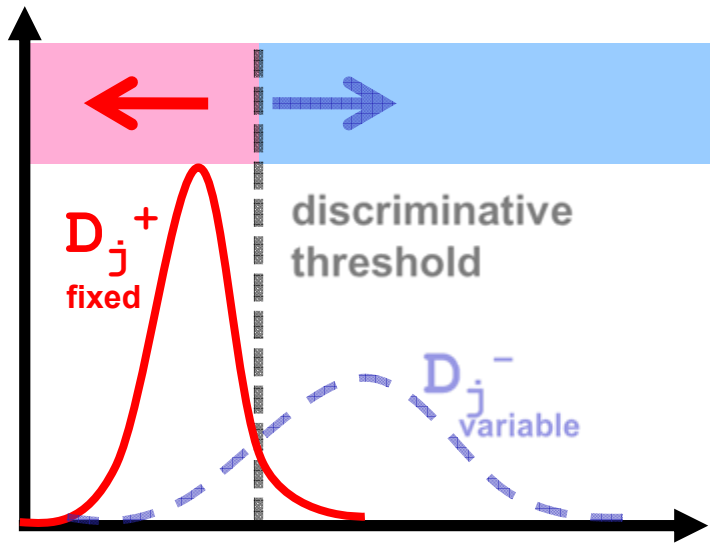

On-line Weak Classifier Update

- Positive distribution D_j^+ is kept fixed
- Negative distribution D_j^- is updated by the current patch
  - generative information used to reduce the amount of label noise

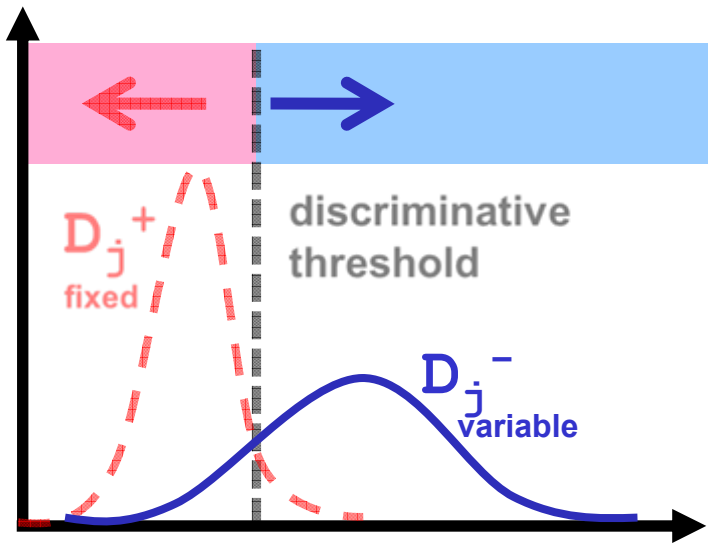

On-line Feature Selection

- Best weak classifier $h_n(\mathbf{x})$ for each selector determined
- During on-line training only negative updates -> only error for negative samples estimated $e_-^{on-line}$
- Estimate combined error

