

Computer Vision 2

WS 2018/19

Part 5 – Tracking by Online Classification

24.10.2018

Prof. Dr. Bastian Leibe

RWTH Aachen University, Computer Vision Group

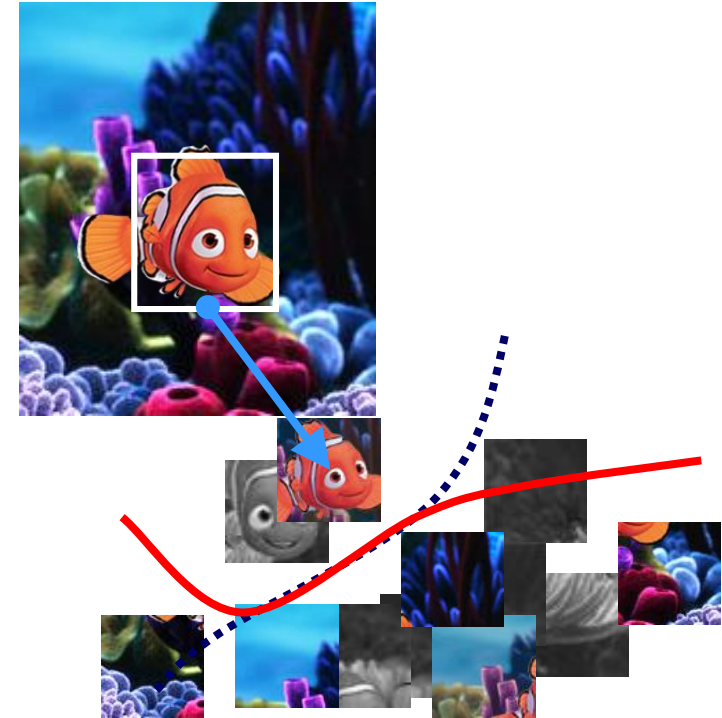
<http://www.vision.rwth-aachen.de>



RWTHAACHEN
UNIVERSITY

Course Outline

- **Single-Object Tracking**
 - Background modeling
 - Template based tracking
 - **Tracking by online classification**
 - Tracking-by-detection
- Bayesian Filtering
- Multi-Object Tracking
- Visual Odometry
- Visual SLAM & 3D Reconstruction
- Deep Learning for Video Analysis



Recap: General LK Image Registration

- Goal

- Find the warping parameters \mathbf{p} that minimize the sum-of-squares intensity difference between the template image $T(\mathbf{x})$ and the warped input image $I(\mathbf{W}(\mathbf{x};\mathbf{p}))$.

- LK formulation

- Formulate this as an optimization problem

$$\arg \min_{\mathbf{p}} \sum_{\mathbf{x}} [I(\mathbf{W}(\mathbf{x}; \mathbf{p})) - T(\mathbf{x})]^2$$

- We assume that an initial estimate of \mathbf{p} is known and iteratively solve for increments to the parameters $\Delta\mathbf{p}$:

$$\arg \min_{\Delta\mathbf{p}} \sum_{\mathbf{x}} [I(\mathbf{W}(\mathbf{x}; \mathbf{p} + \Delta\mathbf{p})) - T(\mathbf{x})]^2$$

Recap: Step-by-Step Derivation

- Key to the derivation
 - Taylor expansion around $\Delta \mathbf{p}$

$$I(\mathbf{W}(\mathbf{x}; \mathbf{p} + \Delta \mathbf{p})) \approx I(\mathbf{W}(\mathbf{x}; \mathbf{p})) + \nabla I \frac{\partial \mathbf{W}}{\partial \mathbf{p}} \Delta \mathbf{p} + \mathcal{O}(\Delta \mathbf{p}^2)$$

$$= I(\mathbf{W}([x, y]; p_1, \dots, p_n))$$

$$+ \begin{bmatrix} \frac{\partial I}{\partial x} & \frac{\partial I}{\partial y} \end{bmatrix} \begin{bmatrix} \frac{\partial W_x}{\partial p_1} & \frac{\partial W_x}{\partial p_2} & \cdots & \frac{\partial W_x}{\partial p_n} \\ \frac{\partial W_y}{\partial p_1} & \frac{\partial W_y}{\partial p_2} & \cdots & \frac{\partial W_y}{\partial p_n} \end{bmatrix} \begin{bmatrix} \Delta p_1 \\ \Delta p_2 \\ \vdots \\ \Delta p_n \end{bmatrix}$$

Gradient

Jacobian

Increment parameters to solve for $\Delta \mathbf{p}$

$$\nabla I$$

$$\frac{\partial \mathbf{W}}{\partial \mathbf{p}}$$



Recap: Inverse Compositional LK Algorithm

- Iterate

- Warp I to obtain $I(\mathbf{W}([x, y]; \mathbf{p}))$

- Compute the error image $T([x, y]) - I(\mathbf{W}([x, y]; \mathbf{p}))$

- Warp the gradient ∇I with $\mathbf{W}([x, y]; \mathbf{p})$

- Evaluate $\frac{\partial \mathbf{W}}{\partial \mathbf{p}}$ at $([x, y]; \mathbf{p})$ (Jacobian)

- Compute steepest descent images $\nabla I \frac{\partial \mathbf{W}}{\partial \mathbf{p}}$

- Compute Hessian matrix $\mathbf{H} = \sum_{\mathbf{x}} \left[\nabla I \frac{\partial \mathbf{W}}{\partial \mathbf{p}} \right]^T \left[\nabla I \frac{\partial \mathbf{W}}{\partial \mathbf{p}} \right]$

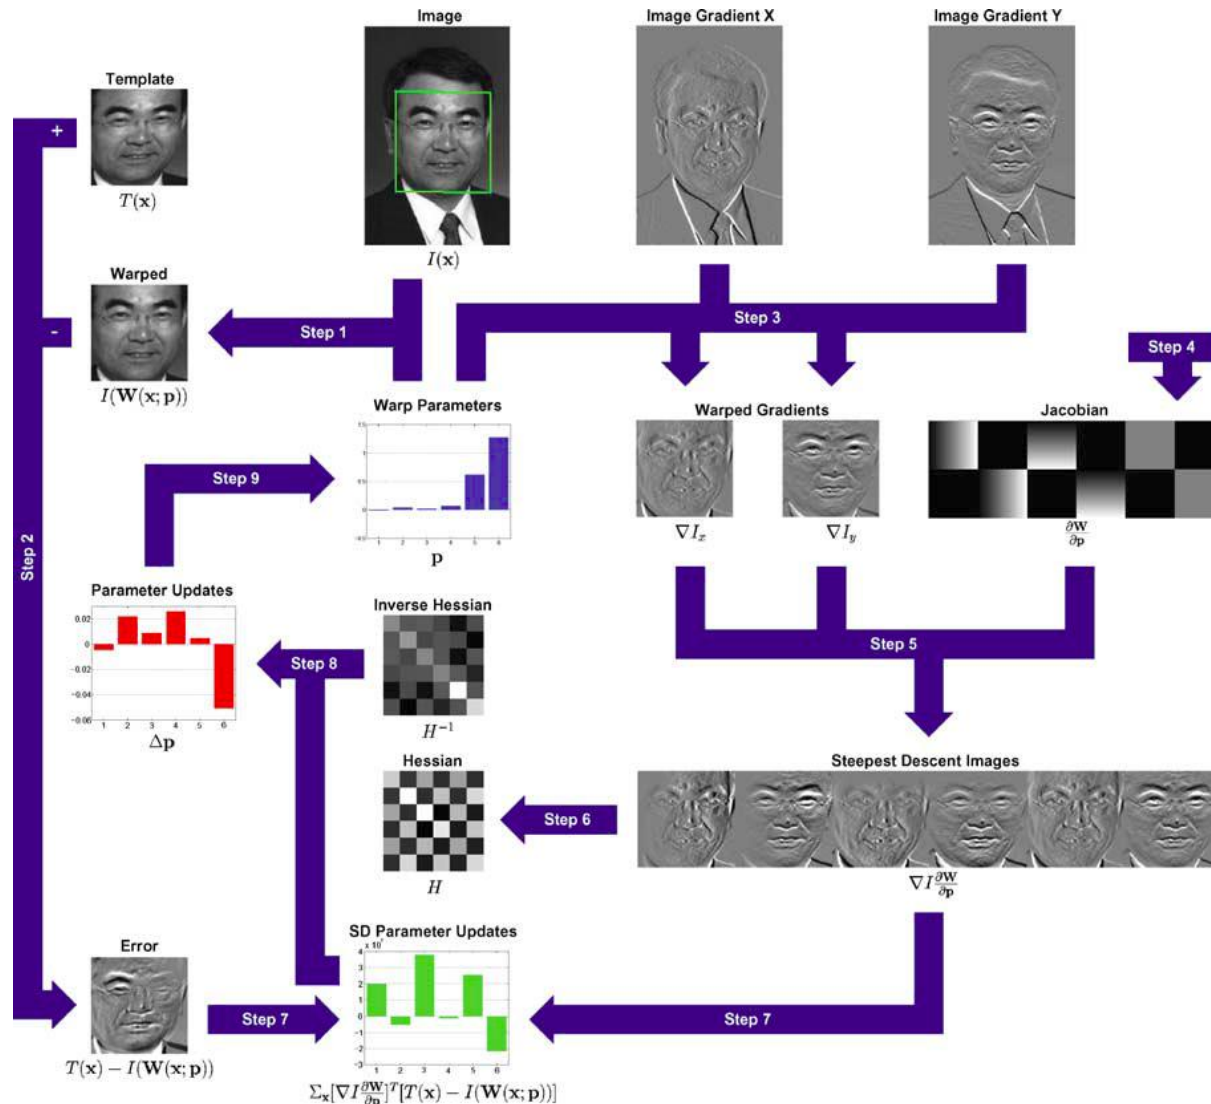
- Compute $\sum_{\mathbf{x}} \left[\nabla I \frac{\partial \mathbf{W}}{\partial \mathbf{p}} \right]^T [T([x, y]) - I(\mathbf{W}([x, y]; \mathbf{p}))]$

- Compute $\Delta \mathbf{p} = \mathbf{H}^{-1} \sum_{\mathbf{x}} \left[\nabla I \frac{\partial \mathbf{W}}{\partial \mathbf{p}} \right]^T [T([x, y]) - I(\mathbf{W}([x, y]; \mathbf{p}))]$

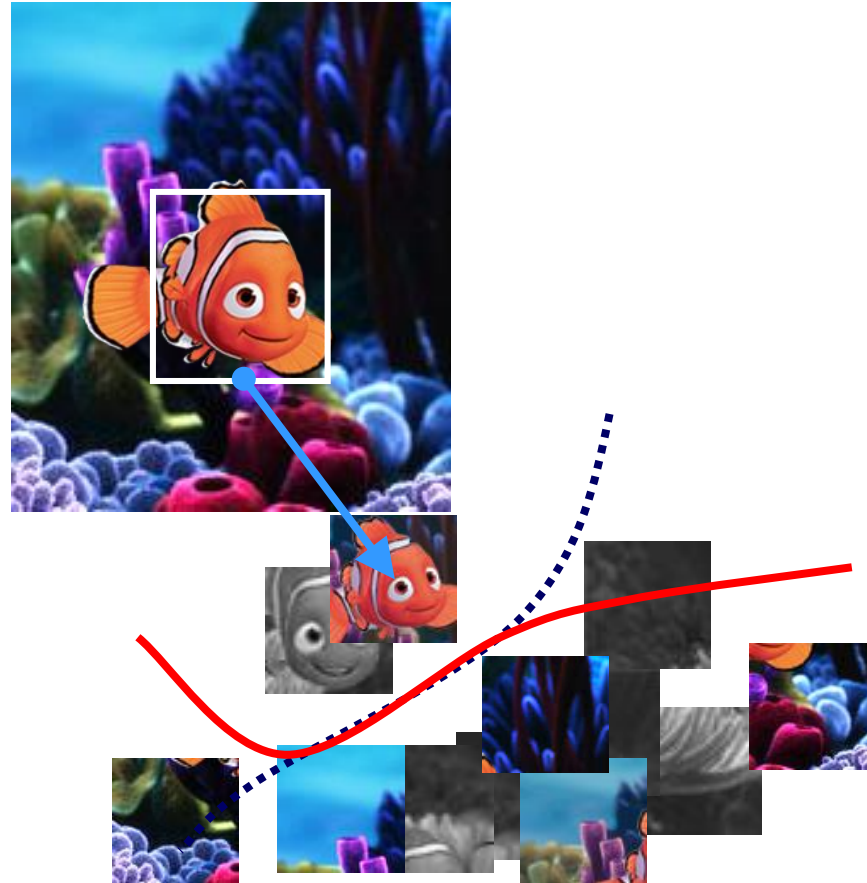
- Update the parameters $\mathbf{p} \leftarrow \mathbf{p} + \Delta \mathbf{p}$

- Until $\Delta \mathbf{p}$ magnitude is negligible

Recap: Inverse Compositional LK Algorithm



Today: Tracking by Online Classification



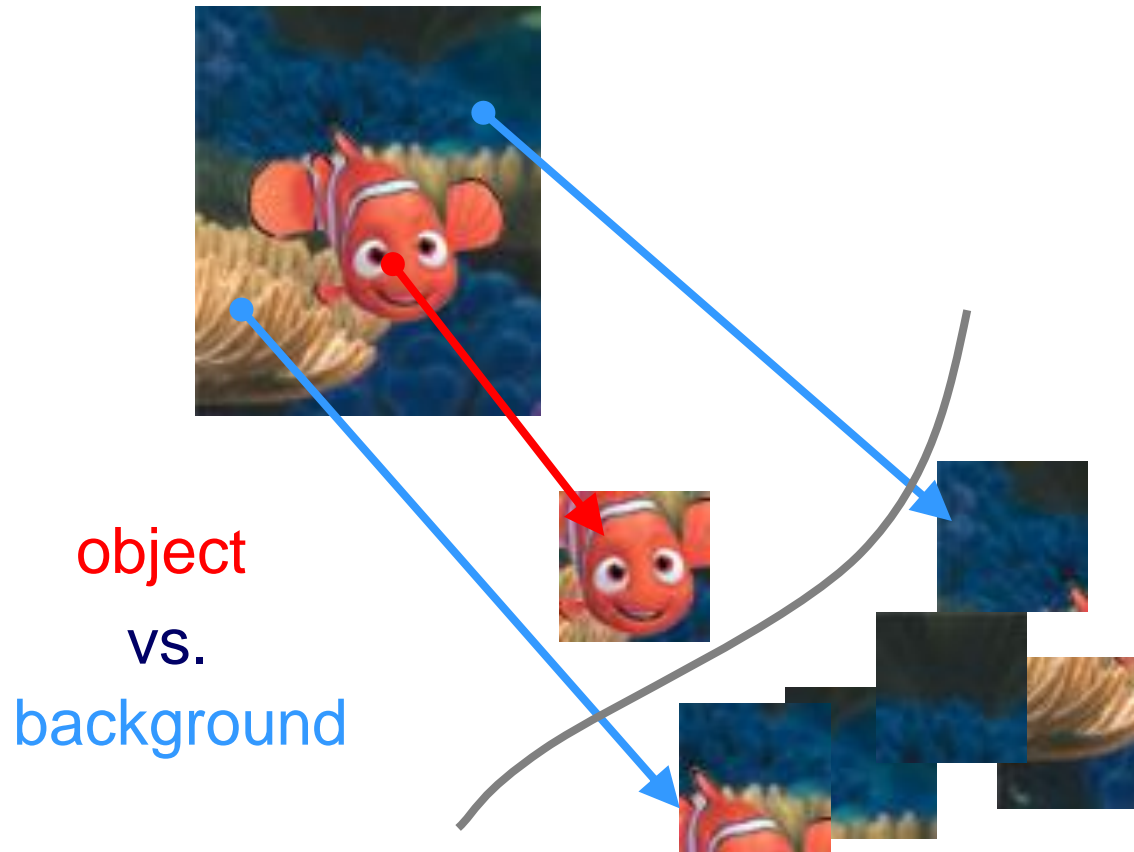
Can Machine Learning solve the problem for us?

Topics of This Lecture

- **Tracking by Online Classification**
 - Motivation
- **Recap: Boosting for Detection**
 - AdaBoost
 - Viola-Jones Detector
- **Extension to Online Classification**
 - Online Boosting
 - Online Feature Selection
 - Results
- **Extensions**
 - Problem: Drift
 - Drift-compensation strategies

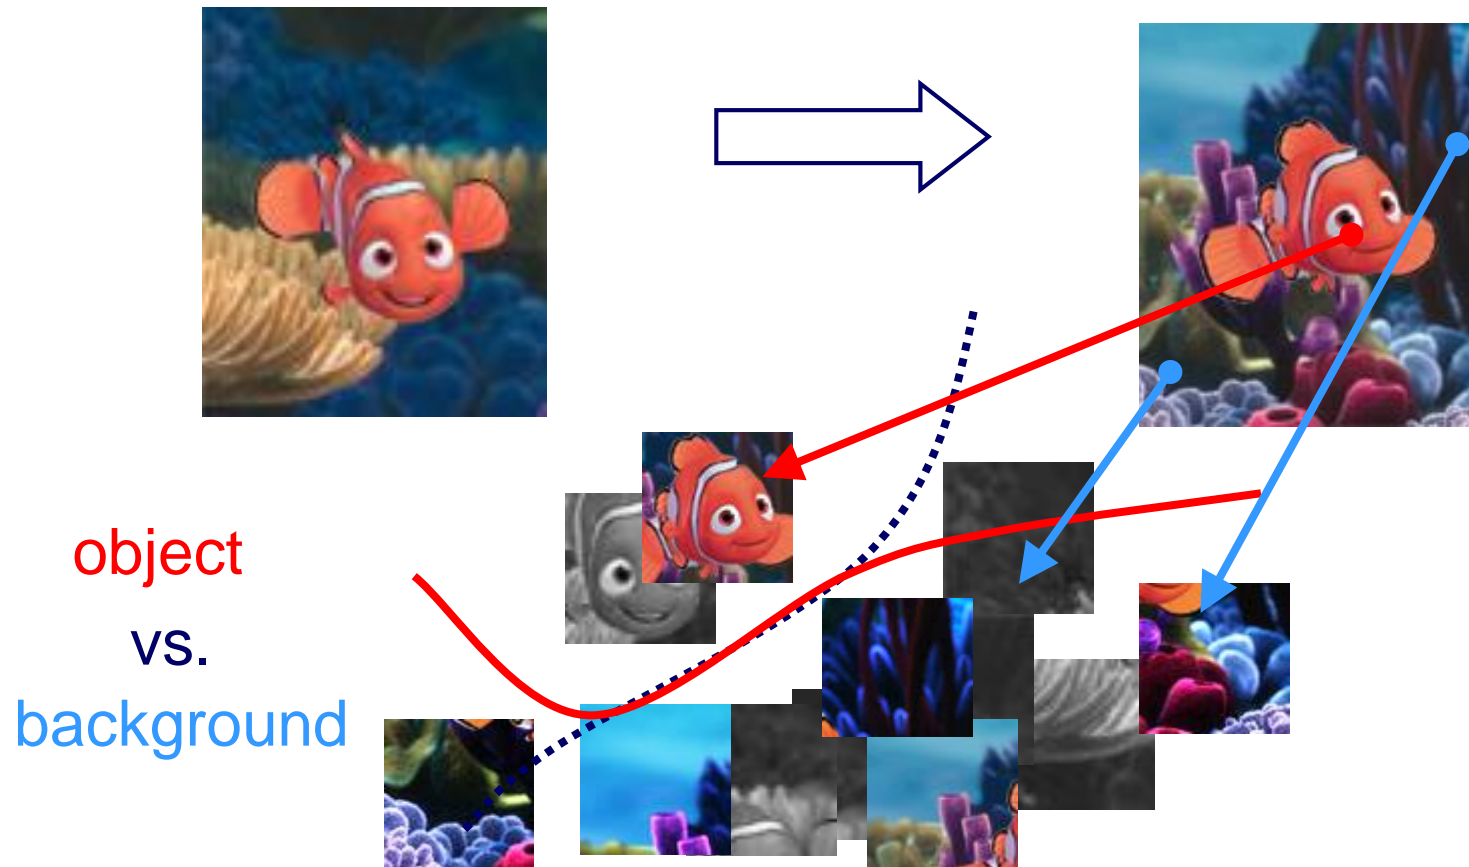
Tracking as Classification

- Tracking as binary classification problem



Tracking as Classification

- Tracking as binary classification problem

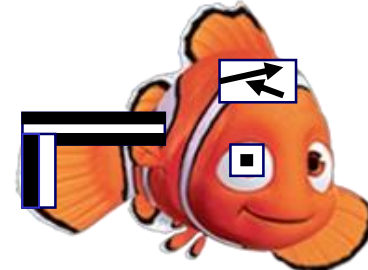


- Handle object and background changes by online updating

Idea: Use Boosting for Feature Selection

Object Detector

Fixed training set
General object detector



$$\text{sign}(\alpha_1 \cdot \text{feature}_1 + \alpha_2 \cdot \text{feature}_2 + \alpha_3 \cdot \text{feature}_3 + \dots)$$

