

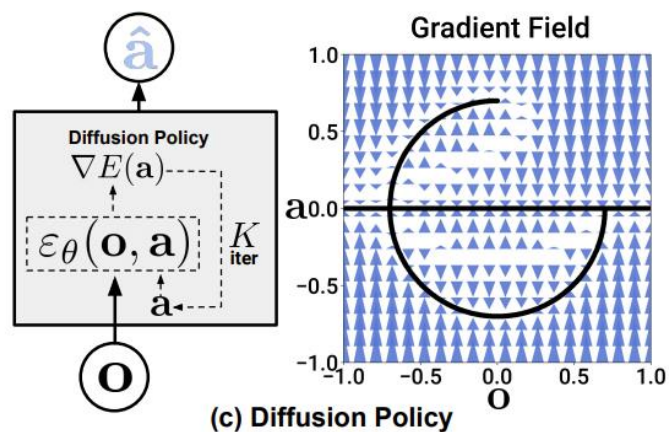
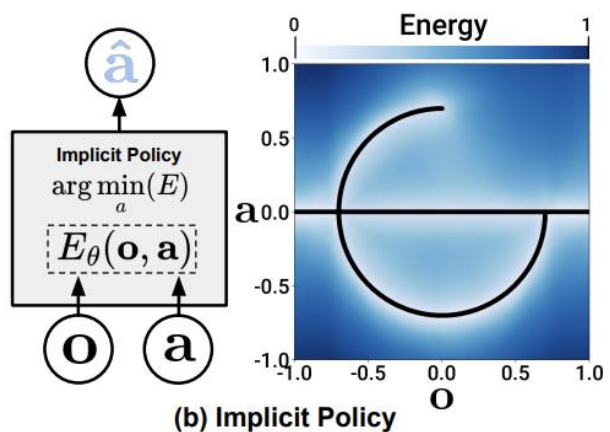
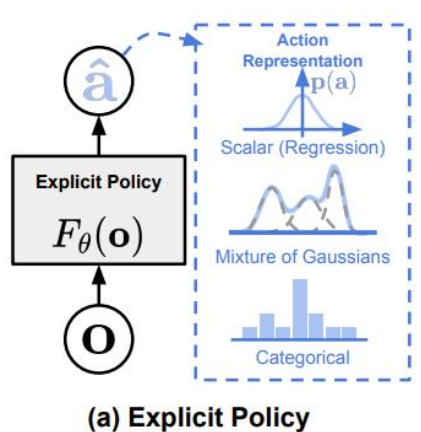


Advanced Behavioral Cloning

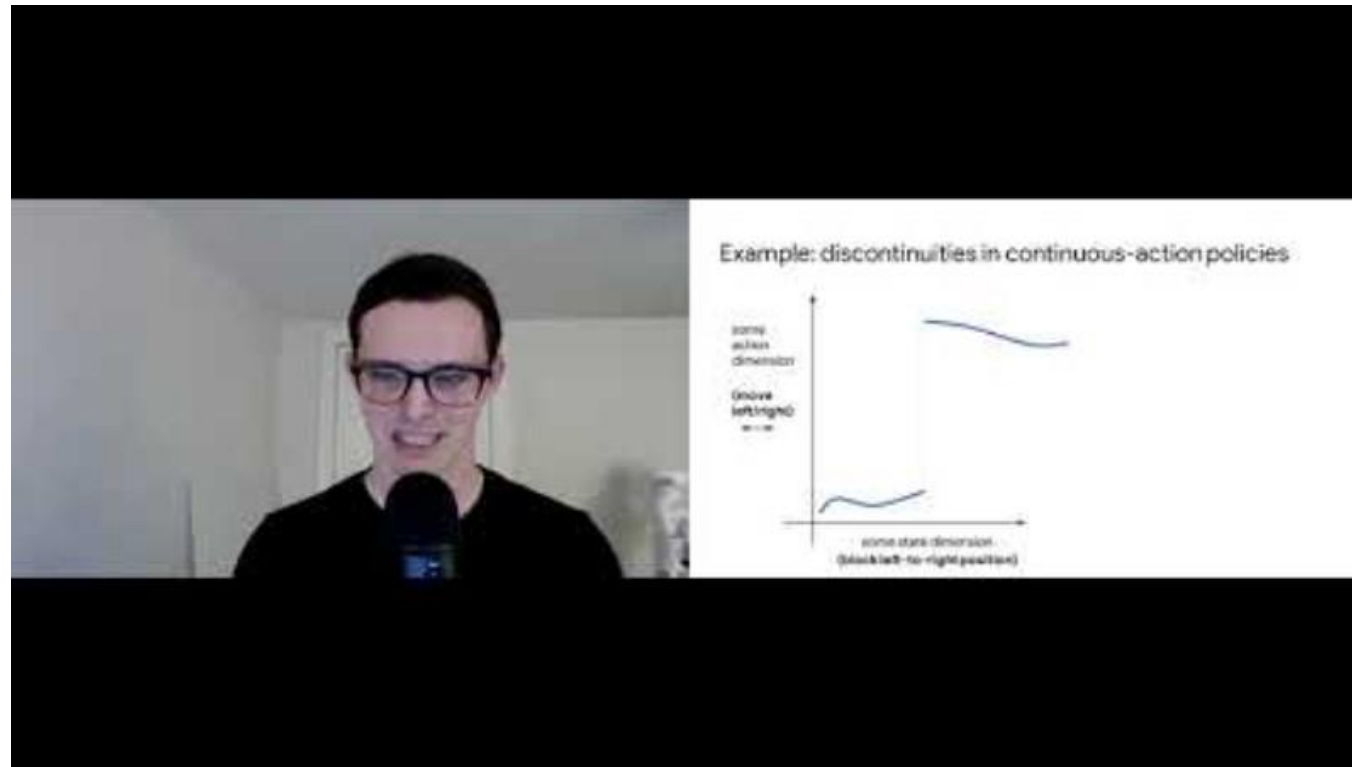


$$D = \{(o, a) \dots\}$$

Instructor: Daniel Brown


Implicit Behavioral Cloning

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



The video frame shows a speaker on the left and a slide on the right. The slide is titled "Example: discontinuities in continuous-action policies". The graph on the slide has a vertical axis labeled "some action dimension" and a horizontal axis labeled "some state dimension (block left-to-right position)". The graph shows two curves: a lower curve on the left and a higher curve on the right, with a vertical jump between them, illustrating a discontinuity in the policy.

