CS540 Introduction to Artificial Intelligence
Lecture 7

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

June 11, 2020

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

P1

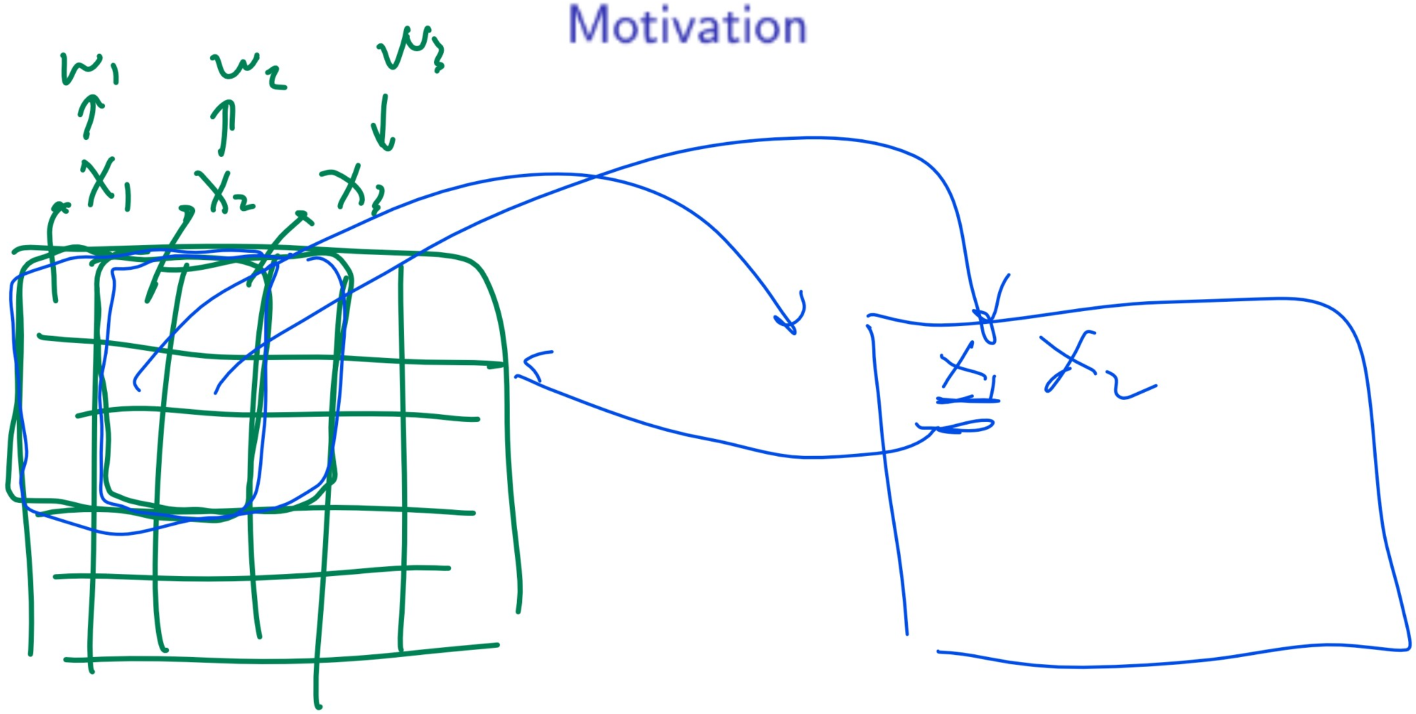

Use info about neighbors.