Boosting for Feature Selection

P. Viola, M. Jones. [Rapid Object Detection using a Boosted Cascade of Simple Features](#). CVPR'01.



Object Tracker

On-line update
Object vs. Background

On-Line Boosting for Feature Selection

H. Grabner, H. Bischof. [On-line Boosting and Vision](#). CVPR'06.

Topics of This Lecture

- Tracking by Online Classification
 - Motivation
- **Recap: Boosting for Detection**
 - AdaBoost
 - Viola-Jones Detector
- Extension to Online Classification
 - Online Boosting
 - Online Feature Selection
 - Results
- Extensions
 - Problem: Drift
 - Drift-compensation strategies

Recap: AdaBoost – “Adaptive Boosting”

- Main idea [Freund & Schapire, 1996]
 - Iteratively select an ensemble of classifiers
 - Reweight misclassified training examples after each iteration to focus training on difficult cases.
- Components
 - $h_m(\mathbf{x})$: “weak” or base classifier
 - Condition: <50% training error over any distribution
 - $H(\mathbf{x})$: “strong” or final classifier
- **AdaBoost:**
 - Construct a strong classifier as a thresholded linear combination of the weighted weak classifiers:

$$H(\mathbf{x}) = \text{sign} \left(\sum_{m=1}^M \alpha_m h_m(\mathbf{x}) \right)$$

Recap: AdaBoost – Algorithm

1. Initialization: Set $w_n^{(1)} = \frac{1}{N}$ for $n = 1, \dots, N$.

2. For $m = 1, \dots, M$ iterations

a) Train a new weak classifier $h_m(\mathbf{x})$ using the current weighting coefficients $\mathbf{W}^{(m)}$ by minimizing the weighted error function

$$J_m = \sum_{n=1}^N w_n^{(m)} I(h_m(\mathbf{x}) \neq t_n)$$

$$I(A) = \begin{cases} 1, & \text{if } A \text{ is true} \\ 0, & \text{else} \end{cases}$$

b) Estimate the weighted error of this classifier on \mathbf{X} :

$$\epsilon_m = \frac{\sum_{n=1}^N w_n^{(m)} I(h_m(\mathbf{x}) \neq t_n)}{\sum_{n=1}^N w_n^{(m)}}$$

c) Calculate a weighting coefficient for $h_m(\mathbf{x})$:

$$\alpha_m = \ln \left\{ \frac{1 - \epsilon_m}{\epsilon_m} \right\}$$

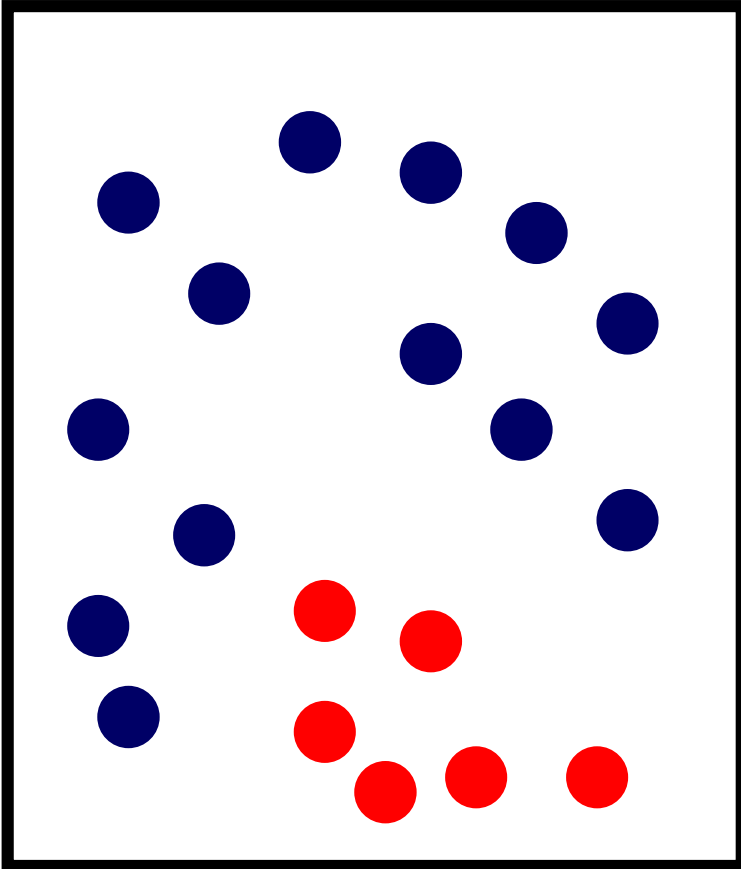
d) Update the weighting coefficients:

$$w_n^{(m+1)} = w_n^{(m)} \exp \{ \alpha_m I(h_m(\mathbf{x}_n) \neq t_n) \}$$

Topics of This Lecture

- **Tracking by Online Classification**
 - Motivation
- **Recap: Boosting for Detection**
 - AdaBoost
 - Viola-Jones Detector
- **Extension to Online Classification**
 - Online Boosting
 - Online Feature Selection
 - Results
- **Extensions**
 - Problem: Drift
 - Drift-compensation strategies

Offline Boosting



Given:

- set of labeled training samples
- weight distribution over them

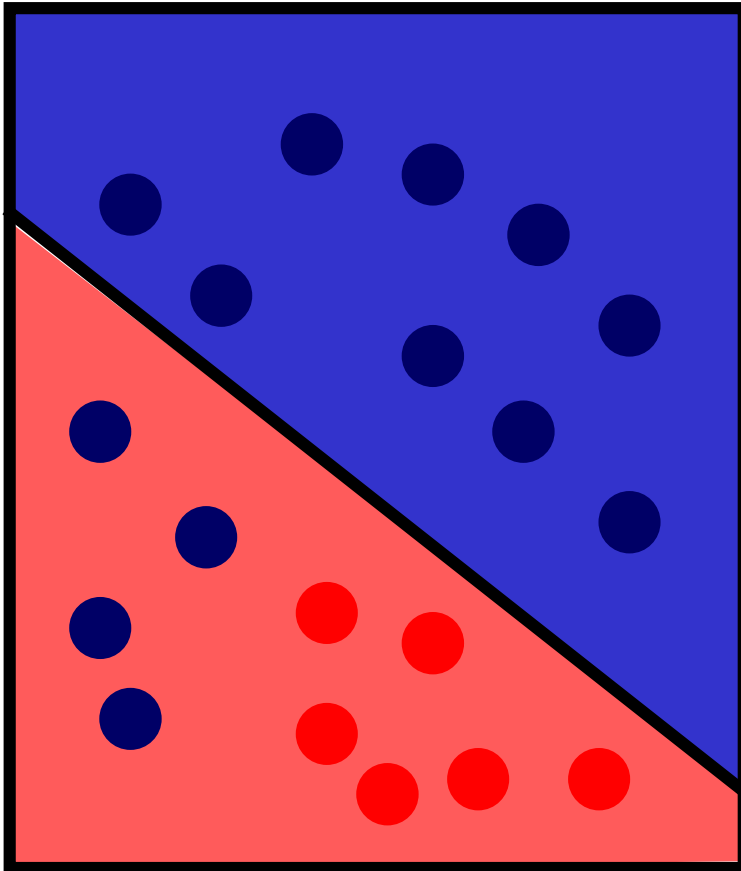
Algorithm:

for n = 1 to N

- train a weak classifier using samples and weight dist.
- calculate error
- calculate weight
- update weight dist.

next

Offline Boosting



Given:

- set of labeled training samples
- weight distribution over them

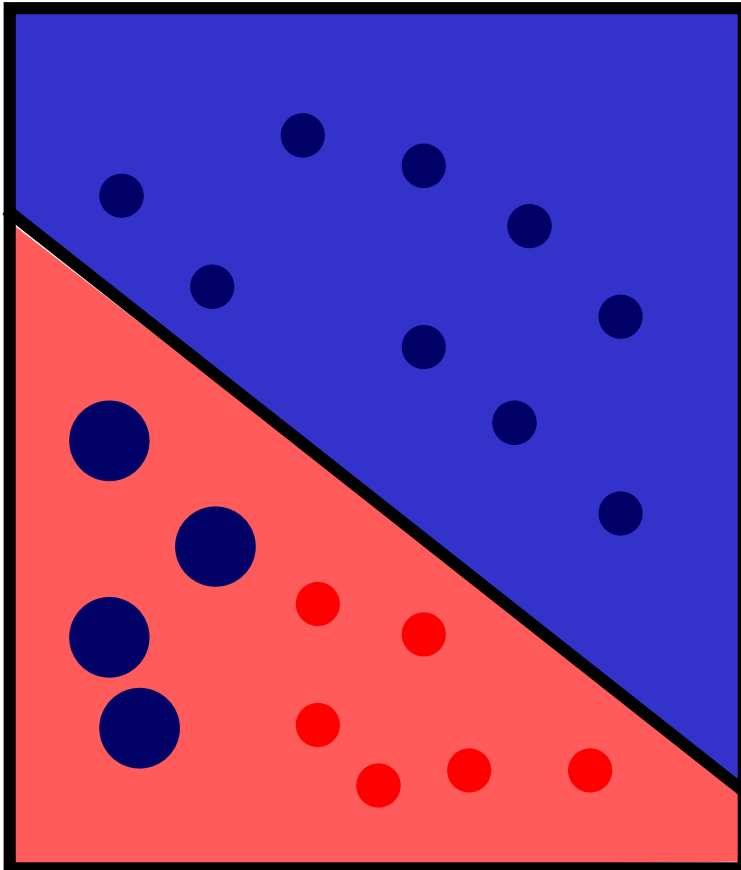
Algorithm:

for n = 1 to N

- train a weak classifier using samples and weight dist.
- calculate error
- calculate weight
- update weight dist.

next

Offline Boosting



$\alpha_1 \cdot$



Given:

- set of labeled training samples
- weight distribution over them

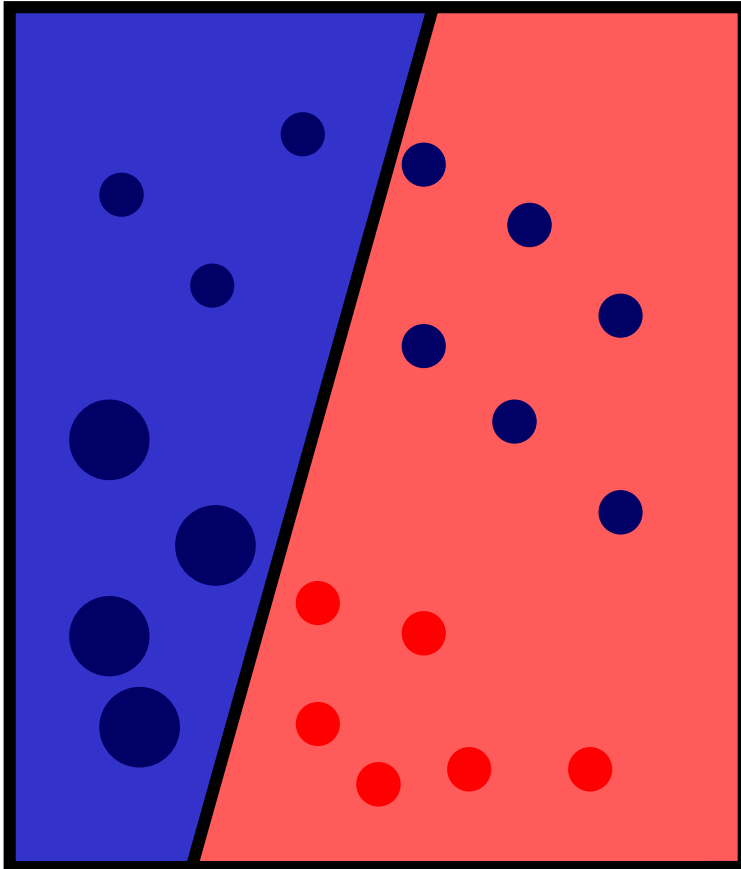
Algorithm:

for n = 1 to N

- train a weak classifier using samples and weight dist.
- calculate error
- calculate weight
- update weight dist.

next

Offline Boosting



Given:

- set of labeled training samples
- weight distribution over them

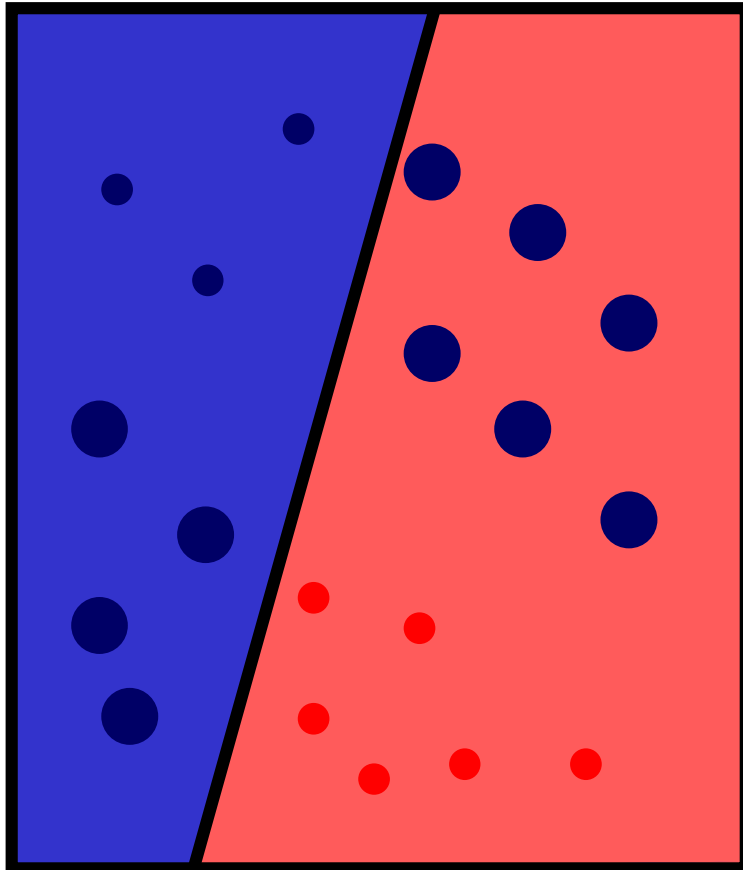
Algorithm:


for n = 1 to N

- train a weak classifier using samples and weight dist.
- calculate error
- calculate weight
- update weight dist.

next

Offline Boosting



$$\alpha_2 \cdot$$


Given:

- set of labeled training samples
- weight distribution over them

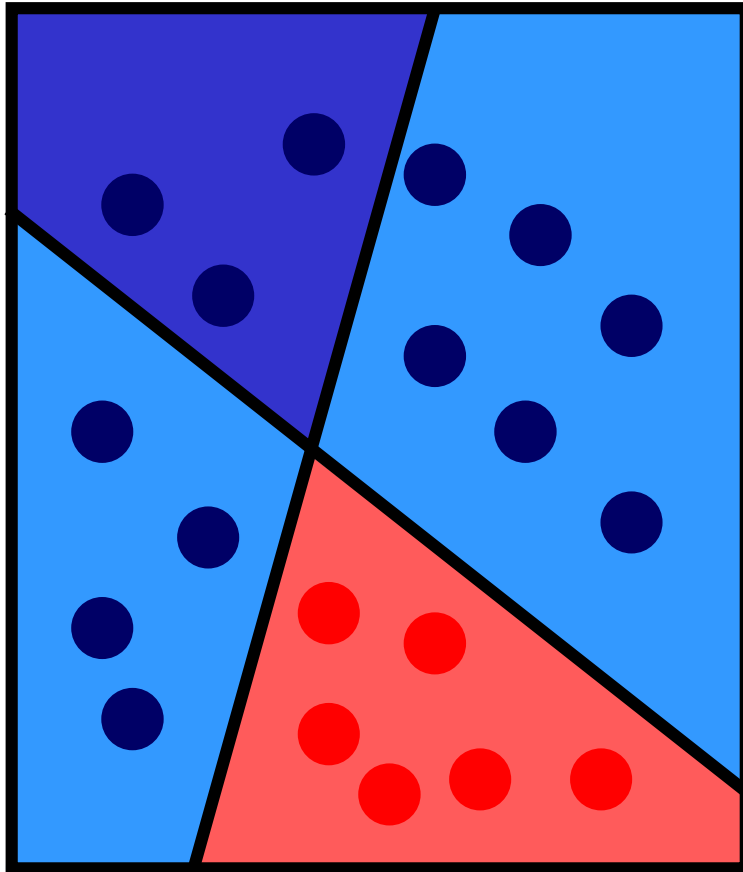
Algorithm:

for $n = 1$ to N

- train a weak classifier using samples and weight dist.
- calculate error
- calculate weight
- update weight dist.

next

Offline Boosting



$$= \alpha_1 \cdot \begin{array}{|c|} \hline \text{dark blue} \\ \hline \text{red} \\ \hline \end{array} + \alpha_2 \cdot \begin{array}{|c|} \hline \text{dark blue} \\ \hline \text{red} \\ \hline \end{array}$$

Given:

- set of labeled training samples
- weight distribution over them

Algorithm:

for $n = 1$ to N

- train a weak classifier using samples and weight dist.
- calculate error
- calculate weight
- update weight dist.

next

Result:

$$h^{strong}(\mathbf{x}) = \text{sign}\left(\sum_{n=1}^N \alpha_n \cdot h_n^{weak}(\mathbf{x})\right)$$

From Offline to Online Boosting

- Goal
 - Formulate the algorithm such that we can present only 1 training sample at a time (and then forget about it).
 - ⇒ Dual problem: instead of keeping all samples and adding weak classifiers, keep a fixed set of weak classifiers and add samples.
- What changes?
 - Updating the classifiers online can be done easily.
 - Many classification approaches can use online updates.
 - Computing the classifier weights is also straightforward if we know the estimated error (which we can compute).