$$D = \left\{ \begin{array}{l} (0, a_1) \dots \\ (0, a_2) \end{array} \right.$$



$E(0, a) \downarrow$
 $E(0, c) \uparrow$

Test time

- BC input $0' \rightarrow a'$
- IBC input $0' \rightarrow E$

$$a' = \arg \min_a E_0(0', a)$$

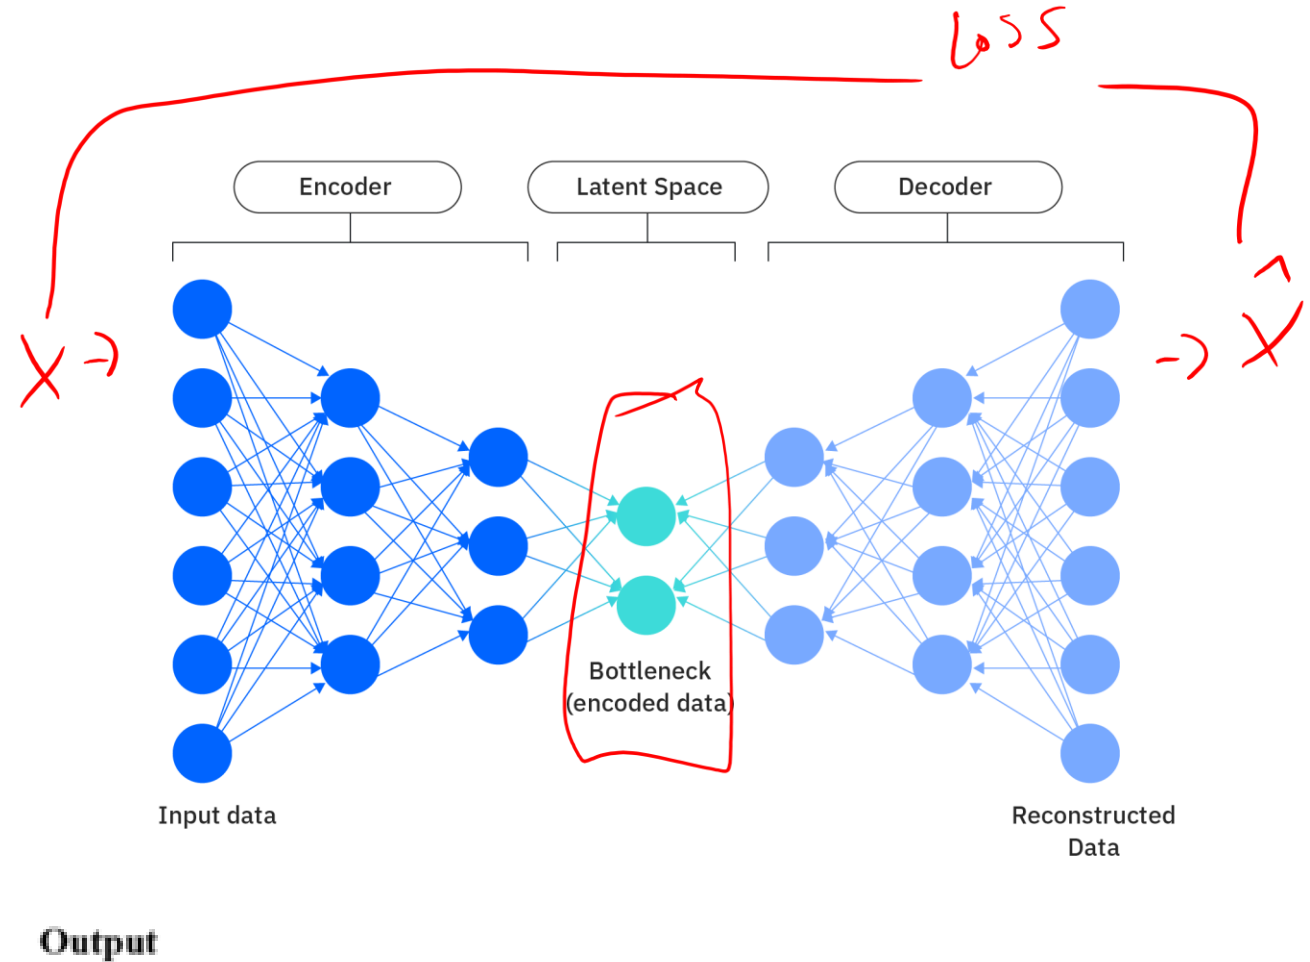



Action Chunking with Transformers (ACT)

- 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

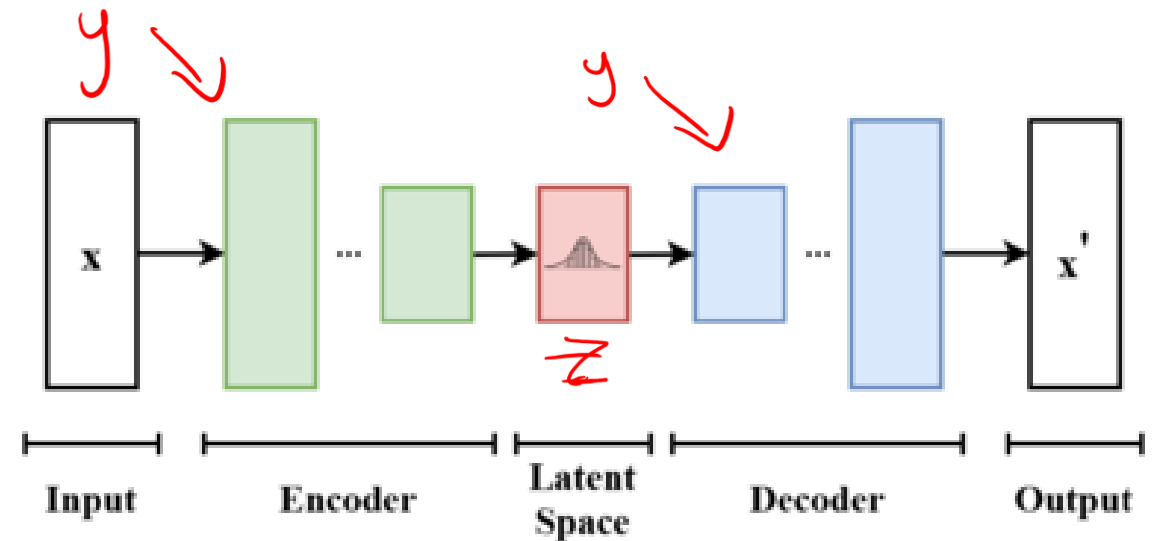


$$\mathcal{L} = \underbrace{\mathbb{E}_{q(z|x)}[\log p(x|z)]}_{\text{encoder}} - \underbrace{D_{\text{KL}}(q(z|x) || p(z))}_{\text{decoder}}$$

prior $N(0, I)$

Conditional Variational Autoencoders (CVAEs)

