CS540 Introduction to Artificial Intelligence
Lecture 7

Young Wu
Based on lecture slides by Jerry Zhu, Yingyu Liang, and Charles Dyer

June 8, 2020

Socrative Test

Admin

- Socrative Student Login: Room CS540E.
- Use Socrative Room CS540 (without the E) for anonymous feedback.
- A: I haven't started P1.
- B: I have started P1.
- C: I have finished part 1.
- D: I have finished P1.
- E: What is P1?

Computer Vision Examples, Part I

Motivation

- Image segmentation
- Image retrieval
- Image colorization
- Image reconstruction
- Image super-resolution
- Image synthesis
- Image captioning

Computer Vision Examples, Part II

Motivation

- Style transfer
- Object tracking
- Visual question answering
- Human pose estimation
- Medical image analysis

Image Features Diagram

Motivation

One Dimensional Convolution

Definition

- The convolution of a vector $x = (x_1, x_2, ..., x_m)$ with a filter $w = (w_{-k}, w_{-k+1}, ... w_{k-1}, w_k)$ is:

$$a = (a_1, a_2, ..., a_m) = x * w$$

$$a_j = \sum_{t=-k}^{k} w_t x_{j-t}, j = 1, 2, ..., m$$

$\overleftrightarrow{w \cdot x}$

- w is also called a kernel (different from the kernel for SVMs).
- The elements that do not exist are assumed to be 0.

Two Dimensional Convolution

Definition

- The convolution of an $m \times m$ matrix X with a $(2k+1) \times (2k+1)$ filter W is:

$$A = X * W$$

$$A_{j,j'} = \sum_{s=-k}^{k} \sum_{t=-k}^{k} W_{s,t} X_{j-s,j'-t}, j, j' = 1, 2, ..., m$$

$W \cdot X$

- The matrix W is indexed by (s, t) for $s = -k, -k+1, ..., k-1, k$ and $t = -k, -k+1, ..., k-1, k$.
- The elements that do not exist are assumed to be 0.