From Offline to Online Boosting

- Main issue
 - Computing the weight distribution for the samples.
 - We do not know a priori the difficulty of a sample!
(Could already have seen the same sample before...)
- Idea of Online Boosting
 - Estimate the importance of a sample by propagating it through a set of weak classifiers.
 - This can be thought of as modeling the information gain w.r.t. the first n classifiers and code it by the importance weight λ for the $n+1$ classifier.
 - Proven [Oza]: Given the same training set, Online Boosting converges to the same weak classifiers as Offline Boosting in the limit of $N \rightarrow \infty$ iterations.

N. Oza and S. Russell. [Online Bagging and Boosting](#).
Artificial Intelligence and Statistics, 2001.

From Offline to Online Boosting

off-line

Given:

- set of labeled training samples

$$\mathcal{X} = \{\langle \mathbf{x}_1, y_1 \rangle, \dots, \langle \mathbf{x}_L, y_L \rangle \mid y_i \pm 1\}$$

- weight distribution over them

$$D_0 = 1/L$$

for $n = 1$ to N

- train a weak classifier using samples and weight dist.

$$h_n^{weak}(\mathbf{x}) = \mathcal{L}(\mathcal{X}, D_{n-1})$$

- calculate error e_n
- calculate weight $\alpha_n = f(e_n)$
- update weight dist. D_n

next

$$h^{strong}(\mathbf{x}) = \text{sign}\left(\sum_{n=1}^N \alpha_n \cdot h_n^{weak}(\mathbf{x})\right)$$

on-line

From Offline to Online Boosting

off-line

Given:

- set of labeled training samples

$$\mathcal{X} = \{\langle \mathbf{x}_1, y_1 \rangle, \dots, \langle \mathbf{x}_L, y_L \rangle \mid y_i \pm 1\}$$

- weight distribution over them

$$D_0 = 1/L$$

for $n = 1$ to N

- train a weak classifier using samples and weight dist.

$$h_n^{weak}(\mathbf{x}) = \mathcal{L}(\mathcal{X}, D_{n-1})$$

- calculate error e_n
- calculate weight $\alpha_n = f(e_n)$
- update weight dist. D_n

next

$$h^{strong}(\mathbf{x}) = \text{sign}\left(\sum_{n=1}^N \alpha_n \cdot h_n^{weak}(\mathbf{x})\right)$$

on-line

Given:

for $n = 1$ to N

next

$$h^{strong}(\mathbf{x}) = \text{sign}\left(\sum_{n=1}^N \alpha_n \cdot h_n^{weak}(\mathbf{x})\right)$$

From Offline to Online Boosting

off-line

Only **one** training example
to **update** the classifier

on-line

Given:

- set of labeled training samples
 $\mathcal{X} = \{\langle \mathbf{x}_1, y_1 \rangle, \dots, \langle \mathbf{x}_L, y_L \rangle \mid y_i \pm 1\}$
- weight distribution over them
 $D_0 = 1/L$

for $n = 1$ to N

- train a weak classifier using samples and weight dist.

$$h_n^{weak}(\mathbf{x}) = \mathcal{L}(\mathcal{X}, D_{n-1})$$

- calculate error e_n
- calculate weight $\alpha_n = f(e_n)$
- update weight dist. D_n

next

$$h^{strong}(\mathbf{x}) = \text{sign}\left(\sum_{n=1}^N \alpha_n \cdot h_n^{weak}(\mathbf{x})\right)$$

Given:

- ONE labeled training sample
 $\langle \mathbf{x}, y \rangle \mid y \pm 1$
- strong classifier to update

for $n = 1$ to N

next

$$h^{strong}(\mathbf{x}) = \text{sign}\left(\sum_{n=1}^N \alpha_n \cdot h_n^{weak}(\mathbf{x})\right)$$

From Offline to Online Boosting

off-line

Update importance for
the current sample

on-line

Given:

- set of labeled training samples
 $\mathcal{X} = \{\langle \mathbf{x}_1, y_1 \rangle, \dots, \langle \mathbf{x}_L, y_L \rangle \mid y_i \pm 1\}$
- weight distribution over them
 $D_0 = 1/L$

for $n = 1$ to N

- train a weak classifier using samples and weight dist.

$$h_n^{weak}(\mathbf{x}) = \mathcal{L}(\mathcal{X}, D_{n-1})$$

- calculate error e_n
- calculate weight $\alpha_n = f(e_n)$
- update weight dist. D_n

next

$$h^{strong}(\mathbf{x}) = \text{sign}\left(\sum_{n=1}^N \alpha_n \cdot h_n^{weak}(\mathbf{x})\right)$$

Given:

- ONE labeled training sample
 $\langle \mathbf{x}, y \rangle \mid y \pm 1$
- strong classifier to update

- initial importance $\lambda = 1$

for $n = 1$ to N

- update importance weight λ

next

$$h^{strong}(\mathbf{x}) = \text{sign}\left(\sum_{n=1}^N \alpha_n \cdot h_n^{weak}(\mathbf{x})\right)$$

From Offline to Online Boosting

off-line

Online update the weak classifier

on-line

Given:

- set of labeled training samples
 $\mathcal{X} = \{\langle \mathbf{x}_1, y_1 \rangle, \dots, \langle \mathbf{x}_L, y_L \rangle \mid y_i \pm 1\}$
- weight distribution over them
 $D_0 = 1/L$

for $n = 1$ to N

- train a weak classifier using samples and weight dist.

$$h_n^{weak}(\mathbf{x}) = \mathcal{L}(\mathcal{X}, D_{n-1})$$

- calculate error e_n
- calculate weight $\alpha_n = f(e_n)$
- update weight dist. D_n

next

$$h^{strong}(\mathbf{x}) = \text{sign}\left(\sum_{n=1}^N \alpha_n \cdot h_n^{weak}(\mathbf{x})\right)$$

Given:

- ONE labeled training sample
 $\langle \mathbf{x}, y \rangle \mid y \pm 1$
- strong classifier to update

- initial importance $\lambda = 1$

for $n = 1$ to N

- update the weak classifier using samples and importance

$$h_n^{weak}(\mathbf{x}) = \mathcal{L}(h_n^{weak}, \langle \mathbf{x}, y \rangle, \lambda)$$

- update importance weight λ

next

$$h^{strong}(\mathbf{x}) = \text{sign}\left(\sum_{n=1}^N \alpha_n \cdot h_n^{weak}(\mathbf{x})\right)$$

From Offline to Online Boosting

off-line

Update errors and weights

on-line

Given:

- set of labeled training samples
 $\mathcal{X} = \{\langle \mathbf{x}_1, y_1 \rangle, \dots, \langle \mathbf{x}_L, y_L \rangle \mid y_i \pm 1\}$
- weight distribution over them
 $D_0 = 1/L$

for $n = 1$ to N

- train a weak classifier using samples and weight dist.

$$h_n^{weak}(\mathbf{x}) = \mathcal{L}(\mathcal{X}, D_{n-1})$$

- calculate error e_n
- calculate weight $\alpha_n = f(e_n)$
- update weight dist. D_n

next

$$h^{strong}(\mathbf{x}) = \text{sign}\left(\sum_{n=1}^N \alpha_n \cdot h_n^{weak}(\mathbf{x})\right)$$

Given:

- ONE labeled training sample
 $\langle \mathbf{x}, y \rangle \mid y \pm 1$
- strong classifier to update

- initial importance $\lambda = 1$

for $n = 1$ to N

- update the weak classifier using samples and importance

$$h_n^{weak}(\mathbf{x}) = \mathcal{L}(h_n^{weak}, \langle \mathbf{x}, y \rangle, \lambda)$$

- update error estimation \hat{e}_n
- update weight $\alpha_n = f(\hat{e}_n)$
- update importance weight λ

next

$$h^{strong}(\mathbf{x}) = \text{sign}\left(\sum_{n=1}^N \alpha_n \cdot h_n^{weak}(\mathbf{x})\right)$$

From Offline to Online Boosting

off-line

Given:

- set of labeled training samples
 $\mathcal{X} = \{\langle \mathbf{x}_1, y_1 \rangle, \dots, \langle \mathbf{x}_L, y_L \rangle \mid y_i \pm 1\}$
- weight distribution over them
 $D_0 = 1/L$

for $n = 1$ to N

- train a weak classifier using samples and weight dist.

$$h_n^{weak}(\mathbf{x}) = \mathcal{L}(\mathcal{X}, D_{n-1})$$

- calculate error e_n
- calculate weight $\alpha_n = f(e_n)$
- update weight dist. D_n

next

$$h^{strong}(\mathbf{x}) = \text{sign}\left(\sum_{n=1}^N \alpha_n \cdot h_n^{weak}(\mathbf{x})\right)$$

on-line

Given:

- ONE labeled training sample
 $\langle \mathbf{x}, y \rangle \mid y \pm 1$
- strong classifier to update

- initial importance $\lambda = 1$

for $n = 1$ to N

- update the weak classifier using samples and importance

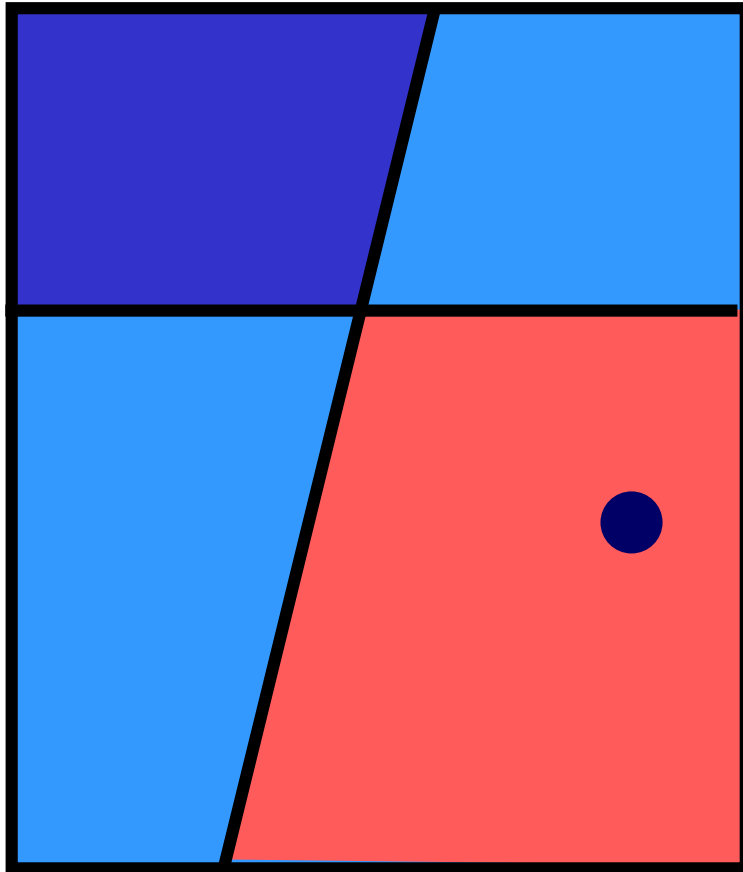
$$h_n^{weak}(\mathbf{x}) = \mathcal{L}(h_n^{weak}, \langle \mathbf{x}, y \rangle, \lambda)$$

- update error estimation \hat{e}_n
- update weight $\alpha_n = f(\hat{e}_n)$
- update importance weight λ

next

$$h^{strong}(\mathbf{x}) = \text{sign}\left(\sum_{n=1}^N \alpha_n \cdot h_n^{weak}(\mathbf{x})\right)$$

Online Boosting



$$\alpha_1 \cdot \begin{array}{|c|} \hline \text{dark blue} \\ \hline \text{red} \\ \hline \end{array} + \alpha_2 \cdot \begin{array}{|c|} \hline \text{light blue} \\ \hline \text{red} \\ \hline \end{array}$$

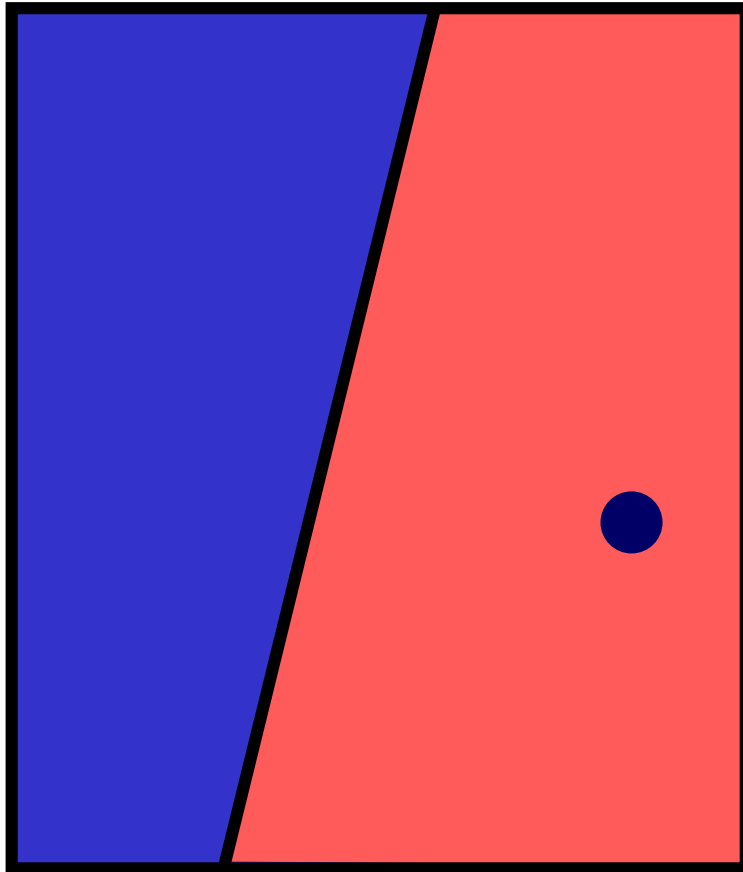
Given:

- ONE labeled training sample
- strong classifier to update

Algorithm:

- initial importance
- for n = 1 to N
- update the weak classifier using sample and importance
 - update error estimation
 - update weight
 - update importance weight
- next

Online Boosting



$$\alpha_1 \cdot \begin{array}{|c|} \hline \text{blue} \\ \hline \end{array} + \alpha_2 \cdot \begin{array}{|c|} \hline \text{red} \\ \hline \end{array}$$

Given:

- ONE labeled training sample
- strong classifier to update

Algorithm:

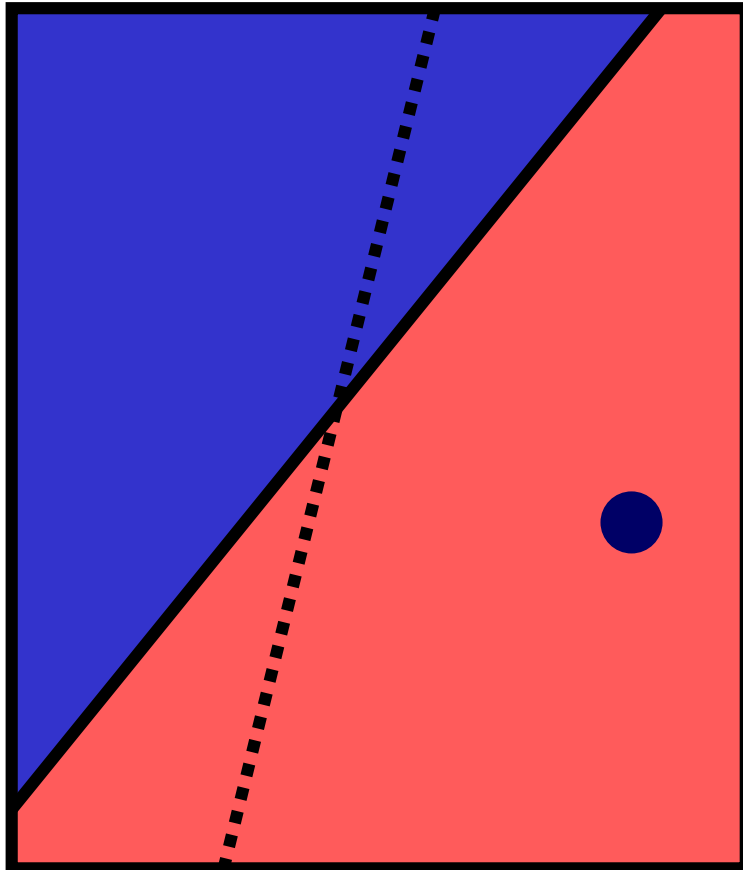
- initial importance

for n = 1 to N

- update the weak classifier using sample and importance
- update error estimation
- update weight
- update importance weight

next

Online Boosting



Given:

- ONE labeled training sample
- strong classifier to update

Algorithm:

- initial importance

for $n = 1$ to N

- update the weak classifier using sample and importance
- update error estimation
- update weight
- update importance weight

next

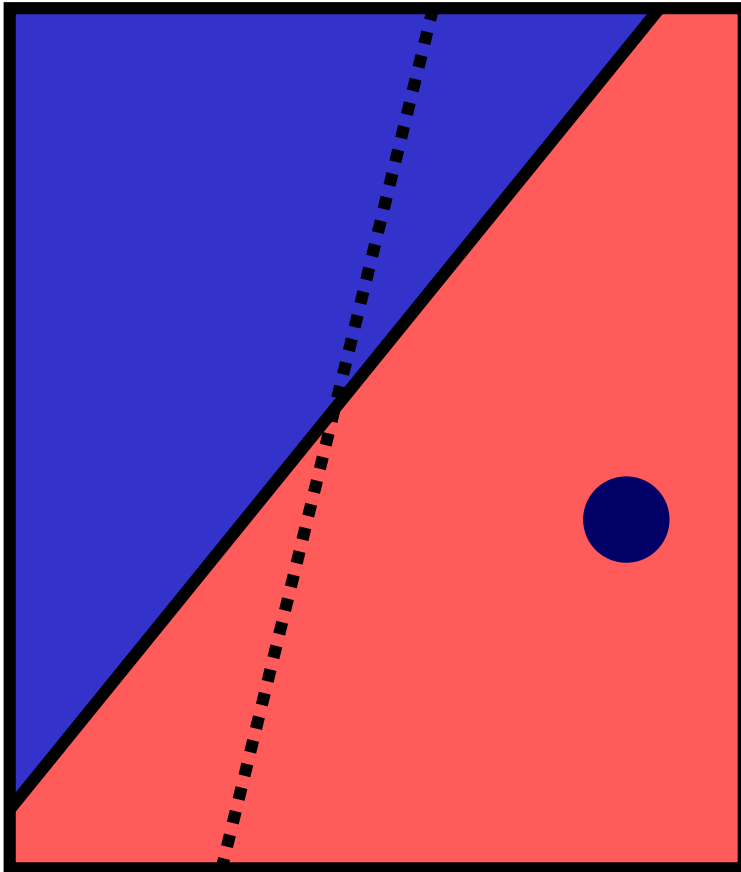
$\alpha_1 \cdot$



$+ \alpha_2 \cdot$



Online Boosting



Given:

- ONE labeled training sample
- strong classifier to update

Algorithm:

- initial importance

for n = 1 to N

- update the weak classifier using sample and importance

- update error estimation

- update weight

- update importance weight

next

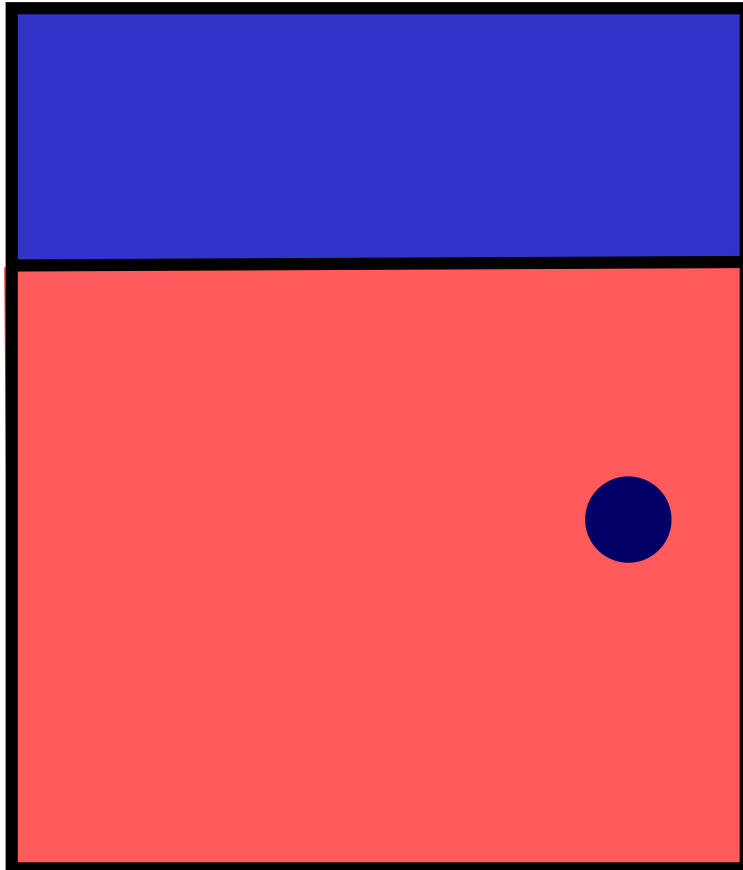
$\alpha_1 \cdot$



$+ \alpha_2 \cdot$



Online Boosting



$$\alpha_1 \cdot \begin{array}{|c|} \hline \text{blue} \\ \hline \text{red} \\ \hline \end{array} + \alpha_2 \cdot \begin{array}{|c|} \hline \text{blue} \\ \hline \text{red} \\ \hline \end{array}$$

Given:

- ONE labeled training sample
- strong classifier to update

Algorithm:

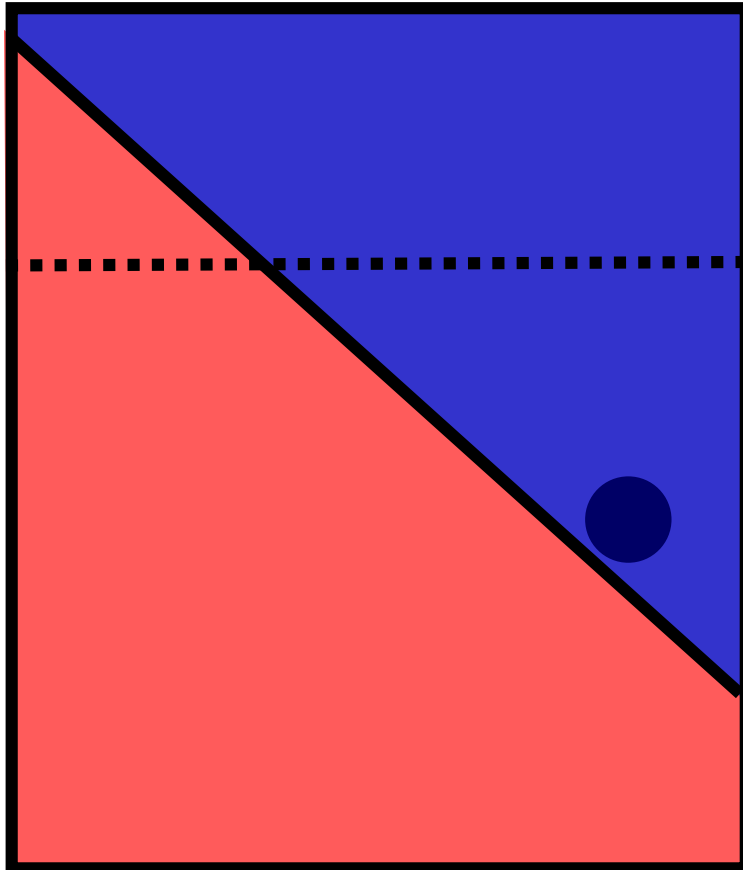
- initial importance

for n = 1 to N

- update the weak classifier using sample and importance
- update error estimation
- update weight
- update importance weight

next

Online Boosting



$$\alpha_1 \cdot \begin{array}{|c|} \hline \text{blue} \\ \hline \text{red} \\ \hline \end{array} + \alpha_2 \cdot \begin{array}{|c|} \hline \text{blue} \\ \hline \text{red} \\ \hline \end{array}$$

Given:

- ONE labeled training sample
- strong classifier to update

Algorithm:

- initial importance

for $n = 1$ to N

- update the weak classifier using sample and importance
- update error estimation
- update weight
- update importance weight

next

Online Boosting



Given:

- ONE labeled training sample
- strong classifier to update

Algorithm:

Converges to the off-line results...

N. Oza and S. Russell. Online Bagging and Boosting.
Artificial Intelligence and Statistics, 2001.

er using

- update importance weight

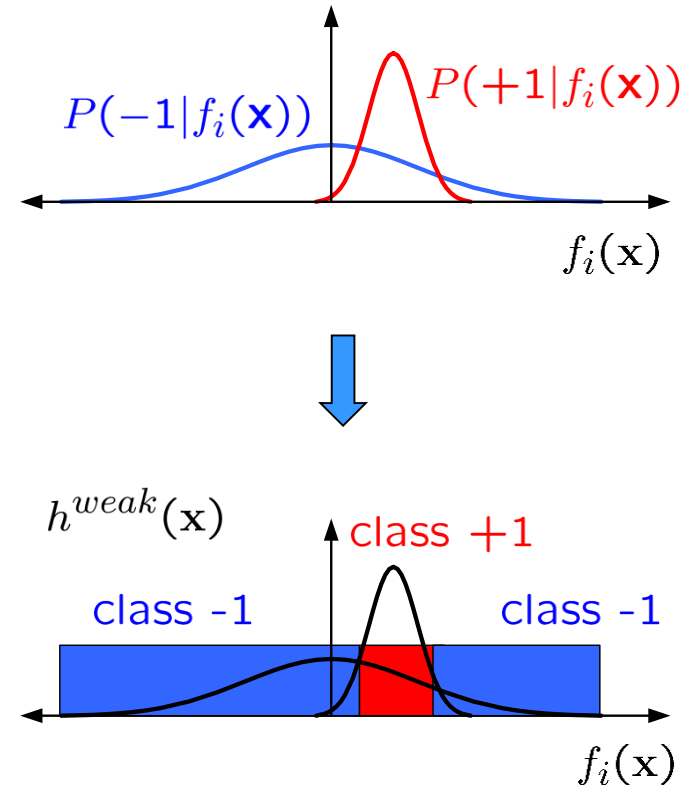
next

$$= \alpha_1 \cdot \begin{array}{|c|} \hline \text{blue} \\ \hline \text{red} \\ \hline \end{array} + \alpha_2 \cdot \begin{array}{|c|} \hline \text{blue} \\ \hline \text{red} \\ \hline \end{array}$$

$$h^{strong}(\mathbf{x}) = \text{sign}\left(\sum_{n=1}^N \alpha_n \cdot h_n^{weak}(\mathbf{x})\right)$$

Online Boosting for Feature Selection

- Each feature corresponds to a weak classifier.
- Features
 - Haar-like wavelets
 - Orientation histograms
 - Locally binary patterns (LBP)
- Fast computation using efficient data structures
 - integral images
 - integral histograms



F. Porikli. [Integral histogram: A fast way to extract histograms in cartesian spaces.](#) CVPR'05.

Online Boosting for Feature Selection

- Introducing “Selector”
 - Selects **one** feature from its local feature pool

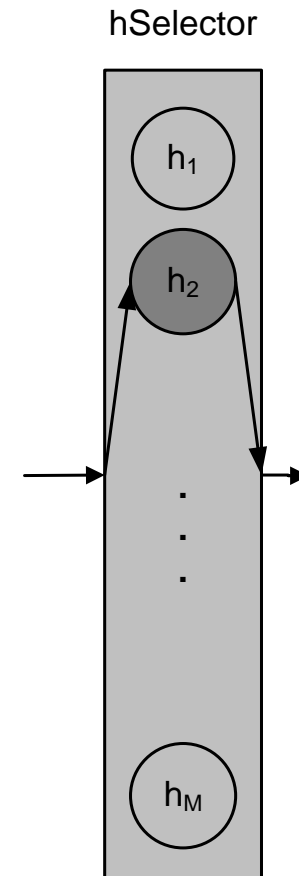
$$\mathcal{H}^{weak} = \{h_1^{weak}, \dots, h_M^{weak}\}$$

$$\mathcal{F} = \{f_1, \dots, f_M\}$$

$$h^{sel}(\mathbf{x}) = h_m^{weak}(\mathbf{x})$$

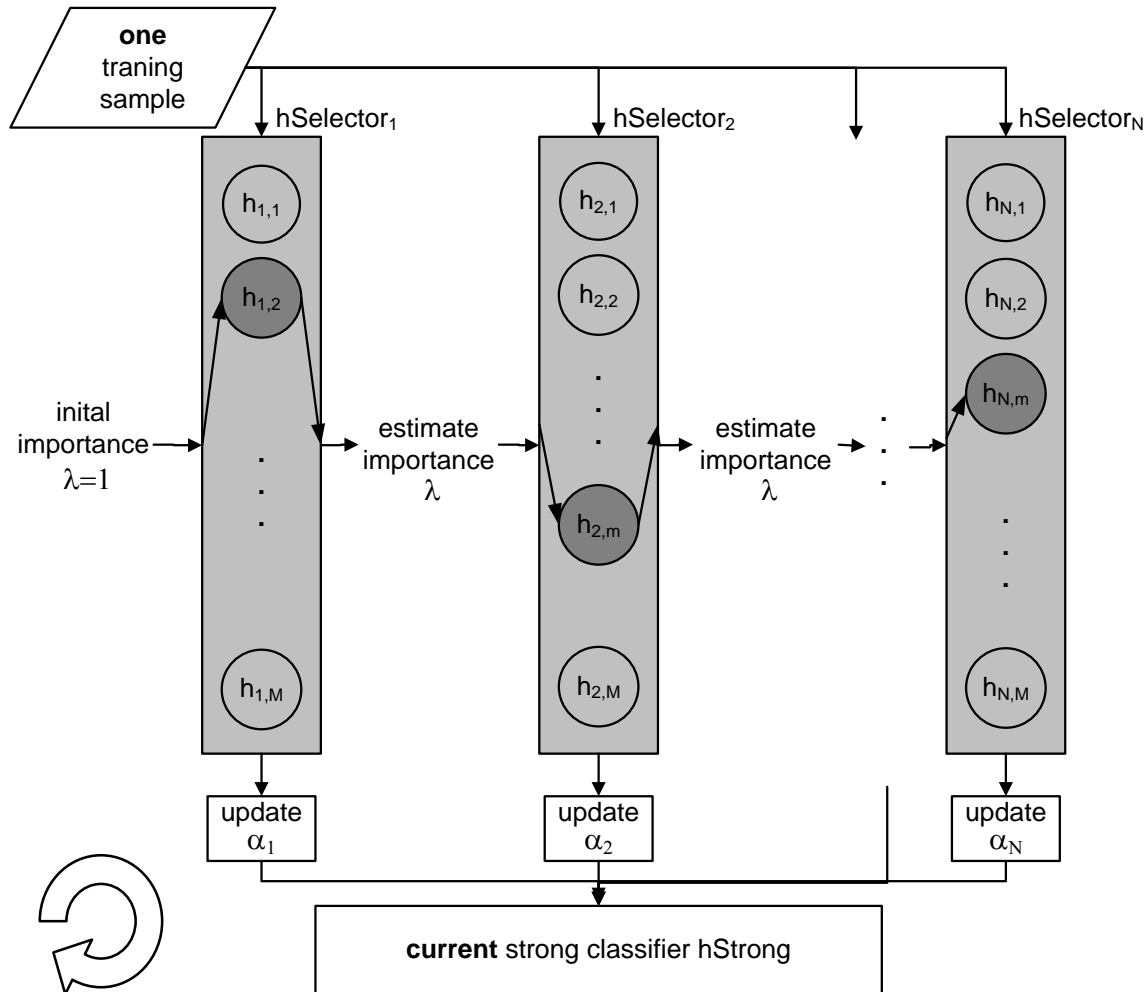
$$m = \arg \min_i e_i$$

On-line boosting is performed on the **Selectors** and not on the weak classifiers directly.



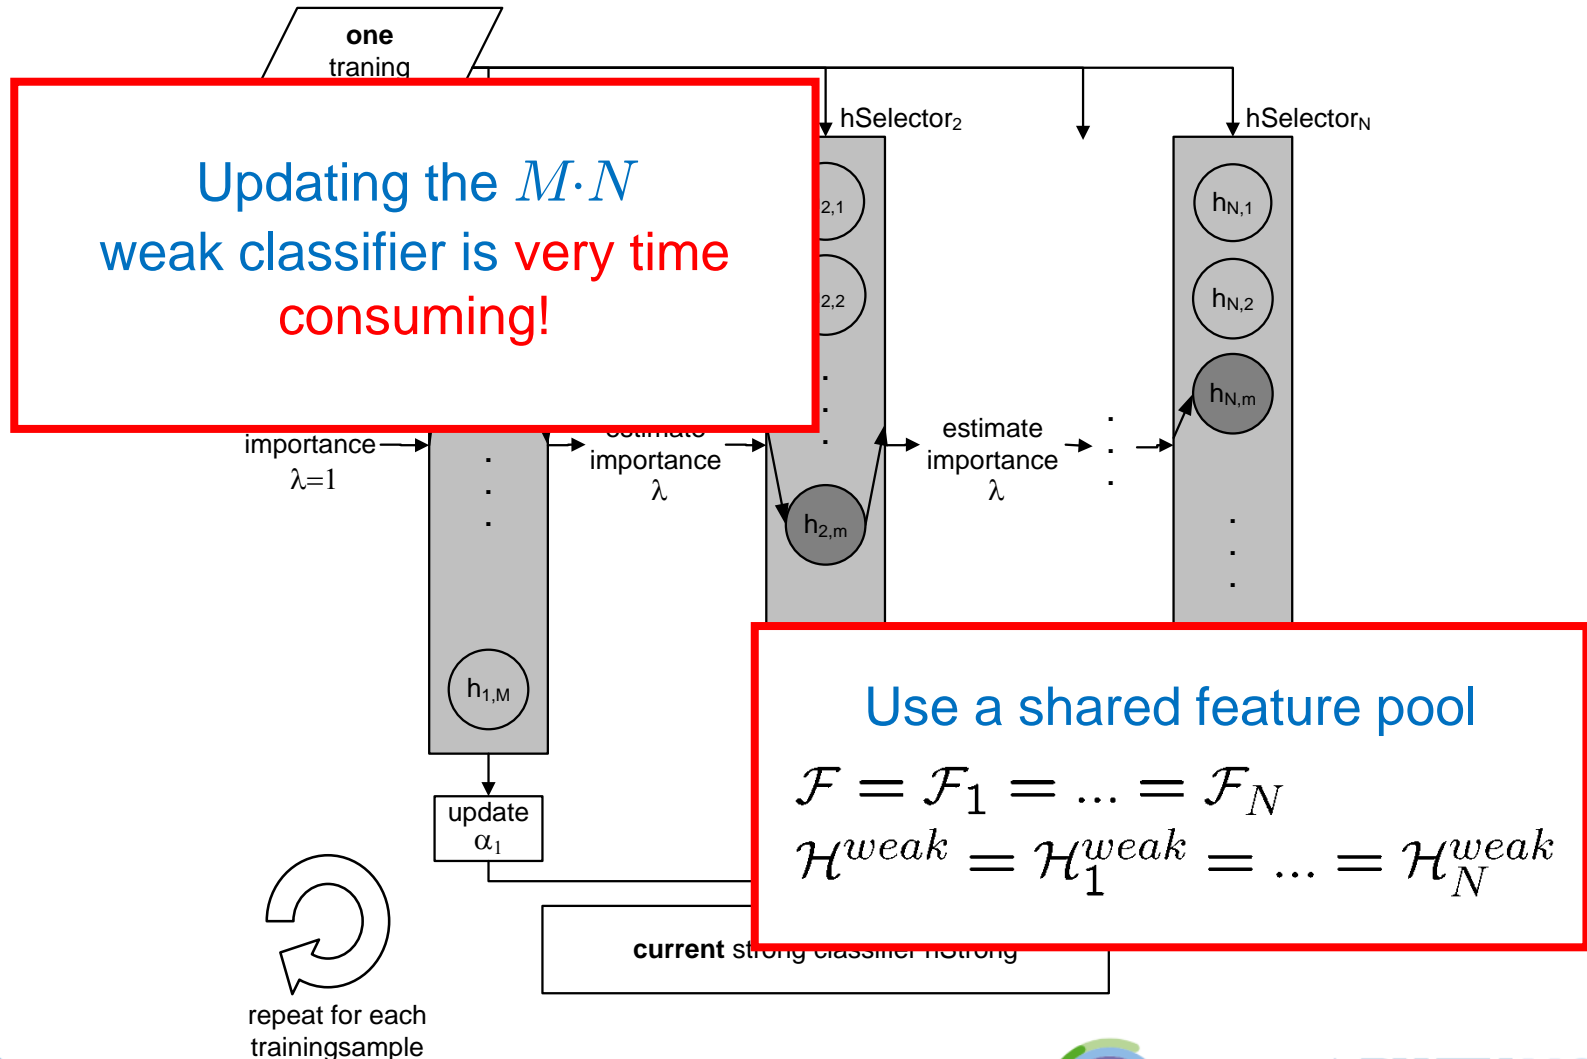
H. Grabner and H. Bischof.
On-line boosting and vision.
CVPR, 2006.

