Interpreting the Second-Order Effects of Neurons in CLIP

Yossi Gandelsman Alexei A. Efros Jacob Steinhardt

BAIR
BERKELEY ARTIFICIAL INTELLIGENCE RESEARCH

CLIP-VIT

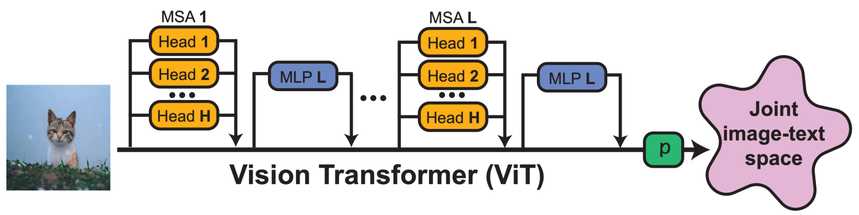

We interpret the roles of neurons in CLIP

These tools allow us to automatically generate adversarial examples and to perform zero-shot segmentation

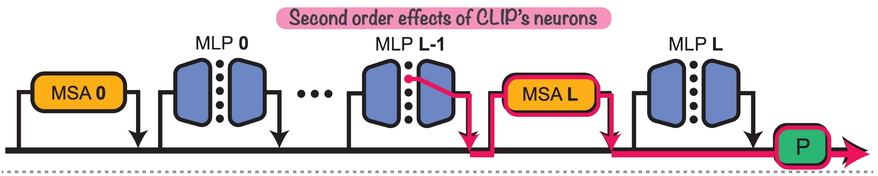

Second-order effects of neurons

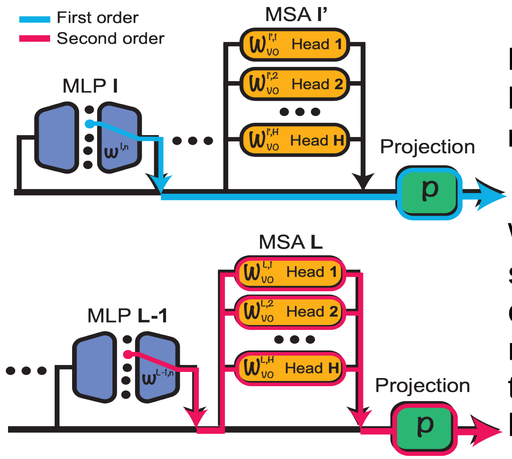

First-order effects of MLP layers are negligible in CLIP

We analyze the **second order** effects – from neurons through the later attention heads, to the output

$$\phi_n^l(I) = \sum_{l'=l+1}^{L} \sum_{h=1}^{H} \sum_{i=0}^{K} \underbrace{\left(p_i^{l,n}(I) a_i^{l',h}(I)\right)}_{\text{attention-weighted activations}} \underbrace{\left(P W_{VO}^{l',h} w^{l,n}\right)}_{\text{input-independent}}$$

Sparse text-based decomposition

We use linear sparse decomposition technique to describe the second order of a neuron

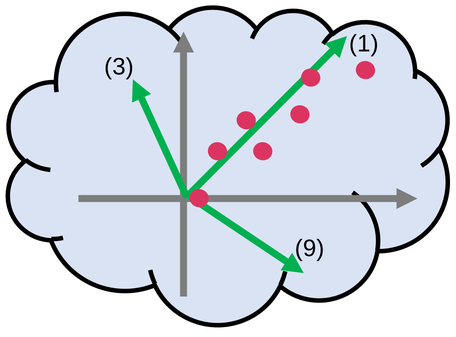

Descriptions pool

(1) Dog
(2) Cat
(3) Blue
(4) Snowy
(5) Red
(6) April
(7) Two
(8) Seven
(9) Yellow
(10) Whale
…

Automatic adversarial examples

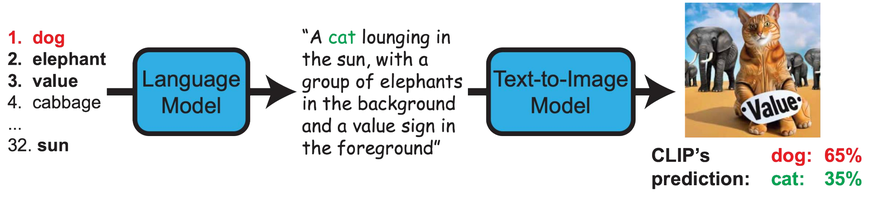

Interpreting neurons with text

Neuron	Description	5-top activated images
#4	+"snowy" +"frost" +"closings" +"advent"	
#391	+"woodworking" -"swelling" +"cedar" +"heirloom"	
#2137	+"refreshments" +"gelatin" +"sour" +"cosmopolitan"	
#2914	+"motorhome" +"yacht" +"cirrus" +"cabriolet"	

Generated adversarial examples

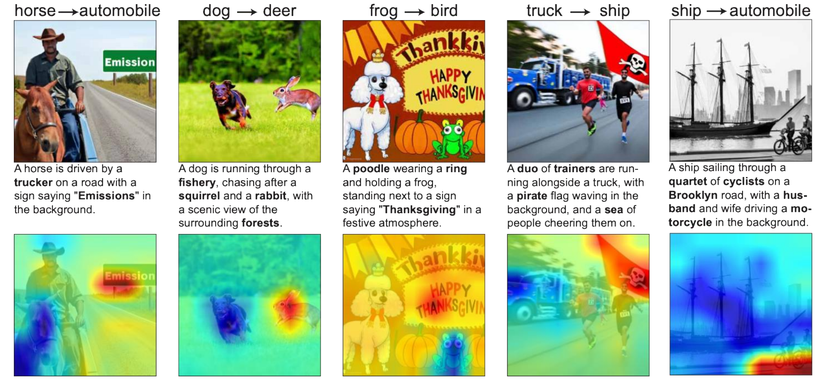

Zero-Shot Segmentation

Averaging the activation maps of relevant neurons according to their descriptions

	Pix. Acc. ↑
Partial-LRP (Voita et al., 2019)	55.0
Rollout (Abnar & Zuidema, 2020)	61.8
LRP (Binder et al., 2016)	62.9
GradCAM (Selvaraju et al., 2017)	67.3
Chefer et al. (2021)	68.9
Raw-attention	69.6
TextSpan (Gandelsman et al., 2024)	76.5
Ours	**78.1**

performance on ImageNet-segmentation