

EECS 442 Discussion

Arash Ushani

November 18, 2015

Announcements

- Project Progress Reports due 11/19

Midterm Student Feedback Session

- Thanks!
- Volume: I'll try to speak louder and repeat questions/answers
- Questions in discussion / explanation of content
- Solutions to MATLAB examples / questions on piazza
- More visuals / details in the slides


Descriptors for Images

- Why use a descriptor?



Histogram of Oriented Gradients

- Key idea: capture the local gradient pattern in image




Histogram of Oriented Gradients

- Key idea: capture the local gradient pattern in image




Compute Gradients

- Find X and Y gradient of image
- Convert this to polar (magnitude and direction of gradient)



Orientation Binning

- Divide image into “cells” and “blocks”
- Each cell contains a small number of pixels, and each block contains a small number of cells
- In each cell, compute a histogram of orientations



Orientation Binning


- Goal: Create a histogram of orientations for each cell
- What are the possible orientations?

Orientation Binning

- Goal: Create a histogram of orientations for each cell
- What are the possible orientations?
0 - 360 (or 0 - 180 if we ignore sign)

Orientation Binning

- Goal: Create a histogram of orientations for each cell
- What are the possible orientations?
0 - 360 (or 0 - 180 if we ignore sign)
- Divide 0 - 180 into N bins (9 is typical)
- Each pixel votes in the histogram according to the magnitude of its gradient
- To address discretization, each bin can vote for multiple pixels based on linear interpolation



Orientation Binning: Example

- Histogram from 0 to 180 with 9 bins (centered at 10, 30, 50, ...)
- Pixel with $\theta = 25, u = 10$

Orientation Binning: Example

- Histogram from 0 to 180 with 9 bins (centered at 10, 30, 50, ...)
- Pixel with $\theta = 25, u = 10$
- Which bins?

Orientation Binning: Example

- Histogram from 0 to 180 with 9 bins (centered at 10, 30, 50, ...)
- Pixel with $\theta = 25, u = 10$
- Which bins? bin 1 (centered at 10) and bin 2 (centered at 30)
- What are the votes for each bin?

Orientation Binning: Example

- Histogram from 0 to 180 with 9 bins (centered at 10, 30, 50, ...)
- Pixel with $\theta = 25$, $u = 10$
- Which bins? bin 1 (centered at 10) and bin 2 (centered at 30)
- What are the votes for each bin?

$$30 - \theta = 5$$

$$\frac{5}{180/N} = \frac{5}{20} = 0.25$$


$$0.25u = 2.5$$

$$0.75u = 7.5$$

- 2.5 to bin 1 (centered at 10), 7.5 to bin 2 (centered at 30)

Blocks

- Group cells into overlapping blocks
- For each 2x2 set of cells, concatenate the four cells histograms together into a single block feature b
- Normalize b



Final Descriptor

- Concatenate all block histograms together into a single descriptor h
- Normalize h
- Threshold all values of h so that all elements of h are below a threshold τ (0.2 is typical)
- Renormalize h
- Why?

Final Descriptor

- Concatenate all block histograms together into a single descriptor h
- Normalize h
- Threshold all values of h so that all elements of h are below a threshold τ (0.2 is typical)
- Renormalize h
- Why? To prevent strong gradients from dominating the whole image

MATLAB

- Go to CTools → Resources → Discussion → 11_18_matlab.zip
- Given an image, compute the HOG descriptor