Online Boosting for Feature Selection

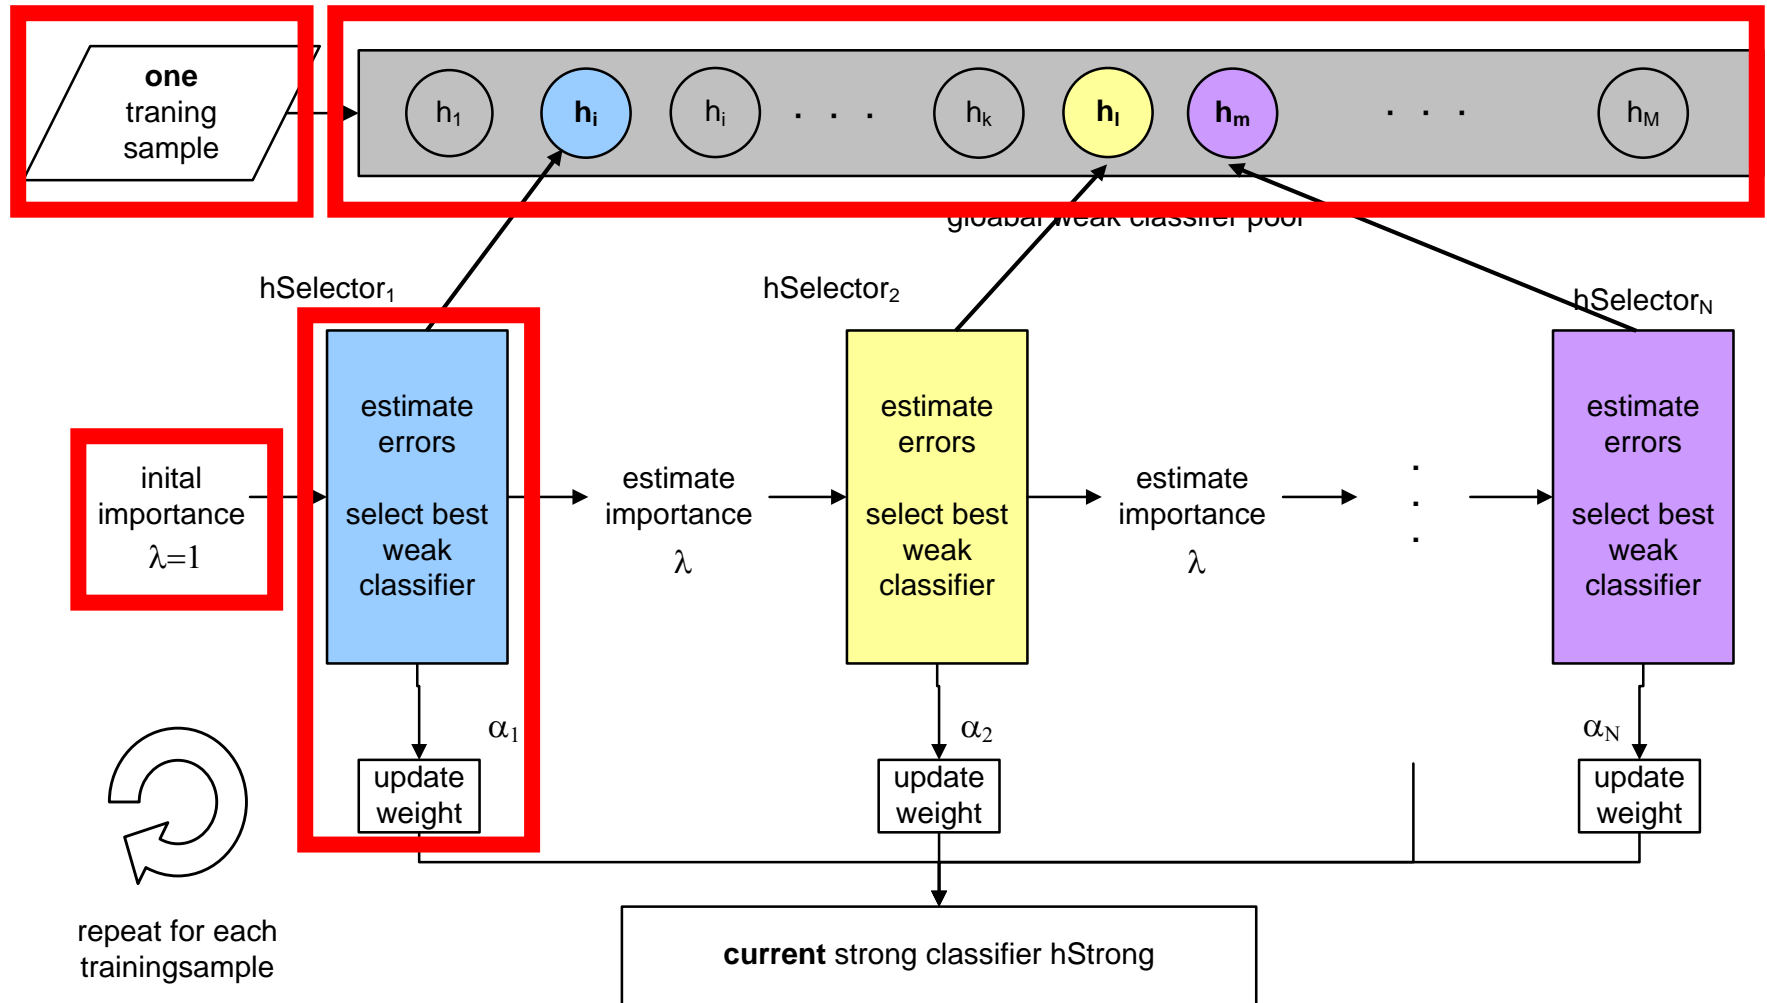


repeat for each
trainingsample

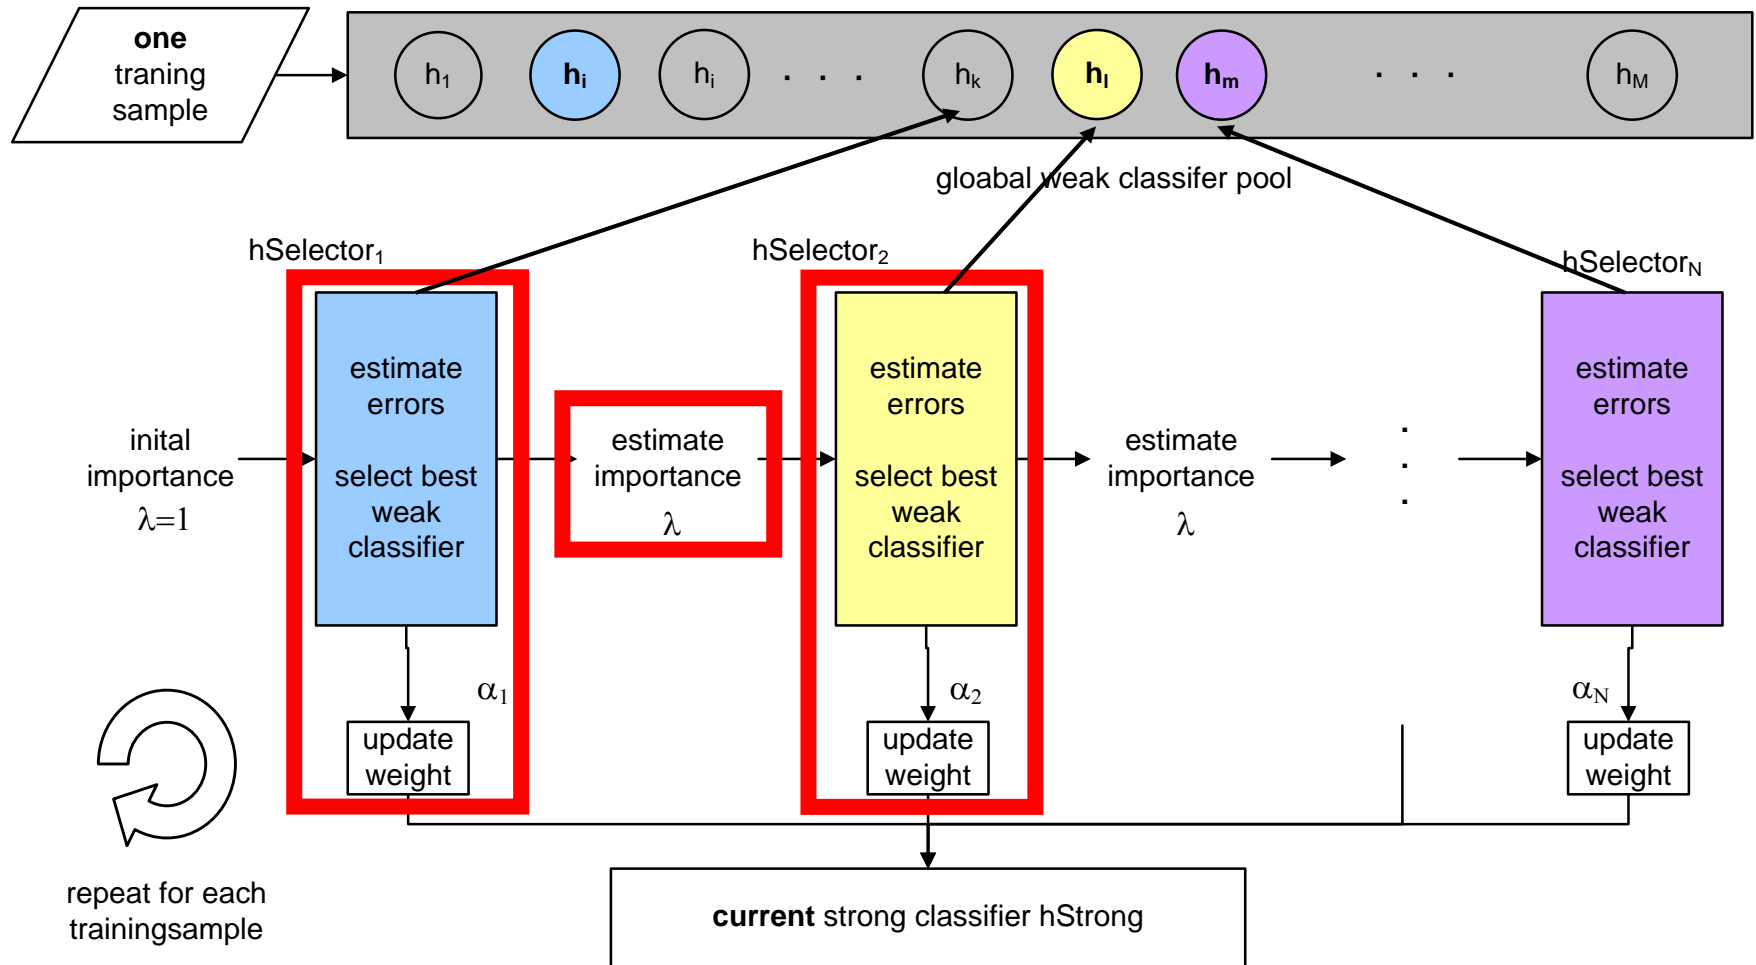
Online Boosting for Feature Selection



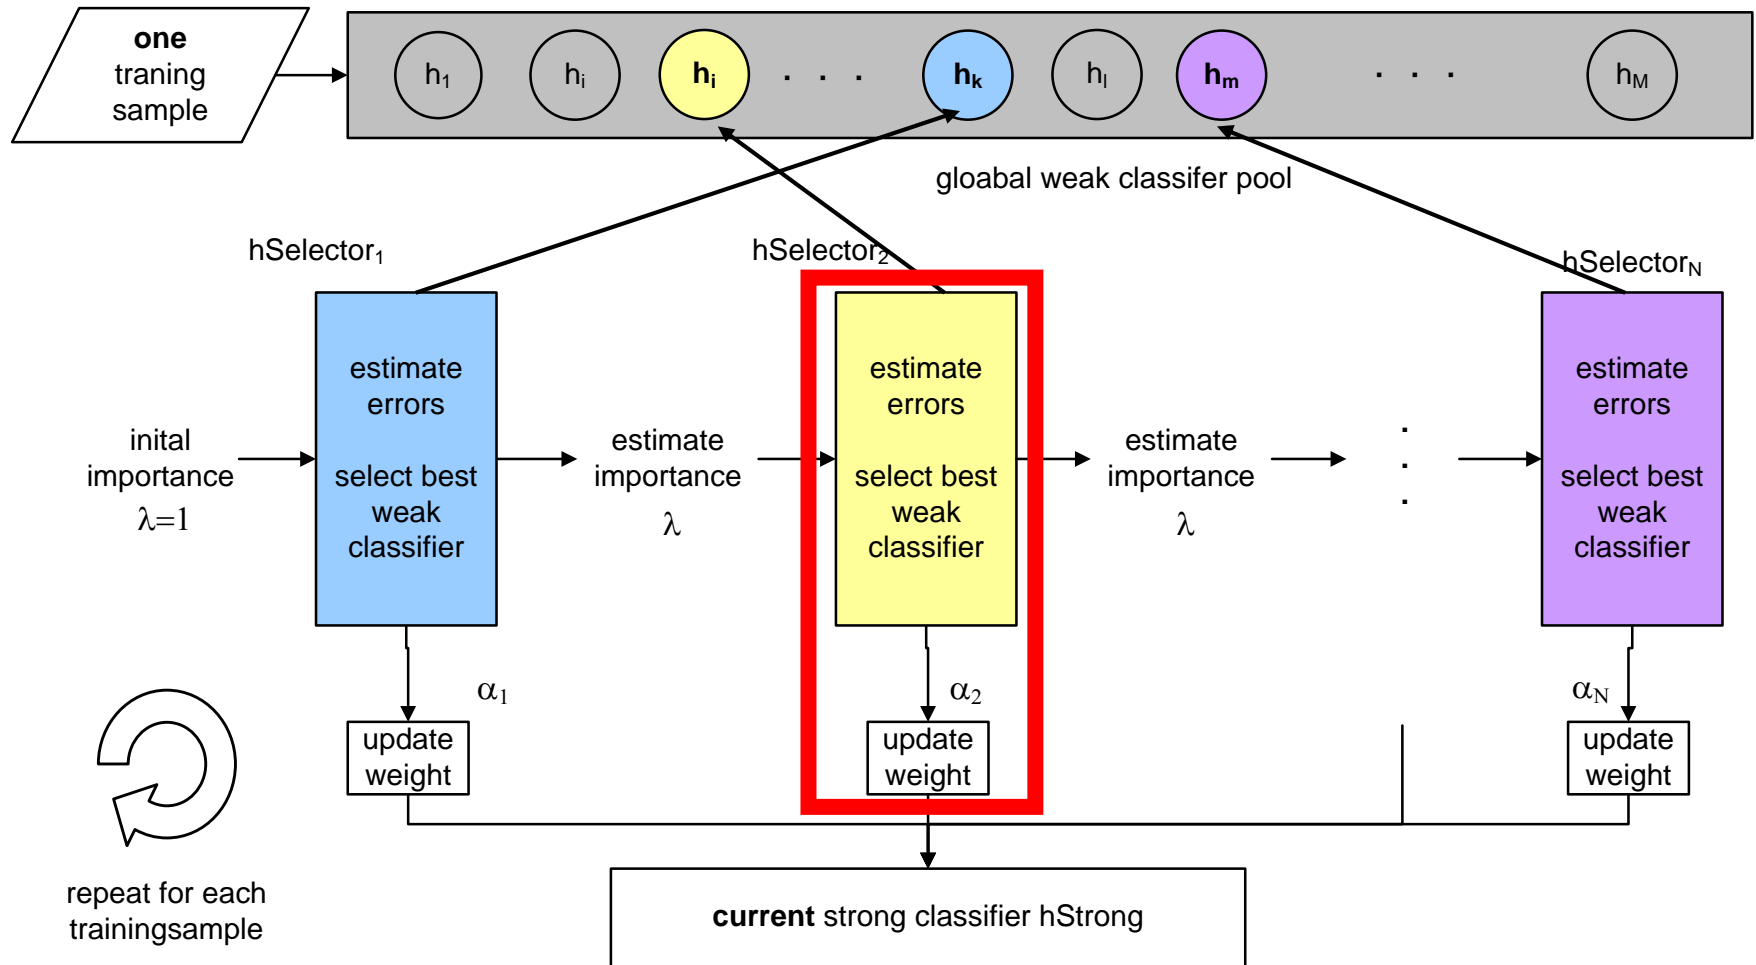
Direct Feature Selection



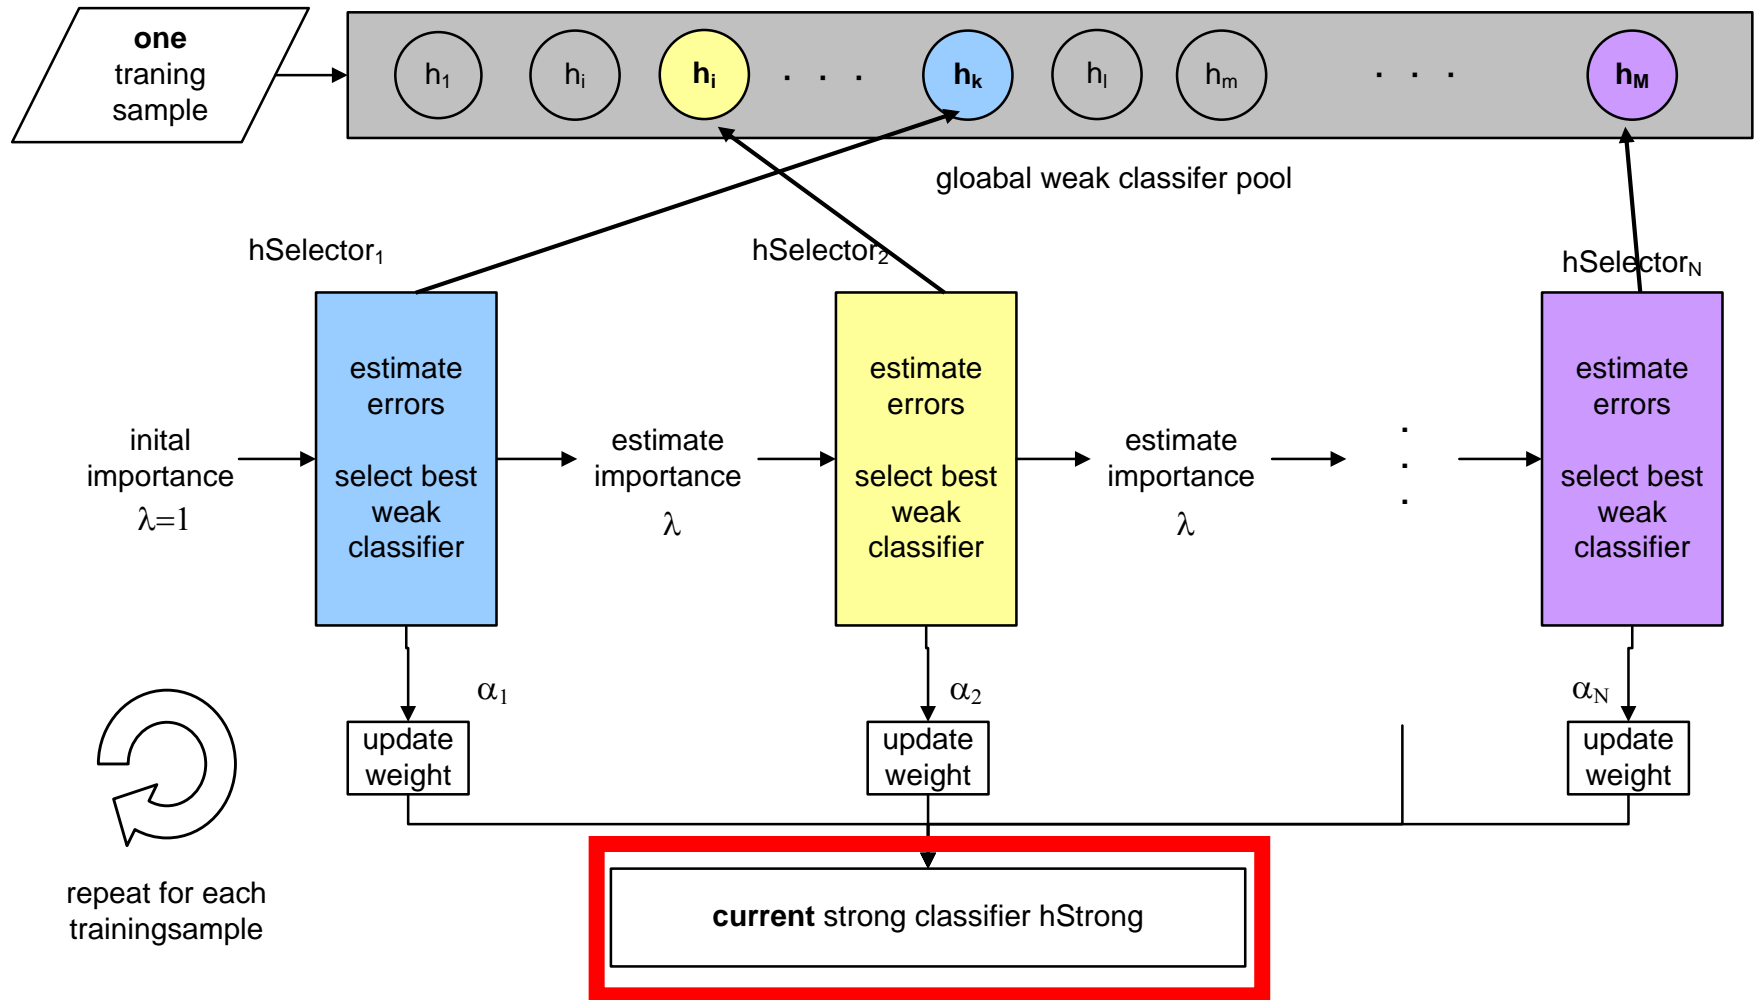
Direct Feature Selection



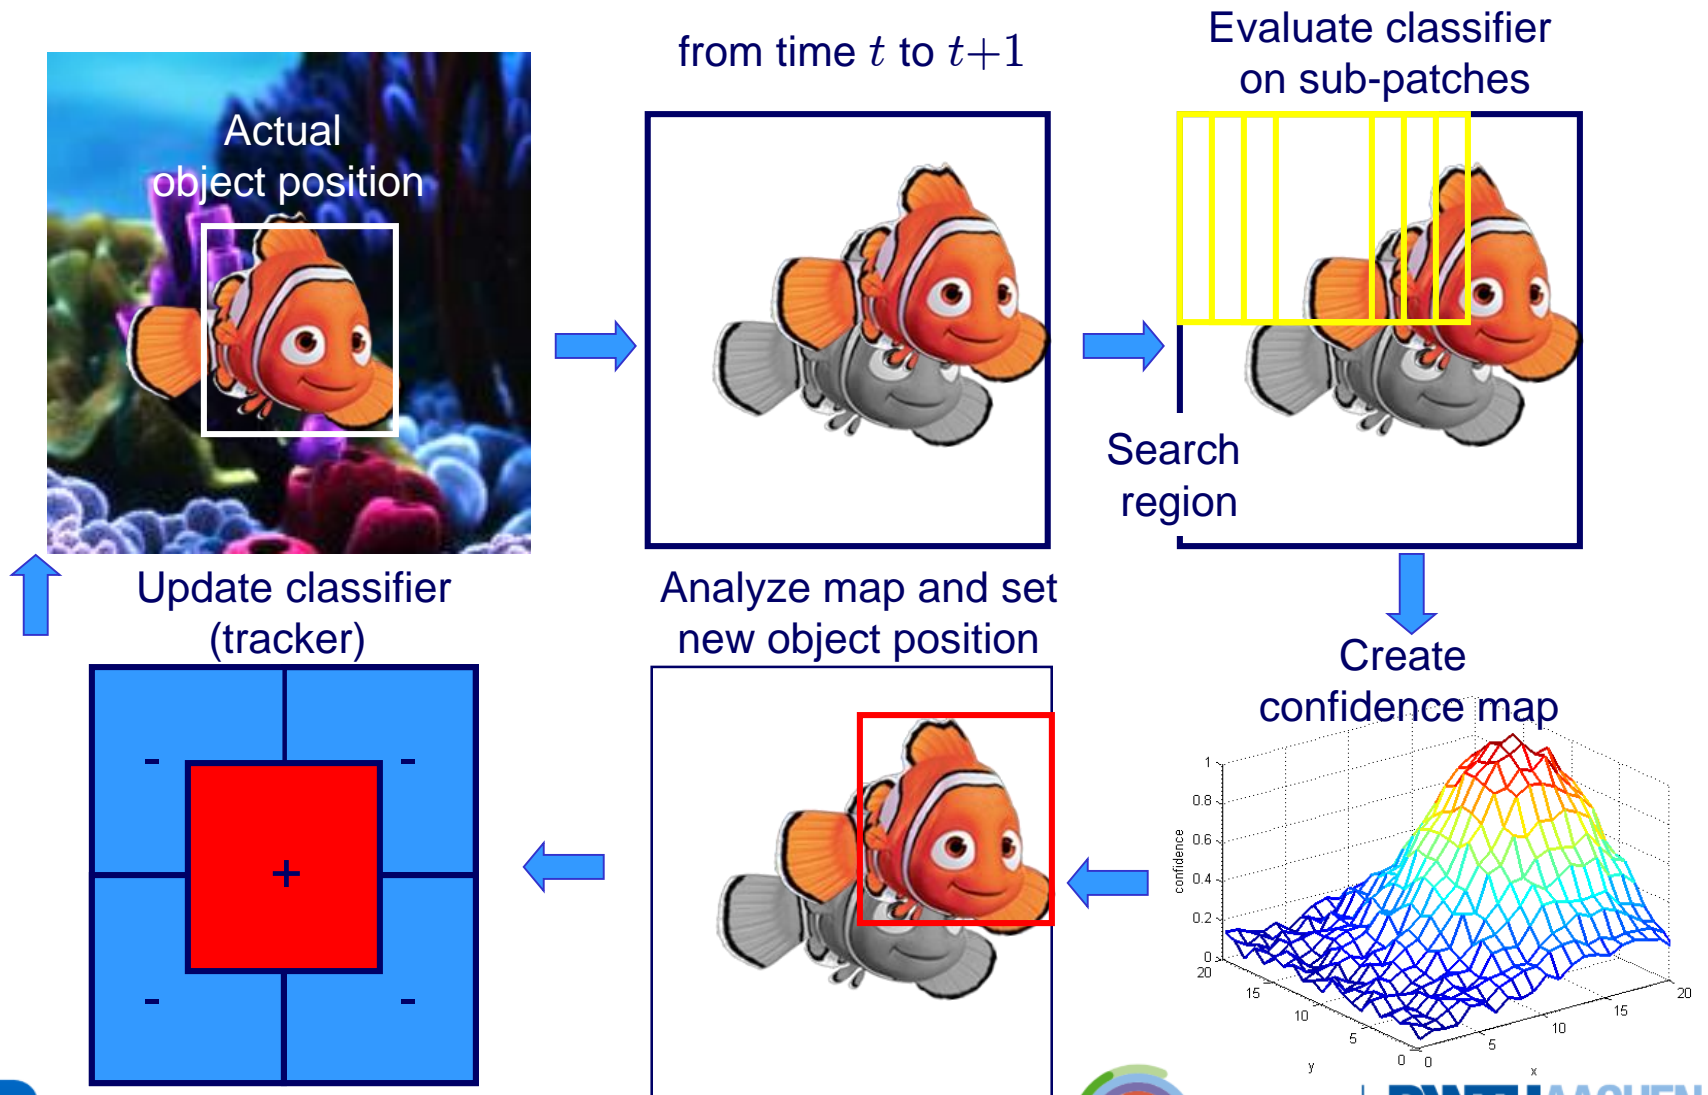
Direct Feature Selection



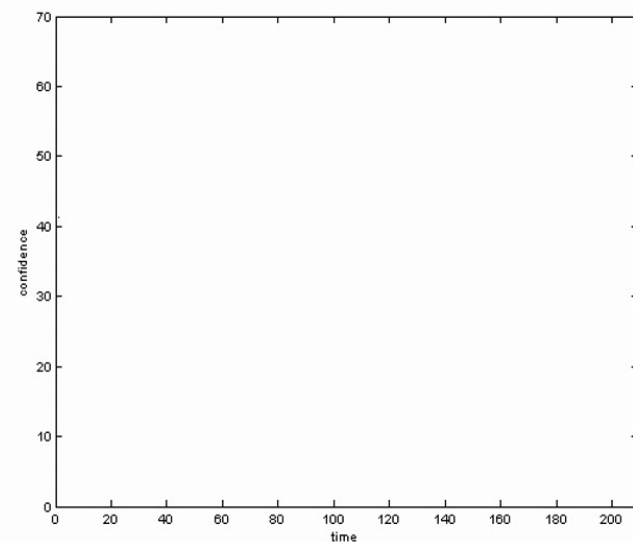
