CS540 Introduction to Artificial Intelligence

Lecture 11

Young Wu

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

June 27, 2022

Coordination Game

Quiz

- You are not allowed to discuss anything about this question in chat. There will be around 10 new questions on the midterm exam. I will post n of them before the exam (this weekend):
- $A : n = 0$.
- $B : n = 1$ if more than 50 percent of you choose B.
- $C : n = 2$ if more than 75 percent of you choose C.
- $D : n = 3$ if more than 95 percent of you choose D.
- $E : n = 0$.
- I will repeat this question a second time. If you fail to coordinate both times, I will not post any of the new questions.

Q2

Exam Date

Admin

Midterm

Q 1

Monday

Alt

- July 11 from 5 : 30 to 8 : 30.
- July 27 (online only, with the other section) from 5 : 30 to 8 : 30.

Wed

- A : I will be available on July 11.
- B : I will be available on July 27.
- C : I am not available on both dates (email me).

Exam Format

Admin

- Similar to $M2$ to $M7, X1$ to $X3$, total of 30 questions.
- No hints, auto-grading will be turned off.
- There could be minor changes to the questions in $M2$ to $M7, X1$ to $X3$, please re-read the questions carefully.
- The last question will ask you for comments.

Exam Formula Sheet

Admin

- Formulas on $W1$ to $W7$ pages: post on Piazza if you want me to add any.
- You are allowed to implement some of these formulas in Excel or a programming language and use them during the exams.
- You are NOT allowed to work with another student.

Review Session

Admin

- July 6 from $5:30$ to $8:30$ on Zoom, go through past exam questions, answer questions.
- July 7 TA Office Hours on Zoom.
- Message me if you would like more office hours.

Image Features Diagram

Motivation

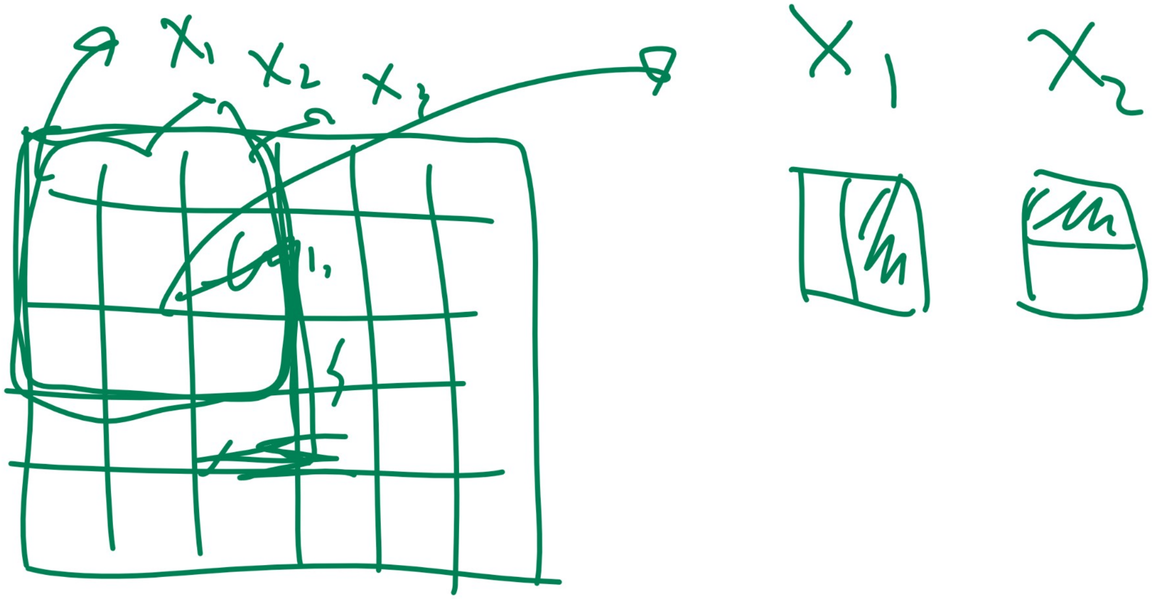

One Dimensional Convolution

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:

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

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

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

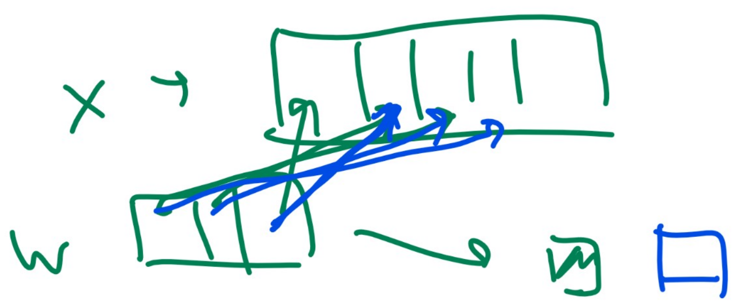

Two Dimensional Convolution

Definition

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

image

$$A = X * W$$

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

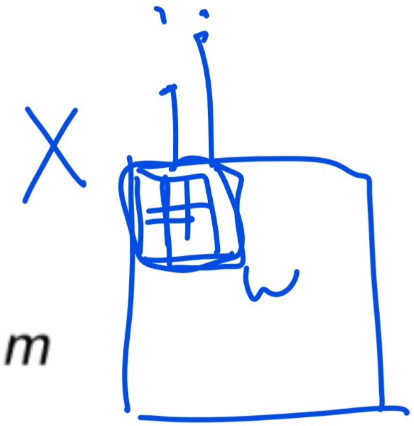

- The matrix W is indexed by (s, t) for $s = -k, -k+1, ..., k-1, k$ and $t = -k, -k+1, ..., k-1, k$.
- The elements that do not exist are assumed to be 0.

Convolution Diagram and Demo

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.

$$\left\{\begin{aligned} \frac{\partial I(s,t)}{\partial s} &\approx \frac{I\left(s+\frac{\varepsilon}{2},t\right)-I\left(s-\frac{\varepsilon}{2},t\right)}{\varepsilon}, \boxed{\varepsilon = 1} \\ \frac{\partial I(s,t)}{\partial t} &\approx \frac{I\left(s,t+\frac{\varepsilon}{2}\right)-I\left(s,t-\frac{\varepsilon}{2}\right)}{\varepsilon}, \varepsilon = 1 \end{aligned}\right.$$

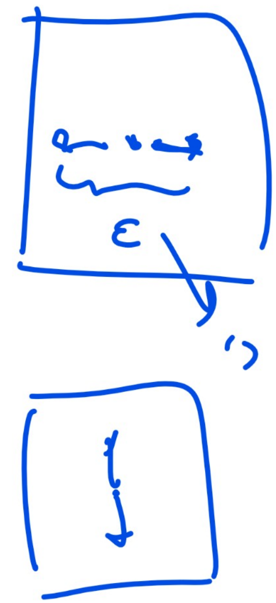

Image Derivative Filters

Definition

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

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

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

Sobel Filter

Definition

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

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

Gradient of Images

Definition

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

$$\nabla_x I = W_x * I, \nabla_y I = W_y * I$$

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

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

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

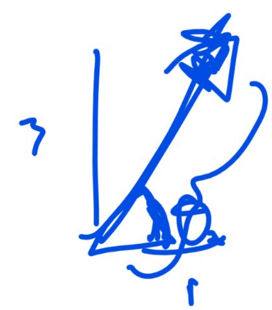

Gradient of Images Demo

Definition

Convolution Example

Quiz

- Find the gradient magnitude and direction for the center cell of the following image. Use the derivative filters $\begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix}$ and $\begin{bmatrix} -1 & 0 & 1 \end{bmatrix}$

$$\begin{bmatrix} 0 & 0 & 0 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

$\begin{pmatrix} -1 \\ 0 \end{pmatrix}$

mag = 1
dir = π

$\begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{pmatrix} \xrightarrow{\text{flip}} \begin{pmatrix} 9 & 8 & 7 \\ 6 & 5 & 4 \\ 3 & 2 & 1 \end{pmatrix}$

$\begin{bmatrix} 0 & -1 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$

flip

$\begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & -1 & 0 \end{bmatrix}$