Convolution Diagram

Definition

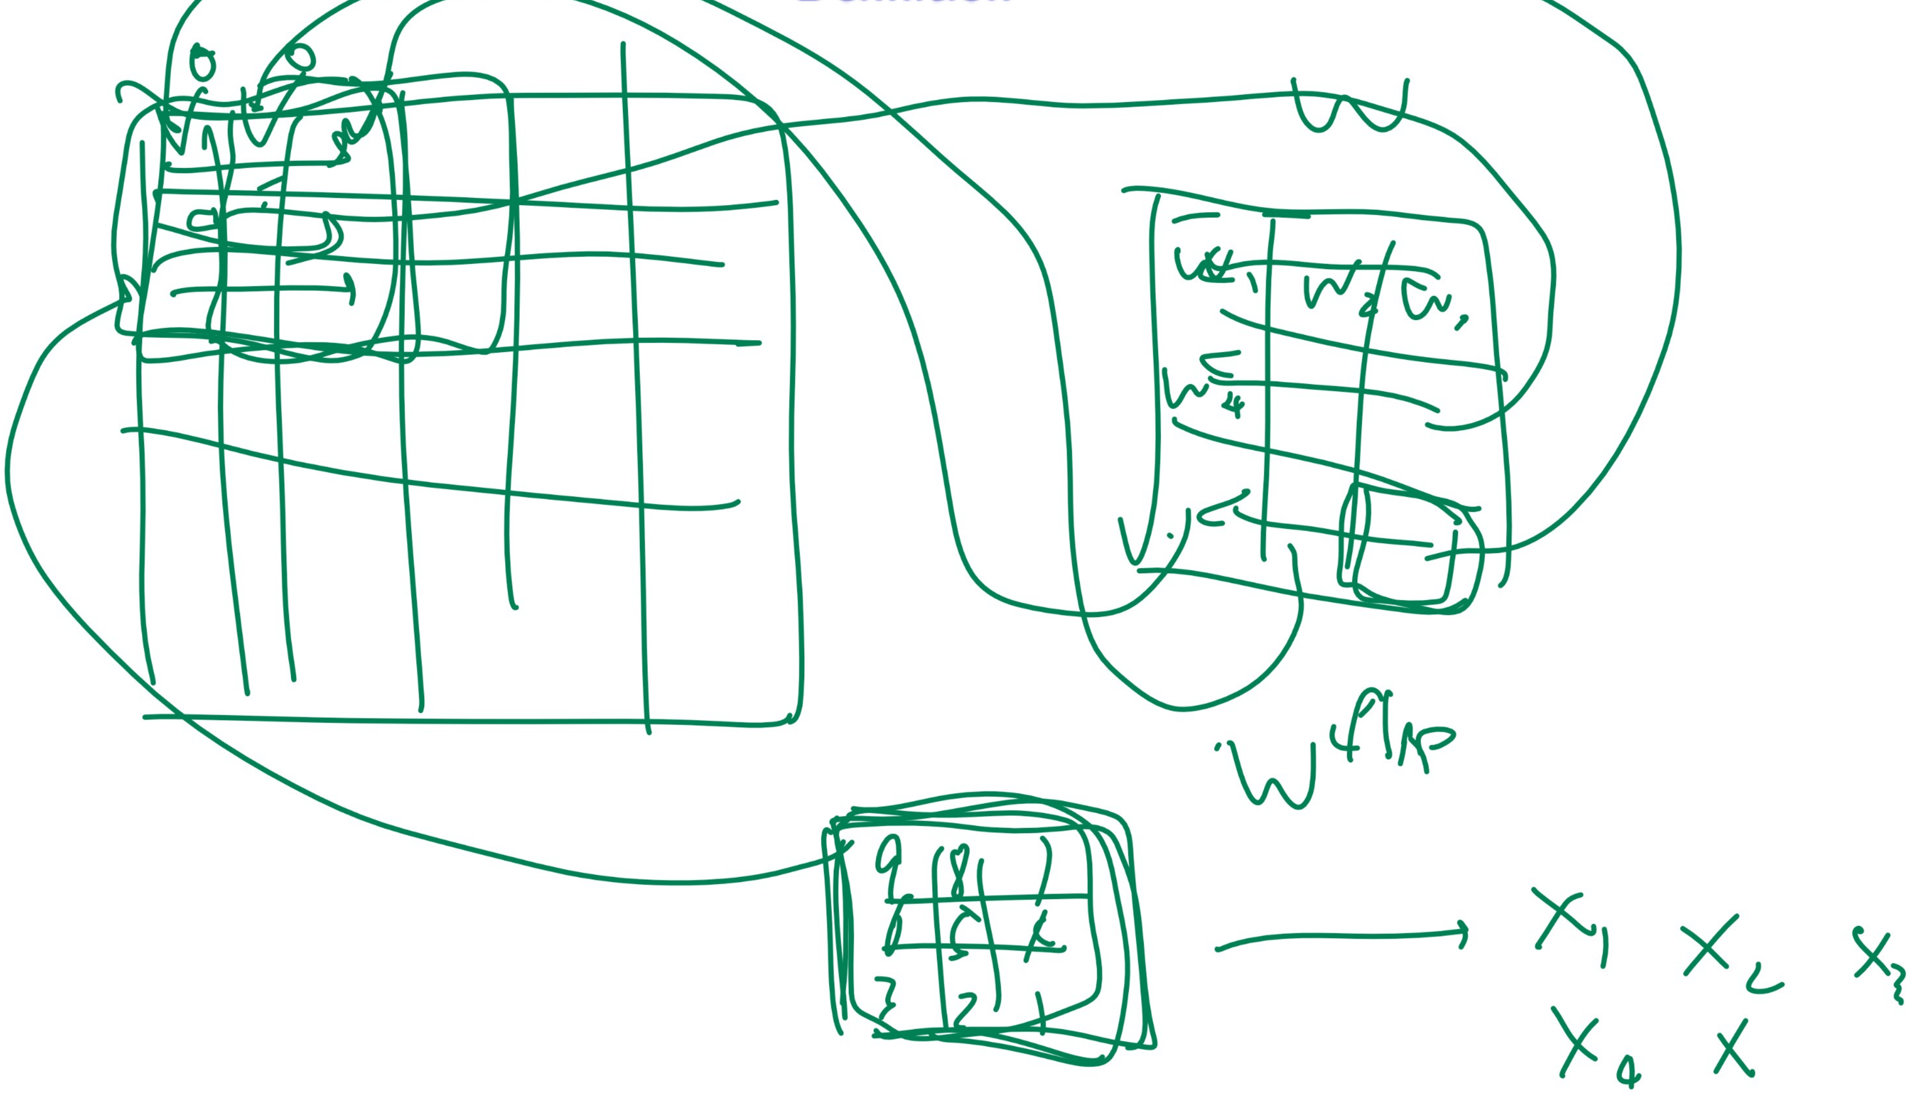

Image Gradient

Definition

- The gradient of an image is defined as the change in pixel intensity due to the change in the location of the pixel.

$$\begin{cases} \dfrac{\partial I(s,t)}{\partial s} \approx \dfrac{I\left(s+\frac{\varepsilon}{2},t\right)-I\left(s-\frac{\varepsilon}{2},t\right)}{\varepsilon}, \varepsilon=1 \\[2ex] \dfrac{\partial I(s,t)}{\partial t} \approx \dfrac{I\left(s,t+\frac{\varepsilon}{2}\right)-I\left(s,t-\frac{\varepsilon}{2}\right)}{\varepsilon}, \varepsilon=1 \end{cases}$$

Image Derivative Filters

Definition

- The gradient can be computed using convolution with the following filters.

$$w_x = \begin{bmatrix} -1 & 0 & 1 \end{bmatrix}, w_y = \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix}$$

1	2	3
4	$\boxed{5}$	6
7	8	9

$\nabla_x I_5 = -2$

$\nabla_y I_5 = -6$

Sobel Filter

Definition

- The Sobel filters also are used to approximate the gradient of an image.

$$W_x = \begin{array}{c} \begin{matrix} -1 & 0 & 1 \end{matrix} \\ \boxed{\begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}} \end{array}, W_y = \boxed{\begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix}}$$

$$\nabla_x I = W_x * I$$

$$\nabla_y I = W_y * I$$

Gradient of Images

Definition

- The gradient of an image I is $(\nabla_x I, \nabla_y I)$.

$$\underline{\nabla_x I} = W_x * I, \underline{\nabla_y I} = W_y * I$$

- The gradient magnitude is G and gradient direction Θ are the following.