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:

$$A = X * W$$

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

- 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

Definition

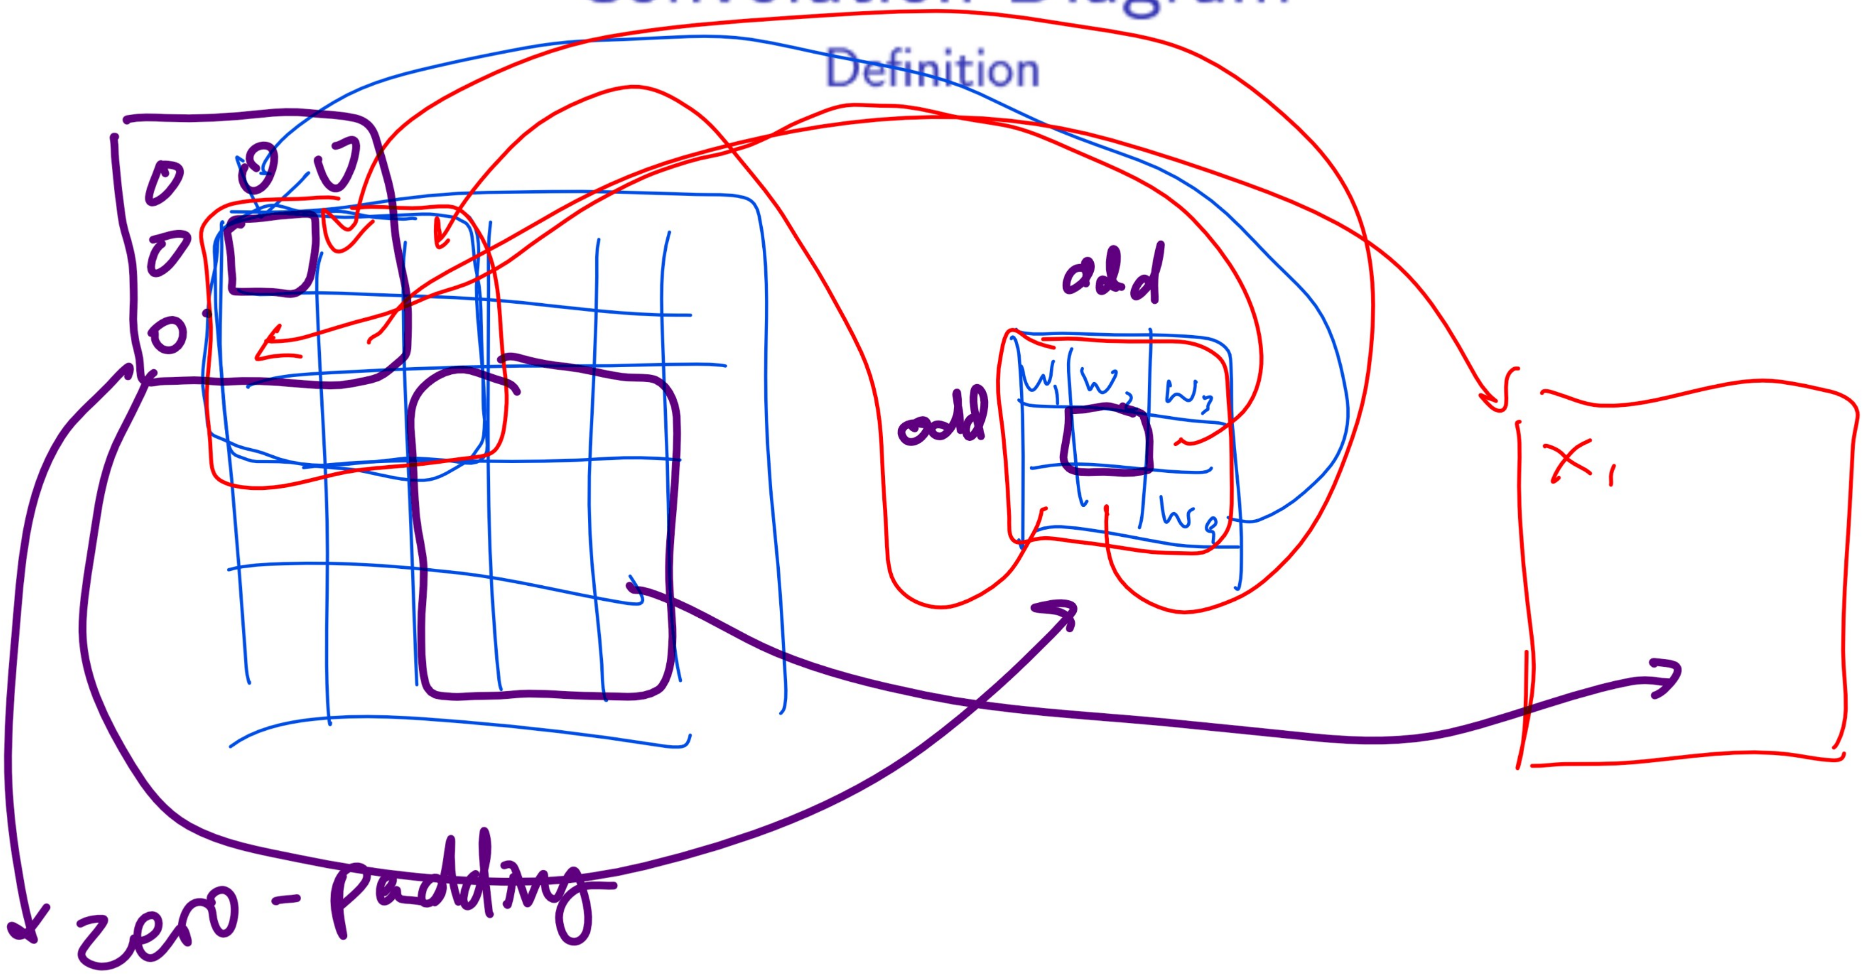

Image Gradient

Definition

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

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

Image Derivative Filters

Definition

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

