Current Trends in Machine Learning

Preparation Meeting

Jürgen Sturm, Rudolph Triebel,
Jan Stühmer, Christian Kerl

What you will learn in the seminar

- Get an overview on current trends in machine learning
- Read and understand scientific publications
- Write a scientific report
- Prepare and give a talk

Important Dates

- First Meeting: 30.10.2013 (today)
  - Fix assignment of papers and date
- Choose your topic until 6.11.2013 (next week, first come first serve!)
- Deadline for the report: 28.02.2014
- Dates for the talks:
  - 8.01.2014
  - 15.01.2014
  - 22.01.2014
  - 29.01.2014

Preparation

- Please do not work on your topic completely alone

→ Meet at least twice with your supervisor

- Recommended schedule
  - 1 month before your talk: Meet your supervisor and discuss paper
  - 1 week before your talk: Meet your supervisor to discuss your slides
  - [optional] after the talk: Feedback of your supervisor regarding the talk
  - 1 week before 28.02.14: Submit a draft of your report

Report and Talk

- Send PDF (not PPTX, not DOC) via email to your supervisor, Latex template available on the web-page
- Recommended length: 6-8 pages
- Required: Minimum 6, Maximum 10 pages
- Language: English or German

Hints for Your Talk

- 20 min. + 5–10 min. for discussion
- Don’t put too much information on one slide
  - 1-2 min. per slide → 10-20 slides
- Recommended structure
  - Introduction, Problem Motivation, Outline
  - Approach
  - Experimental results
  - Discussion
  - Summary of (scientific) contributions

Evaluation Criterions

- Gained expertise in the topic
- Quality of your talk
- Quality of the report
- Active participation in the seminar is required (ask questions, comment talks)

Regular Attendance Is Required

- Attendance at each appointment is necessary
- In case of absence: Medical attest

Papers

Paper title	Supervisor	Student name	Date
What makes Paris look like Paris?	Jürgen Sturm		
LeafSnap: Automatic Plant Specis Identification	Jürgen Sturm		
Active Learning for Level Set Estimation	Jürgen Sturm		
ImageNet Classification with Deep Convolutional Neural Networks	Christian Kerl		
Fast, Accurate Detection of 100,000 Object Classes on a Single Machine	Christian Kerl		
Multipath Sparse Coding Using Hierarchical Matching Pursuit	Christian Kerl		
Decision Tree Fields	Rudolph Triebel		
An Online Boosting Algorithm with Theoretical Justifications	Rudolph Triebel		
Active Learning for Large Multi-Class Problems	Rudolph Triebel		

Overview of available Topics

What makes Paris look like Paris?
[Doersch et al, SIGGRAPH 2012]

- Large repository of geotagged imagery (Google Streetview)
- Find visual elements that are geographically informative
- In which city were these two images taken?

LeafSnap: A Computer Vision System for Automatic Plant Species Identification

[Kumar et al, ECCV 2012]

- Visually identifying plant species using a smartphone
- Segmentation, feature extraction
- Learn a classifier based on curvature

