Computer Vision - Lecture 9

Sliding-Window based Object Detection

26.11.2015

Bastian Leibe
RWTH Aachen
http://www.vision.rwth-aachen.de

leibe@vision.rwth-aachen.de

Course Outline

- Image Processing Basics
- Segmentation
  - Segmentation and Grouping
  - Segmentation as Energy Minimization
- Recognition & Categorization
  - Global Representations
  - Sliding-Window Object Detection
  - Image Classification
- Local Features & Matching
- 3D Reconstruction
- Motion and Tracking

Recap: Appearance-Based Recognition

- Basic assumption
  - Objects can be represented by a set of images (“appearances”).
  - For recognition, it is sufficient to just compare the 2D appearances.
  - No 3D model is needed.

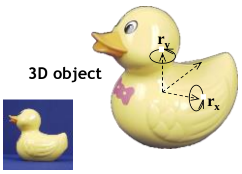

⇒ Fundamental paradigm shift in the 90’s

Recap: Recognition Using Histograms

- Histogram comparison

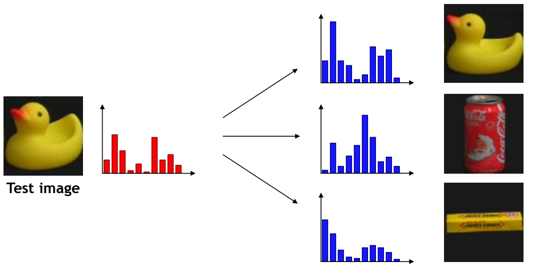

Recap: Comparison Measures

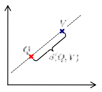

- Vector space interpretation
  - Euclidean distance
  - Mahalanobis distance
- Statistical motivation
  - Chi-square
  - Bhattacharyya
- Information-theoretic motivation
  - Kullback-Leibler divergence, Jeffreys divergence
- Histogram motivation
  - Histogram intersection
- Ground distance
  - Earth Movers Distance (EMD)

Recap: Recognition Using Histograms

- Simple algorithm
  1. Build a set of histograms $H=\{h_i\}$ for each known object
     - More exactly, for each *view* of each object
  2. Build a histogram h_t for the test image.
  3. Compare h_t to each $h_i \in H$
     - Using a suitable comparison measure
  4. Select the object with the best matching score
     - Or reject the test image if no object is similar enough.

“Nearest-Neighbor” strategy

Recap: Multidimensional Representations

- Combination of several descriptors
  - Each descriptor is applied to the whole image.
  - Corresponding pixel values are combined into one feature vector.
  - Feature vectors are collected in multidimensional histogram.

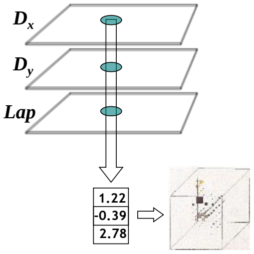

Application: Brand Identification in Video

[Hall, Pellison, Riff, Crowley, 2004]

Application: Brand Identification in Video

[Hall, Pellison, Riff, Crowley, 2004]

Application: Brand Identification in Video

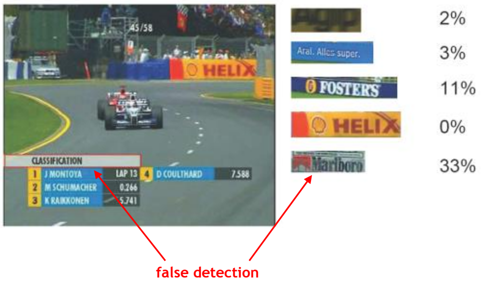

[Hall, Pellison, Riff, Crowley, 2004]

You're Now Ready for First Applications…

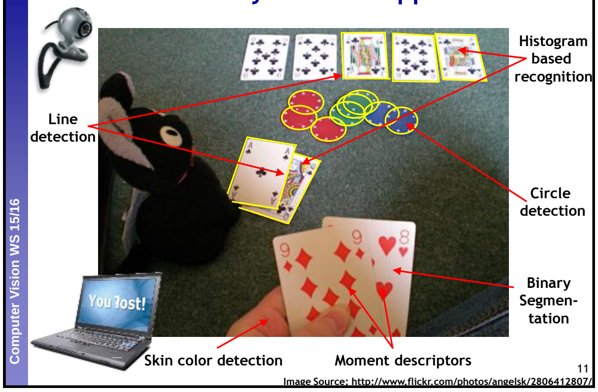

Image Source: http://www.flickr.com/photos/angelsk/2806412807/

Topics of This Lecture

- Object Categorization
  - Problem Definition
  - Challenges
- Sliding-Window based Object Detection
  - Detection via Classification
  - Global Representations
  - Classifier Construction
- Classification with Boosting
  - AdaBoost
  - Viola-Jones Face Detection
- Classification with SVMs
  - Support Vector Machines
  - HOG Detector

Identification vs. Categorization

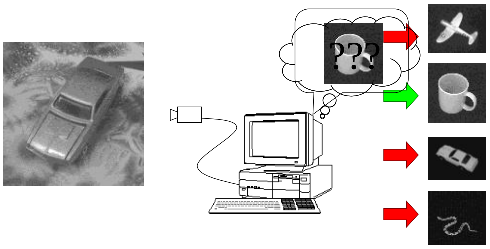

Identification vs. Categorization

- Find *this particular* object

- Recognize ANY car

- Recognize ANY cow

Object Categorization – Potential Applications

There is a wide range of applications, including.

