CS540 Introduction to Artificial Intelligence

Lecture 11

Young Wu

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

July 7, 2022

Image Features Diagram

Motivation

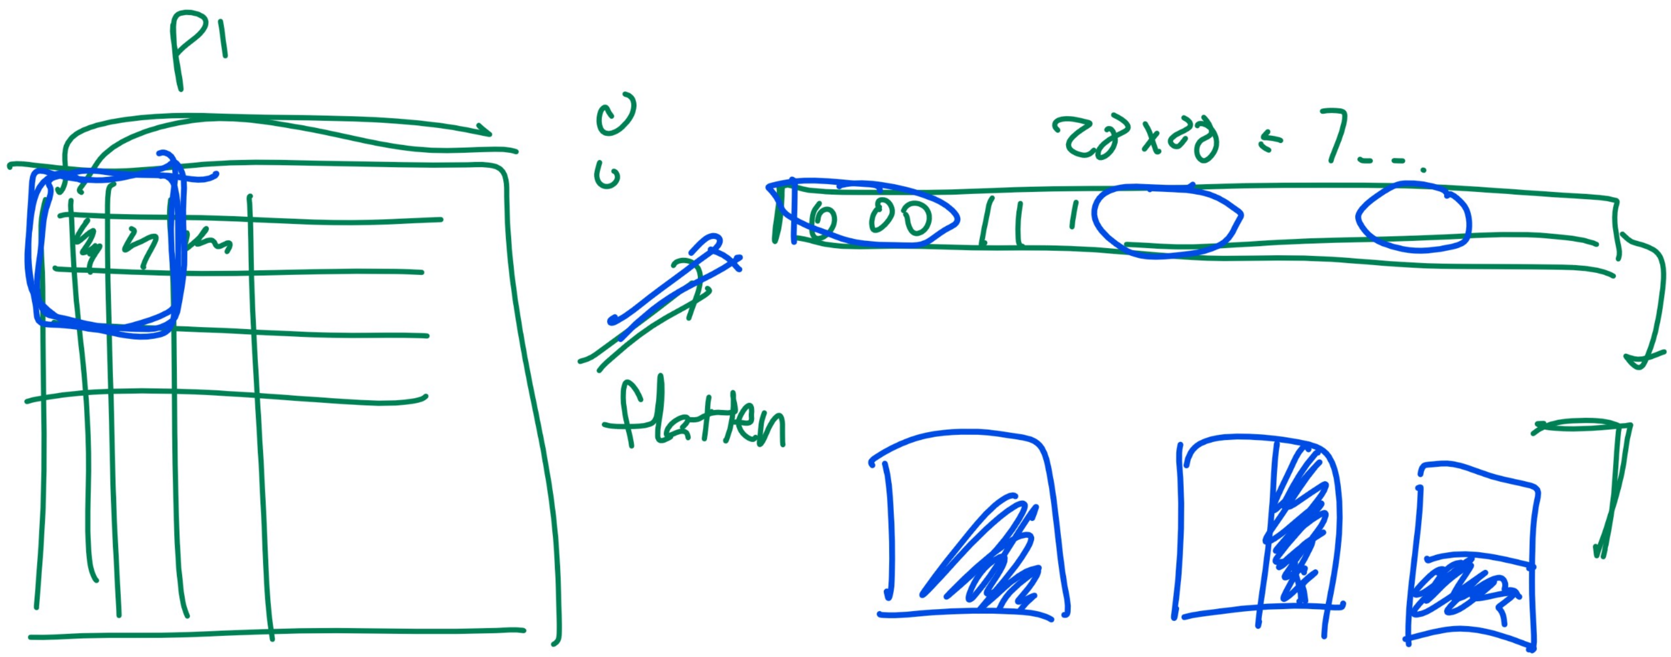

One Dimensional Convolution

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:

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

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

filter = (weights)

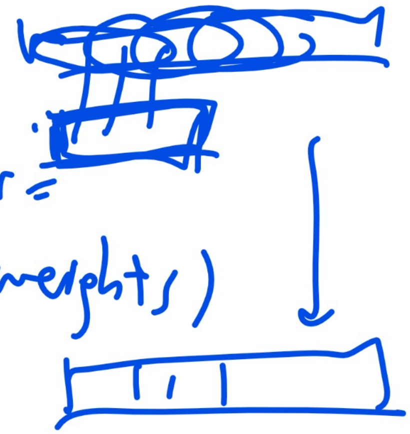

- w is also called a kernel (different from the kernel for SVMs).
- The elements that do not exist are assumed to be 0.

