



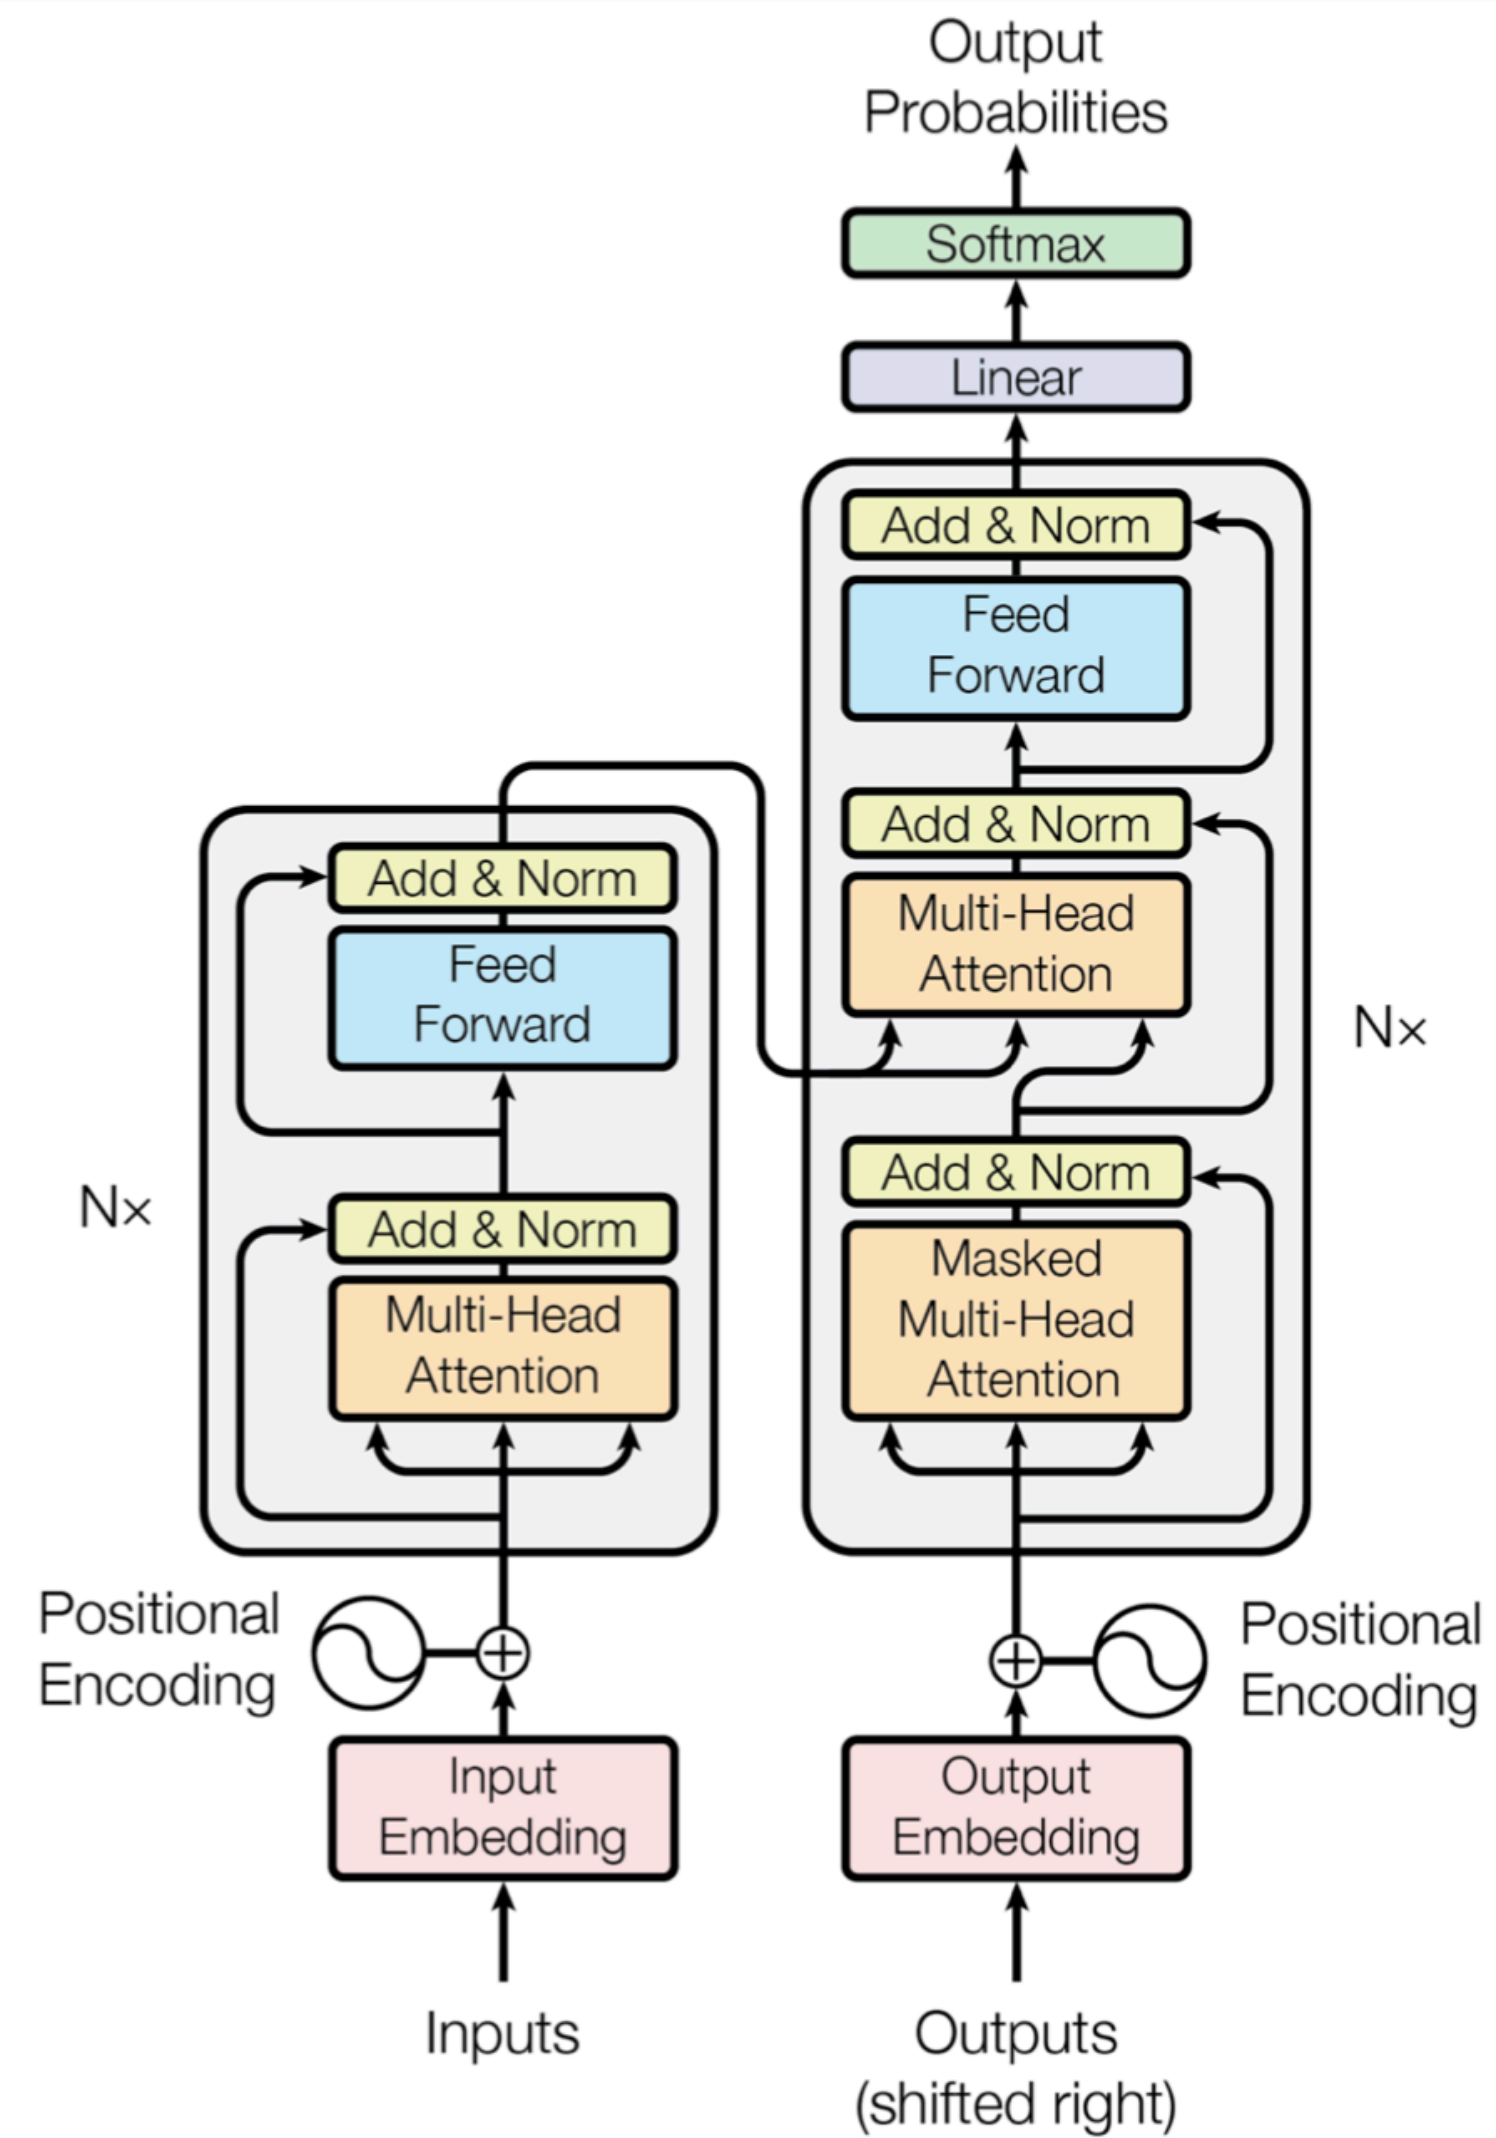
香港科技大學
THE HONG KONG
UNIVERSITY OF SCIENCE
AND TECHNOLOGY