Autonomous robots

Navigation, driver safety

Consumer electronics

Content-based retrieval and analysis for images and videos

Medical image analysis

Slide adapted from Kristen Grauman

How many object categories are there?

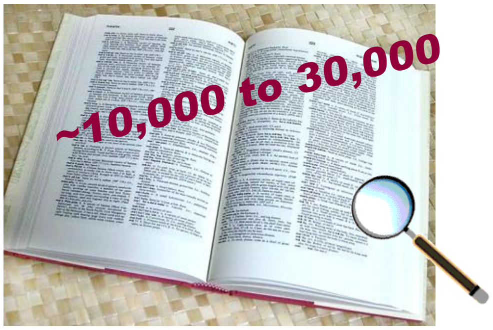

~10,000 to 30,000

Source: Fei-Fei Li, Rob Fergus, Antonio Torralba.

Biederman 1987

~10,000 to 30,000

Challenges: Robustness

Illumination

Object pose

Clutter

Occlusions

Intra-class appearance

Viewpoint

Slide credit: Kristen Grauman

Challenges: Robustness

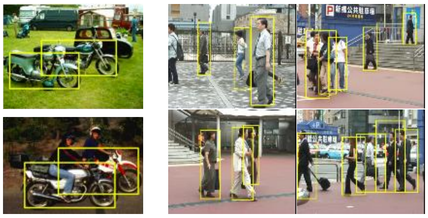

- Detection in crowded, real-world scenes
  - Learn object variability
    - Changes in appearance, scale, and articulation
  - Compensate for clutter, overlap, and occlusion

[Leibe, Seemann, Schiele, CVPR'05]

Topics of This Lecture

- Object Categorization
  - Problem Definition
  - Challenges
- Sliding-Window based Object Detection
  - Detection via Classification
  - Global Representations
  - Classifier Construction
- Classification with Boosting
  - AdaBoost
  - Viola-Jones Face Detection
- Classification with SVMs
  - Support Vector Machines
  - HOG Detector

Detection via Classification: Main Idea

- Basic component: a binary classifier

Car/non-car Classifier

Yes, car.

No, not a car.

Slide credit: Kristen Grauman

Detection via Classification: Main Idea

- If the object may be in a cluttered scene, slide a window around looking for it.

Car/non-car Classifier

- Essentially, this is a brute-force approach with many local decisions.

Slide credit: Kristen Grauman

What *is* a Sliding Window Approach?

- Search over space and scale

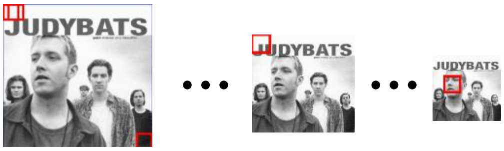

- Detection as subwindow classification problem
- *"In the absence of a more intelligent strategy, any global image classification approach can be converted into a localization approach by using a sliding-window search."*

Detection via Classification: Main Idea

Fleshing out this pipeline a bit more, we need to:

1. Obtain training data
2. Define features
3. Define classifier

Training examples

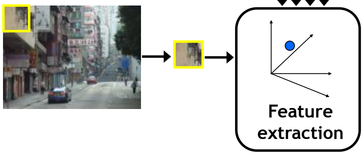

Car/non-car Classifier

Slide credit: Kristen Grauman

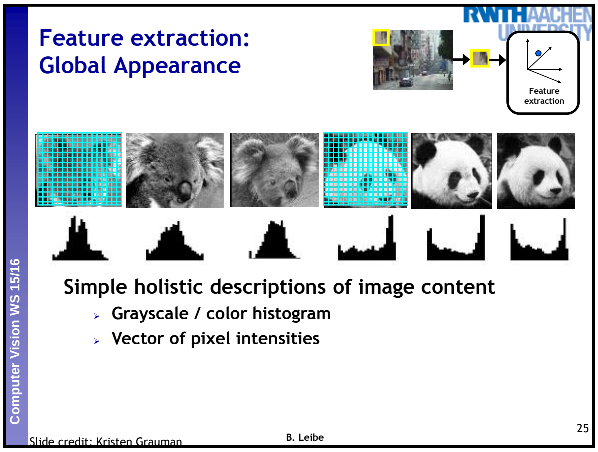

Feature extraction: Global Appearance

Simple holistic descriptions of image content
- Grayscale / color histogram
- Vector of pixel intensities

Slide credit: Kristen Grauman

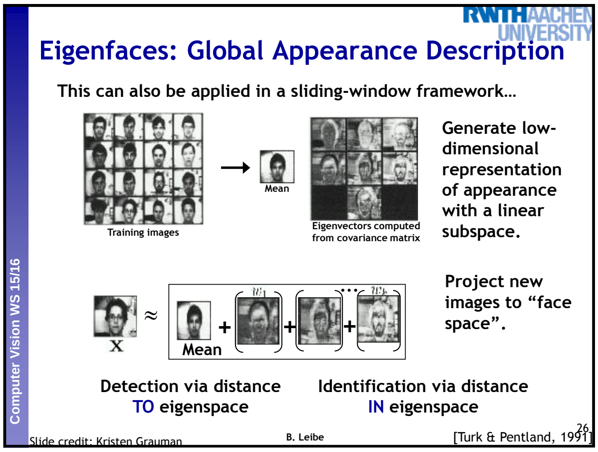

Eigenfaces: Global Appearance Description

This can also be applied in a sliding-window framework…

Generate low-dimensional representation of appearance with a linear subspace.