$$e = \frac{1}{2}\left(e_+^{off-line} + e_-^{on-line}\right)$$

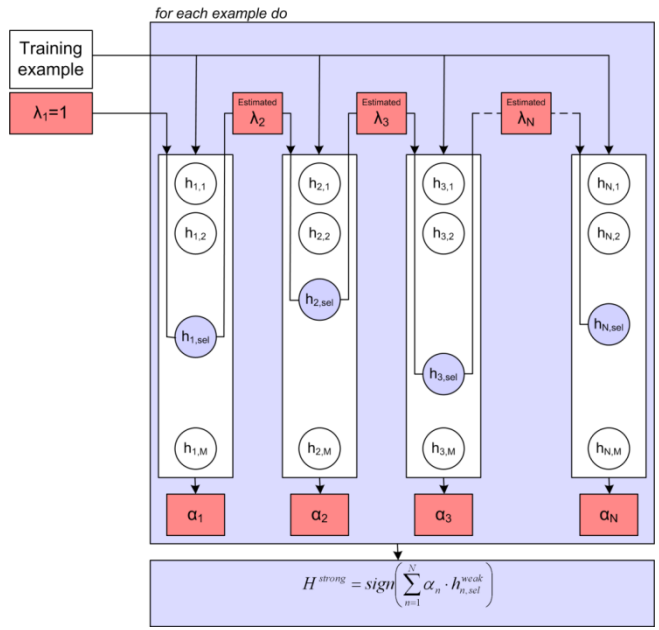

Adaptive Grid-based Detector

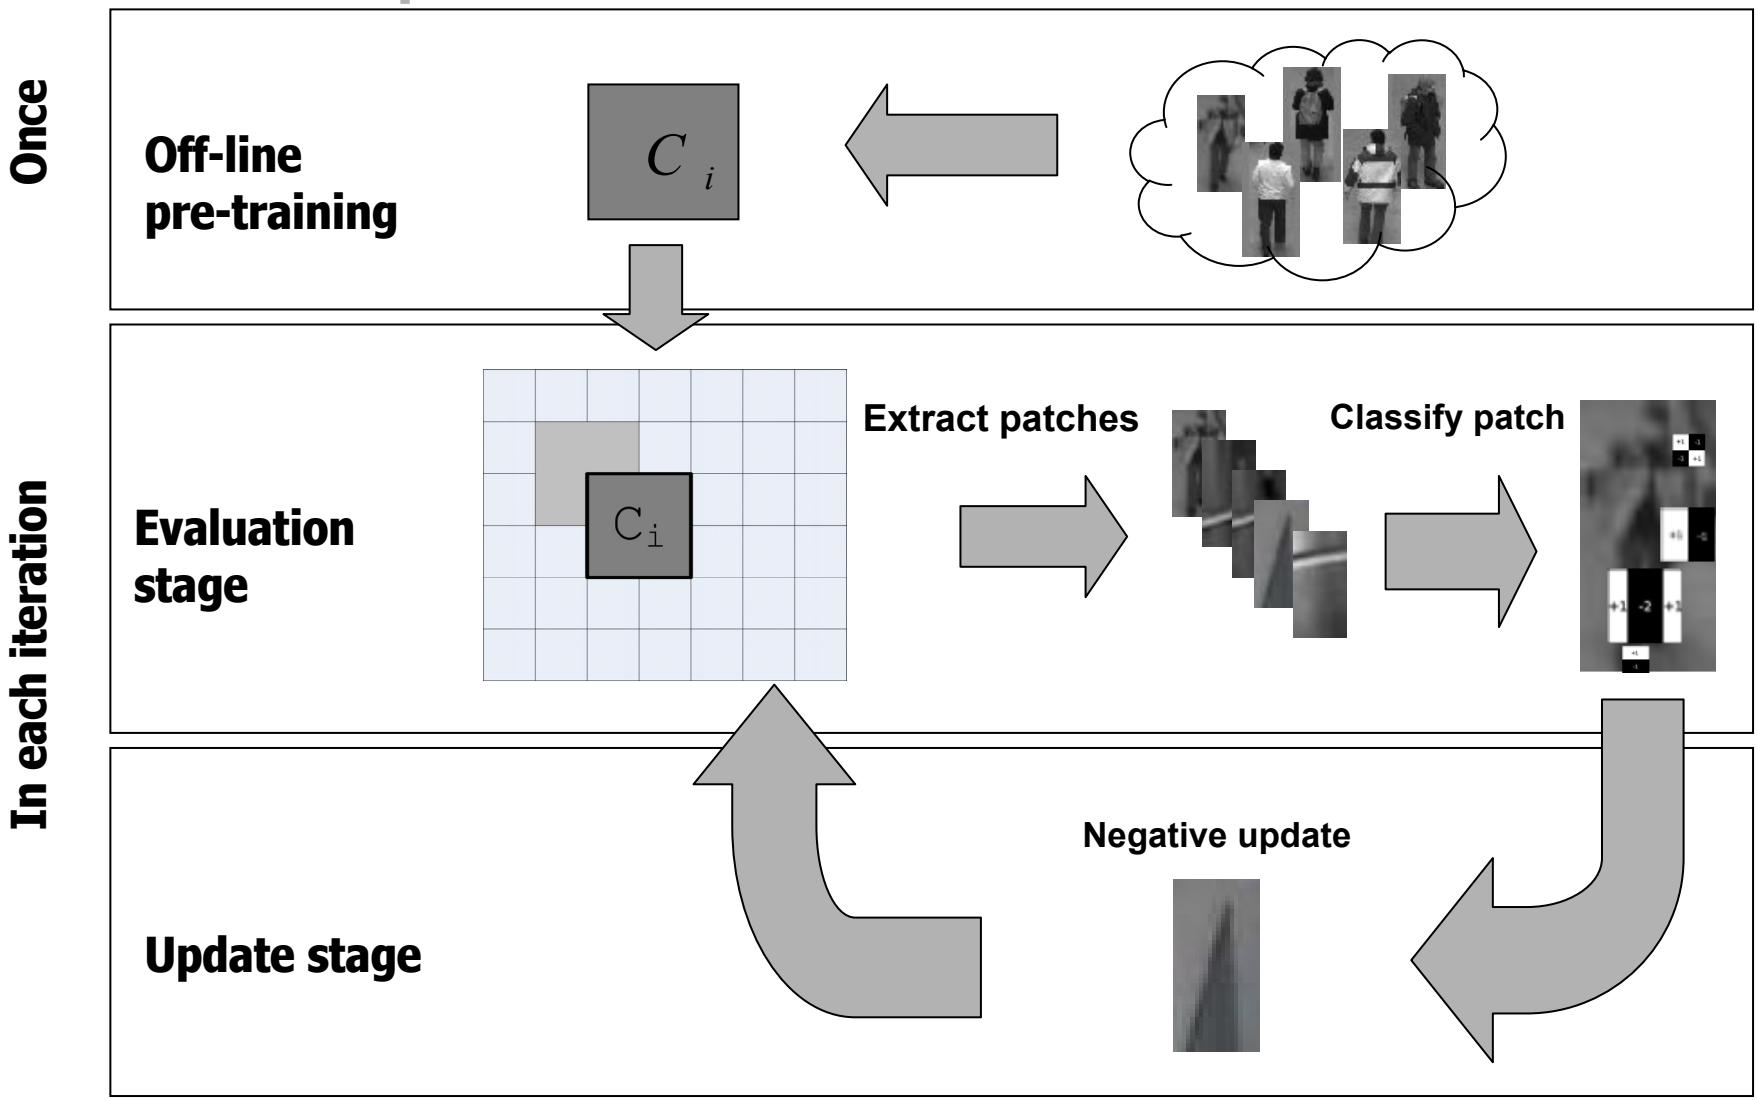

Experiments

- Person detection (publicly available datasets)
  - Caviar dataset [3]
  - PETS 2006 dataset [4]
- Car detection on highways
- Long-term experiment
  - 580.000 updates, running over one week

[3] http://homepages.inf.ed.ac.uk/rbf/CAVIARDATA1
[4] http://www.pets2006.net

Pets 2006 Dataset [4]