COMP 5212
Machine Learning
Lecture 20

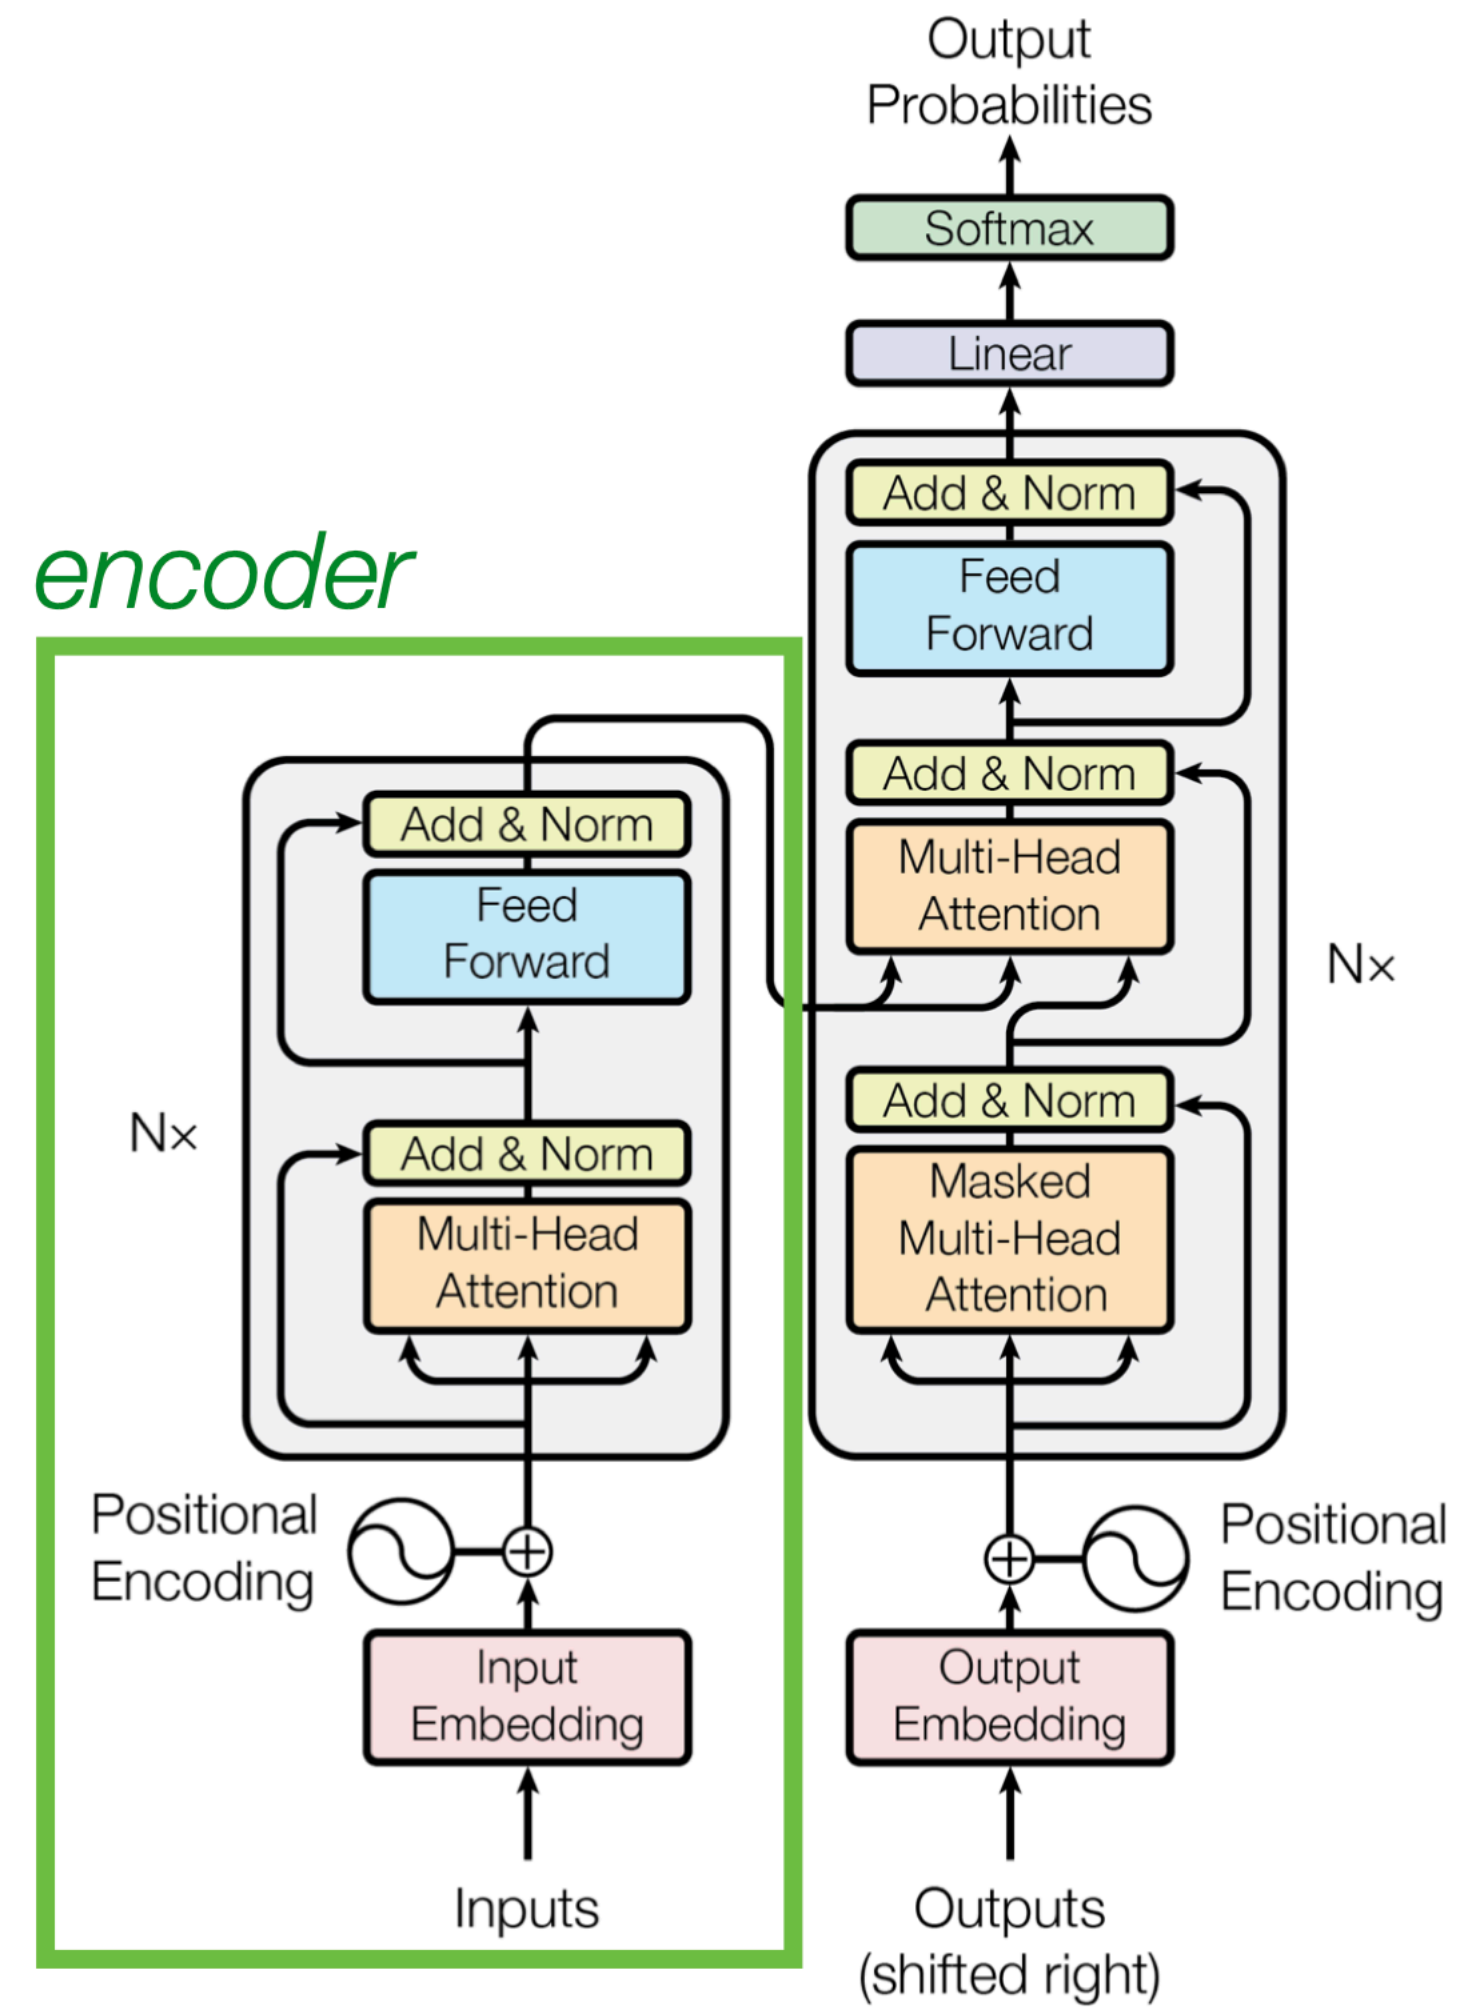
Transformers, VAEs

Junxian He
Apr 19, 2024

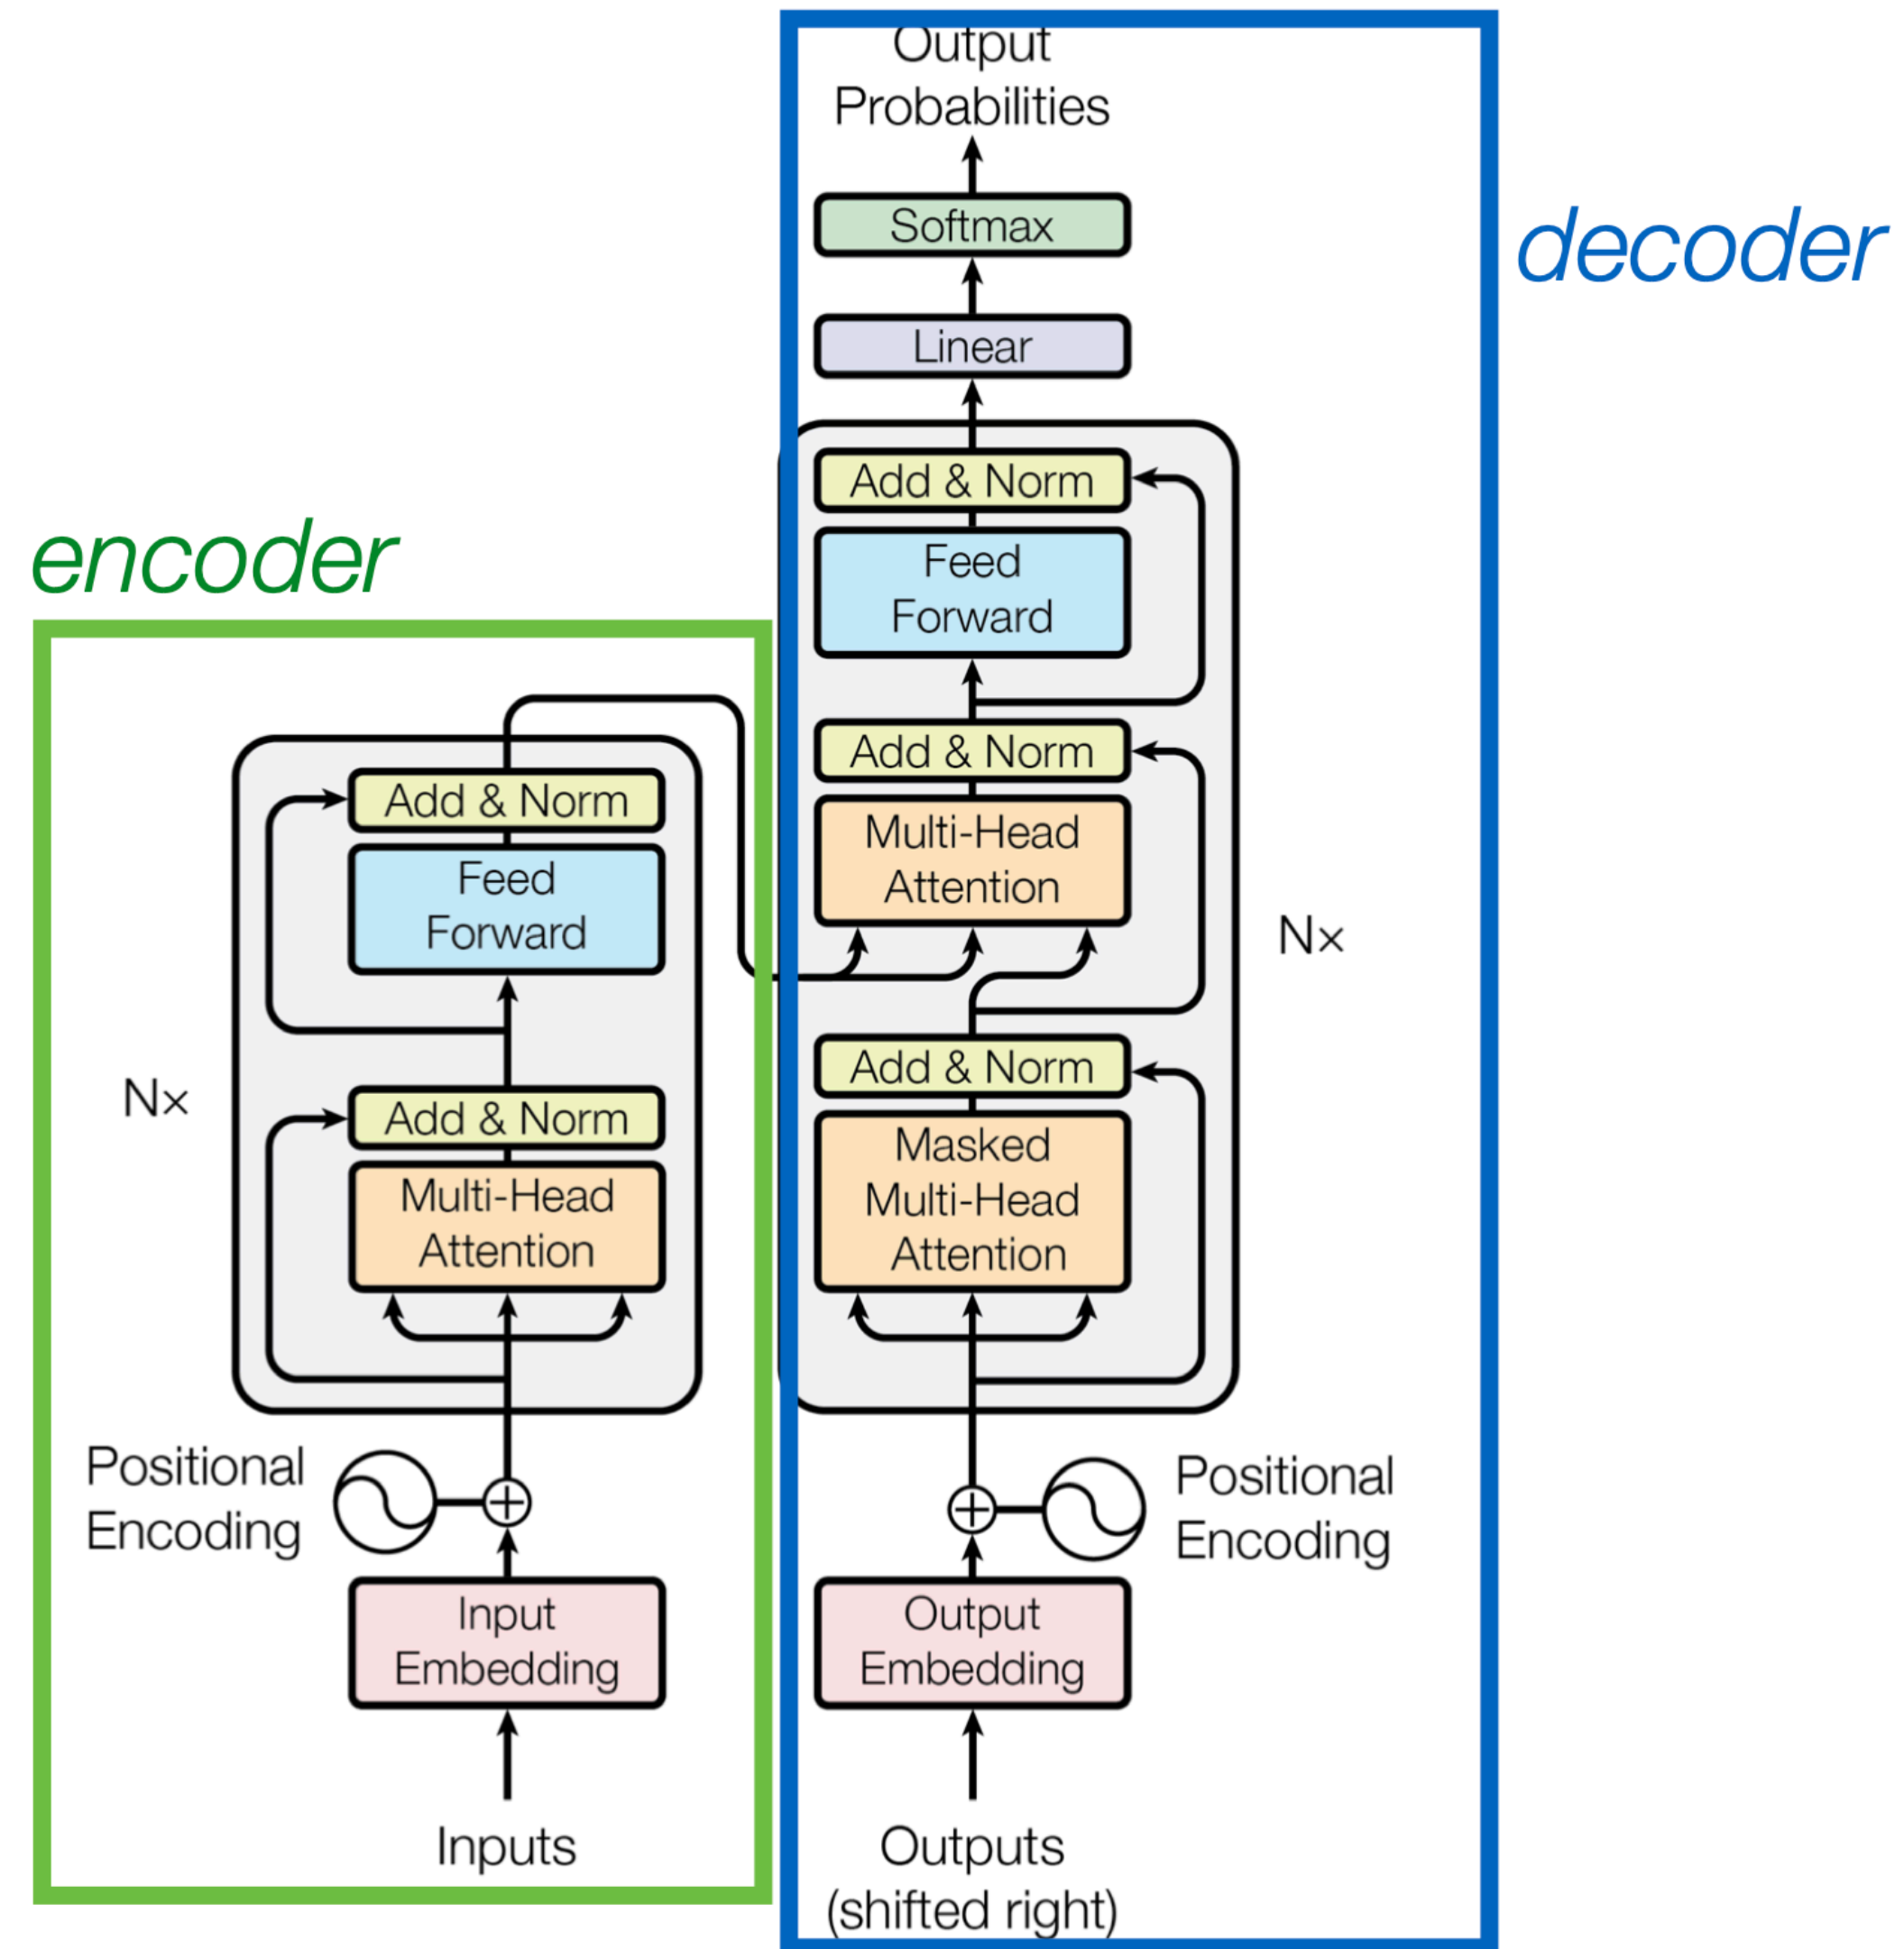