Project new images to “face space”.

Detection via distance TO eigenspace

Identification via distance IN eigenspace

Slide credit: Kristen Grauman

[Turk & Pentland, 1991]

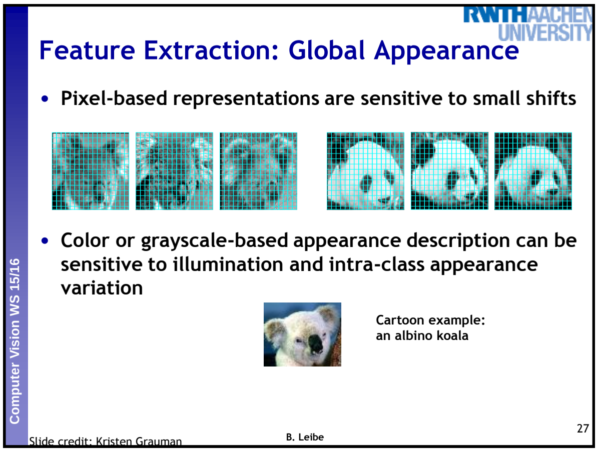

Feature Extraction: Global Appearance

- Pixel-based representations are sensitive to small shifts
- Color or grayscale-based appearance description can be sensitive to illumination and intra-class appearance variation

Cartoon example: an albino koala

Slide credit: Kristen Grauman

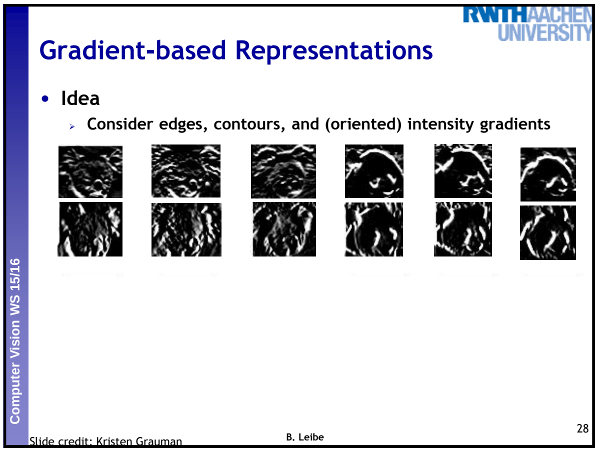

Gradient-based Representations

- Idea
  - Consider edges, contours, and (oriented) intensity gradients

Slide credit: Kristen Grauman

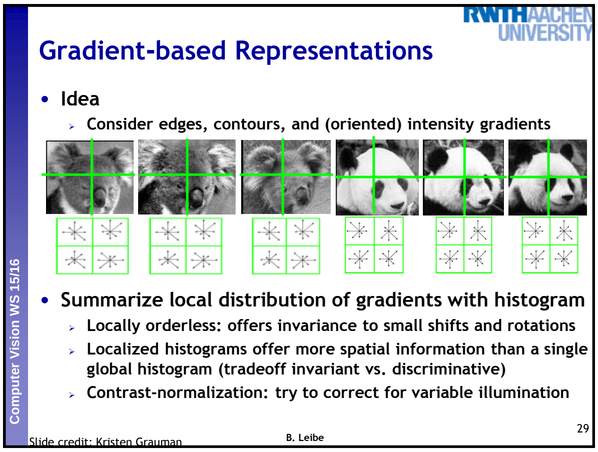

Gradient-based Representations

- Idea
  - Consider edges, contours, and (oriented) intensity gradients
- Summarize local distribution of gradients with histogram
  - Locally orderless: offers invariance to small shifts and rotations
  - Localized histograms offer more spatial information than a single global histogram (tradeoff invariant vs. discriminative)
  - Contrast-normalization: try to correct for variable illumination

Slide credit: Kristen Grauman

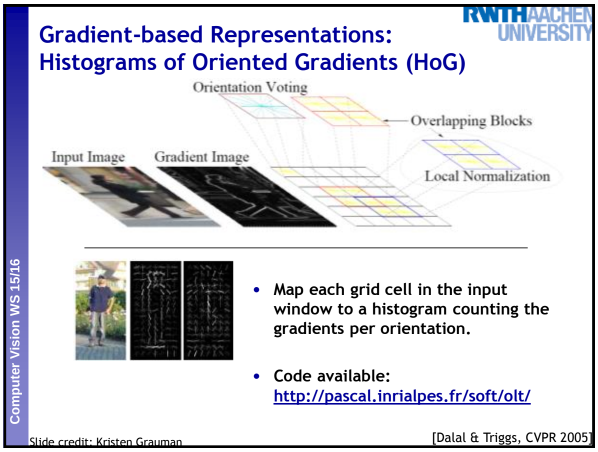

Gradient-based Representations: Histograms of Oriented Gradients (HoG)

- Map each grid cell in the input window to a histogram counting the gradients per orientation.
- Code available: http://pascal.inrialpes.fr/soft/olt/

Slide credit: Kristen Grauman

[Dalal & Triggs, CVPR 2005]