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

1	2	3
4	5	6
7	8	9

$(-2, -6)$

$\nabla_x I(5) = 4 - 6 = -2$

$\nabla_y I(5) = 2 - 8 = -6$

Sobel Filter

Definition

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

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

Gradient of Images

Definition

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

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

$(-2, -4)$

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

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

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

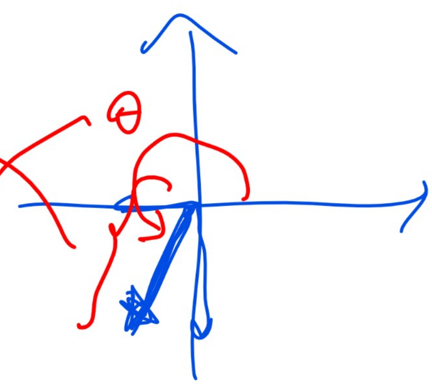

Gradient of Images Demo

Definition

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

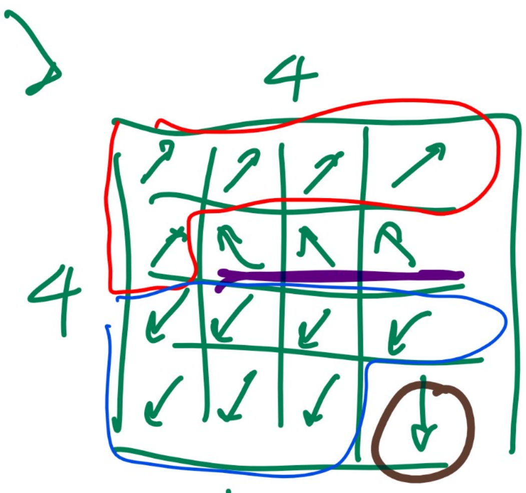

$(\nabla_x I, \nabla_y I)$

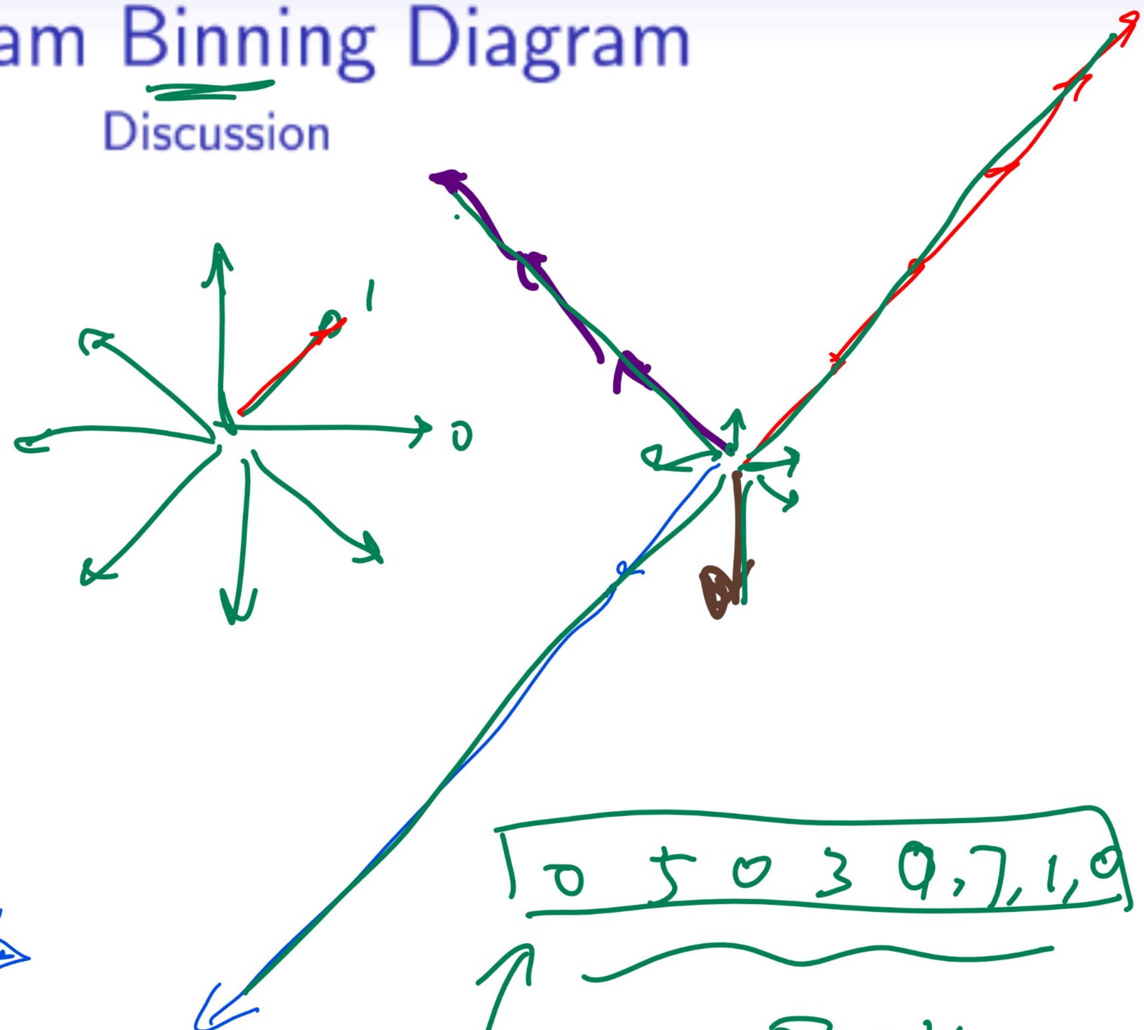

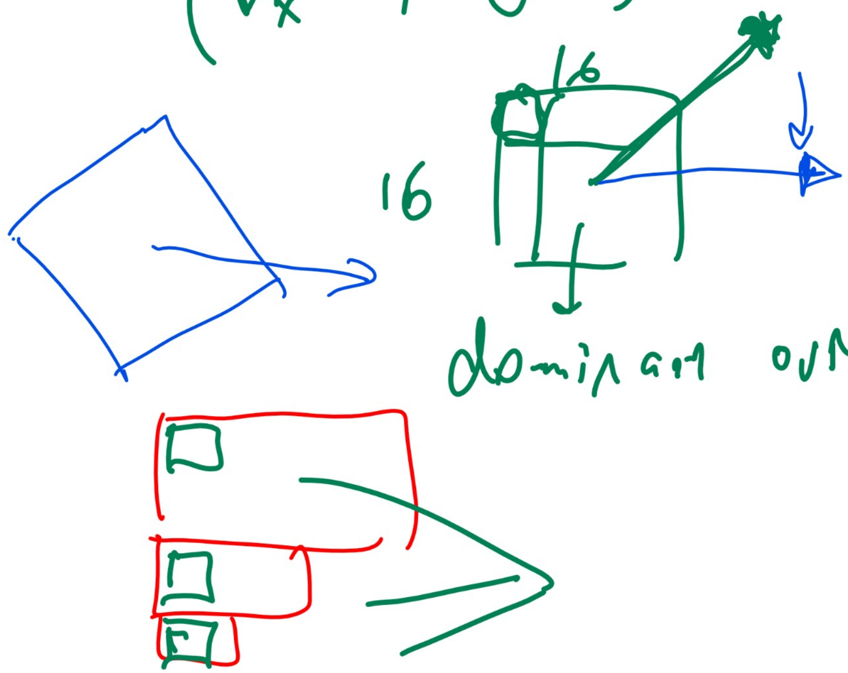

dominant orientation

[0 5 0 3 0, 7, 1, 0]