Transformer



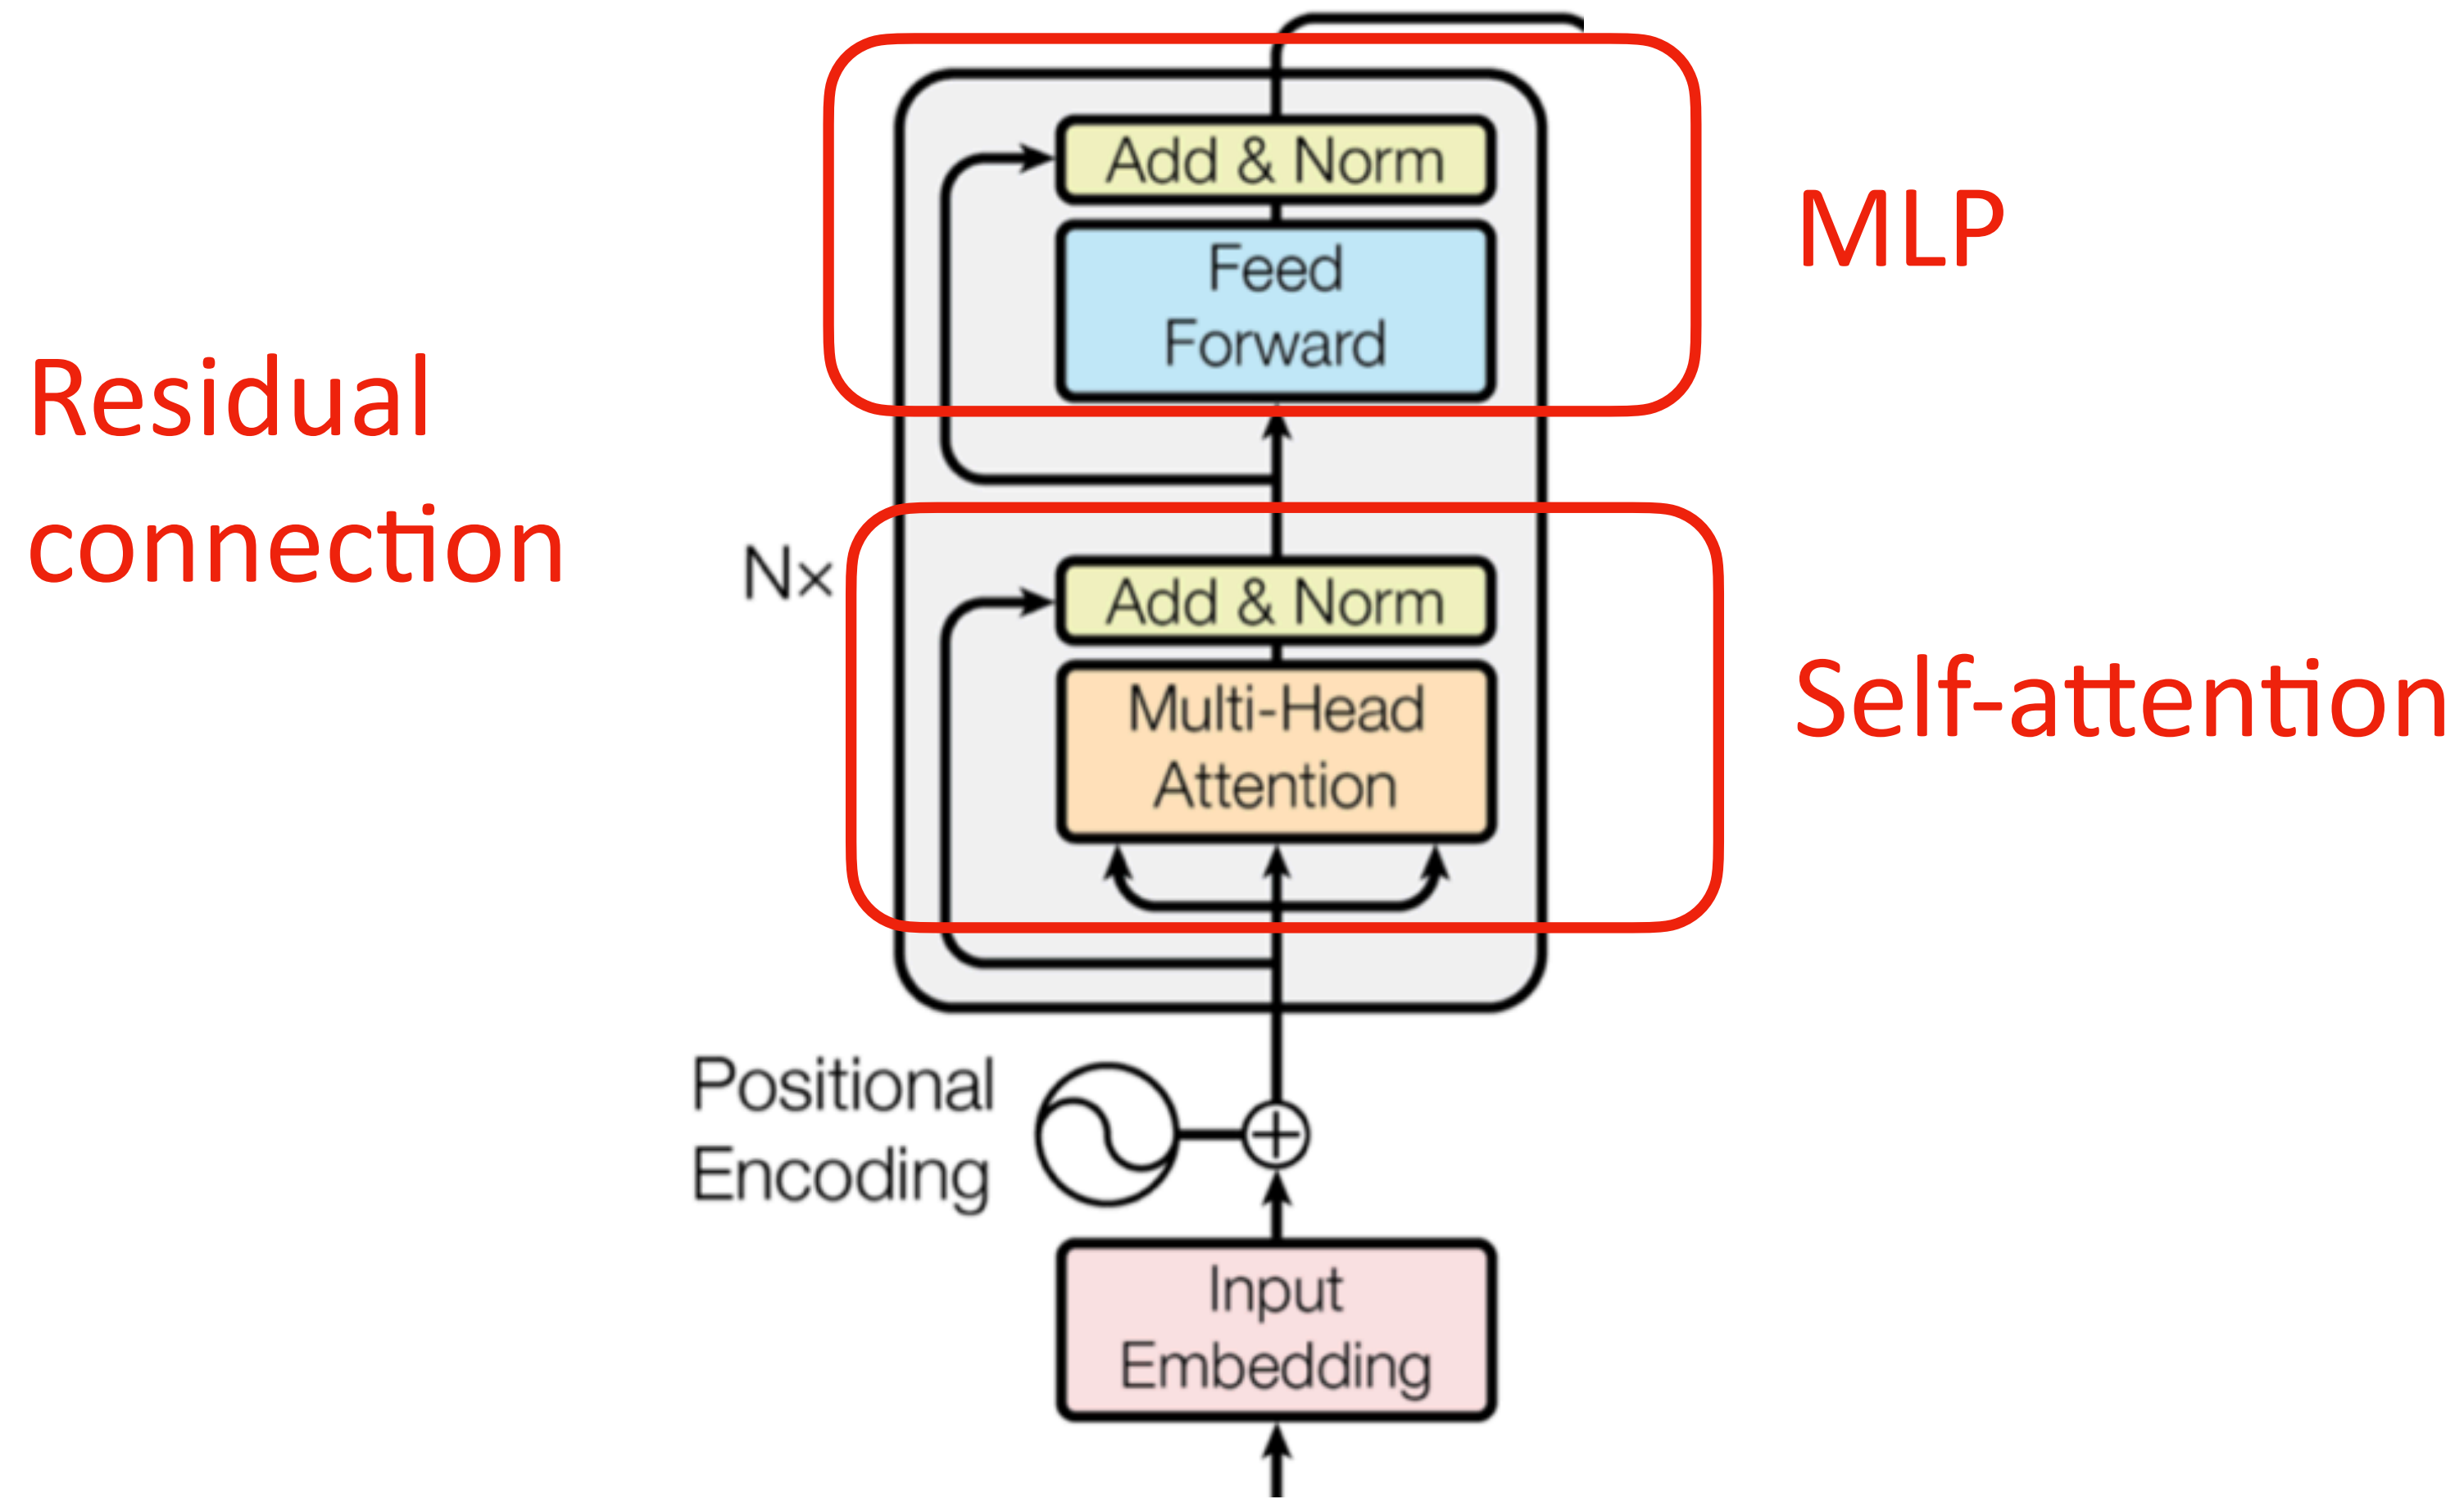
Encoder



Decoder



Transformer Encoder




What is Attention

$$Q \in R^{n \times d} \quad K \in R^{m \times d} \quad