CS 188: Artificial Intelligence

Lecture 24: Computer Vision

Pieter Abbeel – UC Berkeley
Slides adapted from Trevor Darrell (and his sources)

Rough evolution of focus in recognition research

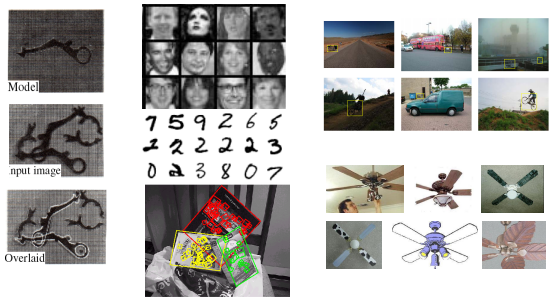

1980s | 1990s to early 2000s | 2000-2010…

Visual Object Recognition Tutorial

Inputs/outputs/assumptions

- What is the **goal**?
  - Say yes/no as to whether an object present in image

  And/or:
  - Determine pose of an object, e.g. for robot to grasp
  - Categorize all objects
  - Forced choice from pool of categories
  - Bounding box on object
  - Full segmentation
  - Build a model of an object category

Scanning windows…

Detection via classification: Main idea

Basic component: a binary classifier

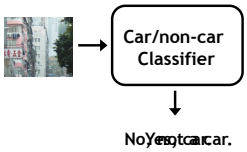

K. Grauman, B. Leibe

Visual Object Recognition Tutorial

Detection via classification: Main idea

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

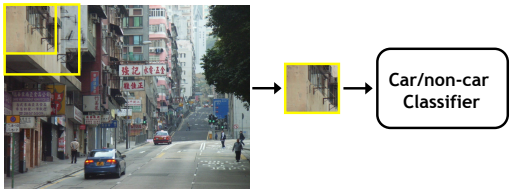

K. Grauman, B. Leibe

Visual Object Recognition Tutorial

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

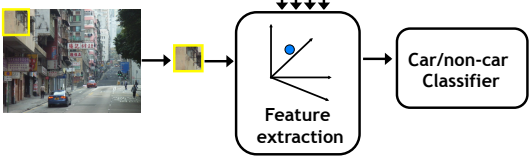

Visual Object Recognition Tutorial

K. Grauman, B. Leibe

Detection via classification: Main idea

- Consider all subwindows in an image
  - Sample at multiple scales and positions (and orientations)
- Make a decision per window:
  - "Does this contain object category X or not?"

Visual Object Recognition Tutorial

K. Grauman, B. Leibe

8

Feature extraction: global appearance

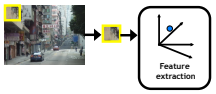

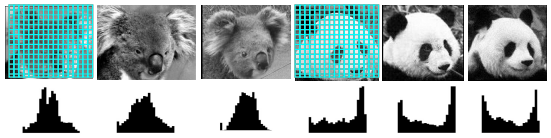

Simple holistic descriptions of image content

- grayscale / color histogram
- vector of pixel intensities

Visual Object Recognition Tutorial

K. Grauman, B. Leibe

Eigenfaces: global appearance description

An early appearance-based approach to face recognition

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

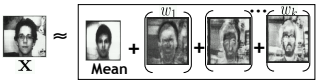

Project new images to "face space".

Recognition via nearest neighbors in face space

Turk & Pentland, 1991

Visual Object Recognition Tutorial

K. Grauman, B. Leibe

Feature extraction: global appearance

- Pixel-based representations sensitive to small shifts

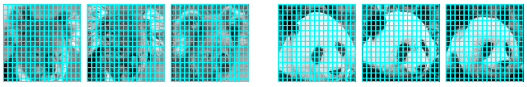

- Color or grayscale-based appearance description can be sensitive to illumination and intra-class appearance variation

Cartoon example: an albino koala

Visual Object Recognition Tutorial

K. Grauman, B. Leibe

Gradient-based representations

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

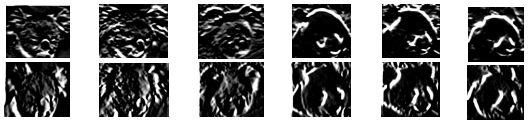

Visual Object Recognition Tutorial

K. Grauman, B. Leibe

HOG

(one of the most widely used features)

Gradient-based representations: Histograms of oriented gradients (HoG)

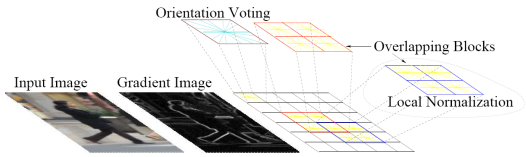

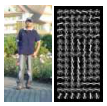

Map each grid cell in the input window to a histogram counting the gradients per orientation.