8×16

$= 128$ features

normalize.. Sum up to 1

HOG

Discussion

faster

MC CV

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

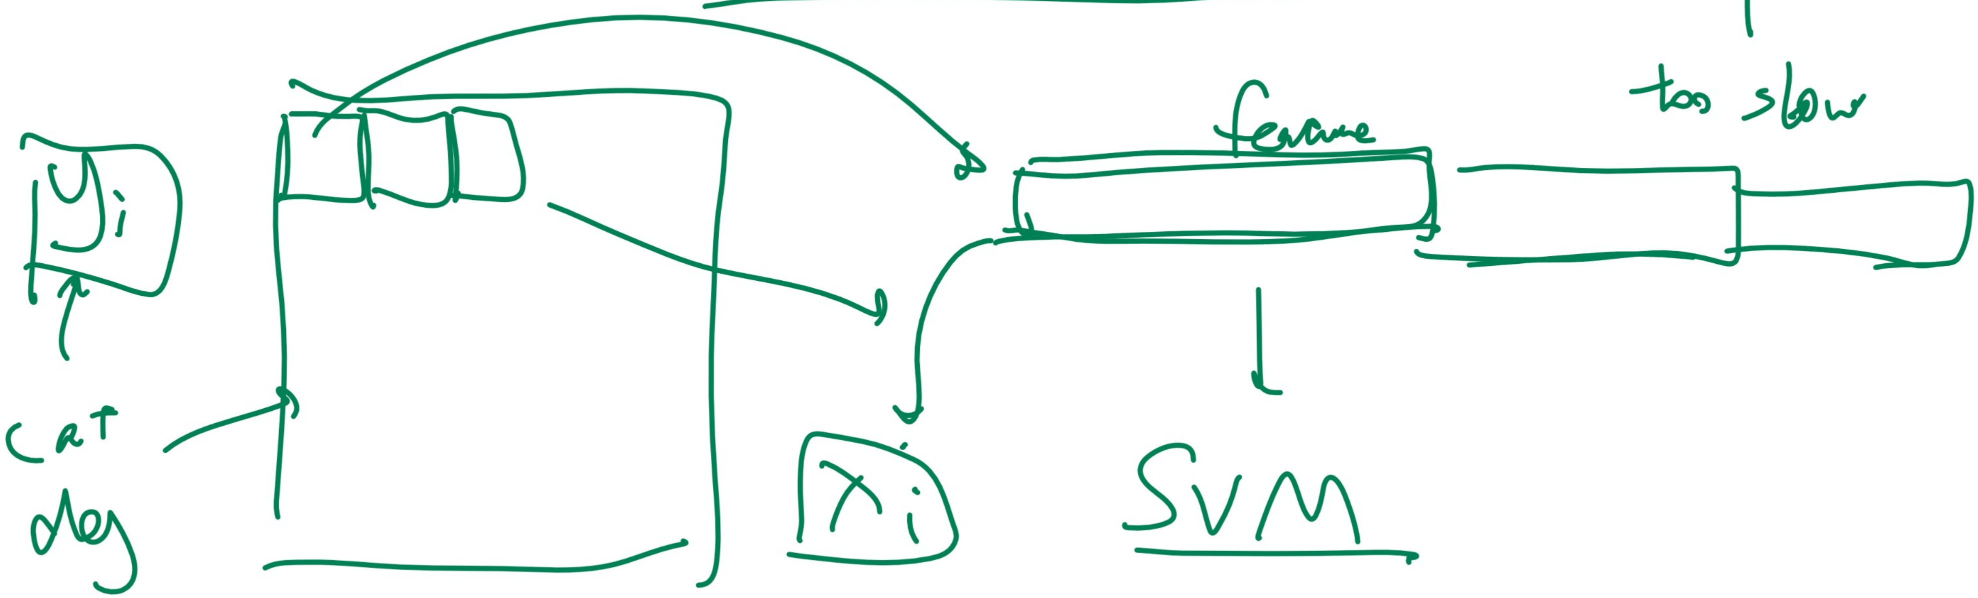

Matching vs Classification Diagram

Discussion

Convolution Example

Quiz

element

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

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

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

$\nabla_x I \quad \nabla_y I$

$(-2, -6)$

$\sqrt{(-2)^2 + (-6)^2}$

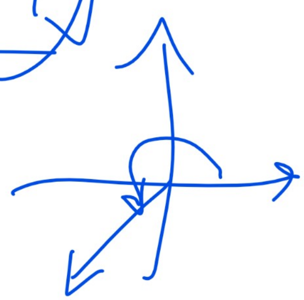

$\arctan\left(\frac{-6}{-2}\right)$

$\mathrm{atan2}(-6, -2)$

Gradient Example

Quiz

Convolution Example 1

Quiz

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

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

flipped

- 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 : $\begin{bmatrix} -1 & 0 & 1 \\ -3 & 0 & 3 \\ -3 & 0 & 3 \end{bmatrix}$

Convolution Example 2

Quiz

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

Q3 select anything

Q4

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

filter

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

flipped

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