V \in R^{m \times d}$$

We have n queries, m (key, value) pairs

Scaled Dot-Product Attention



Q: Query
K: key
V: value

$$\text{Attention weight} = \text{softmax}(QK^T)$$


Dot-products grow large in magnitude

$$\text{Scaled Attention weight} = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right) \quad \text{Shape is } m \times n$$

Attention weight represents the strength to "attend" values V

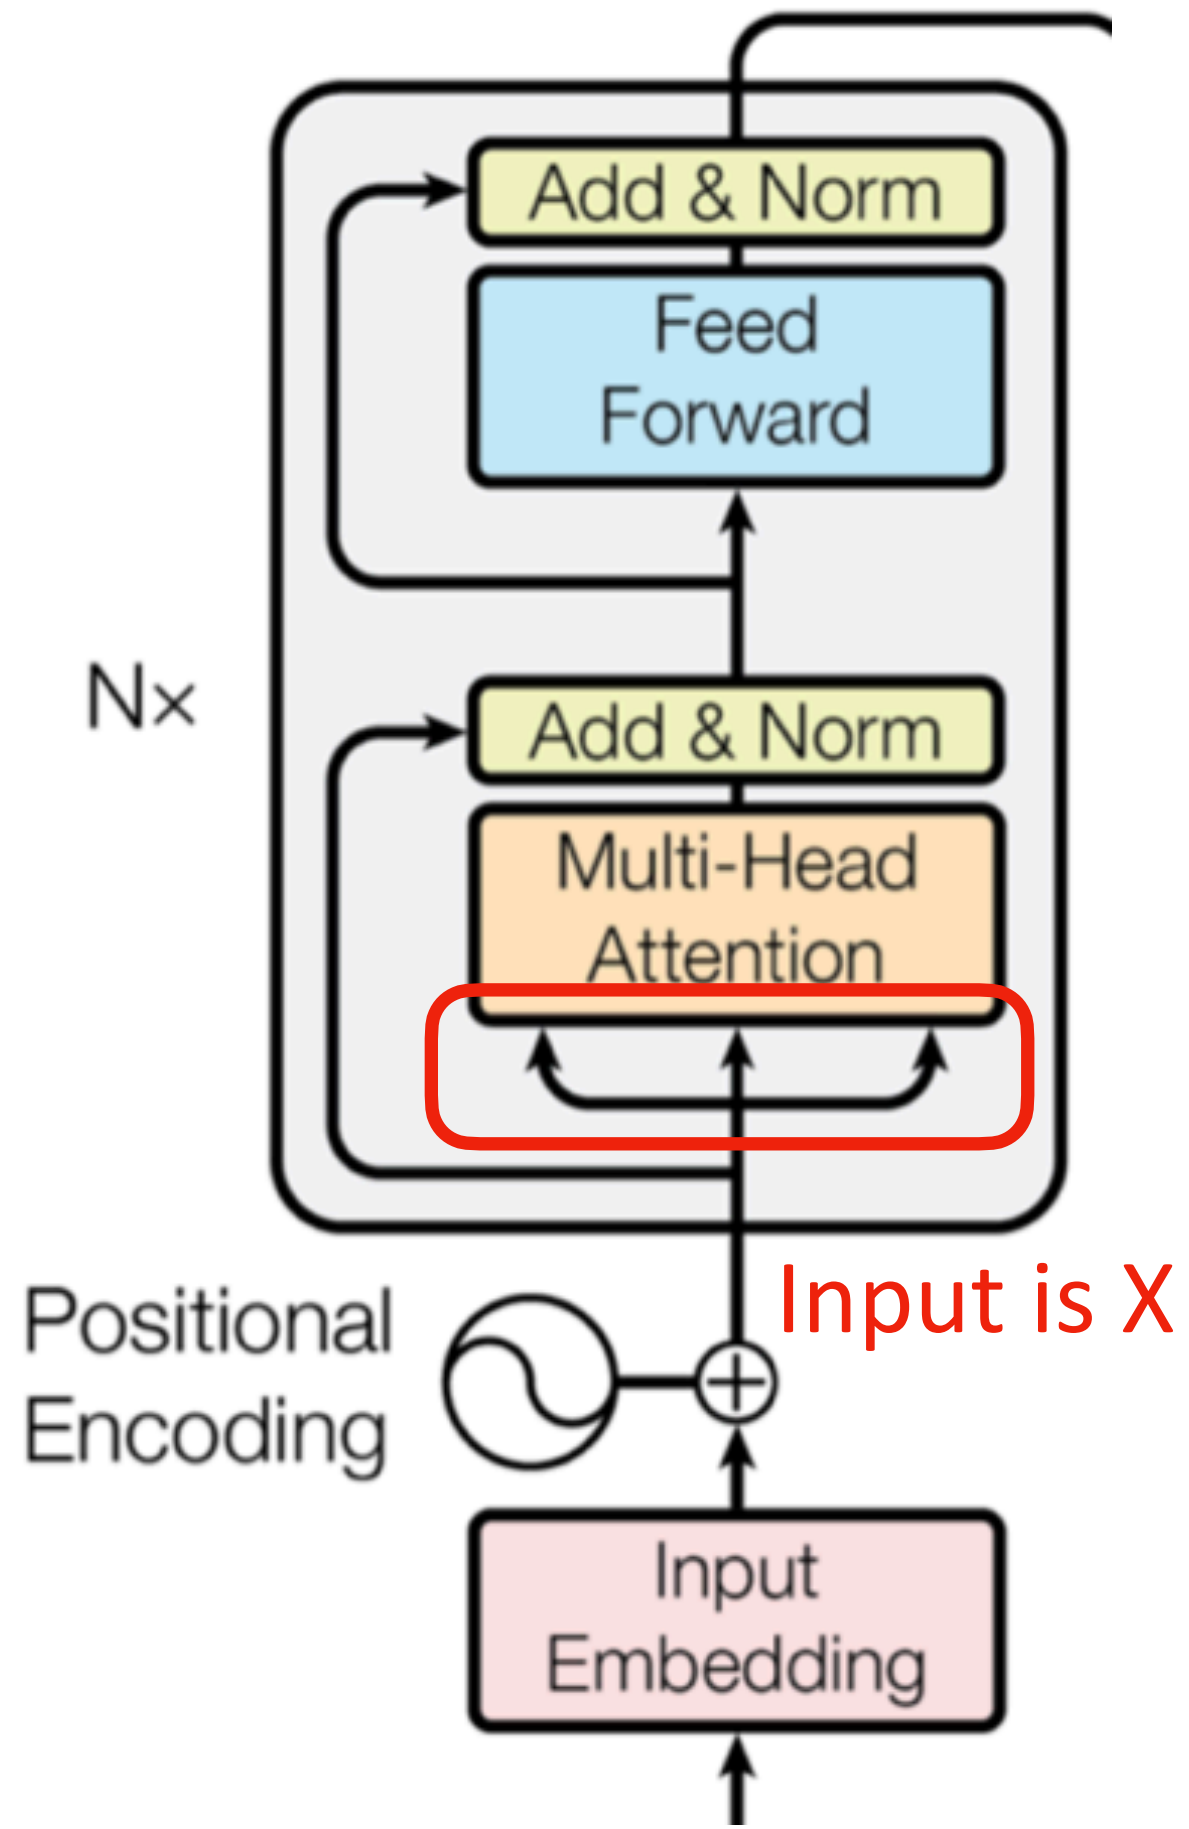
$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