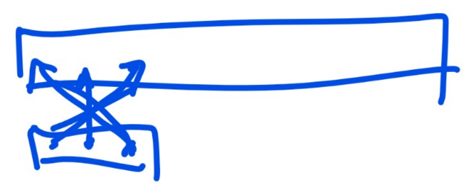

x

Two Dimensional Convolution

Definition

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

$$A = X * W$$

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

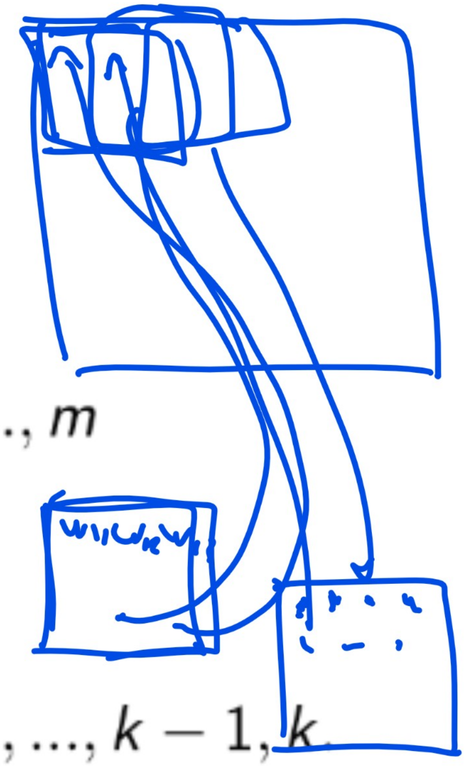

- The matrix W is indexed by (s, t) for $s = -k, -k+1, \ldots, k-1, k$ and $t = -k, -k+1, \ldots, k-1, k$.
- The elements that do not exist are assumed to be 0.

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.

$$\frac{\partial I(s,t)}{\partial s} \approx \frac{I\left(s+\frac{\varepsilon}{2},t\right) - I\left(s-\frac{\varepsilon}{2},t\right)}{\varepsilon}, \varepsilon = 1$$

$$\frac{\partial I(s,t)}{\partial t} \approx \frac{I\left(s,t+\frac{\varepsilon}{2}\right) - I\left(s,t-\frac{\varepsilon}{2}\right)}{\varepsilon}, \varepsilon = 1$$

Image Derivative Filters

Definition

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

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

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 $(\nabla_x I, \nabla_y I)$.

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

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

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

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

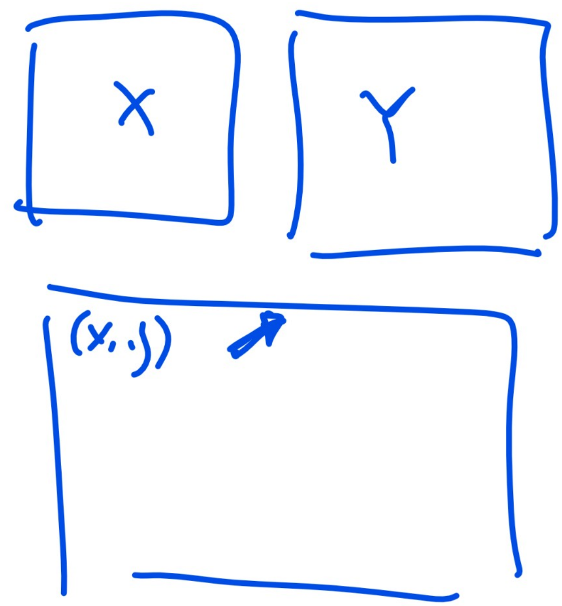

Gradient of Images Demo

Definition

Q1

Pick the least popular choice :