∇_y

$\begin{bmatrix} -1 & -1 & \\ & & \\ & & \end{bmatrix}$

∇_x

$\begin{bmatrix} 0 & & \\ & & \\ & & \end{bmatrix}$

$\begin{bmatrix} 0 & 0 & 0 \\ -1 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix}$

flip

$\begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & -1 \\ 0 & 0 & 0 \end{bmatrix}$

zero padding

Gradient Example

Quiz

Convolution Example 1

Quiz

image filter Q3

$$\begin{bmatrix} 0 & 0 & 0 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} * \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}$$

flip

$$\begin{bmatrix} 1 & 0 & -1 \\ 2 & 0 & -2 \\ 1 & 0 & -1 \end{bmatrix}$$

center pixel of conved image

$= 0 \cdot 1 + 0 \cdot 0 + 0 \cdot -1$
$+ 1 \cdot 2 + 1 \cdot 0 + 1 \cdot -2$
$+ 1 \cdot 1 + 1 \cdot 0 + 1 \cdot -1$

- $A: \begin{bmatrix} -1 & -3 & -3 \\ 0 & 0 & 0 \\ 1 & 3 & 3 \end{bmatrix}$, $B: \begin{bmatrix} -3 & -3 & 3 \\ -4 & -4 & 4 \\ -3 & -3 & 3 \end{bmatrix}$
- $C: \begin{bmatrix} -3 & -4 & -3 \\ -3 & -4 & -3 \\ 3 & 4 & 3 \end{bmatrix}$, $D: \begin{bmatrix} -1 & 0 & 1 \\ -3 & 0 & 3 \\ -3 & 0 & 3 \end{bmatrix}$
- E : I don't understand

Convolution Example 2

Quiz

$$\begin{bmatrix} 0 & 0 & 0 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} * \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix}$$

- $A : \begin{bmatrix} -1 & -3 & -3 \\ 0 & 0 & 0 \\ 1 & 3 & 3 \end{bmatrix}, B : \begin{bmatrix} -3 & -3 & 3 \\ -4 & -4 & 4 \\ -3 & -3 & 3 \end{bmatrix}$
- $C : \begin{bmatrix} -3 & -4 & -3 \\ -3 & -4 & -3 \\ 3 & 4 & 3 \end{bmatrix}, D : \begin{bmatrix} -1 & 0 & 1 \\ -3 & 0 & 3 \\ -3 & 0 & 3 \end{bmatrix}$

SIFT

Discussion

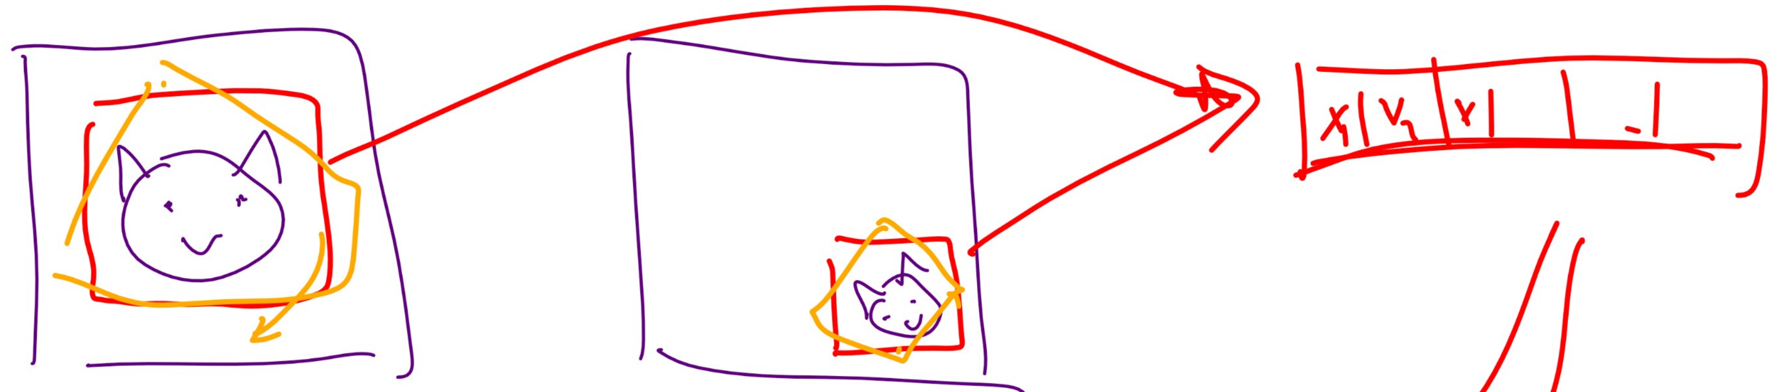

- Scale Invariant Feature Transform (SIFT) features are features that are invariant to changes in the location, scale, orientation, and lighting of the pixels.