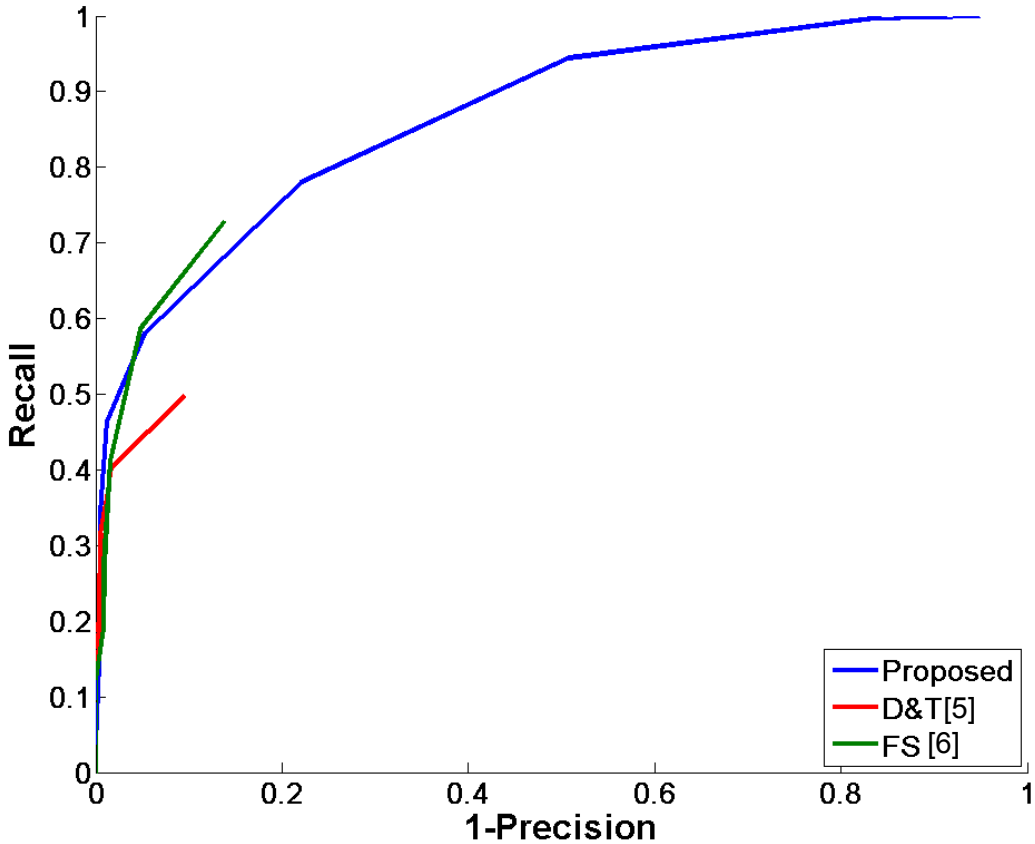

[4] http://www.pets2006.net
[5] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. CVPR'05
[6] P. Felzenszwalb, D. McAllester, and D. Ramanan. A discriminatively trained, multiscale, deformable part model. CVPR'08

Caviar dataset – ShopAssistant2cor [3]

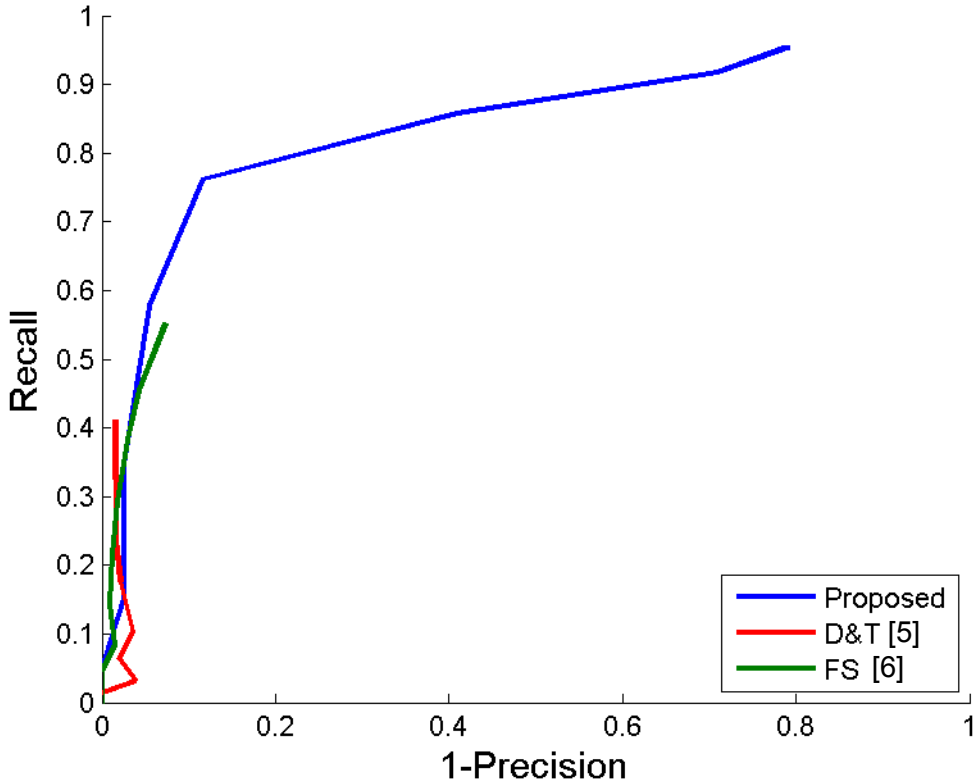

[3] http://homepages.inf.ed.ac.uk/rbf/CAVIARDATA1
[5] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. CVPR'05
[6] P. Felzenszwalb, D. McAllester, and D. Ramanan. A discriminatively trained, multiscale, deformable part model. CVPR'08