A

B

C

D

E

Convolution Example

filter

$$\begin{matrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{matrix} \Rightarrow \begin{matrix} 9 & 8 & 7 \\ 6 & 5 & 4 \\ 3 & 2 & 1 \end{matrix}$$

flipped filter

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

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

flipped filter

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

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

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

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

Gradient Example

Quiz

Convolution Example 1

Quiz

missing entry ⇒ $\boxed{\text{zero-padding}}$

stride 1

no padding

$3\times3 * 3\times3 \Rightarrow 1\times1$

Q2

$$\begin{bmatrix} 0 & 0 & 0 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} * \overbrace{\begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}}^{\text{filter}}$$

filter ⇒ flipped filter

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

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

E. not understand.

3x3

stride 3

$1\cdot2+1\cdot1 = 3$

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

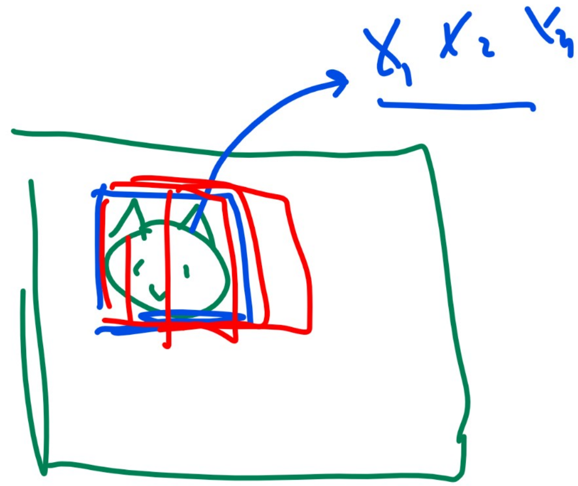

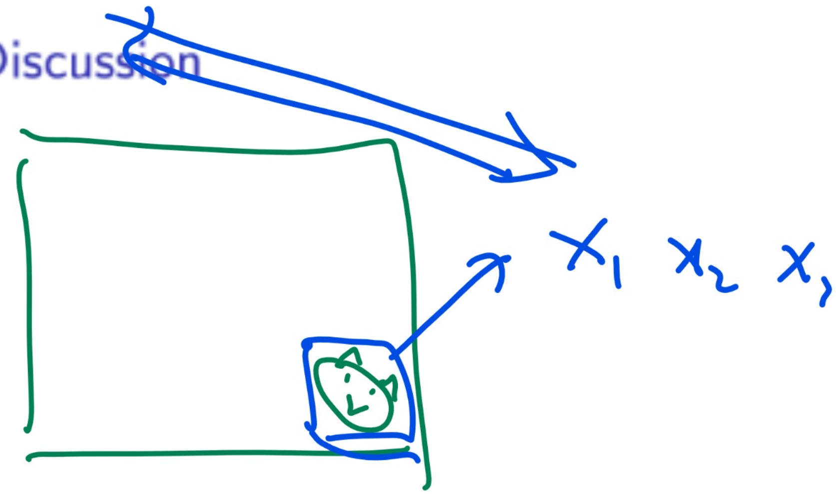

- Scale Invariant Feature Transform (SIFT) features are features that are invariant to changes in the location, scale, orientation, and lighting of the pixels.