- Encoder and decoder both condition on extra info y




- Loss function is reconstruction plus KL divergence

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

Transformers


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

BERT
Encoder
CVAE




GPT
Decoder


Neural language modeling




Neural language modeling

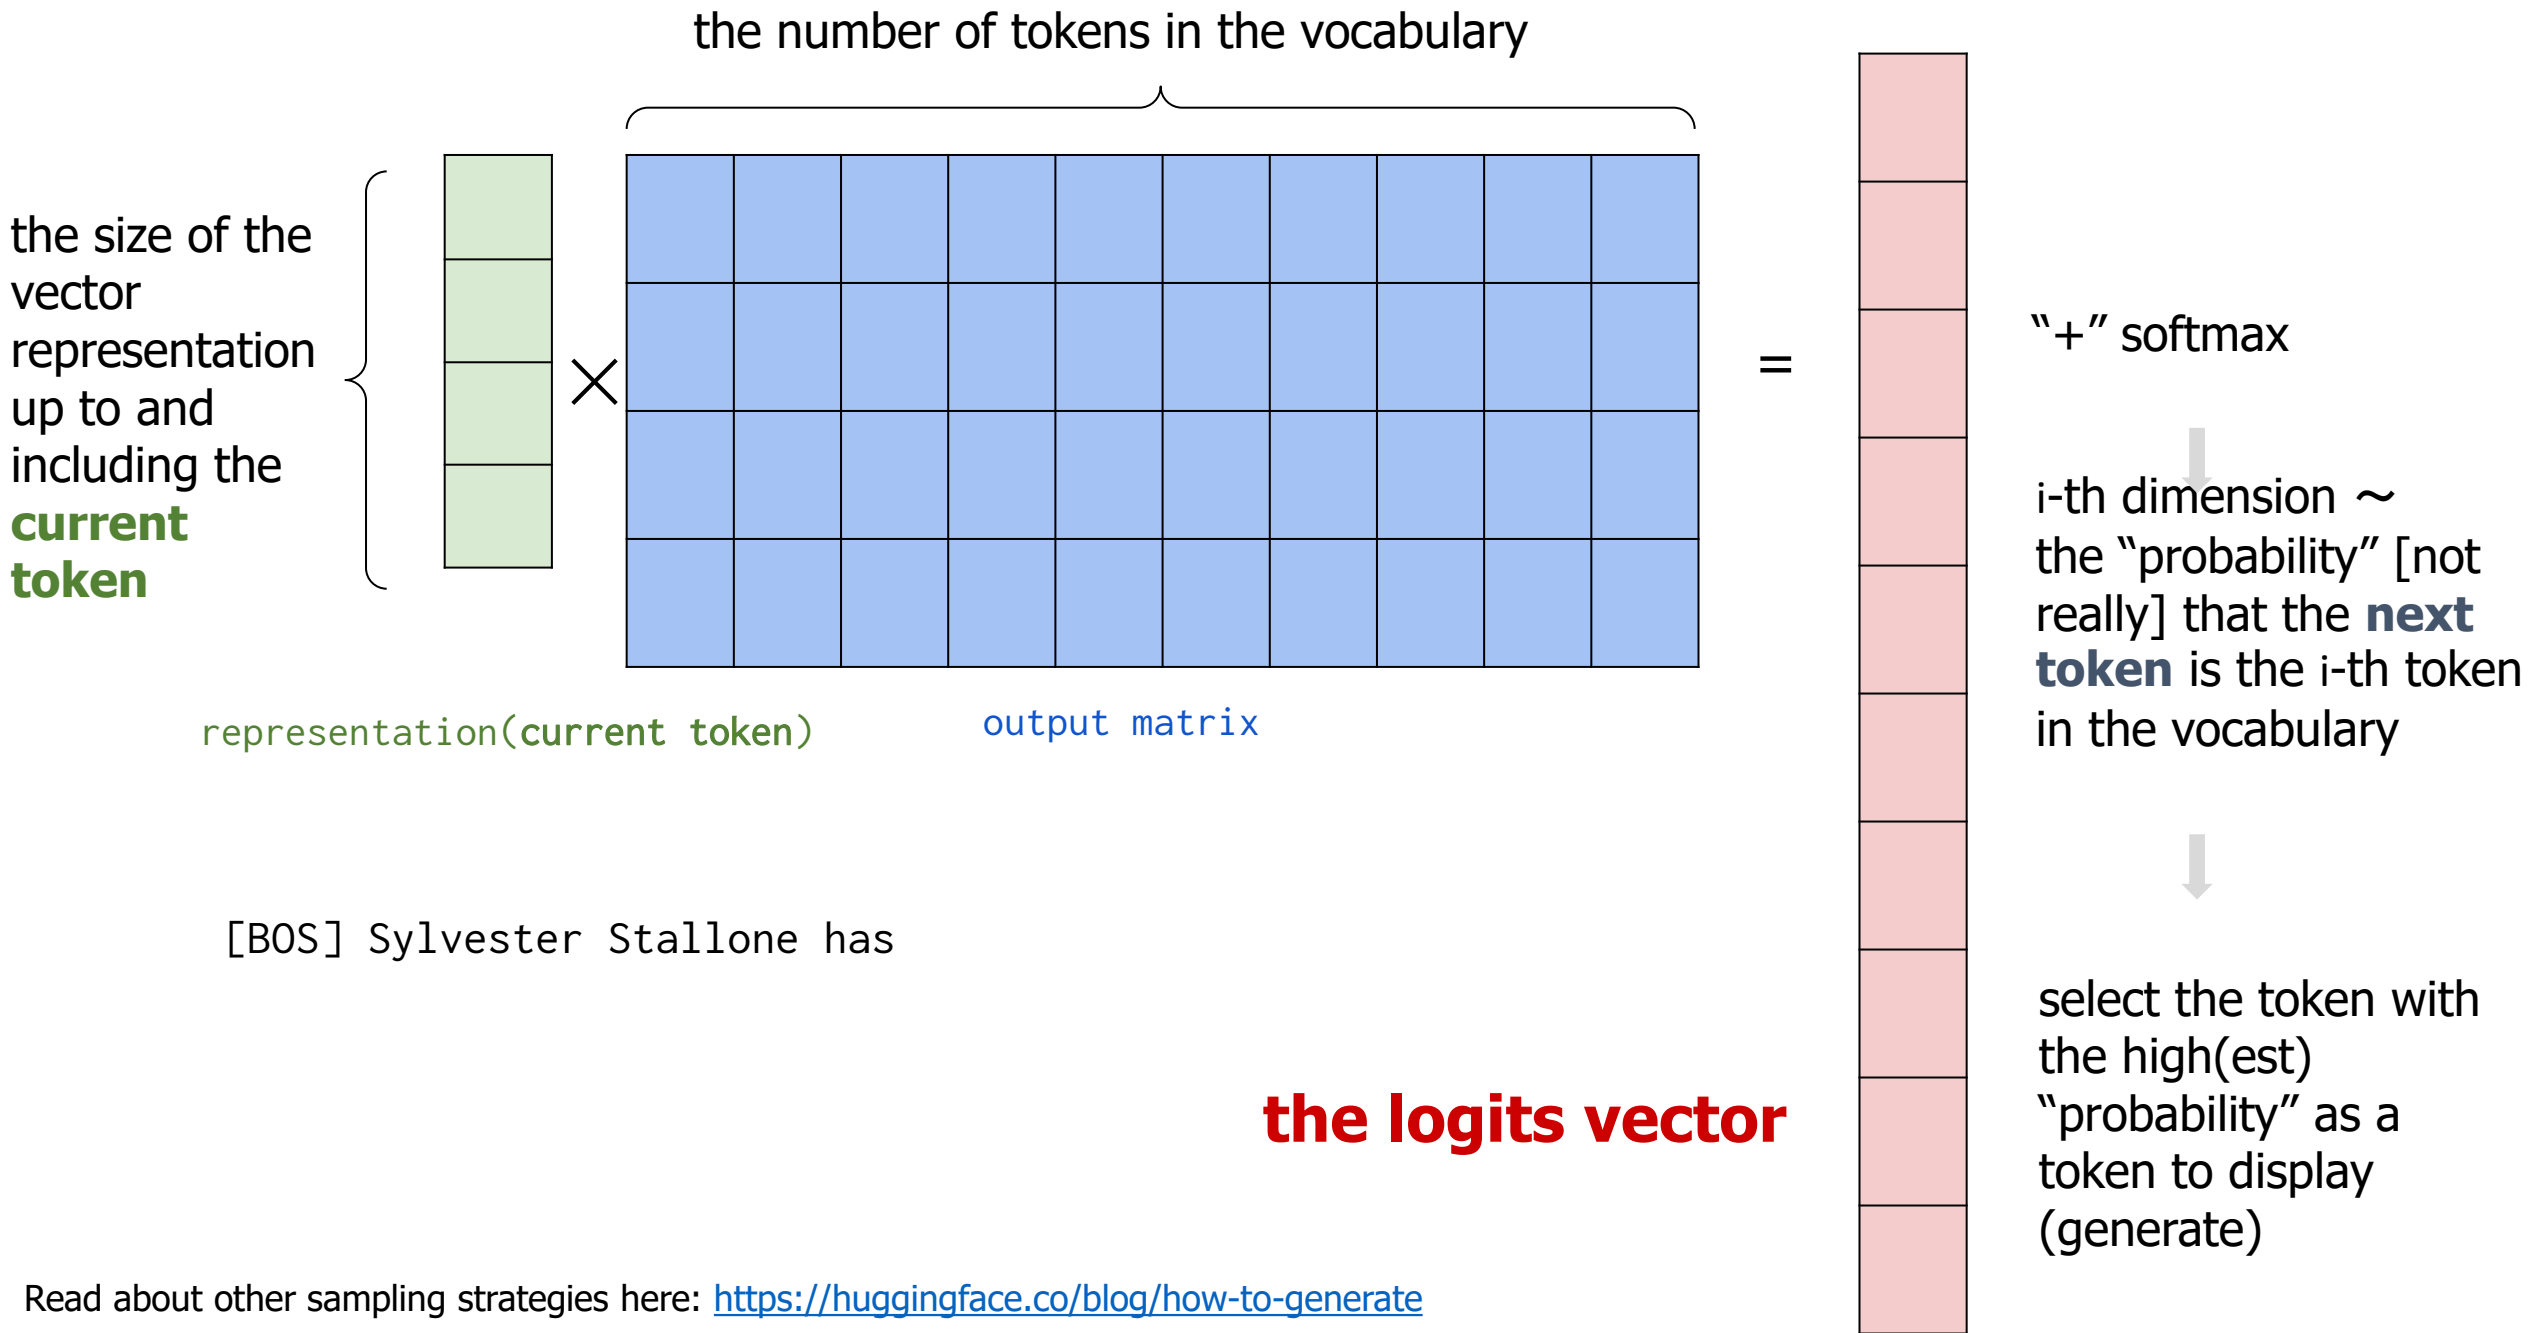


Neural language modeling




Neural language modeling





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)