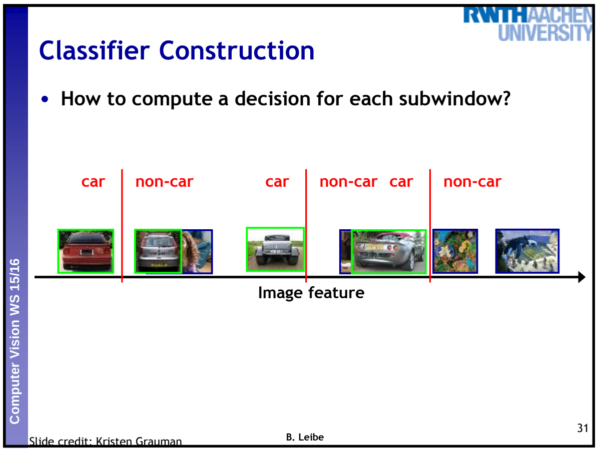

Classifier Construction

- How to compute a decision for each subwindow?

Slide credit: Kristen Grauman

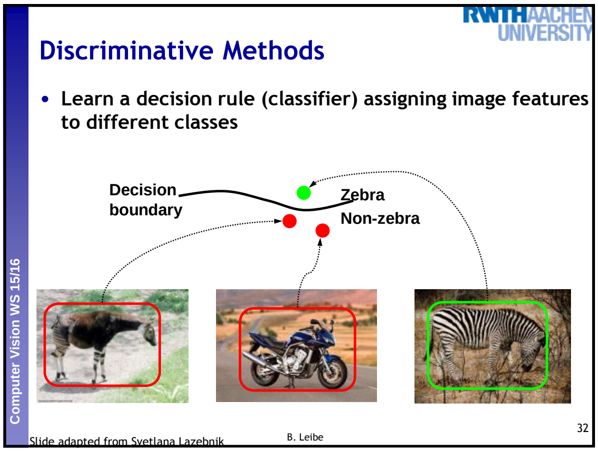

Discriminative Methods

- Learn a decision rule (classifier) assigning image features to different classes

Slide adapted from Svetlana Lazebnik

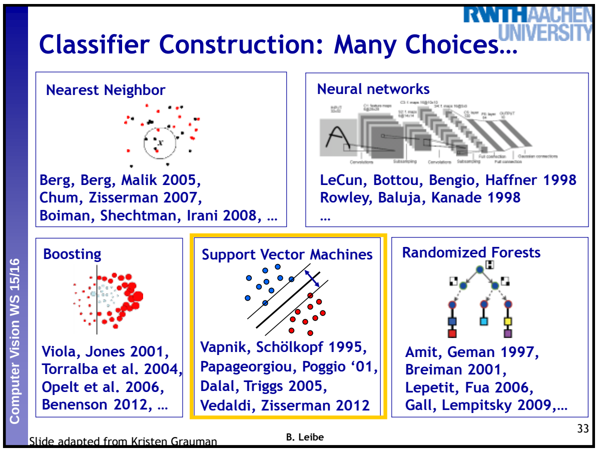

Classifier Construction: Many Choices…

Nearest Neighbor

Berg, Berg, Malik 2005,
Chum, Zisserman 2007,
Boiman, Shechtman, Irani 2008, …

Neural networks

LeCun, Bottou, Bengio, Haffner 1998
Rowley, Baluja, Kanade 1998
…

Boosting

Viola, Jones 2001,
Torralba et al. 2004,
Opelt et al. 2006,
Benenson 2012, …

Support Vector Machines

Vapnik, Schölkopf 1995,
Papageorgiou, Poggio '01,
Dalal, Triggs 2005,
Vedaldi, Zisserman 2012

Randomized Forests

Amit, Geman 1997,
Breiman 2001,
Lepetit, Fua 2006,
Gall, Lempitsky 2009,…

Slide adapted from Kristen Grauman

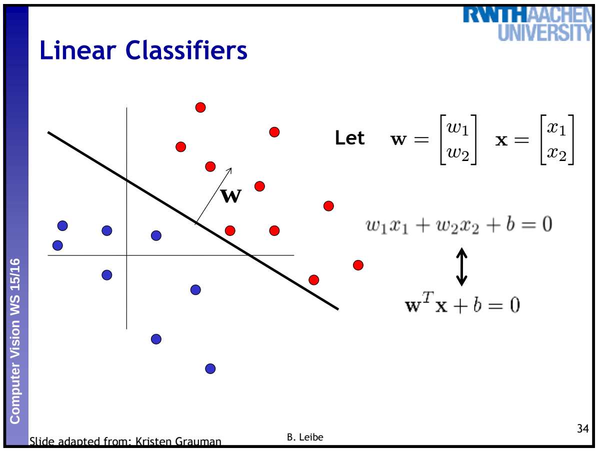

Linear Classifiers

Let $\mathbf{w} = \begin{bmatrix} w_1 \\ w_2 \end{bmatrix}$ $\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$

