






Advanced Behavioral Cloning



Instructor: Daniel Brown

Implicit Behavioral Cloning

- Paper: <https://arxiv.org/abs/2109.00137>
- Video: <https://www.youtube.com/watch?v=QsIGqRUSRzs>





Action Chunking with Transformers (ACT)

- Paper: <https://arxiv.org/pdf/2304.13705>
- Videos: <https://tonyzhaozh.github.io/aloha/>


Variational Autoencoders (VAEs)

- Autoencoders learn latent representations



Variational Autoencoders (VAEs)


- Autoencoders learn latent representations
- VAEs map input into a distribution over latent variables z
- Loss function is reconstruction plus KL divergence



$$\mathcal{L} = \mathbb{E}_{q(z|x)}[\log p(x|z)] - D_{\text{KL}}(q(z|x) || p(z))$$

Conditional Variational Autoencoders (CVAEs)

- Encoder and decoder both condition on extra info y



- Loss function is reconstruction plus KL divergence


$$\mathcal{L} = \mathbb{E}_{q(z|x,y)}[\log p(x|z, \underline{y})] - D_{\text{KL}}(q(z|x, \underline{y}) || p(z|\underline{y}))$$

Transformers


- State of the art ways to ingest and output sequential data.

BERT


Encoder




Neural language modeling




Neural language modeling




Neural language modeling



Neural language modeling



the size of the vector representation up to and including the **current token**



[BOS] Sylvester Stallone has


the logits vector

“+” softmax

i-th dimension ~ the “probability” [not really] that the **next token** is the i-th token in the vocabulary

select the token with the high(est) “probability” as a token to display (generate)

Neural sequence modeling




Problems:


- How do we deal with different length inputs?
- How do we model long-range dependencies?

Recurrent Neural Networks


- Standard RNN





- Long short-term memory (LSTM)





Large Language Models













Input

Thinking


Machines


Embedding


X_1 

X_2 

Queries


q_1 


q_2 




W^Q

Keys


k_1 


k_2 



W^K

Values


v_1 

v_2 




W^V


Perform dot product between query and all keys to get a raw score for each previous word (including current word).




Normalize these scores via a softmax to get a probability distribution. Then return a weighted sum of the values.



Word	Value vector	Score	Value X Score
<S>		0.001	
a		0.3	
robot		0.5	
must		0.002	
obey		0.001	
the		0.0003	
orders		0.005	
given		0.002	
it		0.19	
		Sum:	





Diffusion Policy

- Paper: <https://arxiv.org/pdf/2303.04137v4>
- Videos: <https://diffusion-policy.cs.columbia.edu/>

Denoising Diffusion (high-level)