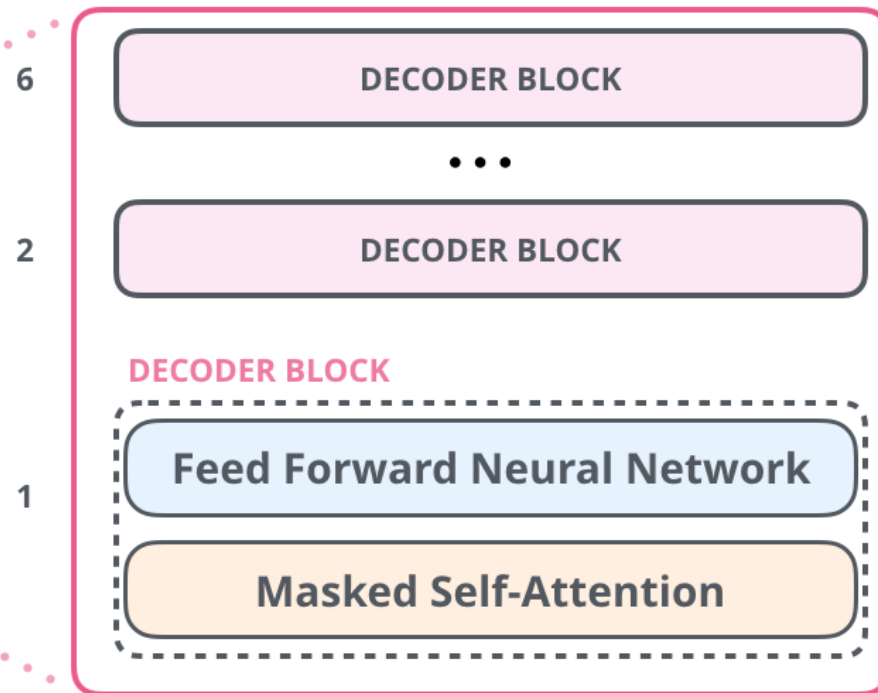

Large Language Models

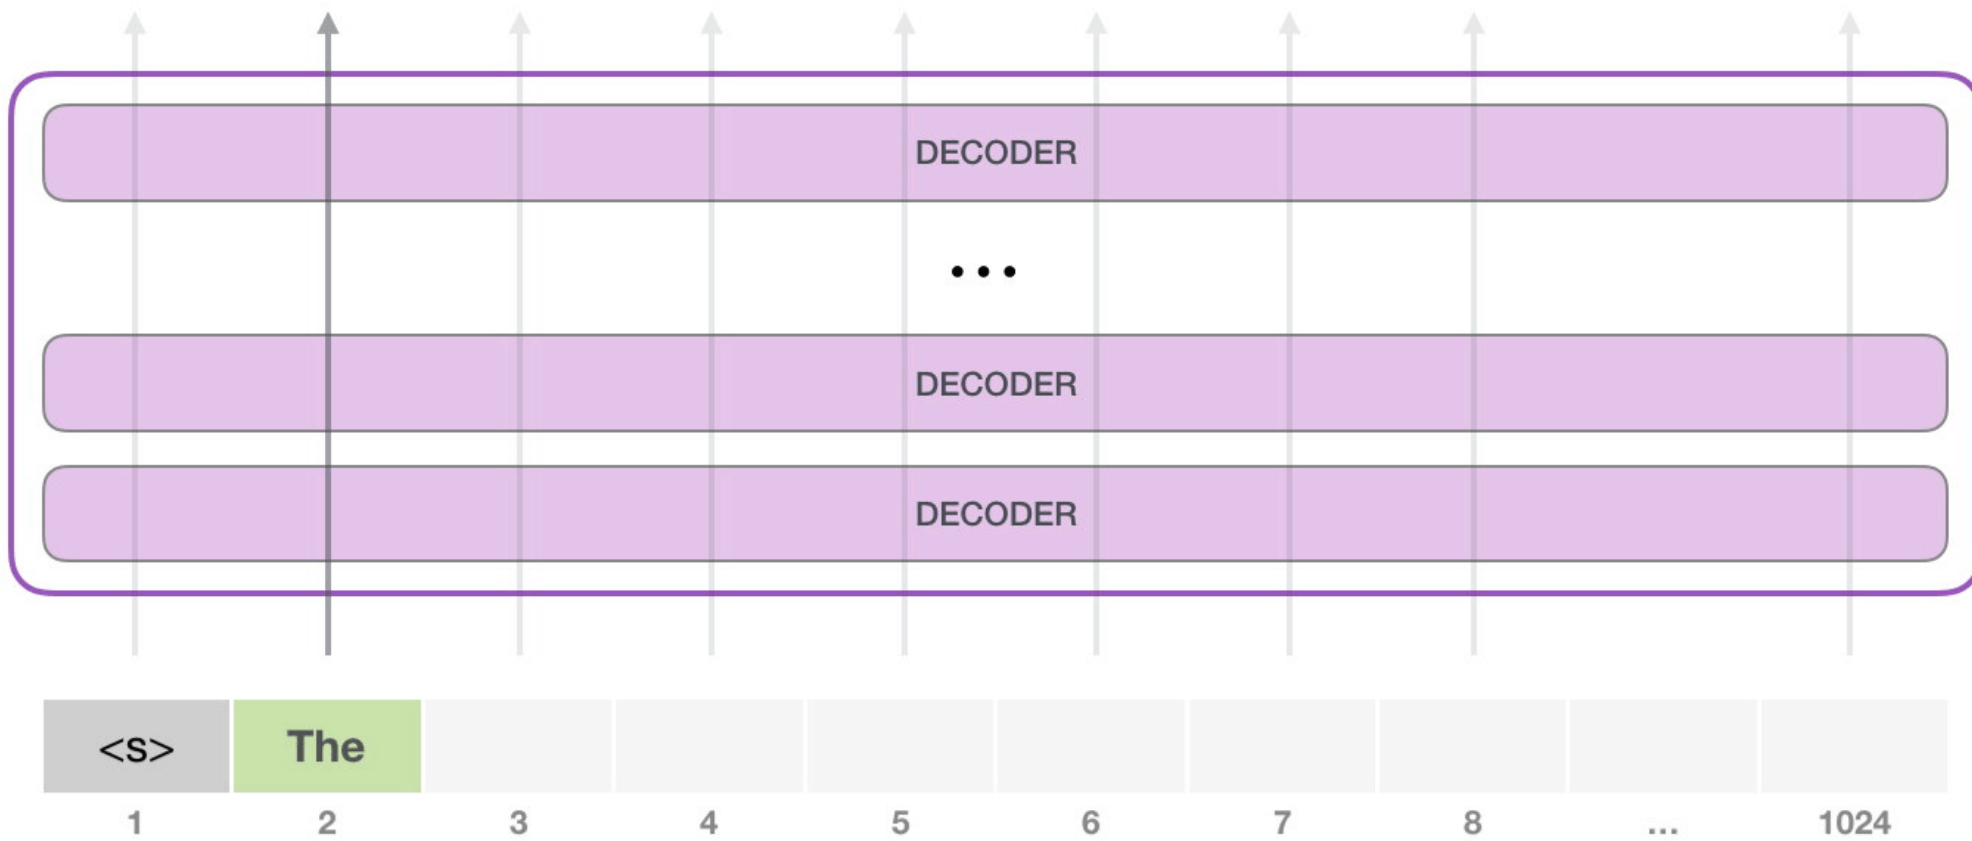
Output




Input











=



Positional encoding for token #1

+



Token embedding of <s>




1

2


...