Car Detection on Highways

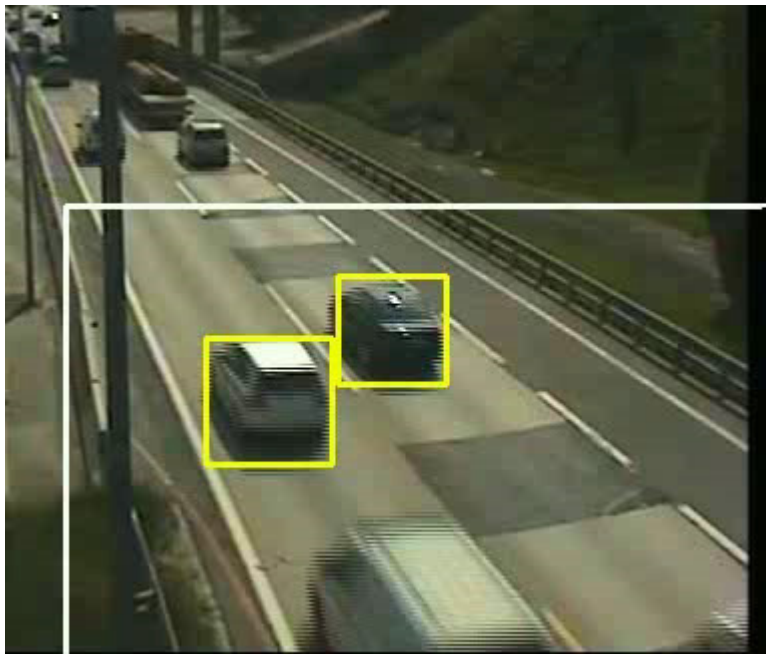

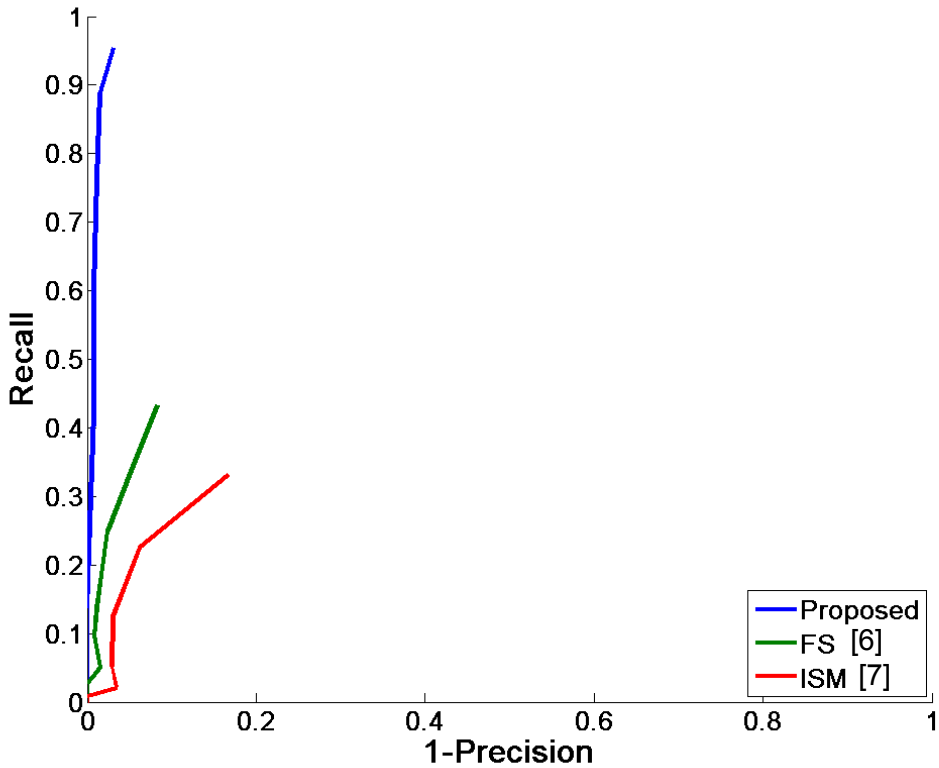

[6] P. Felzenszwalb, D. McAllester, and D. Ramanan. A discriminatively trained, multiscale, deformable part model. CVPR'08
[7] B. Leibe, A. Leonardis, and B. Schiele. Robust object detection with interleaved categorization and segmentation, Journal of Computer Vision 2008