either horizontal or verticle edges,

$$G = \sqrt{\nabla_x^2 + \nabla_y^2}$$

$$\Theta = \arctan\left(\frac{\nabla_y}{\nabla_x}\right)$$

Gradient of Images Demo

Definition

$\checkmark$

SIFT

Discussion

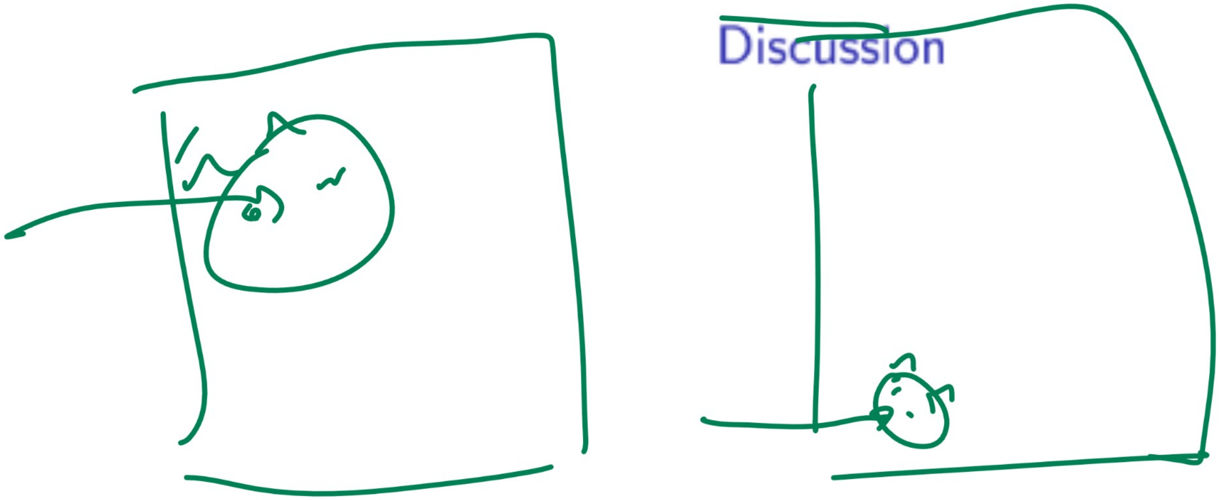

- Scale Invariant Feature Transform (SIFT) features are features that are invariant to changes in the location, scale, orientation, and lighting of the pixels.

Histogram Binning Diagram

Discussion

8 direct

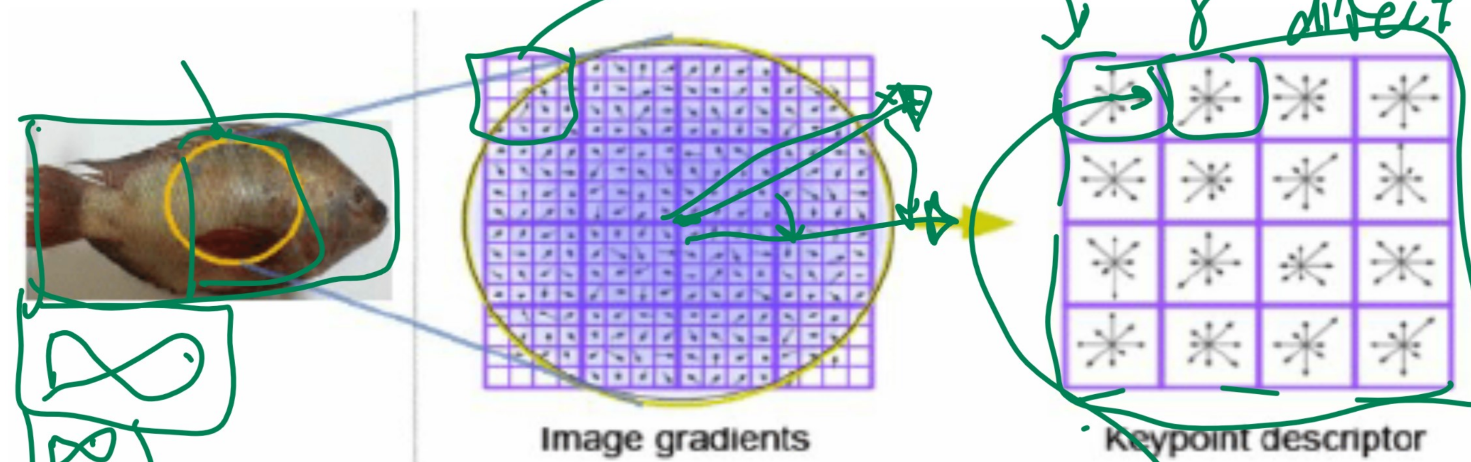

$(\nabla_x I, \nabla_y I)$

$\nabla_y I$

$\nabla_x I$

8 sum of G

16

128 features.

normalize sum 1

SIFT

HOG

Discussion

- Histogram of Oriented Gradients features is similar to SIFT but does not use dominant orientations.

directions

Matching vs Classification Diagram

Discussion