



## Today

- HMM examples
- Support vector machines (SVM)
  - Basic algorithm
  - Kernels
    - Structured input spaces: Pyramid match kernels
  - Multi-class
  - HOG + SVM for person detection
    - Visualizing a feature: Hoggles
- Evaluating an object detector

### Window-based models: Three case studies



Boosting + face detection

Viola & Jones



NN + scene Gist classification

e.g., Hays & Efros



SVM + person detection

e.g., Dalal & Triggs























































![](_page_15_Figure_2.jpeg)

![](_page_16_Figure_1.jpeg)

![](_page_16_Figure_2.jpeg)

# Questions

• What if the data is not linearly separable?

![](_page_17_Figure_3.jpeg)

Lana Lazebnik

![](_page_18_Figure_1.jpeg)

- HMM examples
- Support vector machines (SVM)
  - Basic algorithm
  - Kernels
    - Structured input spaces: Pyramid match kernels
  - Multi-class
  - HOG + SVM for person detection
    - Visualizing a feature: Hoggles
- Evaluating an object detector

![](_page_18_Figure_11.jpeg)

![](_page_19_Figure_1.jpeg)

#### Nonlinear SVMs

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

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

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

$$\sum_{i} \alpha_{i} y_{i} K(\mathbf{x}_{i}, \mathbf{x}) + b$$

# "Kernel trick": Example

2-dimensional vectors  $\mathbf{x} = [x_1 \ x_2]$ ; let  $K(\mathbf{x}_i, \mathbf{x}_j) = (1 + \mathbf{x}_i^T \mathbf{x}_j)^2$ Need to show that  $K(\mathbf{x}_i, \mathbf{x}_j) = \varphi(\mathbf{x}_i)^T \varphi(\mathbf{x}_j)$ :  $K(\mathbf{x}_i, \mathbf{x}_j) = (1 + \mathbf{x}_i^T \mathbf{x}_j)^2$ ,  $= 1 + x_{i1}^2 x_{j1}^2 + 2 x_{i1} x_{j1} x_{i2} x_{j2} + x_{i2}^2 x_{j2}^2 + 2 x_{i1} x_{j1} + 2 x_{i2} x_{j2}$   $= [1 \ x_{i1}^2 \ \sqrt{2} \ x_{i1} x_{i2} \ x_{i2}^2 \ \sqrt{2} x_{i1} \ \sqrt{2} x_{i2}]^T$   $[1 \ x_{j1}^2 \ \sqrt{2} \ x_{j1} x_{j2} \ x_{j2}^2 \ \sqrt{2} x_{j1} \ \sqrt{2} x_{j2}]$   $= \varphi(\mathbf{x}_i)^T \varphi(\mathbf{x}_j)$ , where  $\varphi(\mathbf{x}) = [1 \ x_1^2 \ \sqrt{2} \ x_1 x_2 \ x_2^2 \ \sqrt{2} x_1 \ \sqrt{2} x_2]$ 

![](_page_20_Figure_3.jpeg)

# SVMs for recognition

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

![](_page_21_Figure_7.jpeg)

![](_page_21_Figure_8.jpeg)

# Multi-class SVMs

• Achieve multi-class classifier by combining a number of binary classifiers

#### <u>One vs. all</u>

- 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

#### SVMs: Pros and cons

- Pros
  - · Kernel-based framework is very powerful, flexible
  - · Often a sparse set of support vectors compact at test time
  - Work very well in practice, even with small training sample sizes
- Cons
  - No "direct" multi-class SVM, must combine two-class SVMs
  - Can be tricky to select best kernel function for a problem
  - · Computation, memory
    - During training time, must compute matrix of kernel values for every pair of examples
    - Learning can take a very long time for large-scale problems

### Today

- HMM examples
- Support vector machines (SVM)
  - Basic algorithm
  - Kernels
    - Structured input spaces: Pyramid match kernels
  - Multi-class
  - HOG + SVM for person detection
    - Visualizing a feature: Hoggles
- Evaluating an object detector

![](_page_23_Picture_11.jpeg)

![](_page_24_Picture_1.jpeg)

# Person detection with HoG's & linear SVM's

![](_page_24_Picture_3.jpeg)

• Map each grid cell in the input window to a histogram counting the gradients per orientation.

• Train a linear SVM using training set of pedestrian vs. non-pedestrian windows.

Dalal & Triggs, CVPR 2005

![](_page_25_Picture_1.jpeg)

![](_page_25_Picture_2.jpeg)

![](_page_26_Picture_1.jpeg)

![](_page_26_Picture_2.jpeg)

![](_page_27_Picture_1.jpeg)

![](_page_27_Figure_2.jpeg)

![](_page_28_Picture_1.jpeg)

![](_page_28_Picture_2.jpeg)

![](_page_29_Picture_1.jpeg)

![](_page_29_Figure_2.jpeg)

![](_page_30_Figure_1.jpeg)

![](_page_30_Figure_2.jpeg)

![](_page_31_Figure_1.jpeg)

![](_page_31_Figure_2.jpeg)

Recall: Examples of kernel functions

• Linear:  $K(x_i, x_j) = x_i^T x_j$ 

• Gaussian RBF: 
$$K(x_i, x_j) = \exp(-\frac{\|x_i - x_j\|^2}{2\sigma^2})$$

Histogram intersection:

$$K(x_i, x_j) = \sum_k \min(x_i(k), x_j(k))$$

- · Kernels go beyond vector space data
- Kernels also exist for "structured" input spaces like sets, graphs, trees...

![](_page_32_Picture_8.jpeg)

![](_page_33_Picture_1.jpeg)

![](_page_33_Figure_2.jpeg)

![](_page_34_Figure_1.jpeg)

![](_page_34_Figure_2.jpeg)

![](_page_35_Figure_1.jpeg)

![](_page_35_Figure_2.jpeg)

![](_page_36_Figure_1.jpeg)

![](_page_36_Figure_2.jpeg)

![](_page_37_Figure_1.jpeg)

![](_page_37_Picture_2.jpeg)

#### Summary: This week

- Object recognition as classification task
  - Boosting (face detection ex)
  - Support vector machines and HOG (person detection ex)
    - Pyramid match kernels
    - Hoggles visualization for understanding classifier mistakes
  - Nearest neighbors and global descriptors (scene rec ex)
- Sliding window search paradigm
  - Pros and cons
  - Speed up with attentional cascade
  - Object proposals as alternative to exhaustive search
- HMM examples
- Evaluation
  - Detectors: Intersection over union, precision recall
  - Classifiers: Confusion matrix