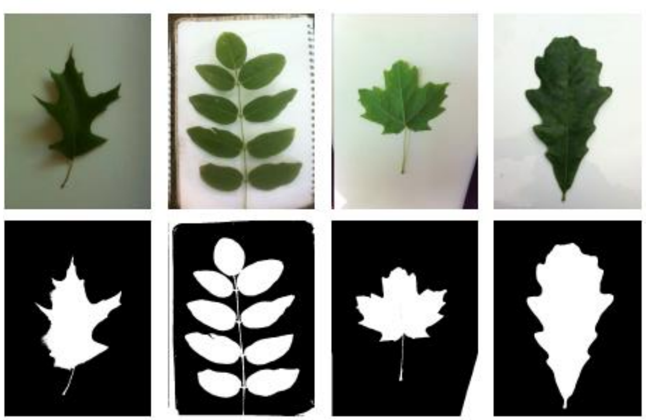

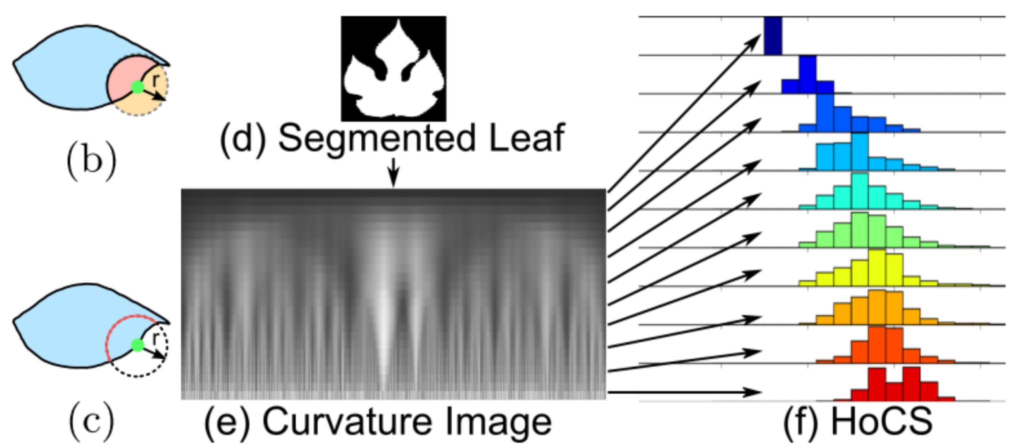

Active Learning for Level Set Estimation
[Gotovos et al., IJCAI 2013]

- Autonomous monitoring of algal population in a lake
- Minimize the number of measurements
- Gaussian process model for classification
- Guide sampling based on GP

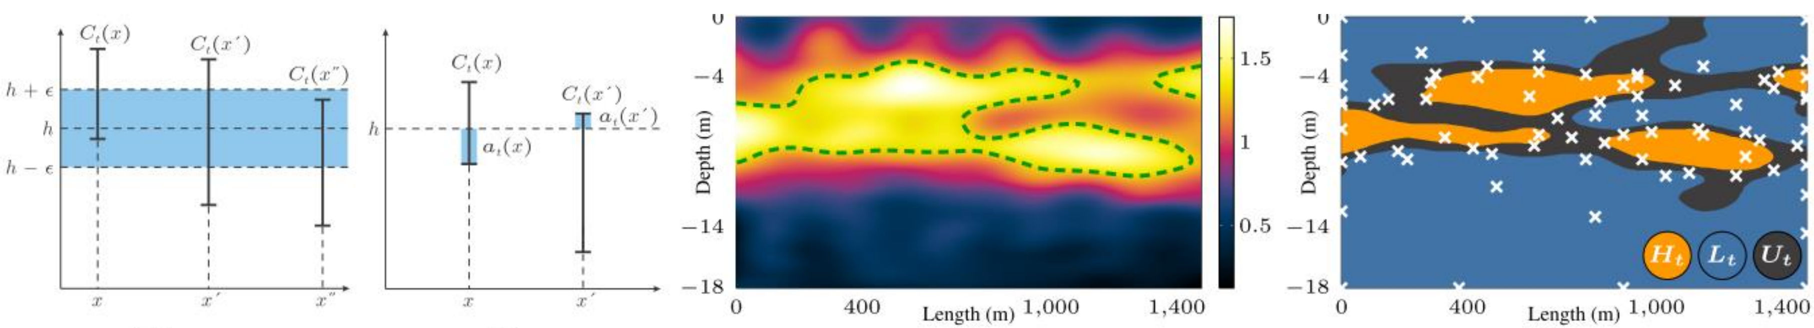

(a) (b) (c) (d)

Decision Tree Fields

- Combination of Conditional Random Fields with Decision Trees
- Efficient training method by parallelization
- Application to occlusion resolution and body-part detection