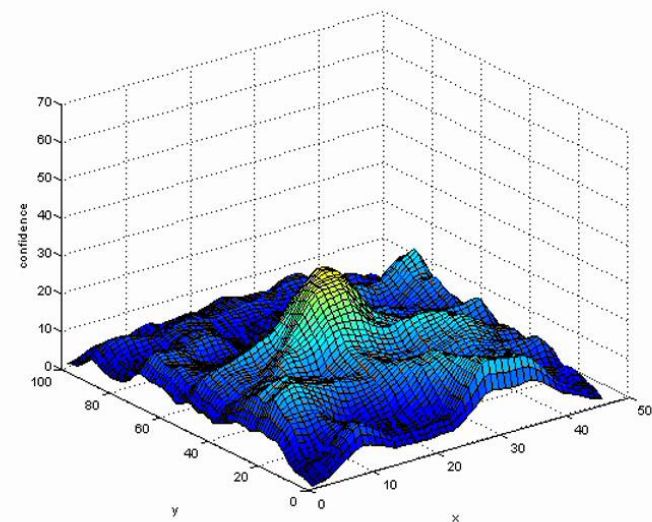
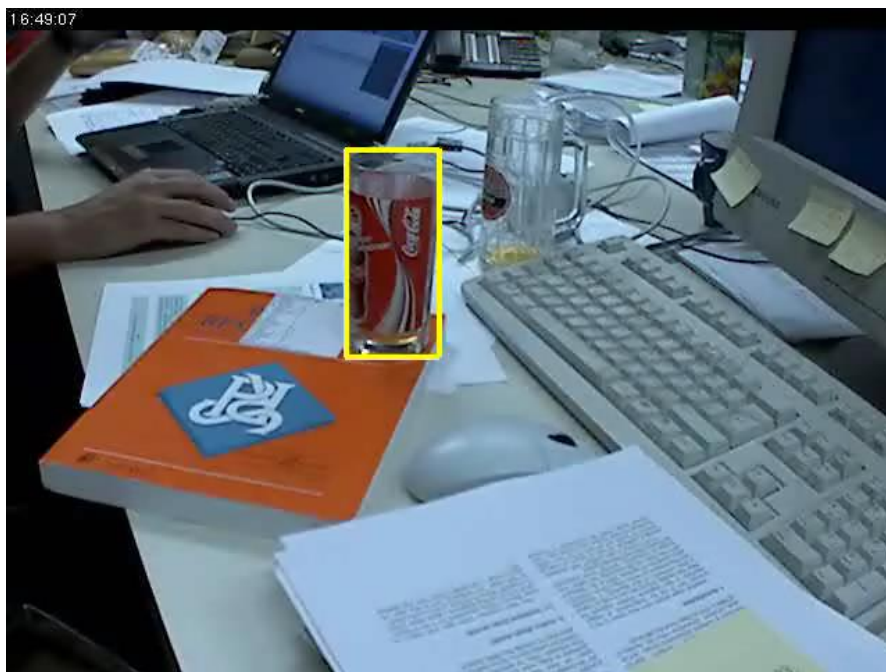
Direct Feature Selection



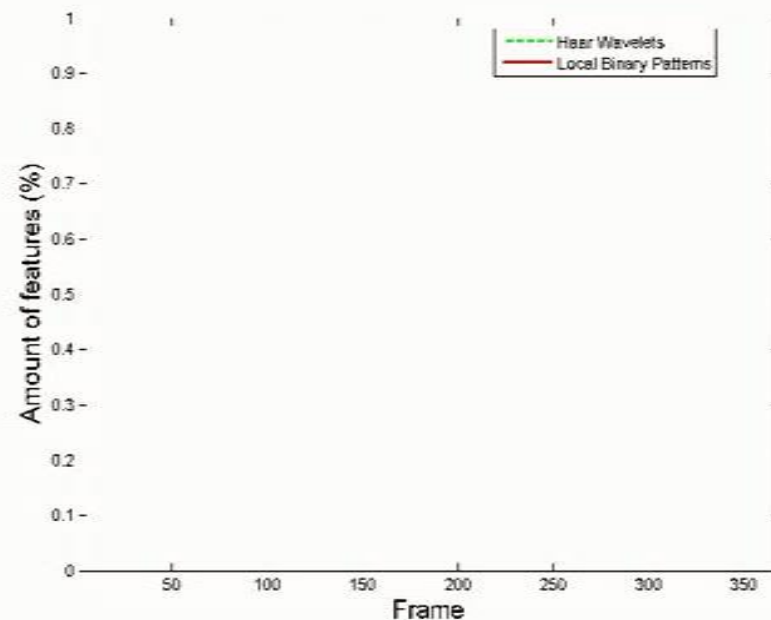
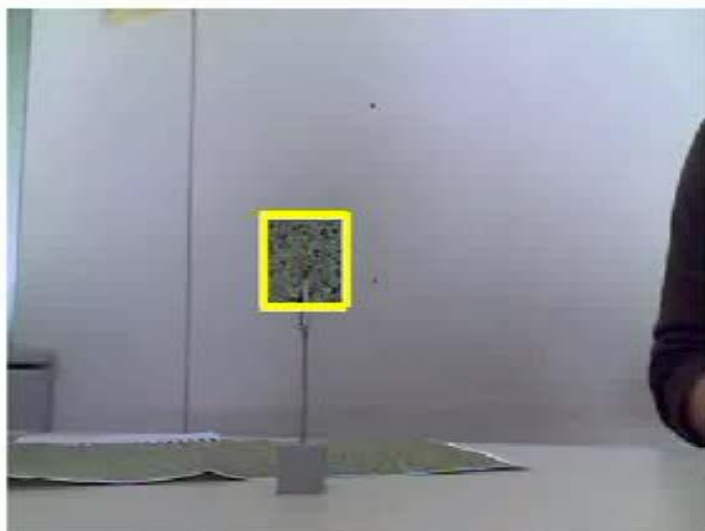
Tracking by Online Classification



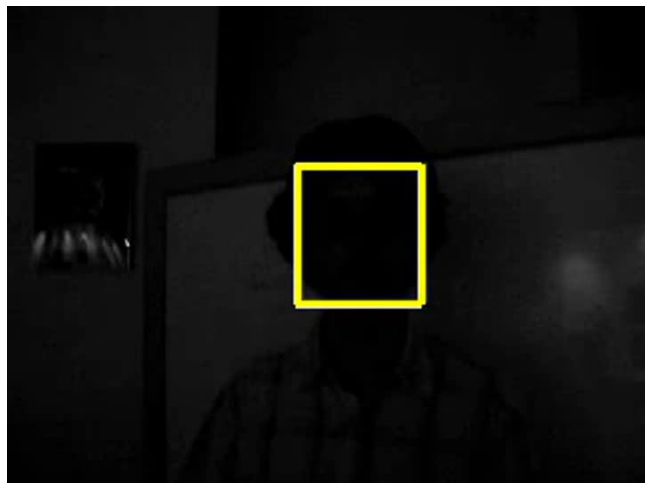
Tracking Results



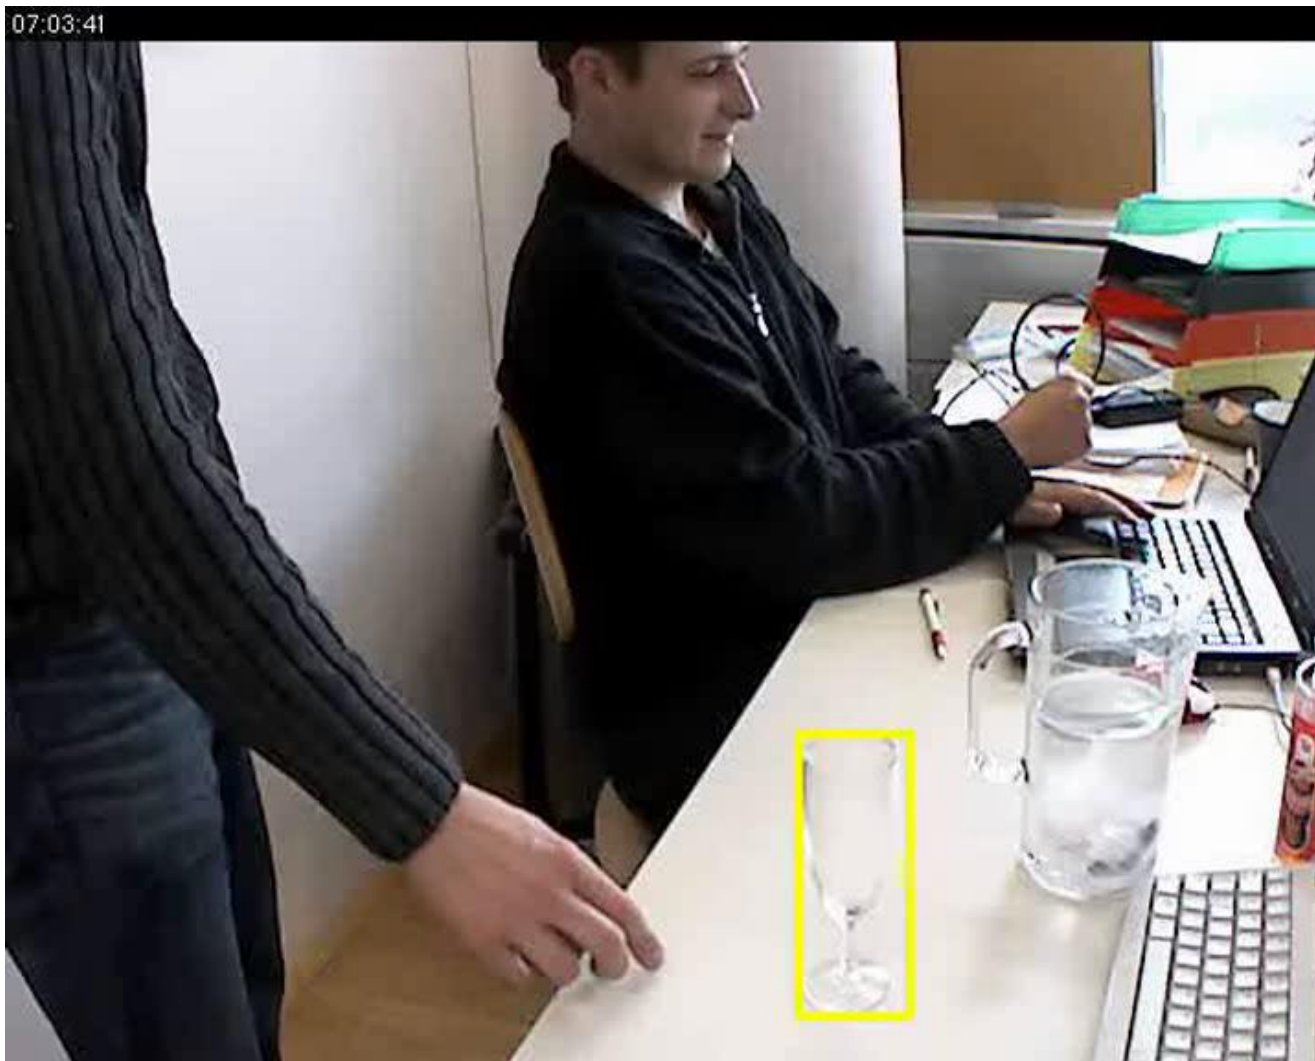
Online Feature Exchange



Additional Tracking Results



“Tracking the Invisible”