$$w_1 x_1 + w_2 x_2 + b = 0$$

$$\updownarrow$$

$$\mathbf{w}^T \mathbf{x} + b = 0$$

Slide adapted from: Kristen Grauman

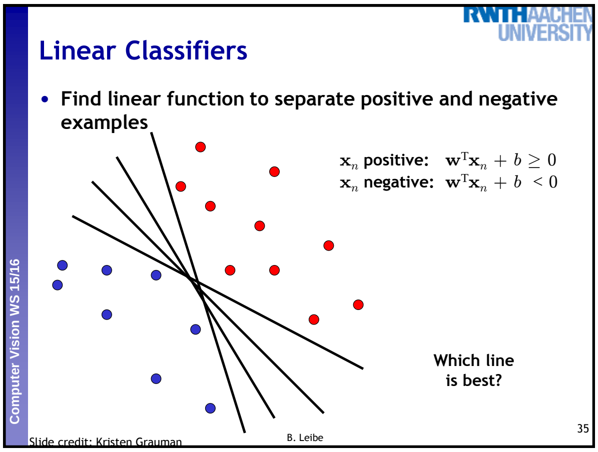

Linear Classifiers

- Find linear function to separate positive and negative examples

$\mathbf{x}_n$ positive: $\mathbf{w}^T\mathbf{x}_n + b \geq 0$

$\mathbf{x}_n$ negative: $\mathbf{w}^T\mathbf{x}_n + b < 0$

Which line is best?

Slide credit: Kristen Grauman

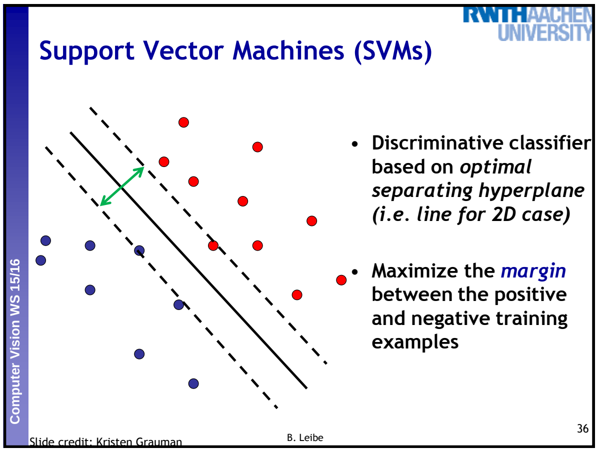

Support Vector Machines (SVMs)

- Discriminative classifier based on *optimal separating hyperplane (i.e. line for 2D case)*
- Maximize the *margin* between the positive and negative training examples

Slide credit: Kristen Grauman

Support Vector Machines

- Want line that maximizes the margin.

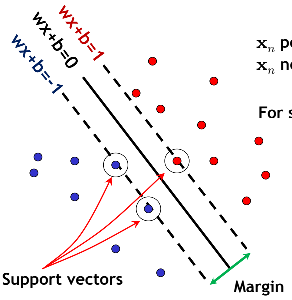

$\mathbf{x}_n$ positive ($t_n = 1$): $\mathbf{w}^\mathrm{T}\mathbf{x}_n + b \geq 1$
$\mathbf{x}_n$ negative ($t_n = -1$): $\mathbf{w}^\mathrm{T}\mathbf{x}_n + b < -1$

For support, vectors, $\mathbf{w}^T\mathbf{x}_n + b = \pm 1$

Quadratic optimization problem

Minimize $\frac{1}{2}\mathbf{w}^T\mathbf{w}$
Subject to $t_n(\mathbf{w}^T\mathbf{x}_n + b) \geq 1$

Packages available for that…

C. Burges, A Tutorial on Support Vector Machines for Pattern Recognition, Data Mining and Knowledge Discovery, 1998

Finding the Maximum Margin Line

- Solution: $\mathbf{w} = \sum_{n=1}^{N} a_n t_n \mathbf{x}_n$

Learned weight

Support vector

C. Burges, A Tutorial on Support Vector Machines for Pattern Recognition, Data Mining and Knowledge Discovery, 1998

Finding the Maximum Margin Line

- Solution: $\mathbf{w} = \sum_{n=1}^{N} a_n t_n \mathbf{x}_n$
- Classification function:

$$f(\mathbf{x}) = \mathrm{sign}(\mathbf{w}^T\mathbf{x} + b)$$
$$= \mathrm{sign}\left(\sum_{n=1}^{N} a_n t_n \boxed{\mathbf{x}_n^T\mathbf{x}} + b\right)$$

If f(x) < 0, classify as neg.,
if f(x) > 0, classify as pos.

  - Notice that this relies on an *inner product* between the test point $\mathbf{x}$ and the support vectors $\mathbf{x}_n$
  - (Solving the optimization problem also involves computing the inner products $\mathbf{x}_n^T\mathbf{x}_m$ between all pairs of training points)

C. Burges, A Tutorial on Support Vector Machines for Pattern Recognition, Data Mining and Knowledge Discovery, 1998

Questions

- What if the features are not 2d?
- What if the data is not linearly separable?
- What if we have more than just two categories?

Slide credit: Kristen Grauman

Questions

- What if the features are not 2d?
  - Generalizes to d-dimensions - replace line with “hyperplane”
- What if the data is not linearly separable?
- What if we have more than just two categories?

Questions

- What if the features are not 2d?
  - Generalizes to d-dimensions - replace line with “hyperplane”
- What if the data is not linearly separable?
  - Non-linear SVMs with special kernels
- What if we have more than just two categories?