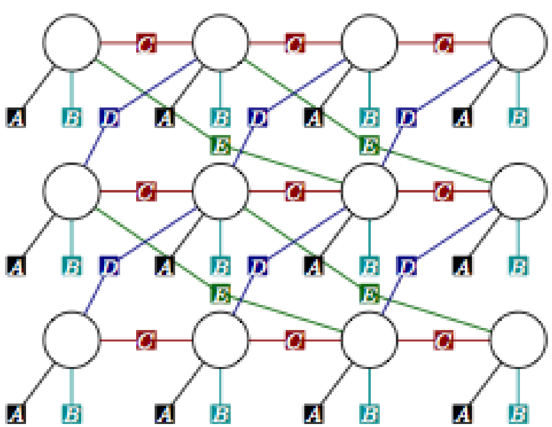

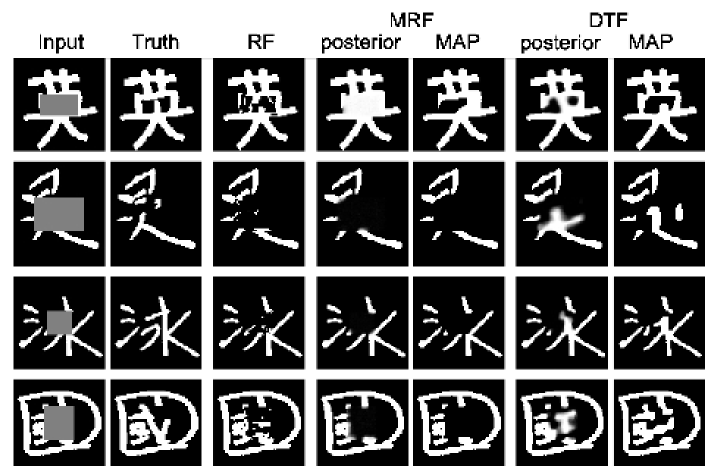

An Online Boosting Algorithm with Theoretical Justifications

- Extends the standard boosting method to online learning
- Bounds proven between online and offline
- Results on standard data sets better than former online boosting

Algorithm 1 Online Boosting with OCP

Input: streaming examples $(x_1, y_1), \ldots, (x_T, y_T)$
parameters $0 < \delta < 1$, $0 \leq \theta < \gamma < \frac{1}{2}$
online weak learner WL
Initialize: $z_t^{(0)} = 0$ for $t \in [T]$
$w_t^{(1)} = 1$ for $t \in [T]$
$\alpha_1^{(i)} = \frac{1}{N}$ for $i \in [N]$
select random $h_1^{(i)} \in \mathcal{H}$ for $i \in [N]$
for $t = 1$ **to** T **do**
  Define $f_t(x) = \sum_{i=1}^{N} \alpha_t^{(i)} h_t^{(i)}(x)$
  Predict with $H_t(x) = \text{sign}\,(f_t(x))$
  if $y_t f_t(x_t) < \theta$ **then**
    $\alpha_{t+1}^{(i)} = \alpha_t^{(i)} + \eta_t y_t h_t^{(i)}(x_t)$, for $i \in [N]$
    project α_{t+1} back into probability simplex
  end if
  for $i = 1$ **to** N **do**
    $h_{t+1}^{(i)} \leftarrow \mathsf{WL}\left(h_t^{(i)}, (x_t, y_t), w_t^{(i)}\right)$
    $z_t^{(i)} = z_t^{(i-1)} + y_t h_t^{(i)}(x_t) - \theta$
    $w_t^{(i+1)} = \min\left\{(1-\gamma)^{z_t^{(i)}/2}, 1\right\}$
  end for
end for

Active Learning for Large Multi-class Problems

- Active Learning for Image Classification
- Based on a probabilistic k-NN method
- Results on standard data set are better than SVM and GP methods