- 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 : $\begin{bmatrix} -1 & 0 & 1 \\ -3 & 0 & 3 \\ -3 & 0 & 3 \end{bmatrix}$

Convolution Example 3

Quiz

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

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

two 1D filter
is faster
one 2D filter

Convolution Example 4

Quiz

Q5

What is the gradient magnitude for the center ~~cell~~ element?

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

I

$$\nabla_x = \begin{bmatrix} -3 & -3 & 3 \\ -4 & -4 & 4 \\ -3 & -3 & 3 \end{bmatrix}, \nabla_y = \begin{bmatrix} -1 & -3 & -3 \\ 0 & 0 & 0 \\ 1 & 3 & 3 \end{bmatrix}$$

- A: 1, B: 2, C: 3, D: 4, E: 5

∇_x ∇_y

$$\left((-4)^2 + (0)^2\right)^{1/2}$$

$I * W_x$

right Sobel filter

$I * W_y$

in Q4

Convolution Example 5

Quiz

Q6

What is the gradient direction bin for the center cell?

$$\nabla_x = \begin{bmatrix} -3 & -3 & 3 \\ -4 & -4 & 4 \\ -3 & -3 & 3 \end{bmatrix}, \nabla_y = \begin{bmatrix} -1 & -3 & -3 \\ 0 & 0 & 0 \\ 1 & 3 & 3 \end{bmatrix}$$

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

$\theta = \pi$

- A: $\left(0, \frac{\pi}{2}\right]$, B: $\left(\frac{\pi}{2}, \pi\right]$, C: $\left(\pi, 3\frac{\pi}{2}\right]$, D: $\left(3\frac{\pi}{2}, 2\pi\right]$

don't use arctan

atan2

$-\frac{\pi}{2}, \frac{\pi}{2}$

(a, b] [a, b]