Slide credit: Kristen Grauman

Non-Linear SVMs: Feature Spaces

- General idea: The original input space can be mapped to some higher-dimensional feature space where the training set is separable:

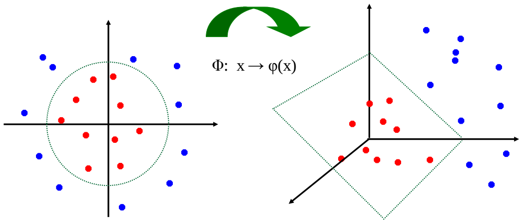

More on that in the Machine Learning lecture…

Slide from Andrew Moore's tutorial: http://www.autonlab.org/tutorials/svm.html

---

Nonlinear SVMs

- *The kernel trick*: instead of explicitly computing the lifting transformation $\varphi(\mathbf{x})$, define a kernel function K such that

$$K(\mathbf{x}_i,\mathbf{x}_j) = \varphi(\mathbf{x}_i) \cdot \varphi(\mathbf{x}_j)$$

- This gives a nonlinear decision boundary in the original feature space:

$$\sum_n a_n t_n K(\mathbf{x}_n, \mathbf{x}) + b$$

C. Burges, A Tutorial on Support Vector Machines for Pattern Recognition, Data Mining and Knowledge Discovery, 1998

Some Often-Used Kernel Functions

- Linear:

$$K(x_i,x_j)= x_i^T x_j$$

- Polynomial of power p:

$$K(x_i,x_j)= (1+ x_i^T x_j)^p$$

- Gaussian (Radial-Basis Function):

$$K(\mathbf{x}_i,\mathbf{x}_j) = \exp(-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{2\sigma^2})$$

Slide from Andrew Moore's tutorial: http://www.autonlab.org/tutorials/svm.html

---

Questions

- What if the features are not 2d?
  - Generalizes to d-dimensions - replace line with "hyperplane"
- What if the data is not linearly separable?
  - Non-linear SVMs with special kernels
- What if we have more than just two categories?

Slide credit: Kristen Grauman

Multi-Class SVMs

- Achieve multi-class classifier by combining a number of binary classifiers
- One vs. all
  - Training: learn an SVM for each class vs. the rest
  - Testing: apply each SVM to test example and assign to it the class of the SVM that returns the highest decision value
- One vs. one
  - Training: learn an SVM for each pair of classes
  - Testing: each learned SVM "votes" for a class to assign to the test example

Slide credit: Kristen Grauman

SVMs for Recognition

1. Define your representation for each example.
2. Select a kernel function.
3. Compute pairwise kernel values between labeled examples
4. Given this "kernel matrix" to SVM optimization software to identify support vectors & weights.
5. To classify a new example: compute kernel values between new input and support vectors, apply weights, check sign of output.

Slide credit: Kristen Grauman

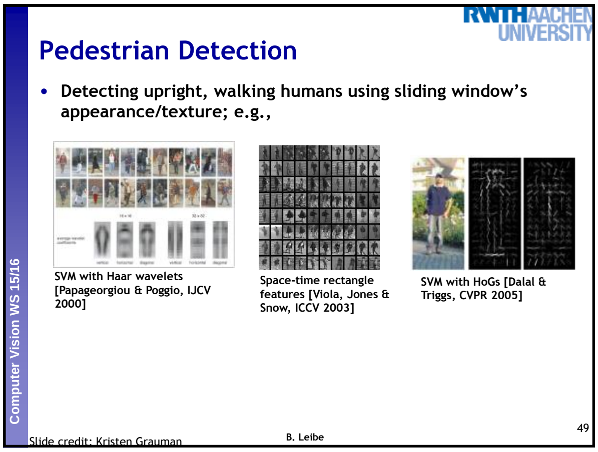

Pedestrian Detection

- Detecting upright, walking humans using sliding window's appearance/texture; e.g.,

SVM with Haar wavelets [Papageorgiou & Poggio, IJCV 2000]

Space-time rectangle features [Viola, Jones & Snow, ICCV 2003]

SVM with HoGs [Dalal & Triggs, CVPR 2005]

Slide credit: Kristen Grauman

B. Leibe

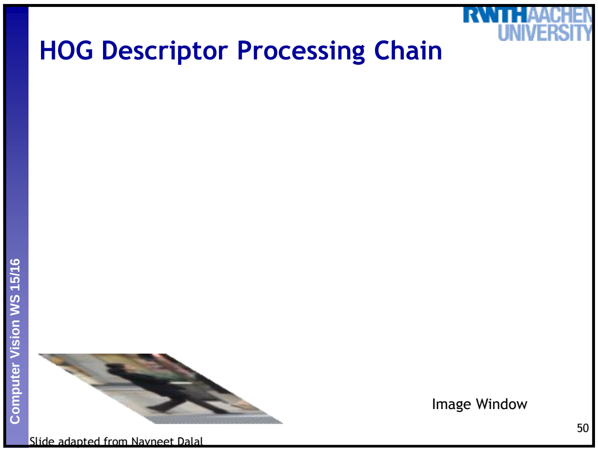

HOG Descriptor Processing Chain

Slide adapted from Navneet Dalal

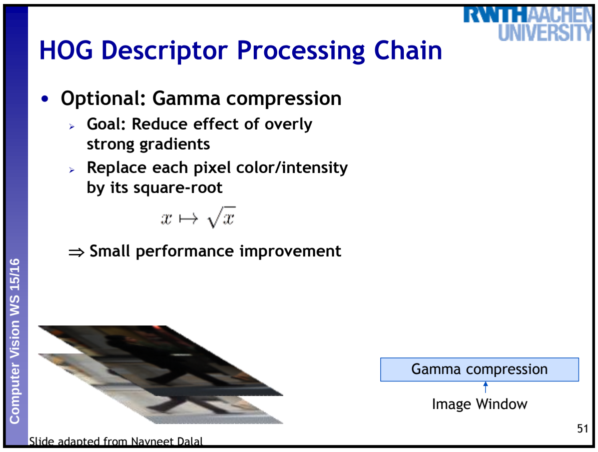

HOG Descriptor Processing Chain

- Optional: Gamma compression
  - Goal: Reduce effect of overly strong gradients
  - Replace each pixel color/intensity by its square-root