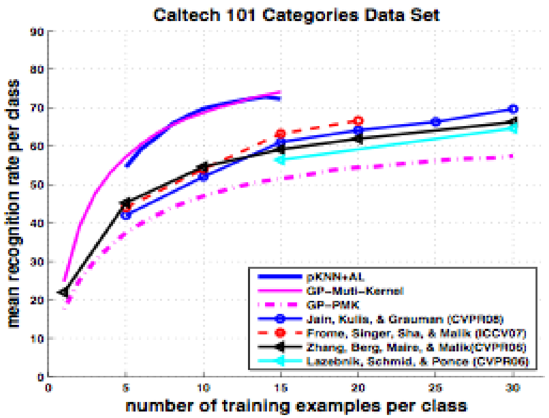

ImageNet Classification with Deep Convolutional Neural Networks
[Krizhevsky et al, NIPS 2012]

- Classification of objects in images
- Trains huge multi-layered convolutional neural networks

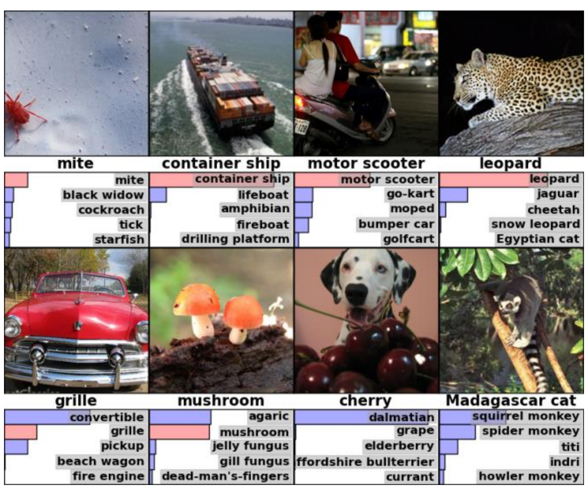

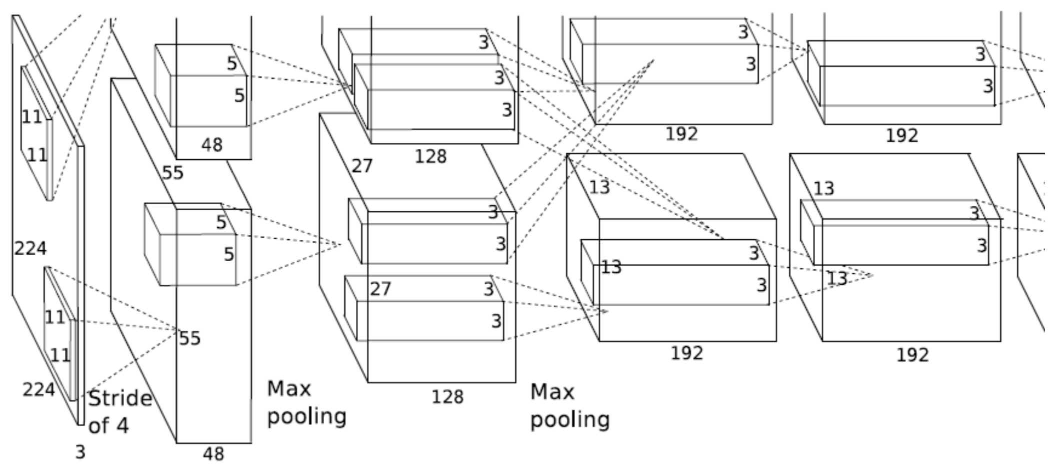

Multipath Sparse Coding Using Hierarchical Matching Pursuit
[Bo et al, CVPR 2013]

- Classification of objects in images
- Uses sparse coding to learn image features
- Layered architecture