1024

Token Embeddings



Positional Encodings







Input




Thinking

Machines

Embedding





Queries






$W^Q x_1 = z_1$

Keys




W^K

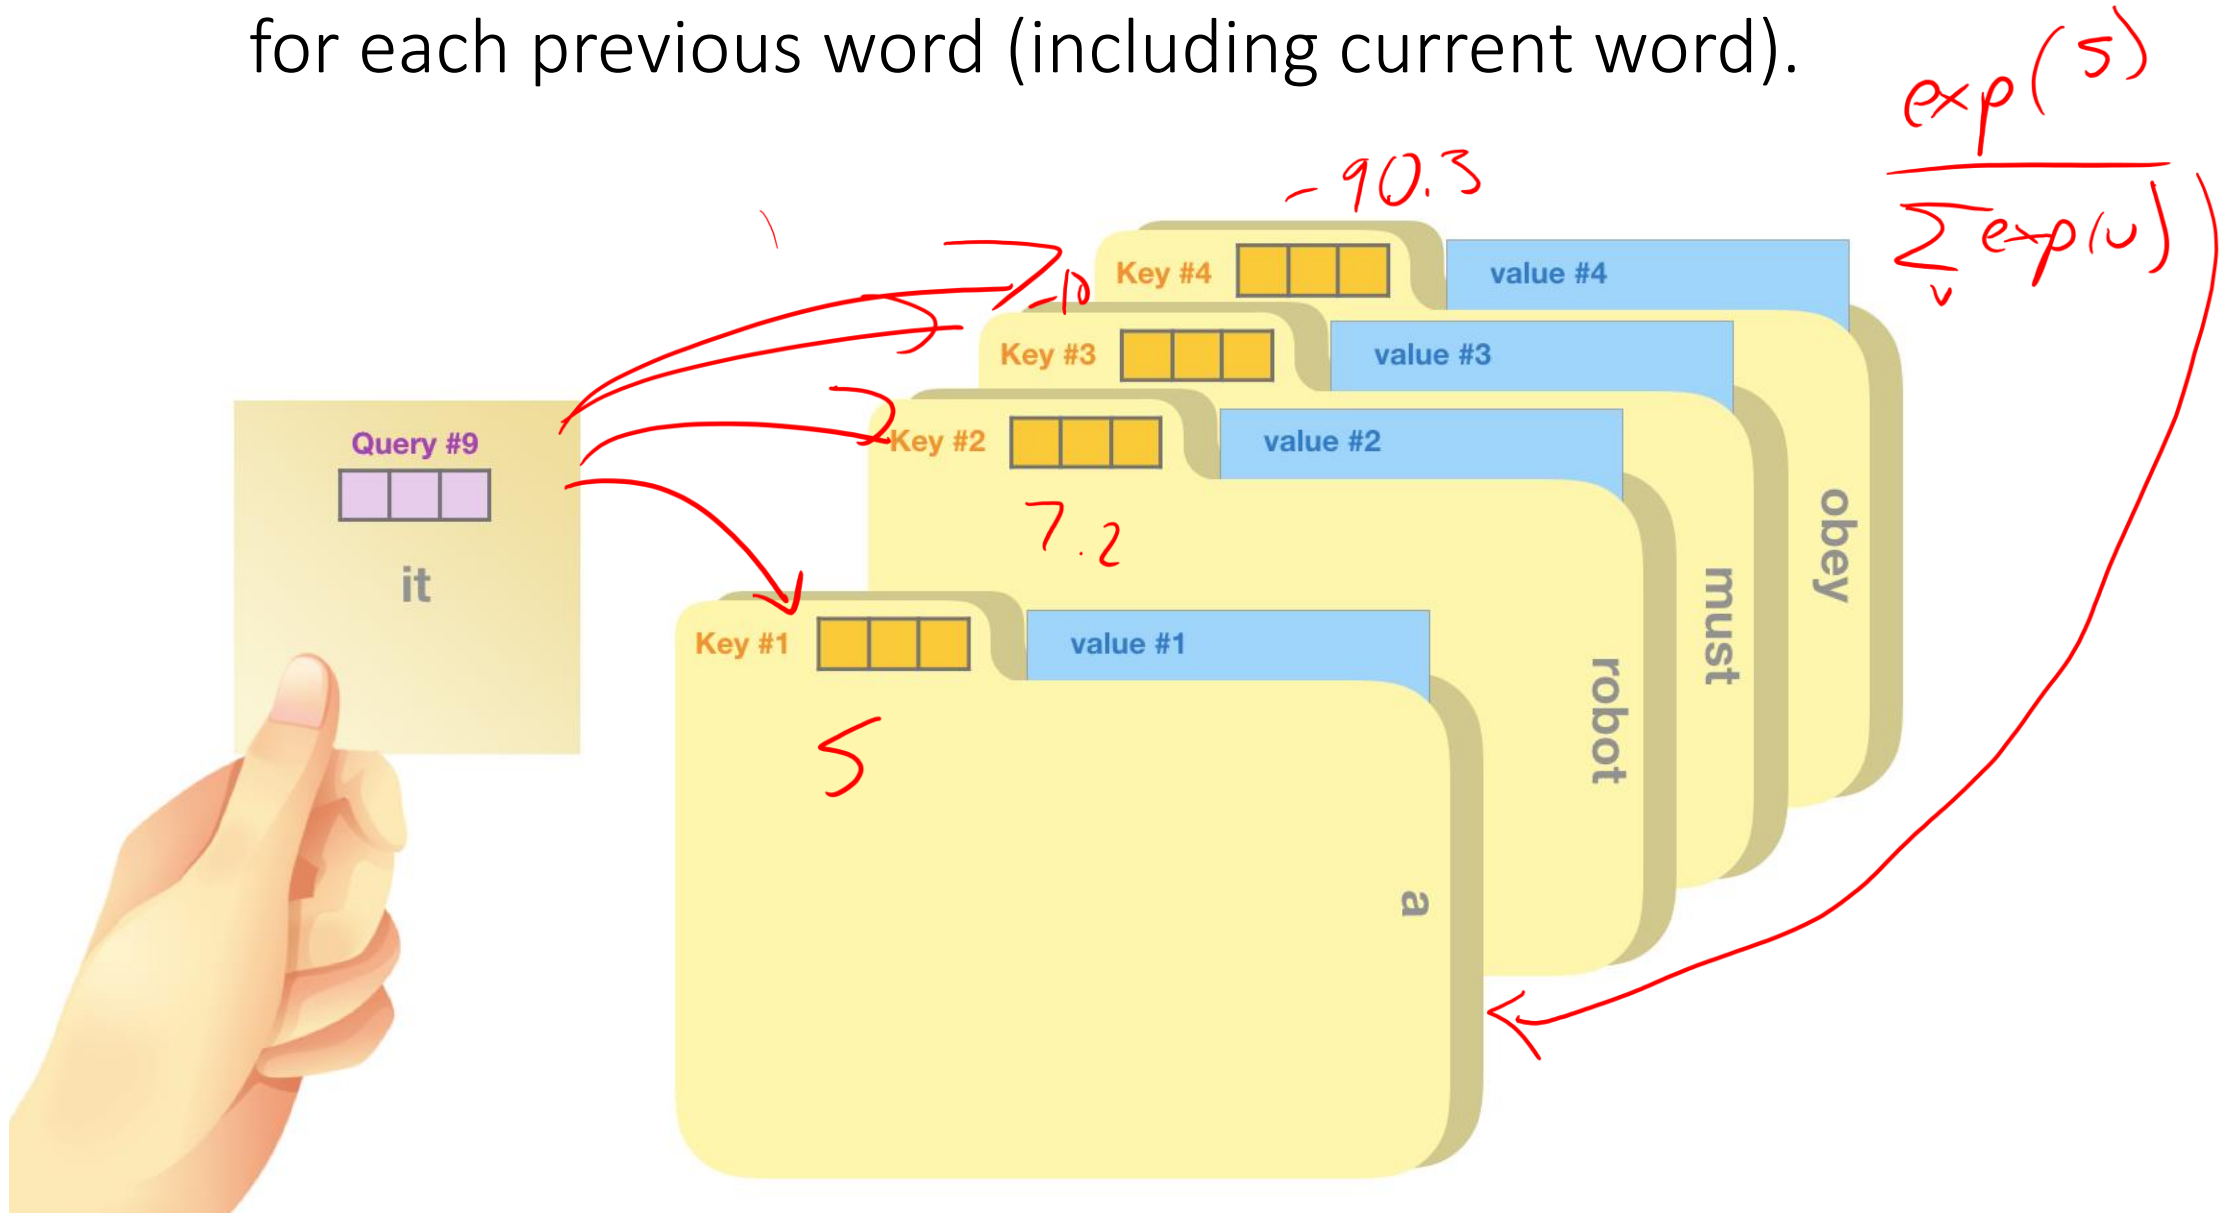