P1 :
in
slides

$$\boxed{C = \frac{1}{2}(y_i - a_i)^2}$$

CE

$\downarrow$

$\underline{\underline{D_w C}}$

$D_w C = \quad (a_{ij})(1 - a_{ij})$

M5Q8:

1 2 3 4 5

monster

$$\underline{H(Y)} - H(Y \mid X)$$

$$1,\ 0 \rightarrow -1\log 1 - 0 \log 0$$

$$H(1, 0) = 0$$

$$\underbrace{H\left(\tfrac{3}{5}, \tfrac{2}{5}\right)}_{-\frac{3}{5}\log\frac{3}{5} - \frac{2}{5}\log\frac{2}{5}} - \frac{1}{5}\boxed{H(Y \mid X = 1)} - \frac{4}{5}\underline{H(Y \mid X = 0)}$$

$$H(Y \mid X = 1) = 0$$

$$H(Y \mid X = 0) = H\left(\tfrac{1}{2}, \tfrac{1}{2}\right)$$

$$-\frac{1}{2}\log\frac{1}{2} - \frac{1}{2}\log\frac{1}{2}$$

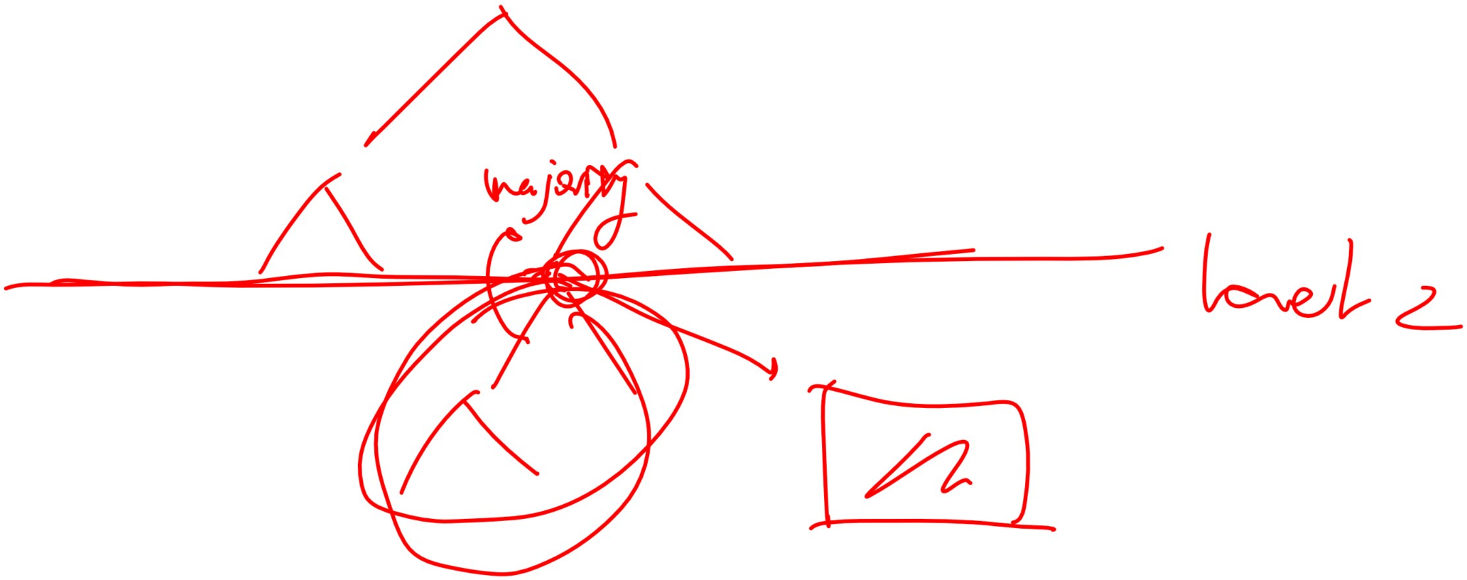
majority
level 2