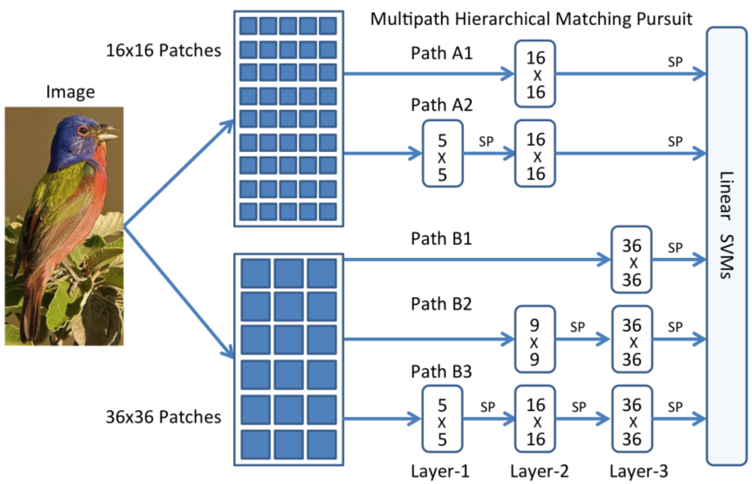

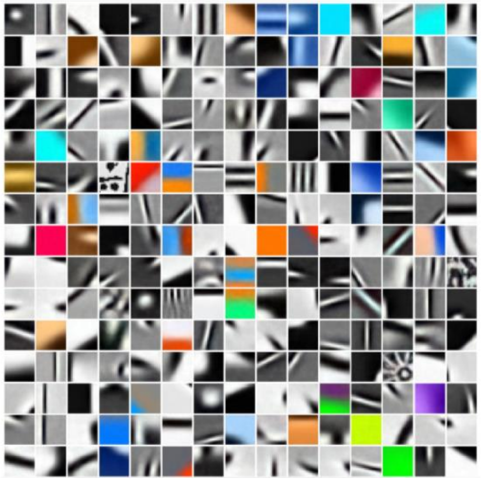

Fast, Accurate Detection of 100,000 Object Classes on a Single Machine

[Dean et al, CVPR 2013]

- Detection + classification of objects in images
- Uses Deformable Part Model per object class
- Accelerates detection through hashing

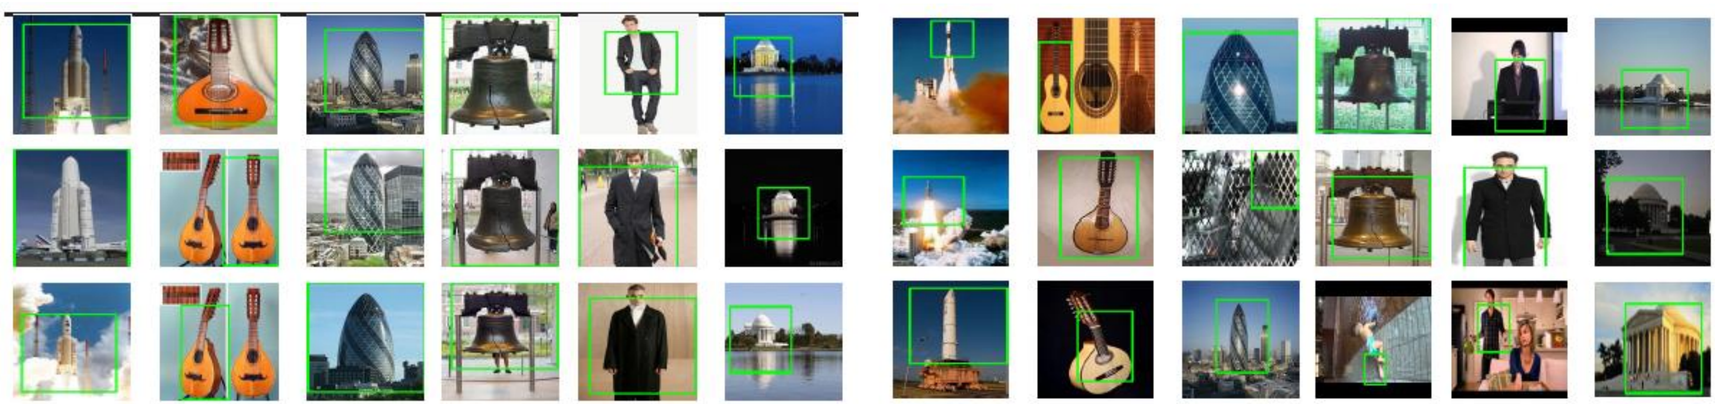

Enjoy the seminar!