Values



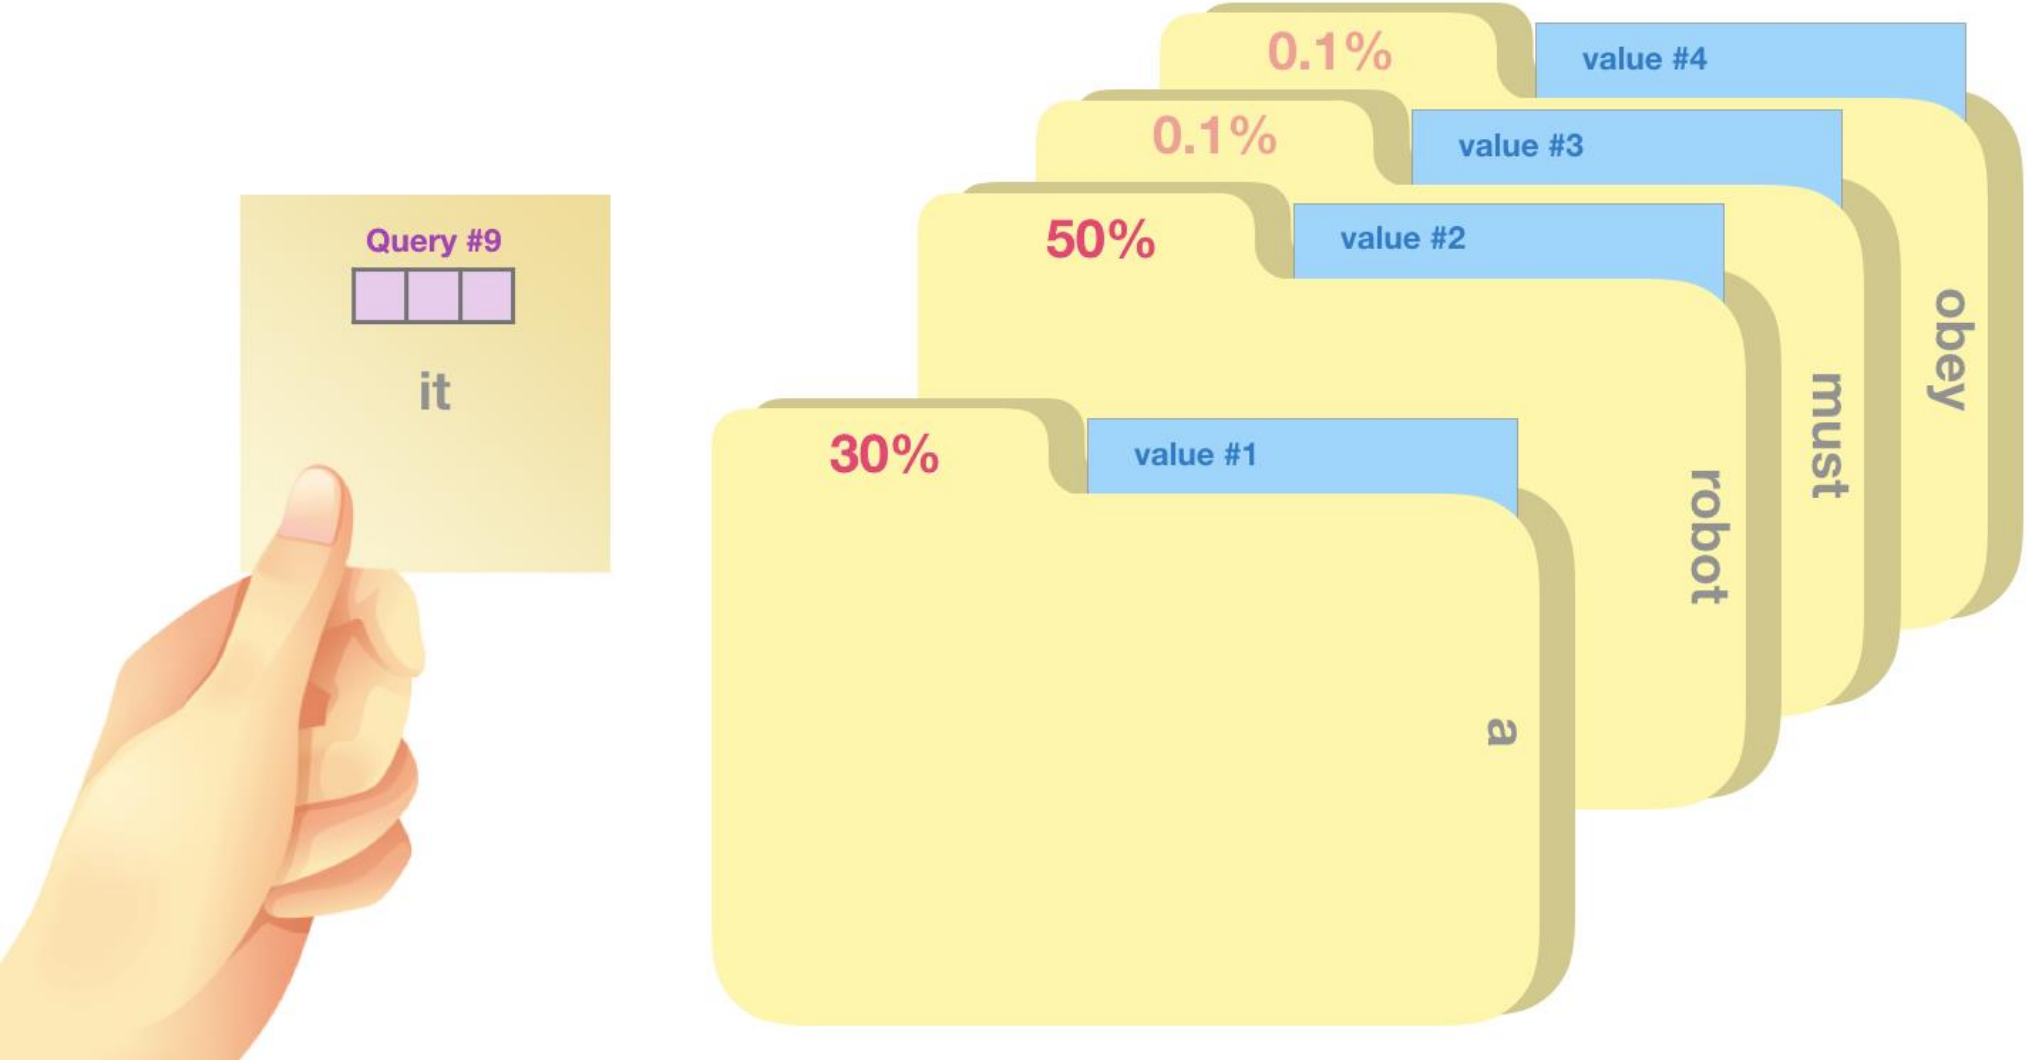
W^V


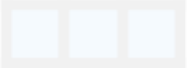





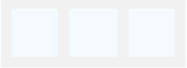



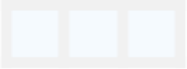