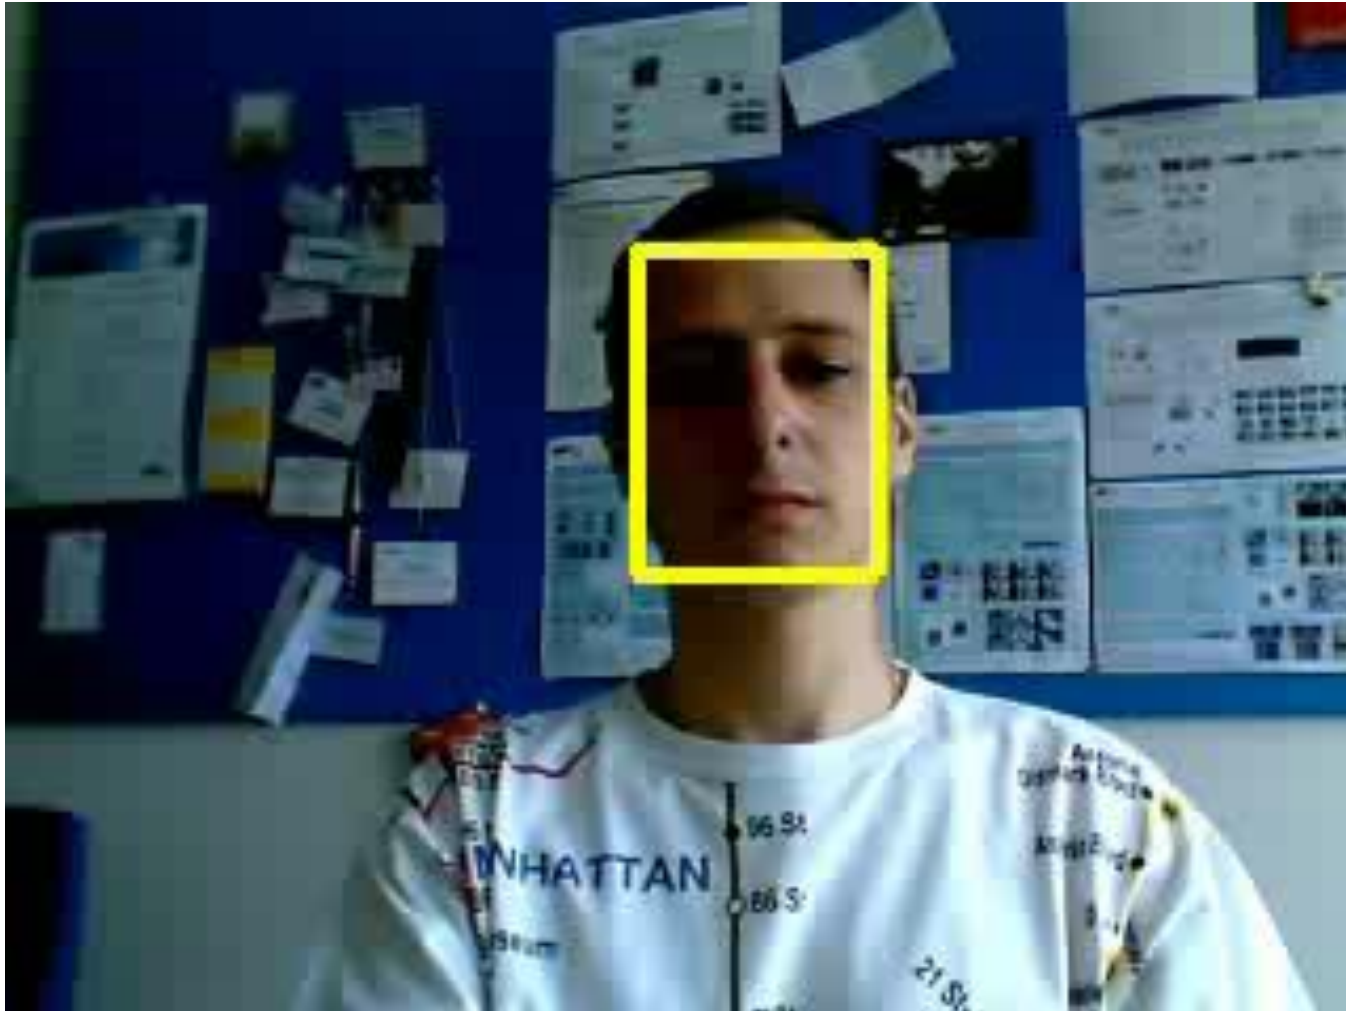
Summary: Tracking by Online Classification

- Interpret tracking as a classification problem
 - Continuously updating a classifier which discriminates the object from the background.
- Online Boosting
 - Adaptation of AdaBoost to process 1 training sample at a time.
 - Process sample by fixed set of classifiers to compute its importance weight.
 - Converges to the same result as Offline Boosting.
- Online Boosting for Feature Selection
 - Perform Boosting on Selectors instead of weak classifiers.
 - Each Selector chooses from a pool of weak classifiers.
 - Selected features and voting weights change over time.
 - Shared feature pool for real-time processing.

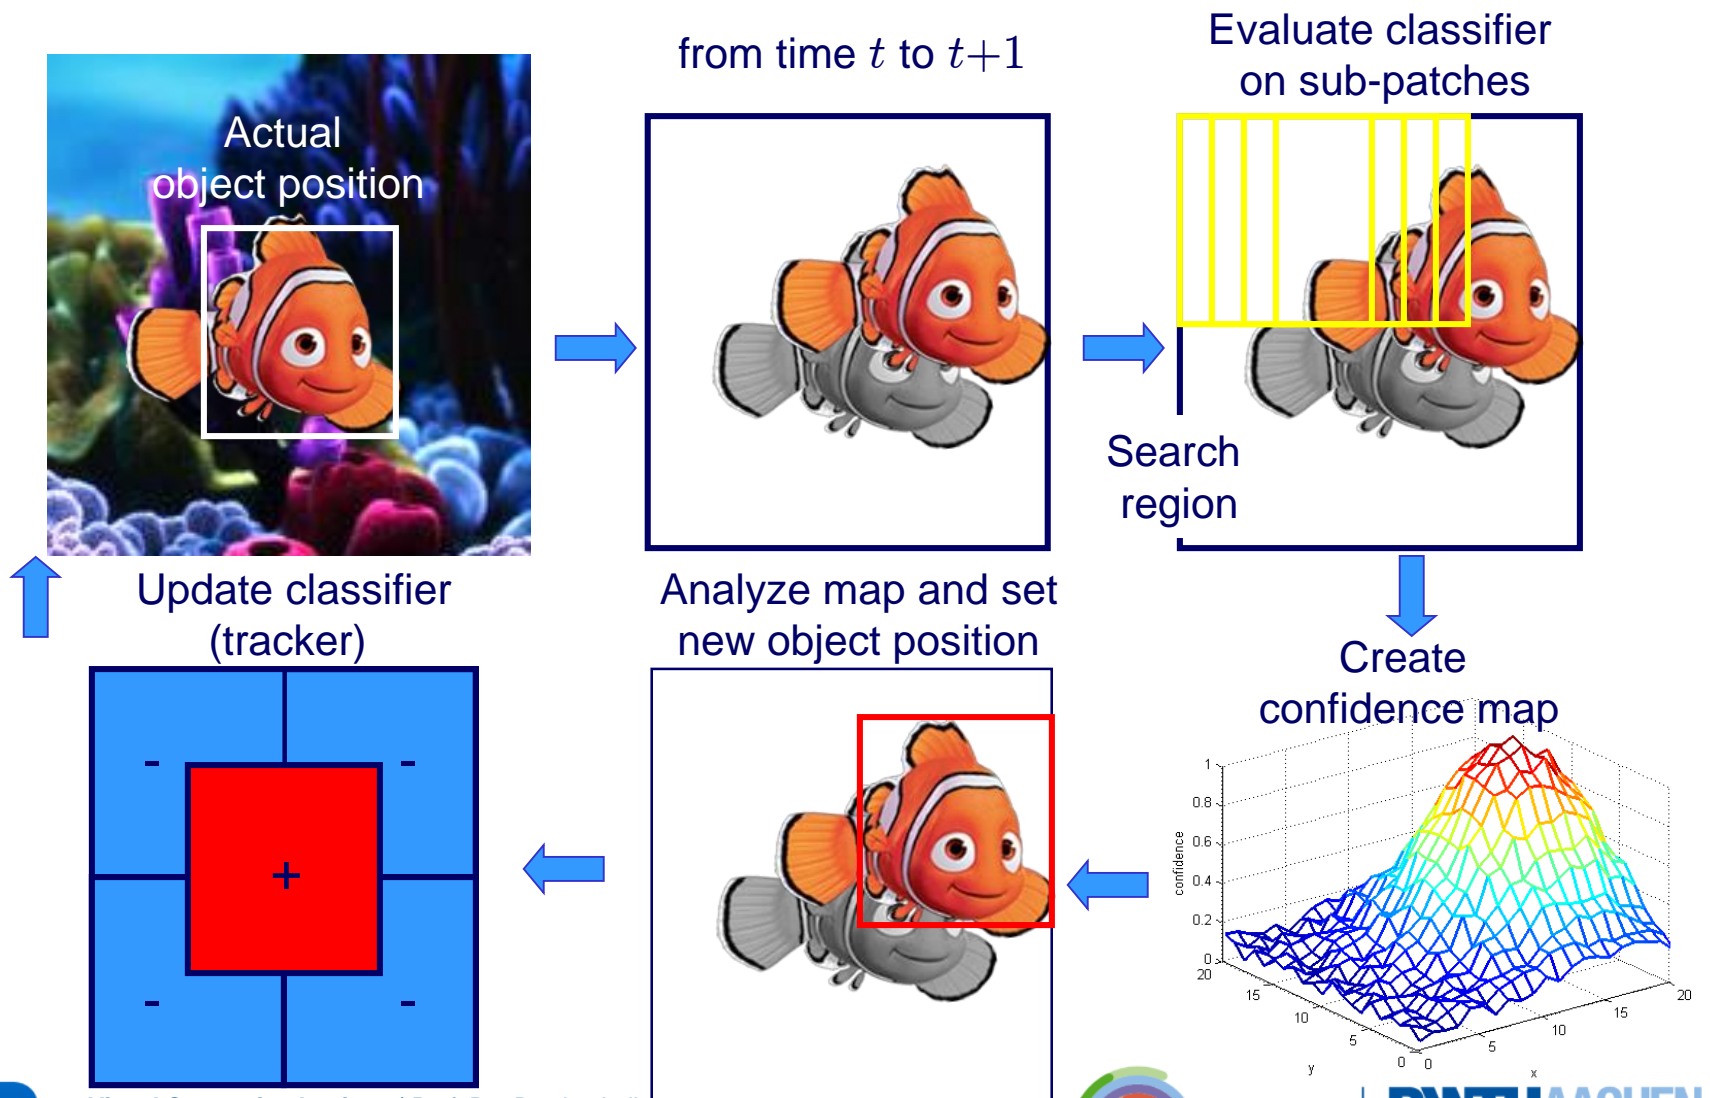
Topics of This Lecture

- Tracking by Online Classification
 - Motivation
- Recap: Boosting for Detection
 - AdaBoost
 - Viola-Jones Detector
- Extension to Online Classification
 - Online Boosting
 - Online Feature Selection
 - Results
- **Extensions**
 - Problem: Drift
 - Drift-compensation strategies

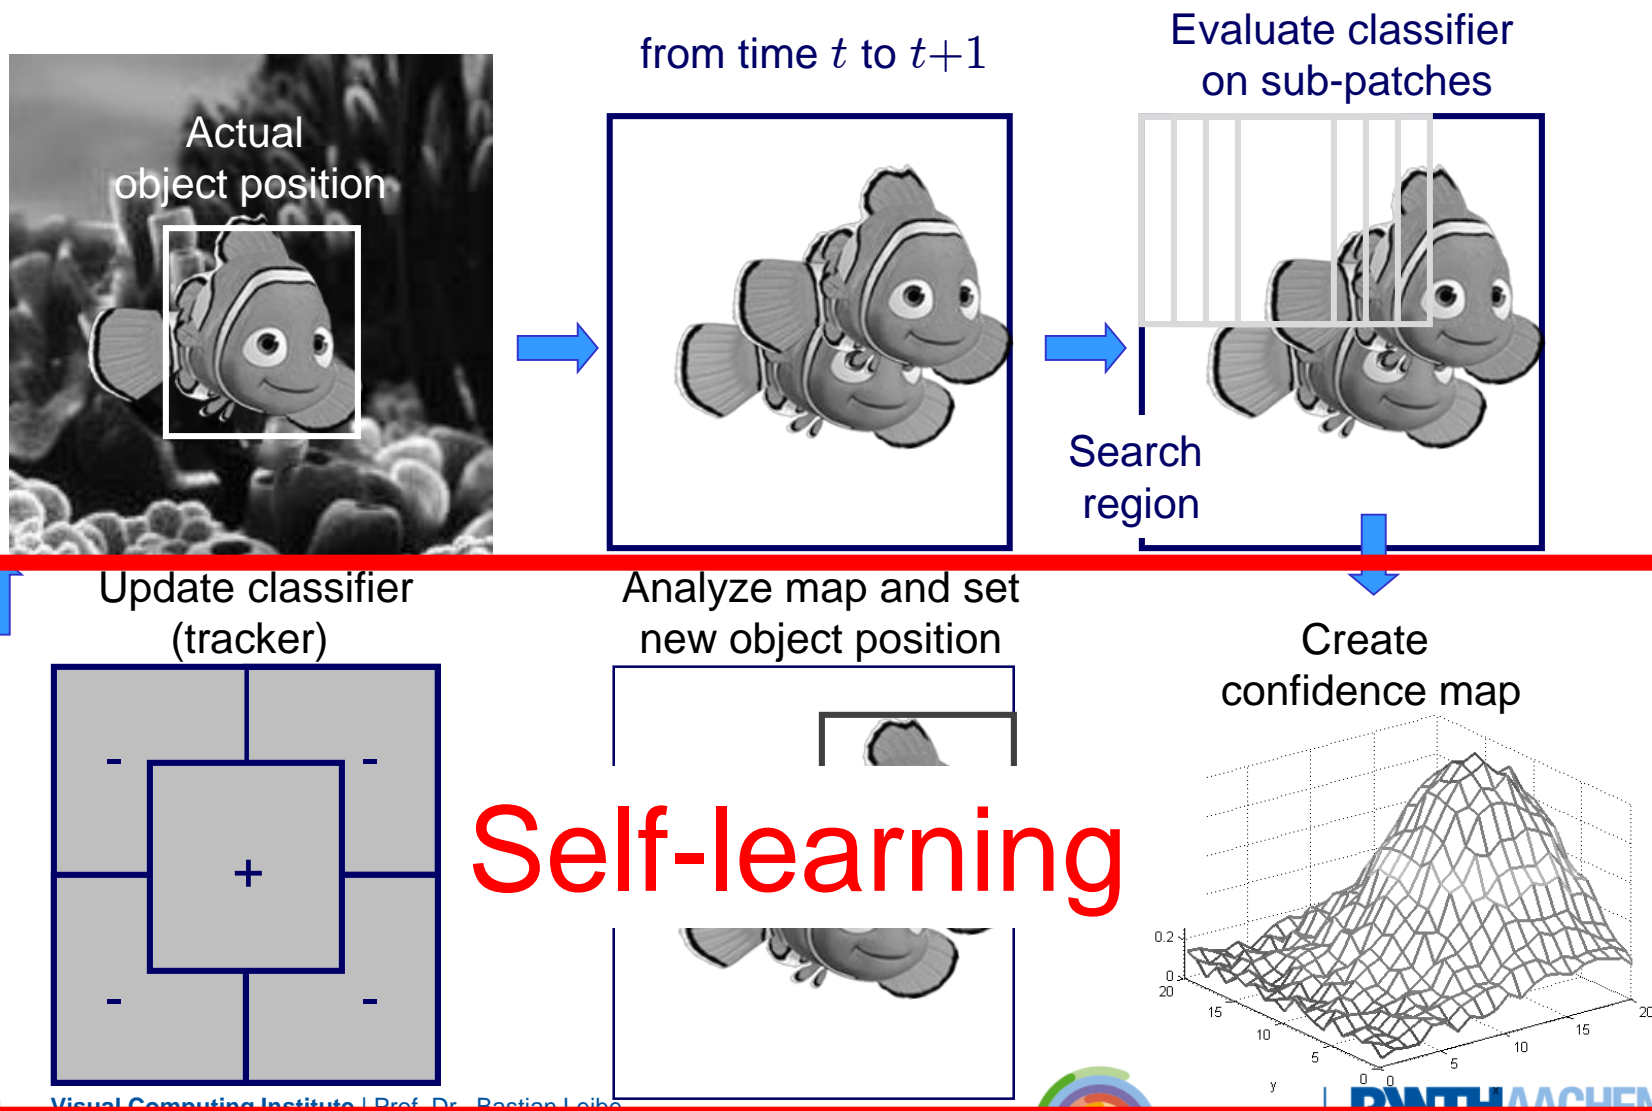
When Does It Fail...