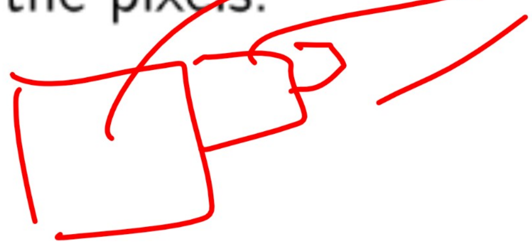

Histogram Binning Diagram

Discussion

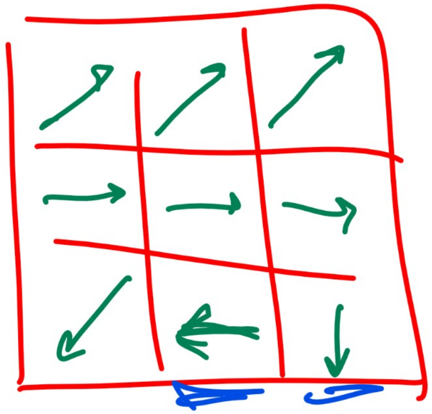

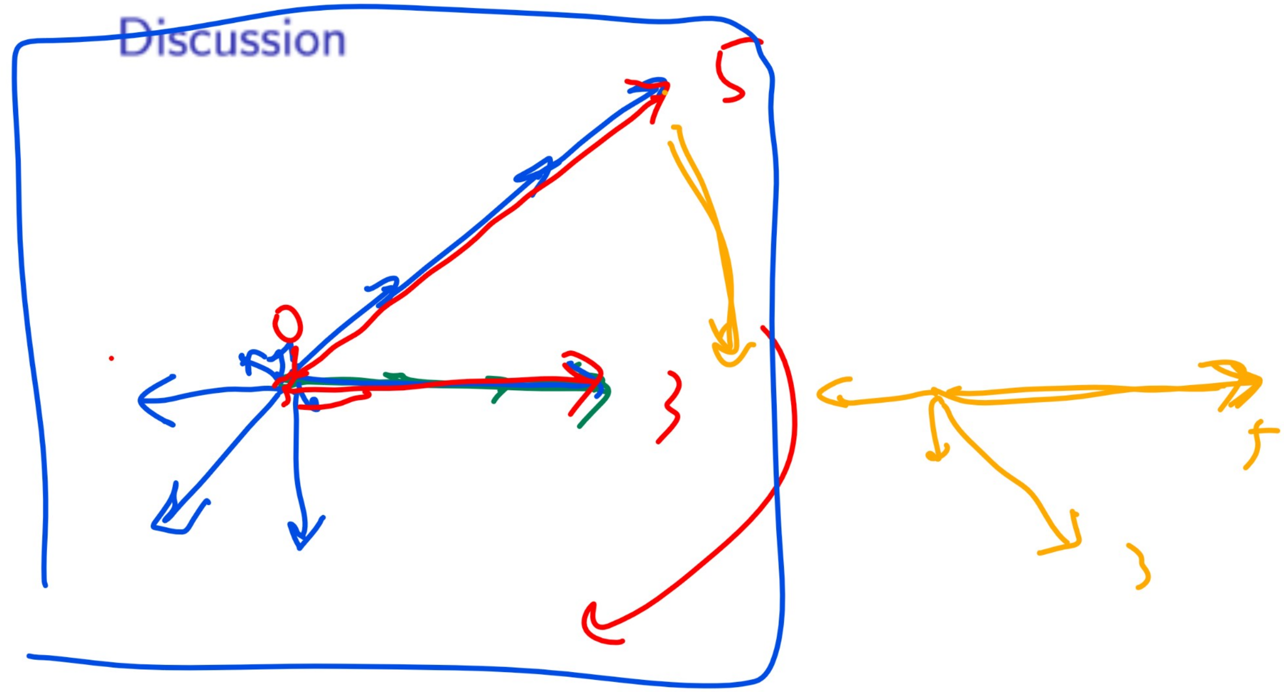

0, 5, 3, 0, . . .

HOG

Discussion

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

Classification

Discussion

- SIFT features are not often used in training classifiers and more often used to match the objects in multiple images.
- HOG features are usually computed for every cell in the image and used as features (in place of pixel intensities) in classification algorithms such as SVM.

6:50