


# CS5670: Computer Vision

## Feature descriptors




# Reading


- Szeliski (1<sup>st</sup> edition): 4.1

# Local features: main components

1) Detection: Identify the interest points




**2) Description: Extract vector feature descriptor surrounding each interest point.**



$$\mathbf{x}_1 = [x_1^{(1)}, \dots, x_d^{(1)}]$$

$$\mathbf{x}_2 = [x_1^{(2)}, \dots, x_d^{(2)}]$$


3) Matching: Determine correspondence between descriptors in two views



# Feature descriptors

We know how to detect good points

Next question: **How to match them?**




**Answer:** Come up with a *descriptor* for each point, find similar descriptors between the two images

# Feature descriptors

We know how to detect good points

Next question: **How to match them?**



Lots of possibilities

- Simple option: match square windows around the point
- State of the art approach: SIFT
  - David Lowe, UBC <http://www.cs.ubc.ca/~lowe/keypoints/>


# Invariance vs. discriminability

- Invariance:
  - Descriptor shouldn't change even if image is transformed
- Discriminability:
  - Descriptor should be highly unique for each point

# Image transformations revisited

- Geometric

**Rotation**



**Scale**



- Photometric

**Intensity change**



# Invariant descriptors

- We looked at invariant / equivariant **detectors**
- Most feature descriptors are also designed to be invariant to:
  - Translation, 2D rotation, scale
- They can usually also handle
  - Limited 3D rotations (SIFT works up to about 60 degrees)
  - Limited affine transforms (some are fully affine invariant)
  - Limited illumination/contrast changes



# How to achieve invariance

Need both of the following:

1. Make sure your detector is invariant
2. Design an invariant feature descriptor
  - Simplest descriptor: a single 0
    - What's this invariant to?
  - Next simplest descriptor: a square, axis-aligned 5x5 window of pixels
    - What's this invariant to?
  - Let's look at some better approaches...

# Rotation invariance for feature descriptors

- Find dominant orientation of the image patch
  - E.g., given by  $\mathbf{x}_{\max}$ , the eigenvector of  $\mathbf{H}$  corresponding to  $\lambda_{\max}$  (the *larger* eigenvalue)
  - Or (better) simply the orientation of the (smoothed) gradient
  - Rotate the patch according to this angle




Figure by Matthew Brown


# Multiscale Oriented PatcheS descriptor

Take 40x40 square window around detected feature

- Scale to 1/5 size (using prefiltering)
- Rotate to horizontal
- Sample 8x8 square window centered at feature
- Intensity normalize the window by subtracting the mean, dividing by the standard deviation in the window (why?)



# Detections at multiple scales



*Figure 1. Multi-scale Oriented Patches (MOPS) extracted at five pyramid levels from one of the Matier images. The boxes show the feature orientation and the region from which the descriptor vector is sampled.*

# Scale Invariant Feature Transform

Basic idea:

- Take 16x16 square window around detected feature
- Compute edge orientation (angle of the gradient -  $90^\circ$ ) for each pixel
- Throw out weak edges (threshold gradient magnitude)
- Create histogram of surviving edge orientations






Image gradients




angle histogram



# SIFT descriptor

Full version

- Divide the 16x16 window into a 4x4 grid of cells (2x2 case shown below)
- Compute an orientation histogram for each cell
- 16 cells \* 8 orientations = 128 dimensional descriptor






# Properties of SIFT

Extraordinarily robust matching technique

- Can handle changes in viewpoint (up to about 60 degree out of plane rotation)
- Can handle significant changes in illumination (sometimes even day vs. night (below))
- Pretty fast—hard to make real-time, but can run in <1s for moderate image sizes
- Lots of code available



# SIFT Example




**868 SIFT features**



# Other descriptors

- HOG: Histogram of Gradients (HOG)
  - Dalal/Triggs
  - Sliding window, pedestrian detection





- FREAK: Fast Retina Keypoint
  - Perceptually motivated
  - Can run in real-time; used in Visual SLAM
- LIFT: Learned Invariant Feature Transform
  - Learned via deep learning

<https://arxiv.org/abs/1603.09114>


**Questions?**

# Summary

- Keypoint detection: repeatable and distinctive
  - Corners, blobs, stable regions
  - Harris, DoG
- Descriptors: robust and selective
  - spatial histograms of orientation
  - SIFT and variants are typically good for stitching and recognition
  - But, need not stick to one



# Which features match?



**Questions?**