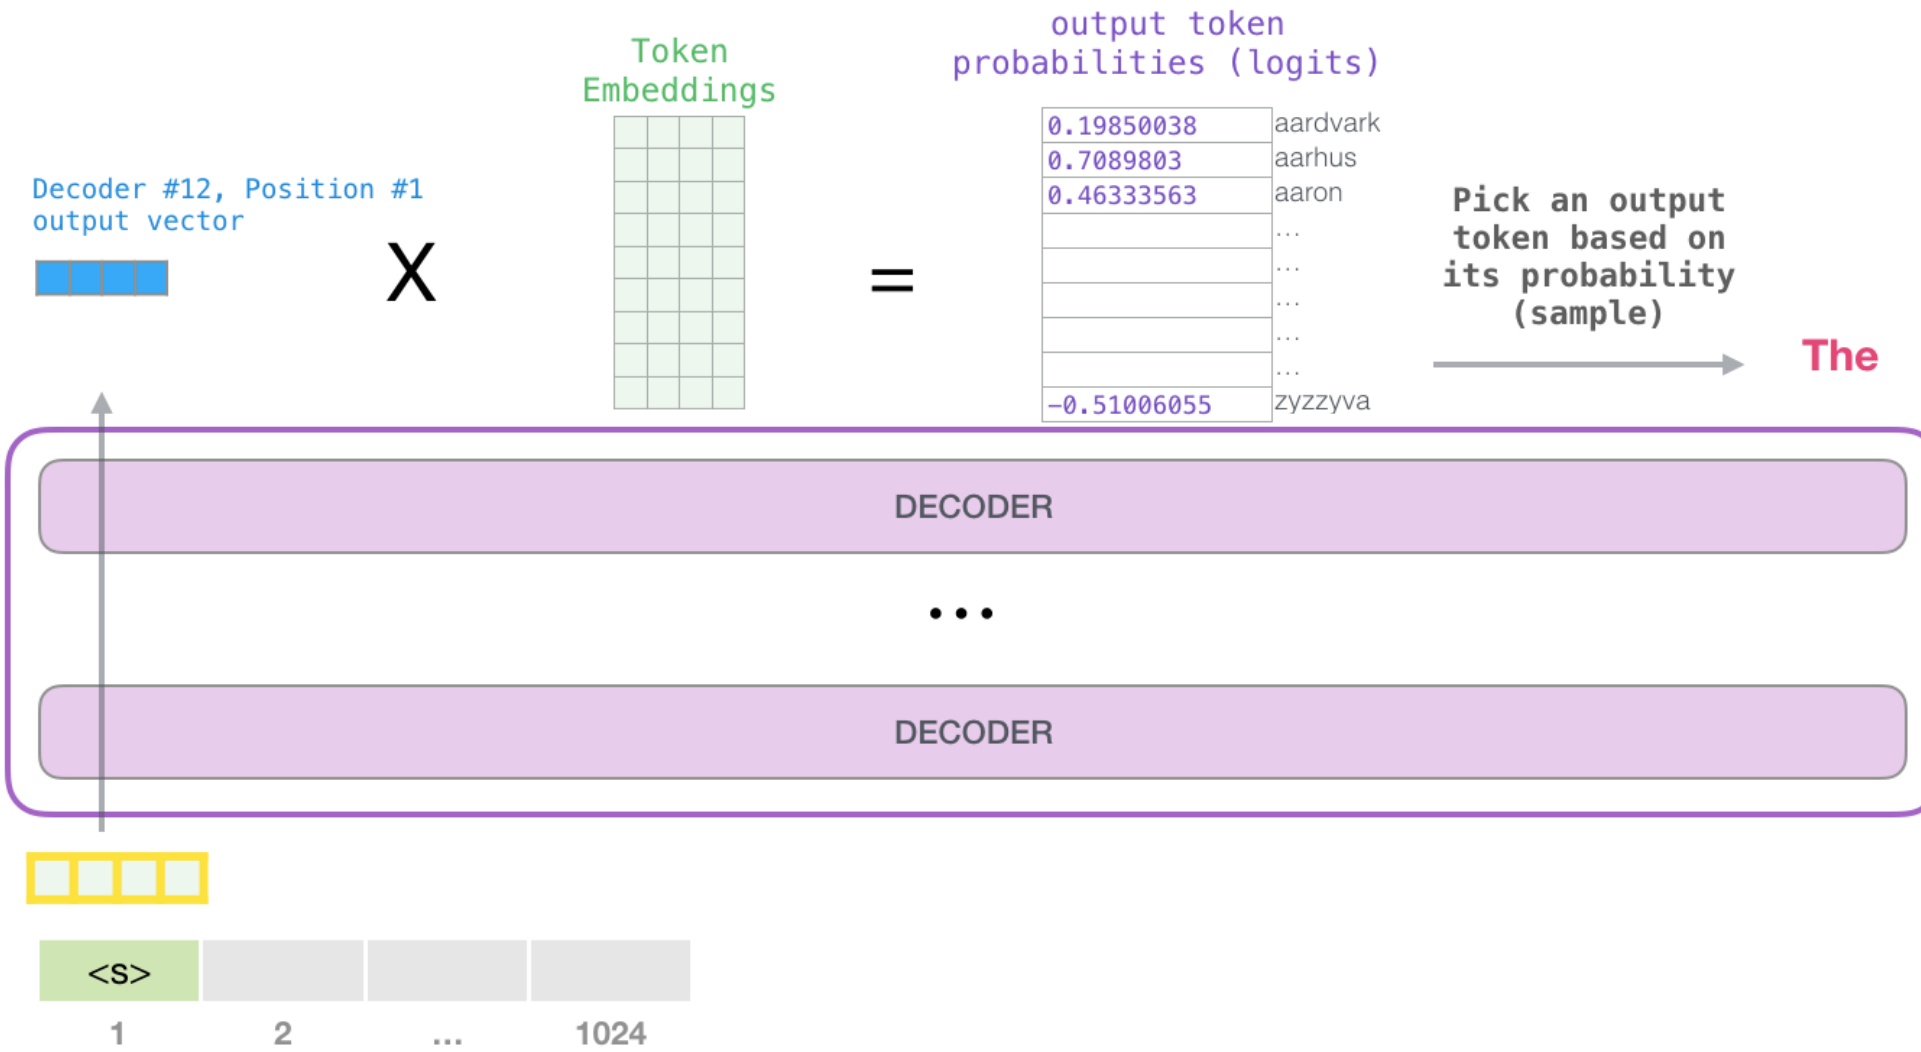
Perform dot product between query and all keys to get a raw score for each previous word (including current word).