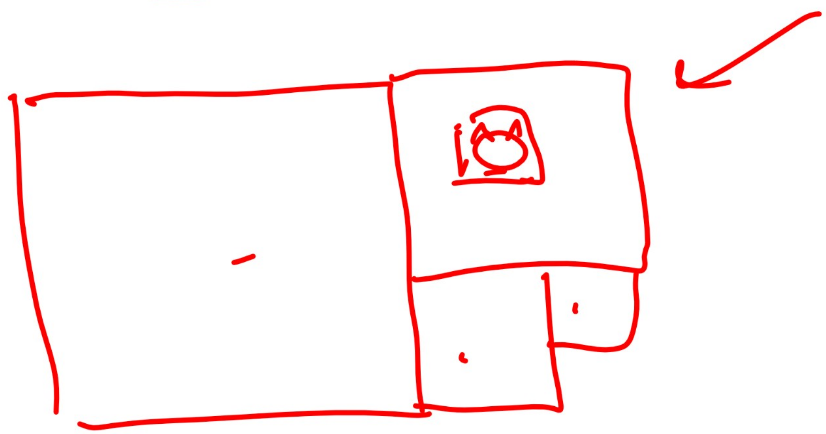

Histogram Binning Diagram

Discussion

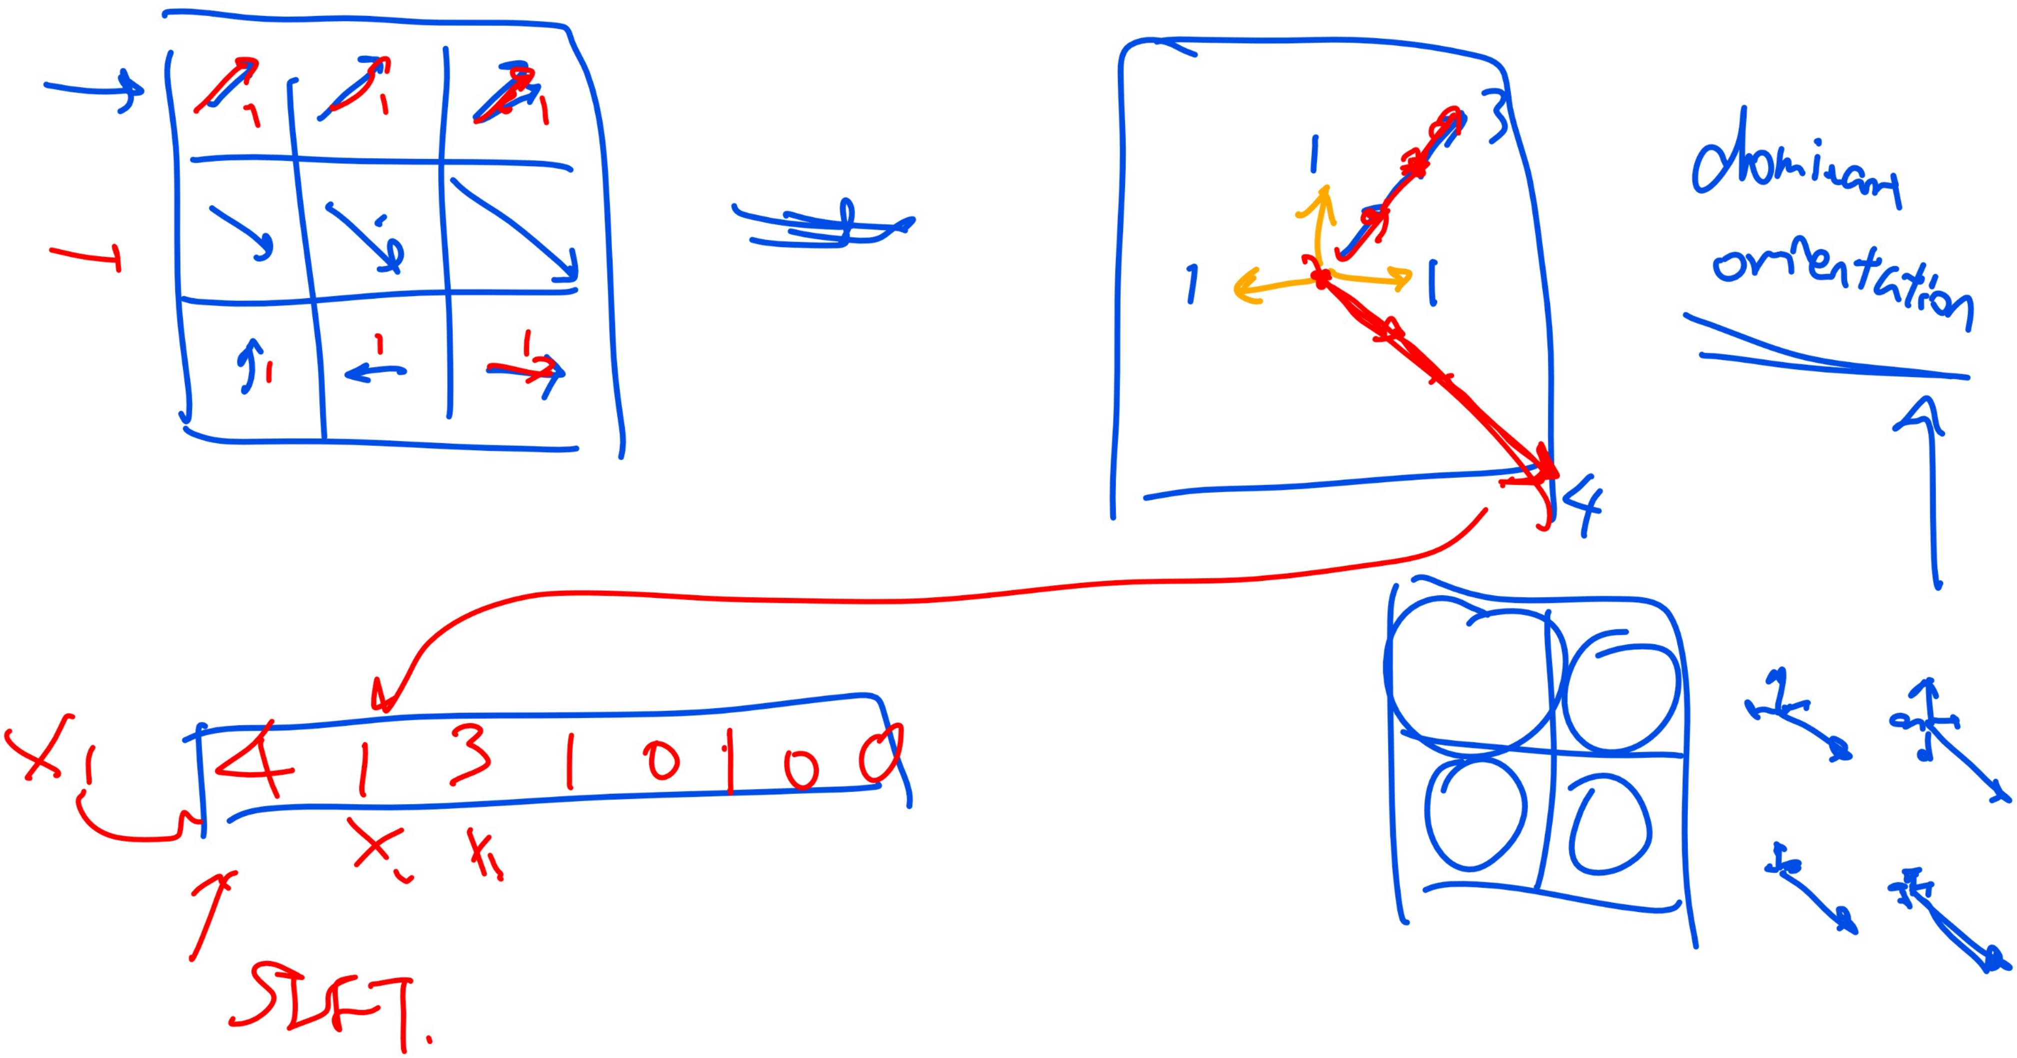

HOG

Discussion

9 bins

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

Classification

Discussion

feature engineering

- 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.