Long-term Experiment

- Learning 24 hours a day, 7 days a week
- 580.000 updates

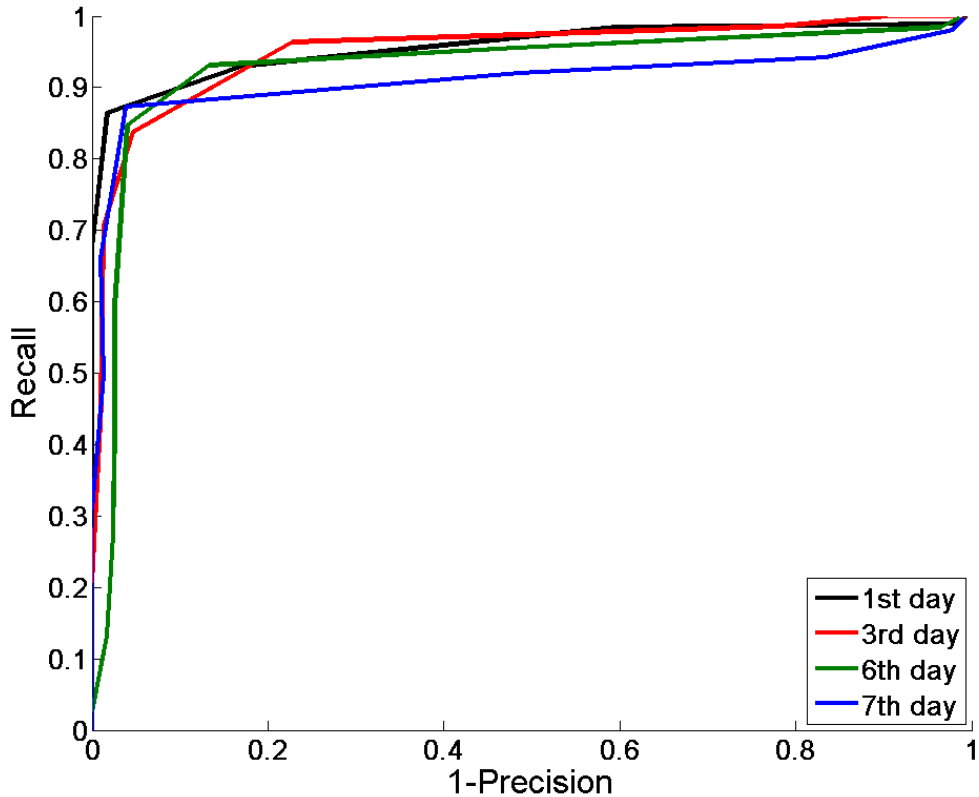

Demo

Conclusion

- Robust object detection for static cameras running stable 24 hours a day, 7 days a week
- Adaptive system which is still robust:
  - Two different representations
  - Fixed update strategies
- Small and compact classifiers enable real-time object detection

Thanks for your attention!

Off-line Pre-training for Feature Selection

Randomly selected features

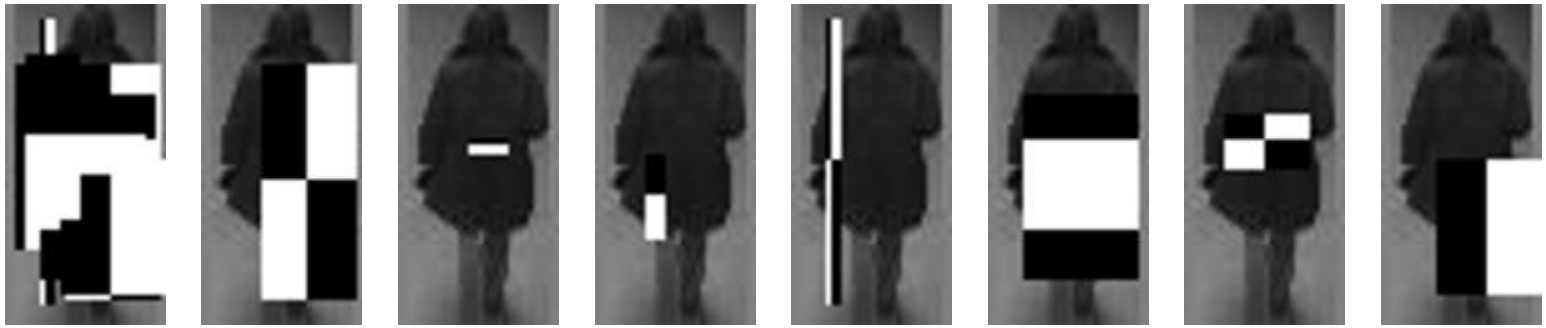

Off-line pre-trained

Classifier Grids

- Compact classifiers due to simplified problem