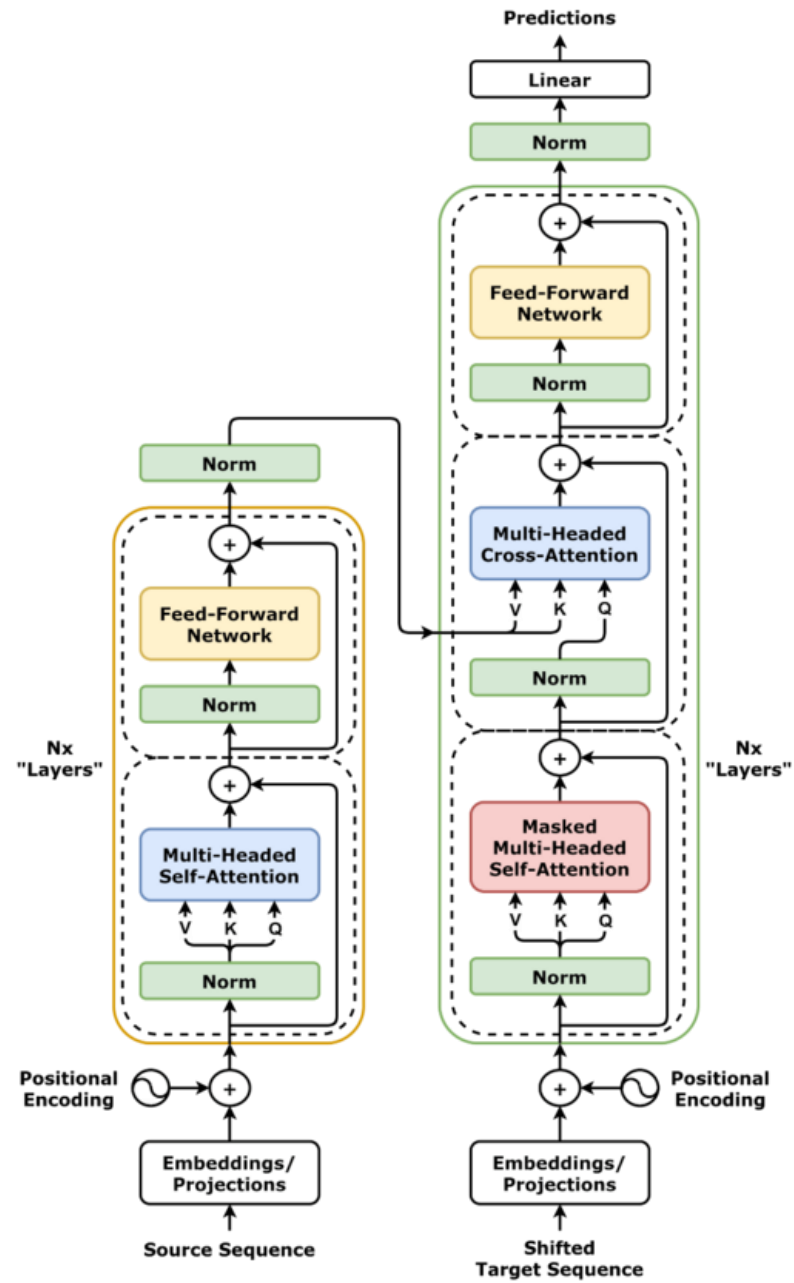


Normalize these scores via a softmax to get a probability distribution. Then return a weighted sum of the values.



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)

$$\alpha_t = 0 \rightarrow \mathcal{N}(0, I)$$

$$P(x_{t+1} | x_t) = \mathcal{N}(x_t | \sqrt{\alpha_t} x_t, (1 - \alpha_t) I)$$

↑ noise schedule

x_0

x_1

Fixed forward diffusion process



Data



Noise



Generative reverse denoising process

$$t \rightarrow \square \rightarrow \hat{\epsilon}$$

$$x_t \rightarrow \square \rightarrow \hat{\epsilon}$$

$$MSE = \min_{\theta} \|\hat{\epsilon}_{\theta}(t, x_t) - \epsilon_t\|^2$$

$$x_{T-1} = x_T - \hat{\epsilon}_T + \gamma$$

Gradient descent