$$x \mapsto \sqrt{x}$$

⇒ Small performance improvement

Slide adapted from Navneet Dalal

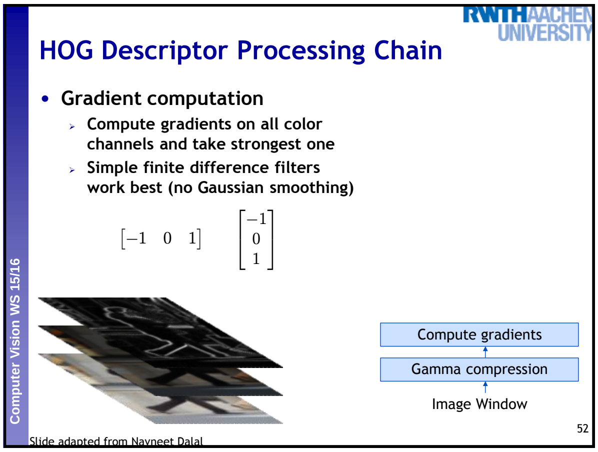

HOG Descriptor Processing Chain

- Gradient computation
  - Compute gradients on all color channels and take strongest one
  - Simple finite difference filters work best (no Gaussian smoothing)

$$\begin{bmatrix} -1 & 0 & 1 \end{bmatrix} \qquad \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix}$$

Slide adapted from Navneet Dalal

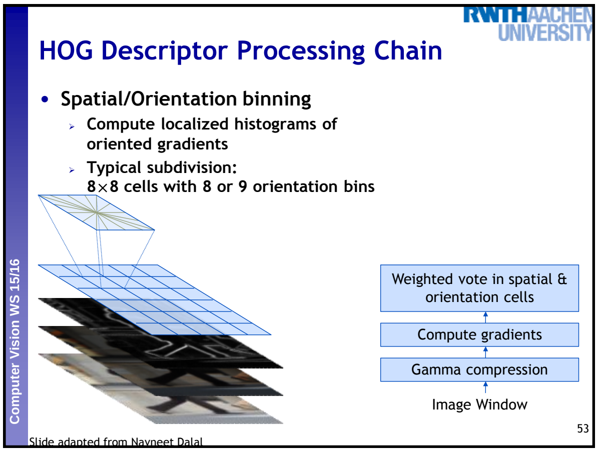

HOG Descriptor Processing Chain

- Spatial/Orientation binning
  - Compute localized histograms of oriented gradients
  - Typical subdivision: 8×8 cells with 8 or 9 orientation bins

Slide adapted from Navneet Dalal

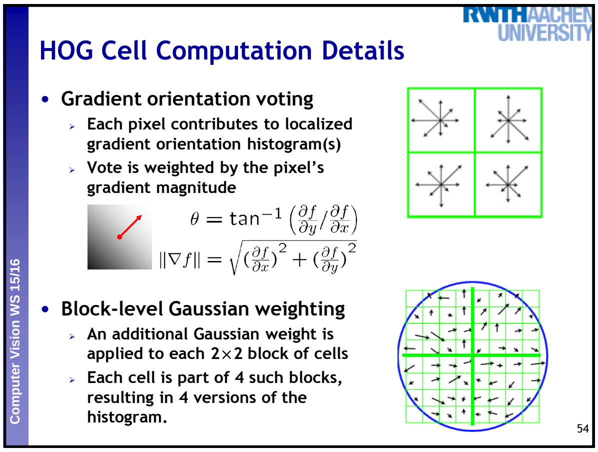

HOG Cell Computation Details

- Gradient orientation voting
  - Each pixel contributes to localized gradient orientation histogram(s)
  - Vote is weighted by the pixel's gradient magnitude

$$\theta = \tan^{-1}\left(\frac{\partial f}{\partial y} \Big/ \frac{\partial f}{\partial x}\right)$$

$$\|\nabla f\| = \sqrt{\left(\frac{\partial f}{\partial x}\right)^2 + \left(\frac{\partial f}{\partial y}\right)^2}$$

- Block-level Gaussian weighting
  - An additional Gaussian weight is applied to each 2×2 block of cells
  - Each cell is part of 4 such blocks, resulting in 4 versions of the histogram.