Why Does It Fail?

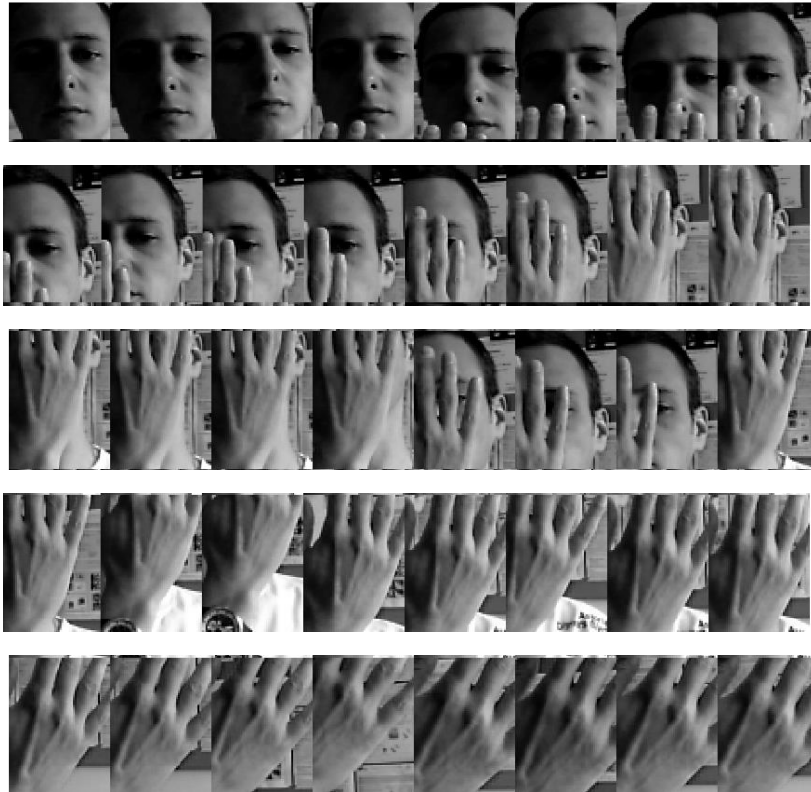


Why Does It Fail?

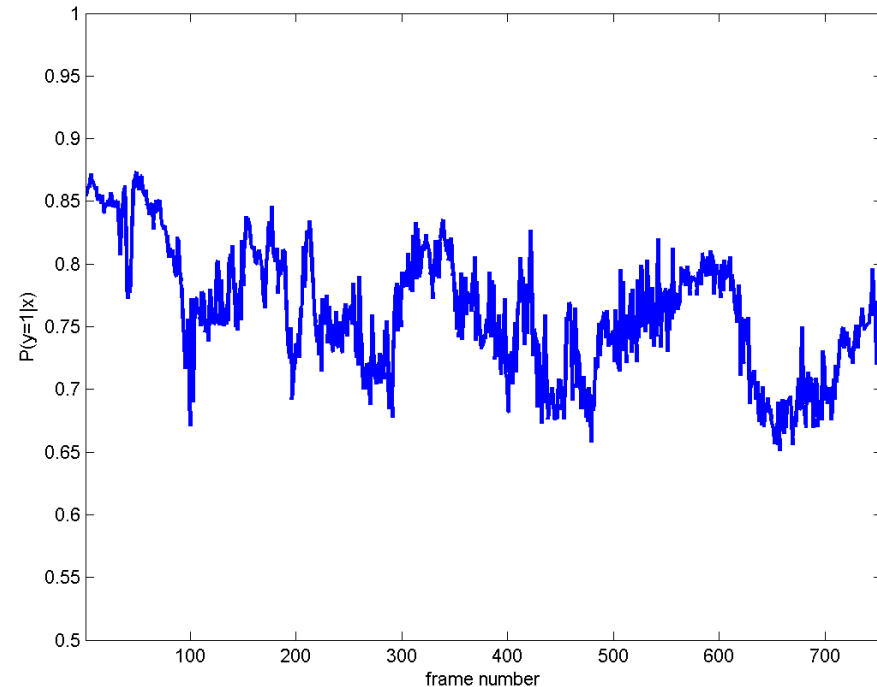


Drifting Due to Self-Learning Policy

Tracked Patches



Confidence

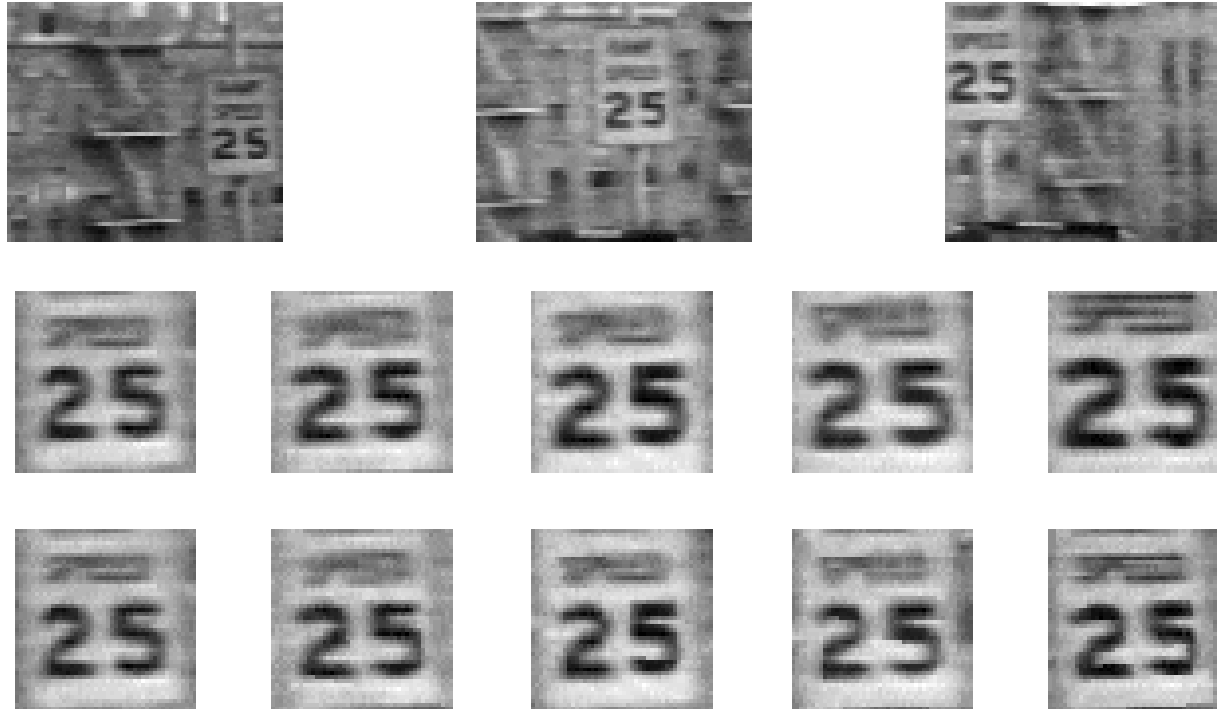


⇒ Not only does it drift, it also remains confident about it!

Self-Learning and Drift

- Drift
 - Major problem in all adaptive or self-learning trackers.
 - Difficulty: distinguish “allowed” appearance change due to lighting or viewpoint variation from “unwanted” appearance change due to drifting.
 - Cannot be decided based on the tracker confidence!
 - Since the confidence is always dependent on the learned model
 - Model may already be affected by drift when the confidence is measured.
 - Several approaches have been proposed to address this.

Strategy 1: Match Against Initialization



- Used mostly in low-level trackers (e.g., KLT)
 - Advantage: robustly catches drift
 - Disadvantage: cannot follow appearance changes

J. Shi and C. Tomasi. [Good Features to Track](#). CVPR 1994.

Strategy 2: Semi-Supervised Learning

Object Detector

Our approach

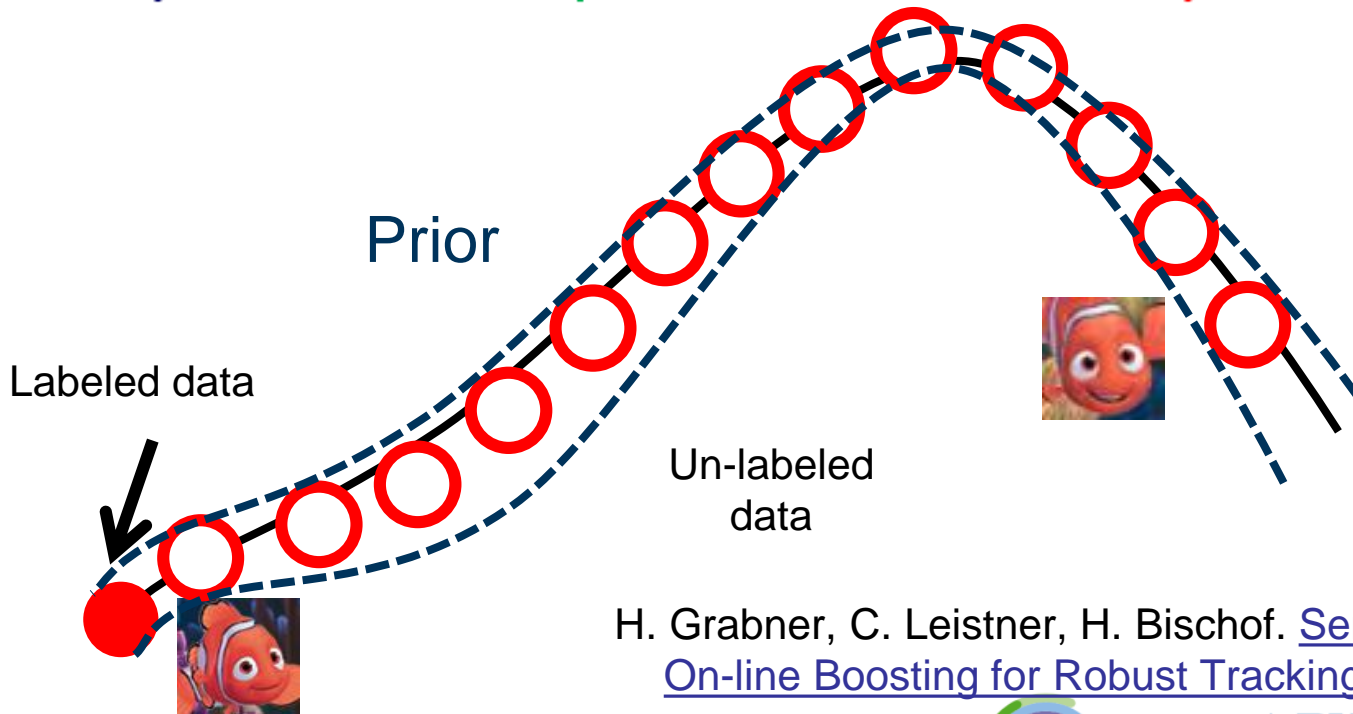
Object Tracker



Fixed Training set
General object detector

Fixed Prior for updating an
Adaptive on-line classifier

On-line update
Object vs. Background

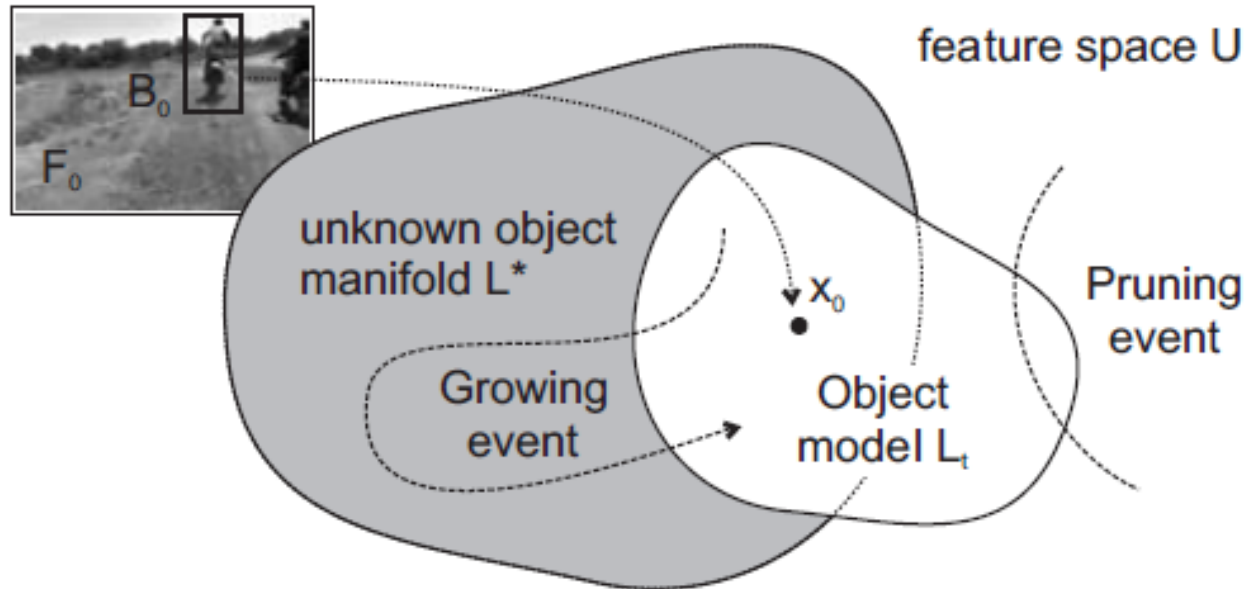


H. Grabner, C. Leistner, H. Bischof. [Semi-Supervised On-line Boosting for Robust Tracking](#). ECCV'08.

Tracking despite Occlusions



Strategy 3: Using Additional Cues

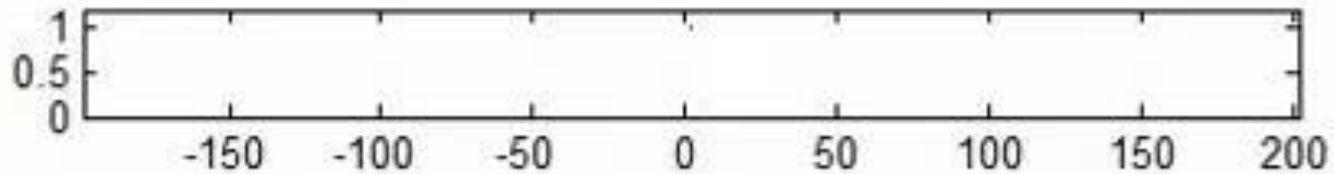


- Tracking-Learning-Detection
 - Combination of KLT and Tracking-by-Detection
 - Use a KLT tracker as additional cue to generate confident (positive and negative) training examples.
 - Learn an object detector on the fly using Online Random Ferns.

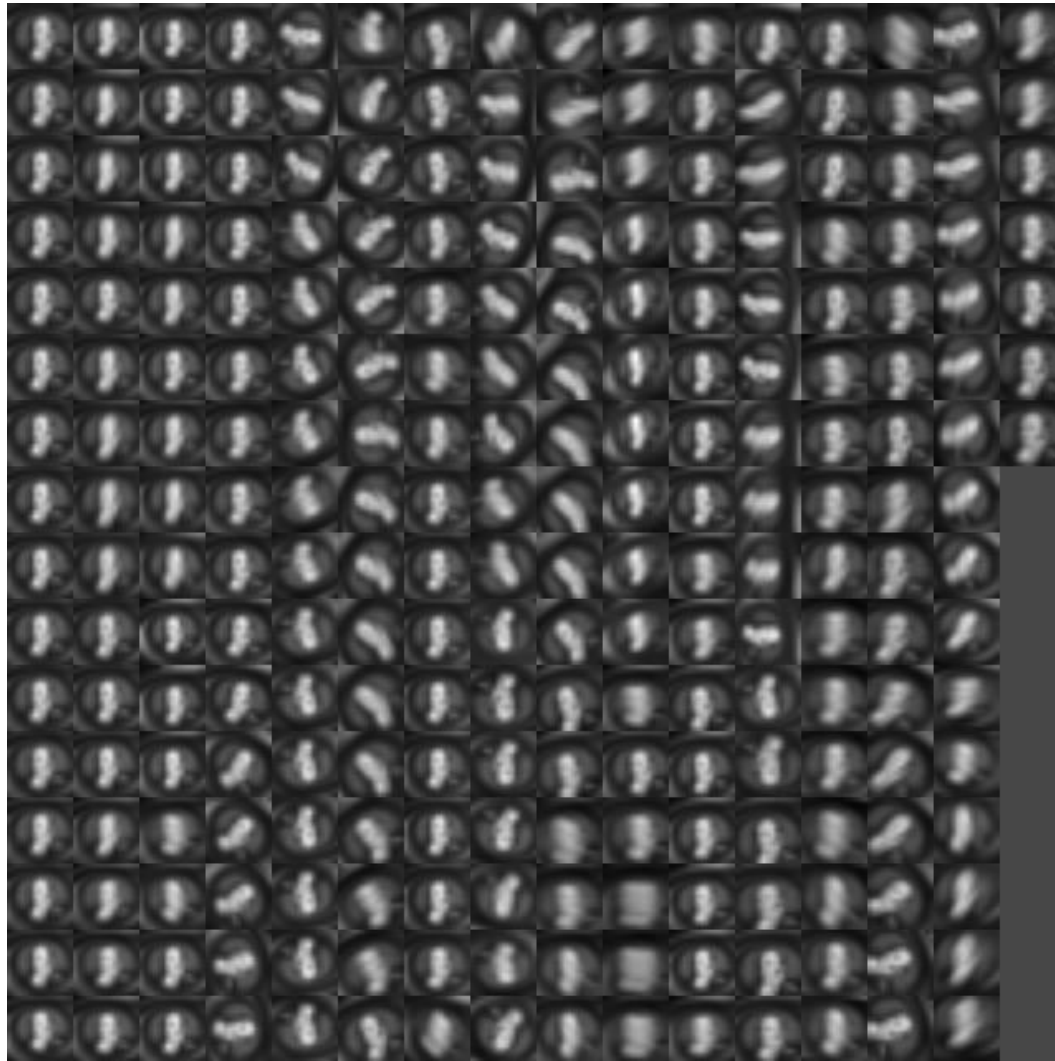
Z. Kalal, K. Mikolajczyk, J. Matas. [Tracking-Learning-Detection](#). PAMI 2011.

TLD Results

2



Accumulated Training Examples



TLD Results



References and Further Reading

- The original Online AdaBoost paper
 - N. Oza and S. Russell. [Online Bagging and Boosting](#). Artificial Intelligence and Statistics, 2001.
- Online Boosting for Tracking
 - H. Grabner, H. Bischof. [On-line Boosting and Vision](#). CVPR'06.
- Semi-Supervised Boosting
 - H. Grabner, C. Leistner, H. Bischof. [Semi-Supervised On-line Boosting for Robust Tracking](#). ECCV'08.
- Tracking-Learning-Detection
 - Z. Kalal, K. Mikolajczyk, J. Matas. [Tracking-Learning-Detection](#). PAMI 2011.