Code available: http://pascal.inrialpes.fr/soft/olt/

Dalal & Triggs, CVPR 2005

Slide credit: Dalal, Triggs, P. Barnum

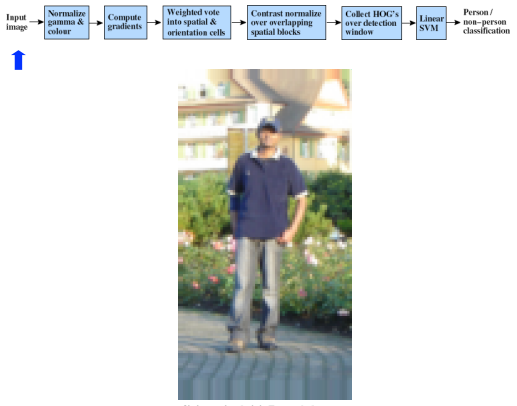

Slide credit: Dalal, Triggs, P. Barnum

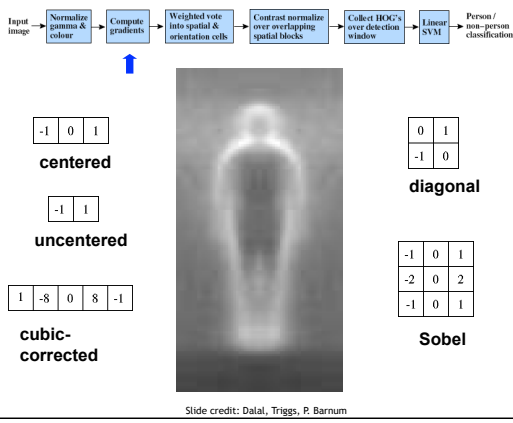

Slide credit: Dalal, Triggs, P. Barnum

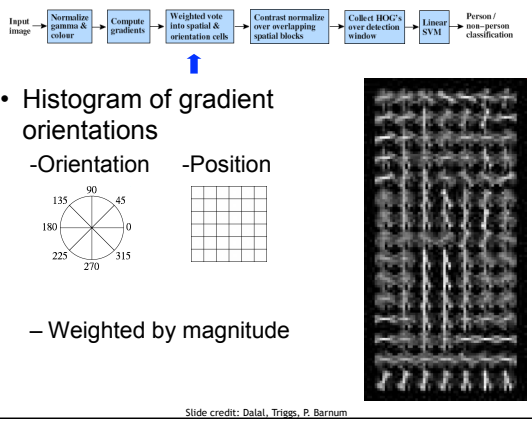

Slide credit: Dalal, Triggs, P. Barnum

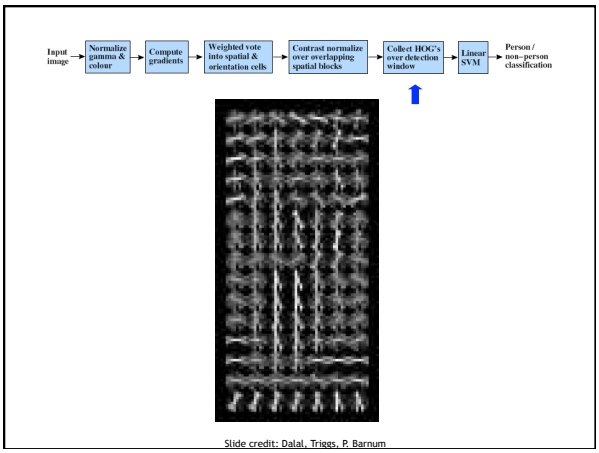

Slide credit: Dalal, Triggs, P. Barnum

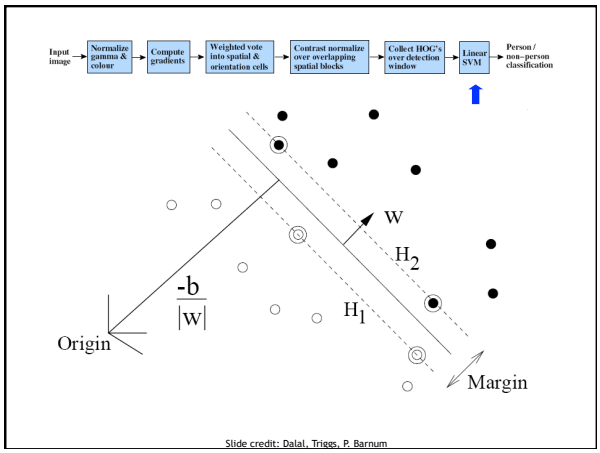

Slide credit: Dalal, Triggs, P. Barnum

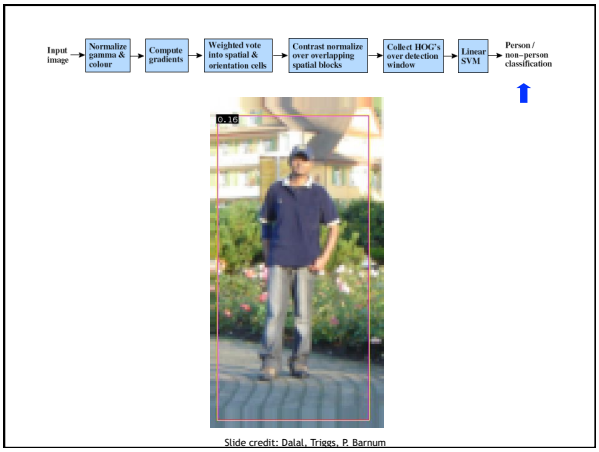

Slide credit: Dalal, Triggs, P. Barnum

Slide credit: Dalal, Triggs, P. Barnum