HOG Cell Computation Details (2)

- Important for robustness: Tri-linear interpolation
  - Each pixel contributes to (up to) 4 neighboring cell histograms
  - Weights are obtained by bilinear interpolation in image space:

$$h(x_1,y_1) \leftarrow w \cdot \left(1 - \frac{x-x_1}{x_2-x_1}\right)\left(1 - \frac{y-y_1}{y_2-y_1}\right)$$

$$h(x_1,y_2) \leftarrow w \cdot \left(1 - \frac{x-x_1}{x_2-x_1}\right)\left(\frac{y-y_1}{y_2-y_1}\right)$$

$$h(x_2,y_1) \leftarrow w \cdot \left(\frac{x-x_1}{x_2-x_1}\right)\left(1 - \frac{y-y_1}{y_2-y_1}\right)$$

$$h(x_2,y_2) \leftarrow w \cdot \left(\frac{x-x_1}{x_2-x_1}\right)\left(\frac{y-y_1}{y_2-y_1}\right)$$

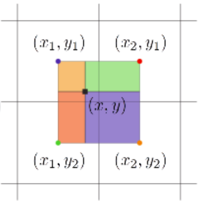

  - Contribution is further split over (up to) 2 neighboring orientation bins via linear interpolation over angles.

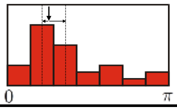

HOG Descriptor Processing Chain

- 2-Stage contrast normalization
  - L2 normalization, clipping, L2 normalization

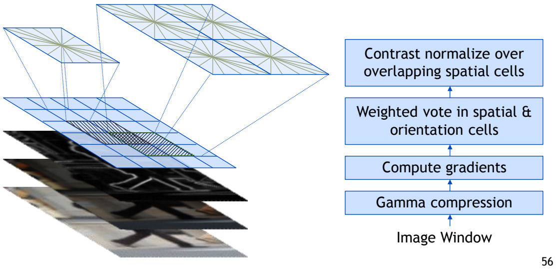

Contrast normalize over overlapping spatial cells

Weighted vote in spatial & orientation cells

Compute gradients

Gamma compression

Image Window

Slide adapted from Navneet Dalal

HOG Descriptor Processing Chain

- Feature vector construction
  - Collect HOG blocks into vector

[..., ..., ..., ...]

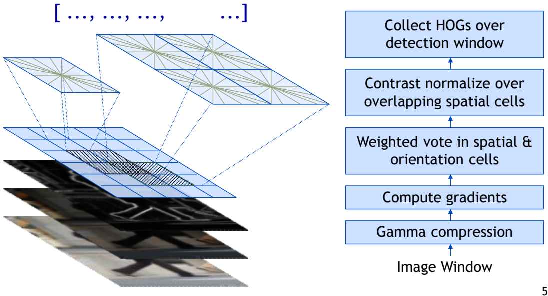

Collect HOGs over detection window

Contrast normalize over overlapping spatial cells

Weighted vote in spatial & orientation cells

Compute gradients

Gamma compression

Image Window

Slide adapted from Navneet Dalal

HOG Descriptor Processing Chain

- SVM Classification
  - Typically using a linear SVM

[..., ..., ..., ...]

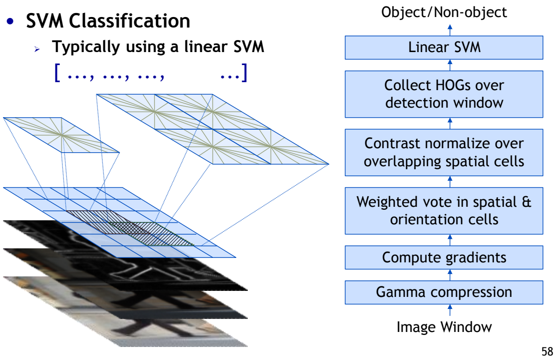

Object/Non-object

Linear SVM

Collect HOGs over detection window

Contrast normalize over overlapping spatial cells

Weighted vote in spatial & orientation cells

Compute gradients

Gamma compression

Image Window

Slide adapted from Navneet Dalal

Pedestrian Detection with HOG

- Train a pedestrian template using a linear SVM
- At test time, convolve feature map with template

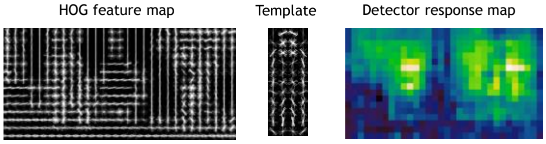

N. Dalal and B. Triggs, Histograms of Oriented Gradients for Human Detection, CVPR 2005

Slide credit: Svetlana Lazebnik

Non-Maximum Suppression

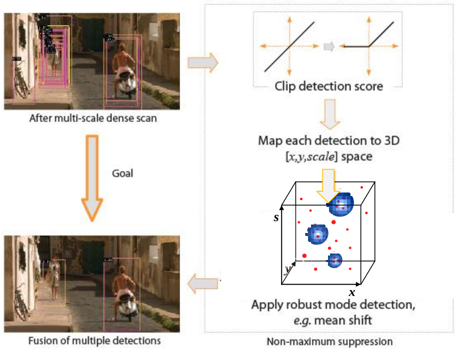

Image source: Navneet Dalal, PhD Thesis

Pedestrian detection with HoGs & SVMs

- Navneet Dalal, Bill Triggs, Histograms of Oriented Gradients for Human Detection, CVPR 2005

Slide credit: Kristen Grauman

References and Further Reading

- Read the HOG paper
  - N. Dalal, B. Triggs,
    Histograms of Oriented Gradients for Human Detection,
    CVPR, 2005.
- HOG Detector
  - Code available: http://pascal.inrialpes.fr/soft/olt/