

Class Plan

Recap



Reminders

Interpretability


1. Grad CAM
2. Bounding Box
3. Segmentation

Convolutions

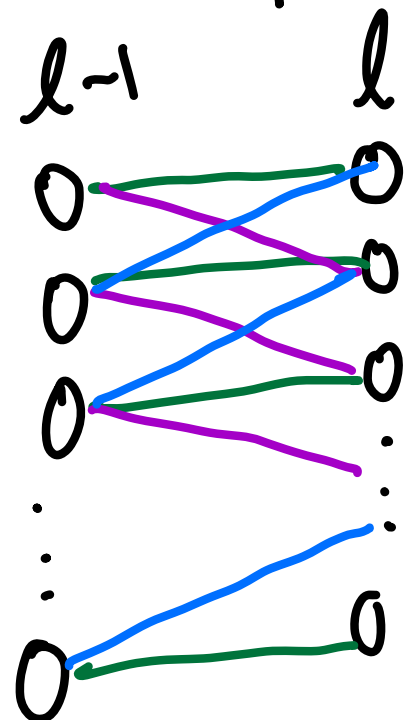
- images
 - audio
 - language
 - video
- we lots of images




3D



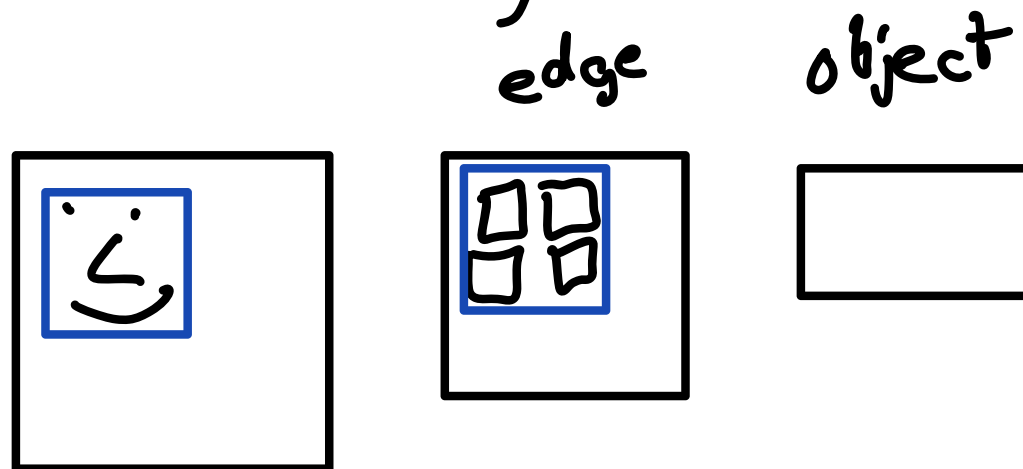
↳ computational



↳ Shift Invariance



↳ Locality

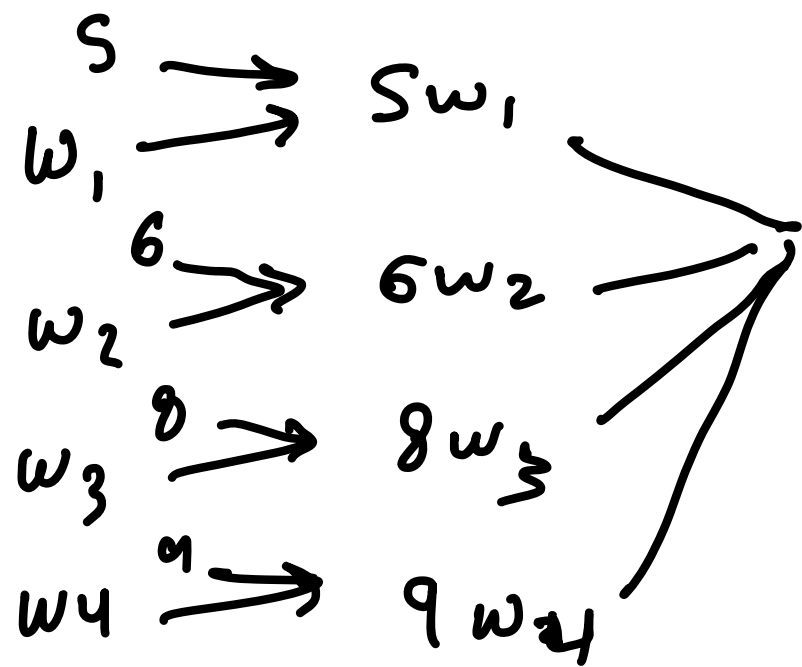


1	2	3
4	5	6
7	8	9

w_1	w_2
w_3	w_4

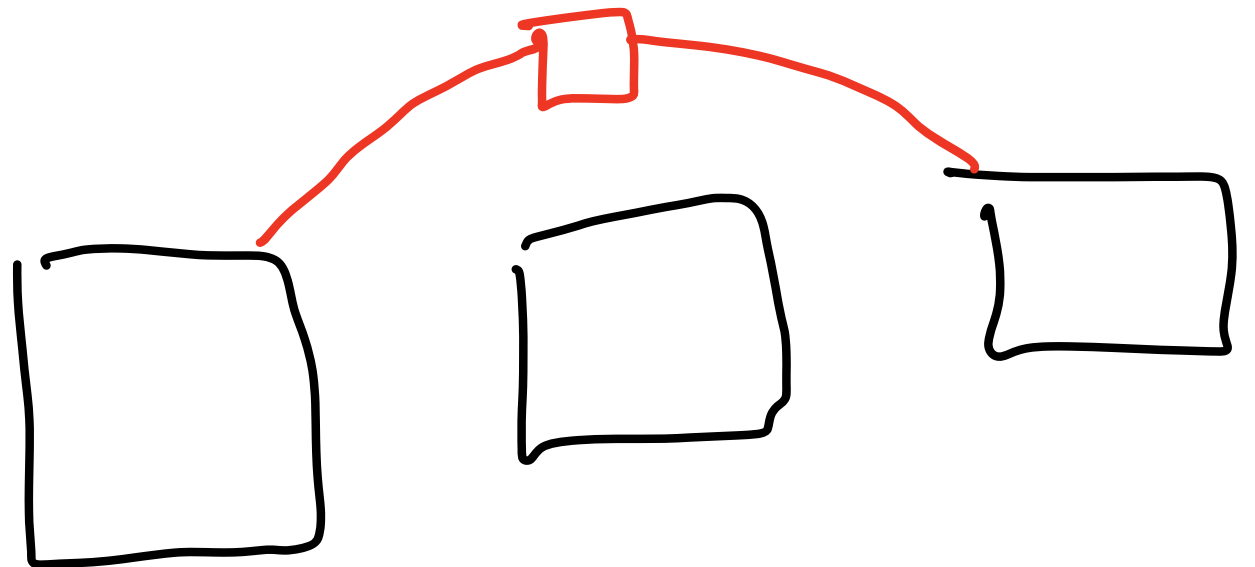
||

	$2w_1 + 3w_2$ $+5w_3 + 6w_4$
	$5w_1 + 6w_2$ $+8w_3 + 9w_4$





Residual Network

0 0 0
 0 0 0
 \vdots \vdots \vdots
 0 0 0



Reminders

- No office hours today
- Google form 23/26
1 point per day
- Homework due Spm Fri.
15 points
↳ 24 hours
- Self grade due Spm Mon.
For each problem
 - 1) ?
 - 2) ?


Object Detection

What is in the image?


and Where is it?

Approach #1: Model Explanation


Gradient Class Activation Mapping



...



channels



probabilities

classes

						j
3	1	-1	7	.5	2	π
			2			

A^k

Outputs class c

"importance kth channel"


$$\alpha_k^c = \sum_i \sum_j \frac{\partial y_k^c}{\partial A_{ij}^k}$$

$$L_{ij}^c = \text{Relu} \left(\sum_k \alpha_k^c A_{ij}^k \right)$$


Grad-CAM for "Cat"



Grad-CAM for "Dog"



Approach #2: Bounding Boxes



Represent with four numbers


- two corners each x, y
- center, height, width

How good is a bounding box?



Jaccard similarity A, B

$$\frac{|A \cap B|}{|A \cup B|} = \frac{\text{green blob}}{\text{green blob} + \text{purple blob}}$$

- classification



- classify each candidate box



Approach #3: Segmentation