Q, K, V



What are Q, K, V in the transformer

Self-Attention



8

$$X \times W^Q = Q$$

$$X \times W^K = K$$

$$X \times W^V = V$$


Query, key, and value are from the same input, thus it is called "self"-attention

$$\text{softmax} \left(\frac{Q \times K^T}{\sqrt{d_k}} \right) V = Z$$

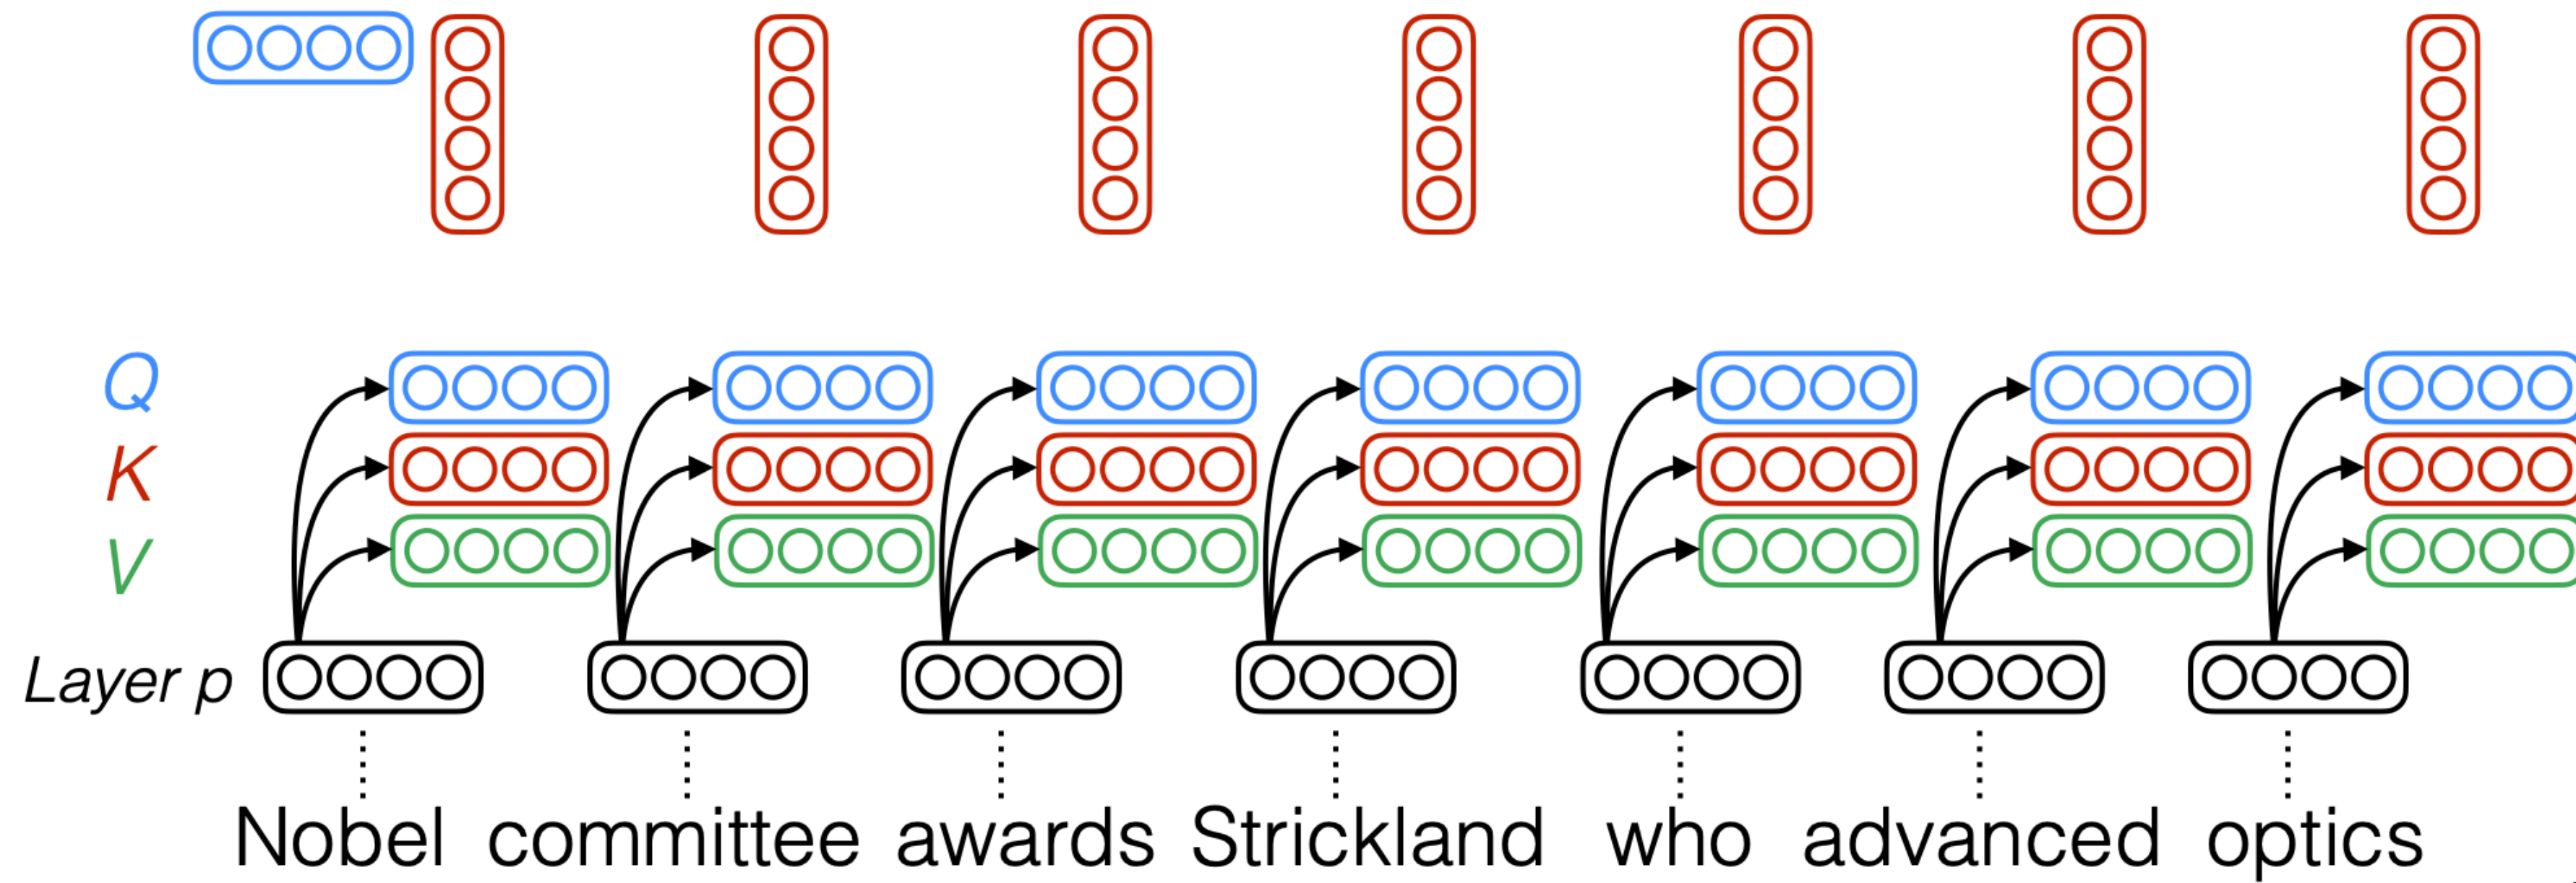
8

Self-Attention

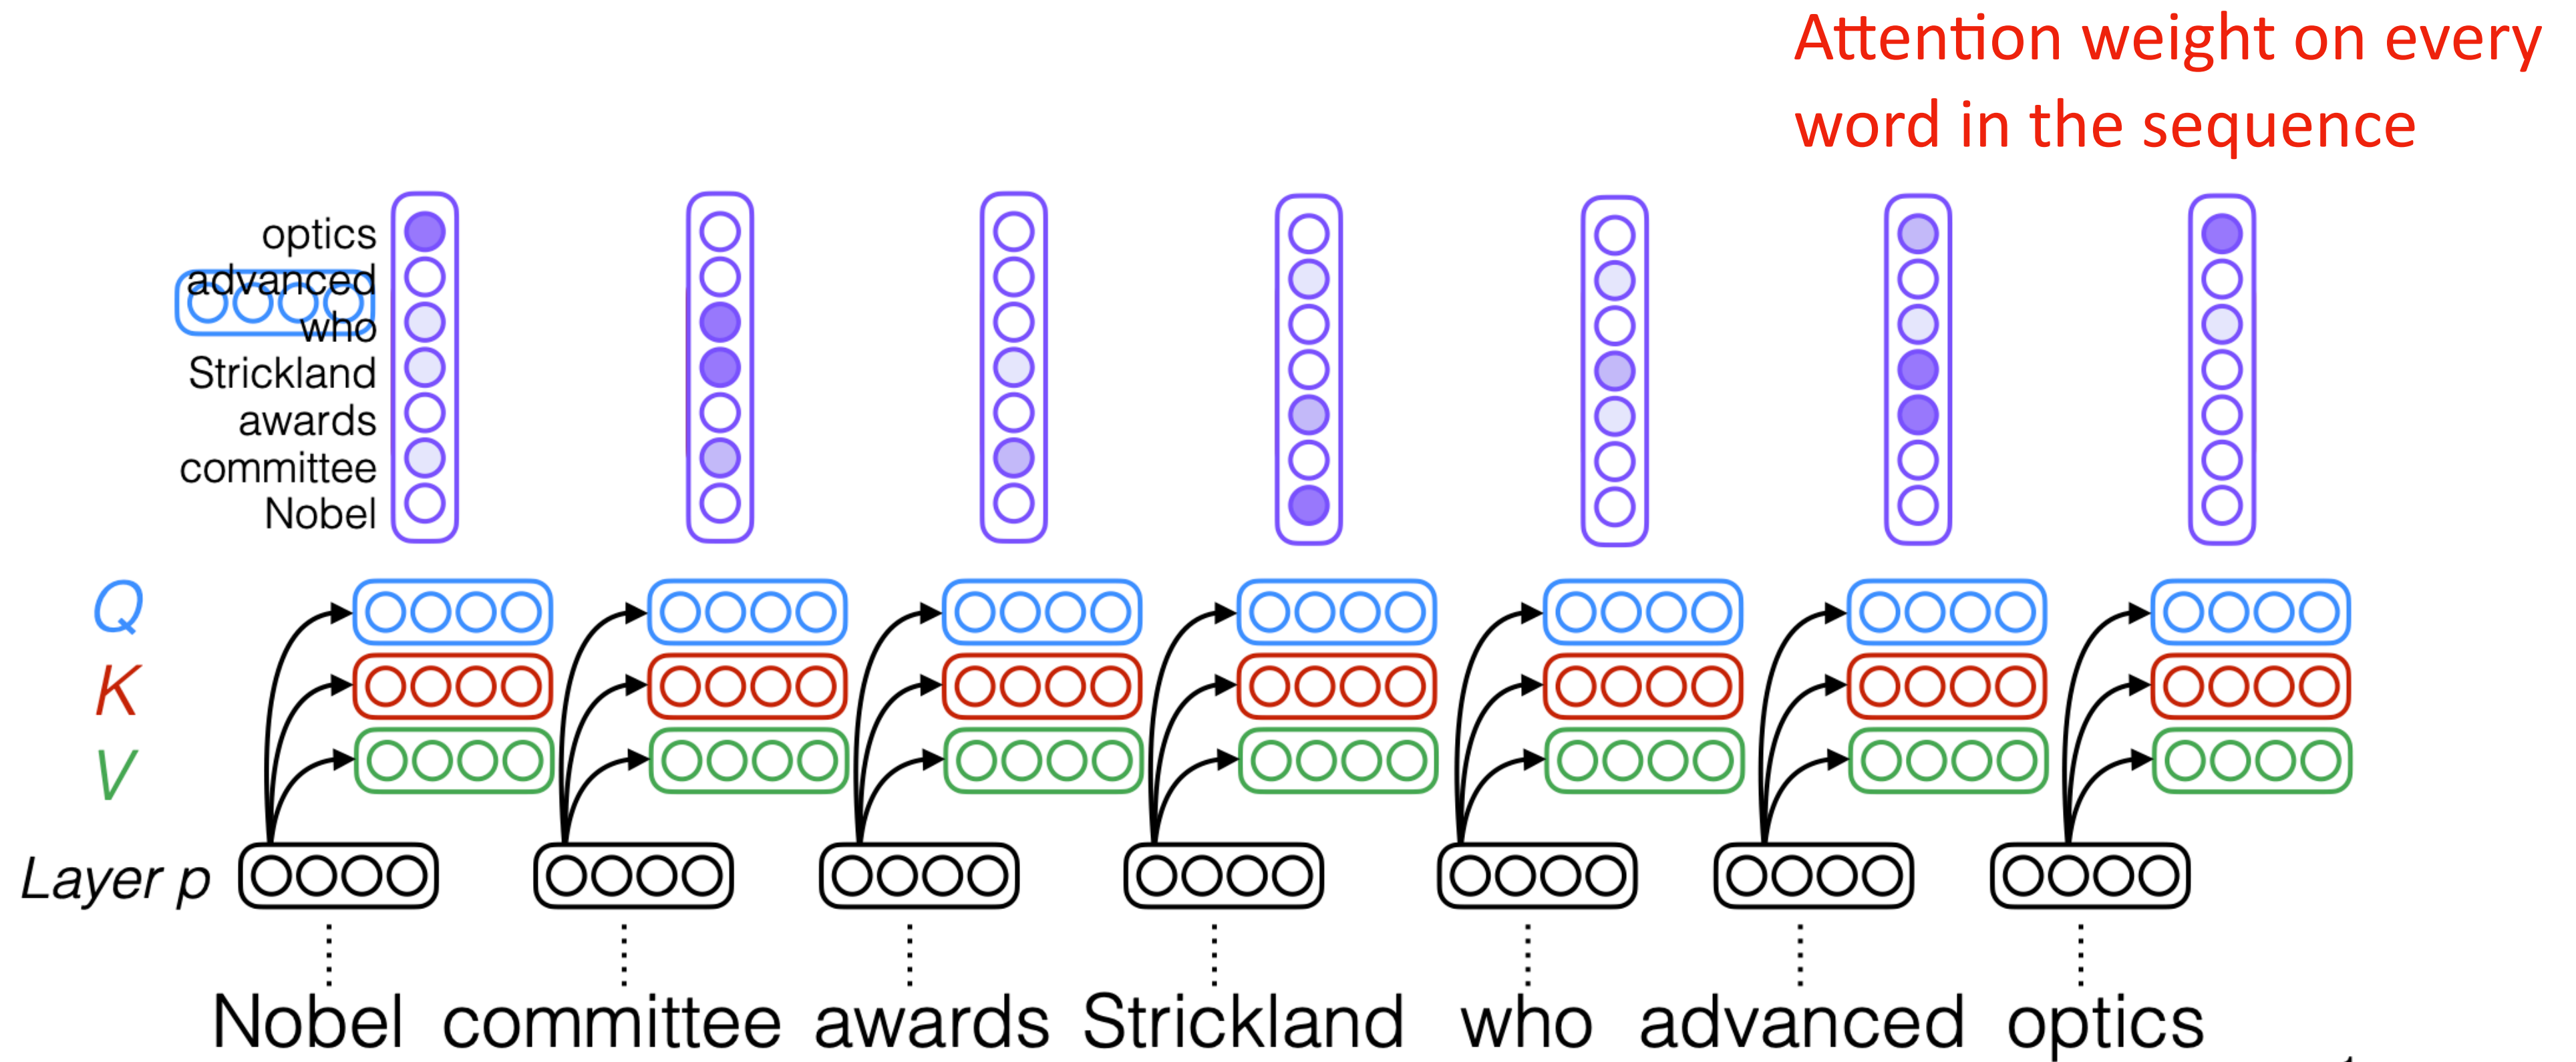
At each step, the attention computation attends to all steps in the input example



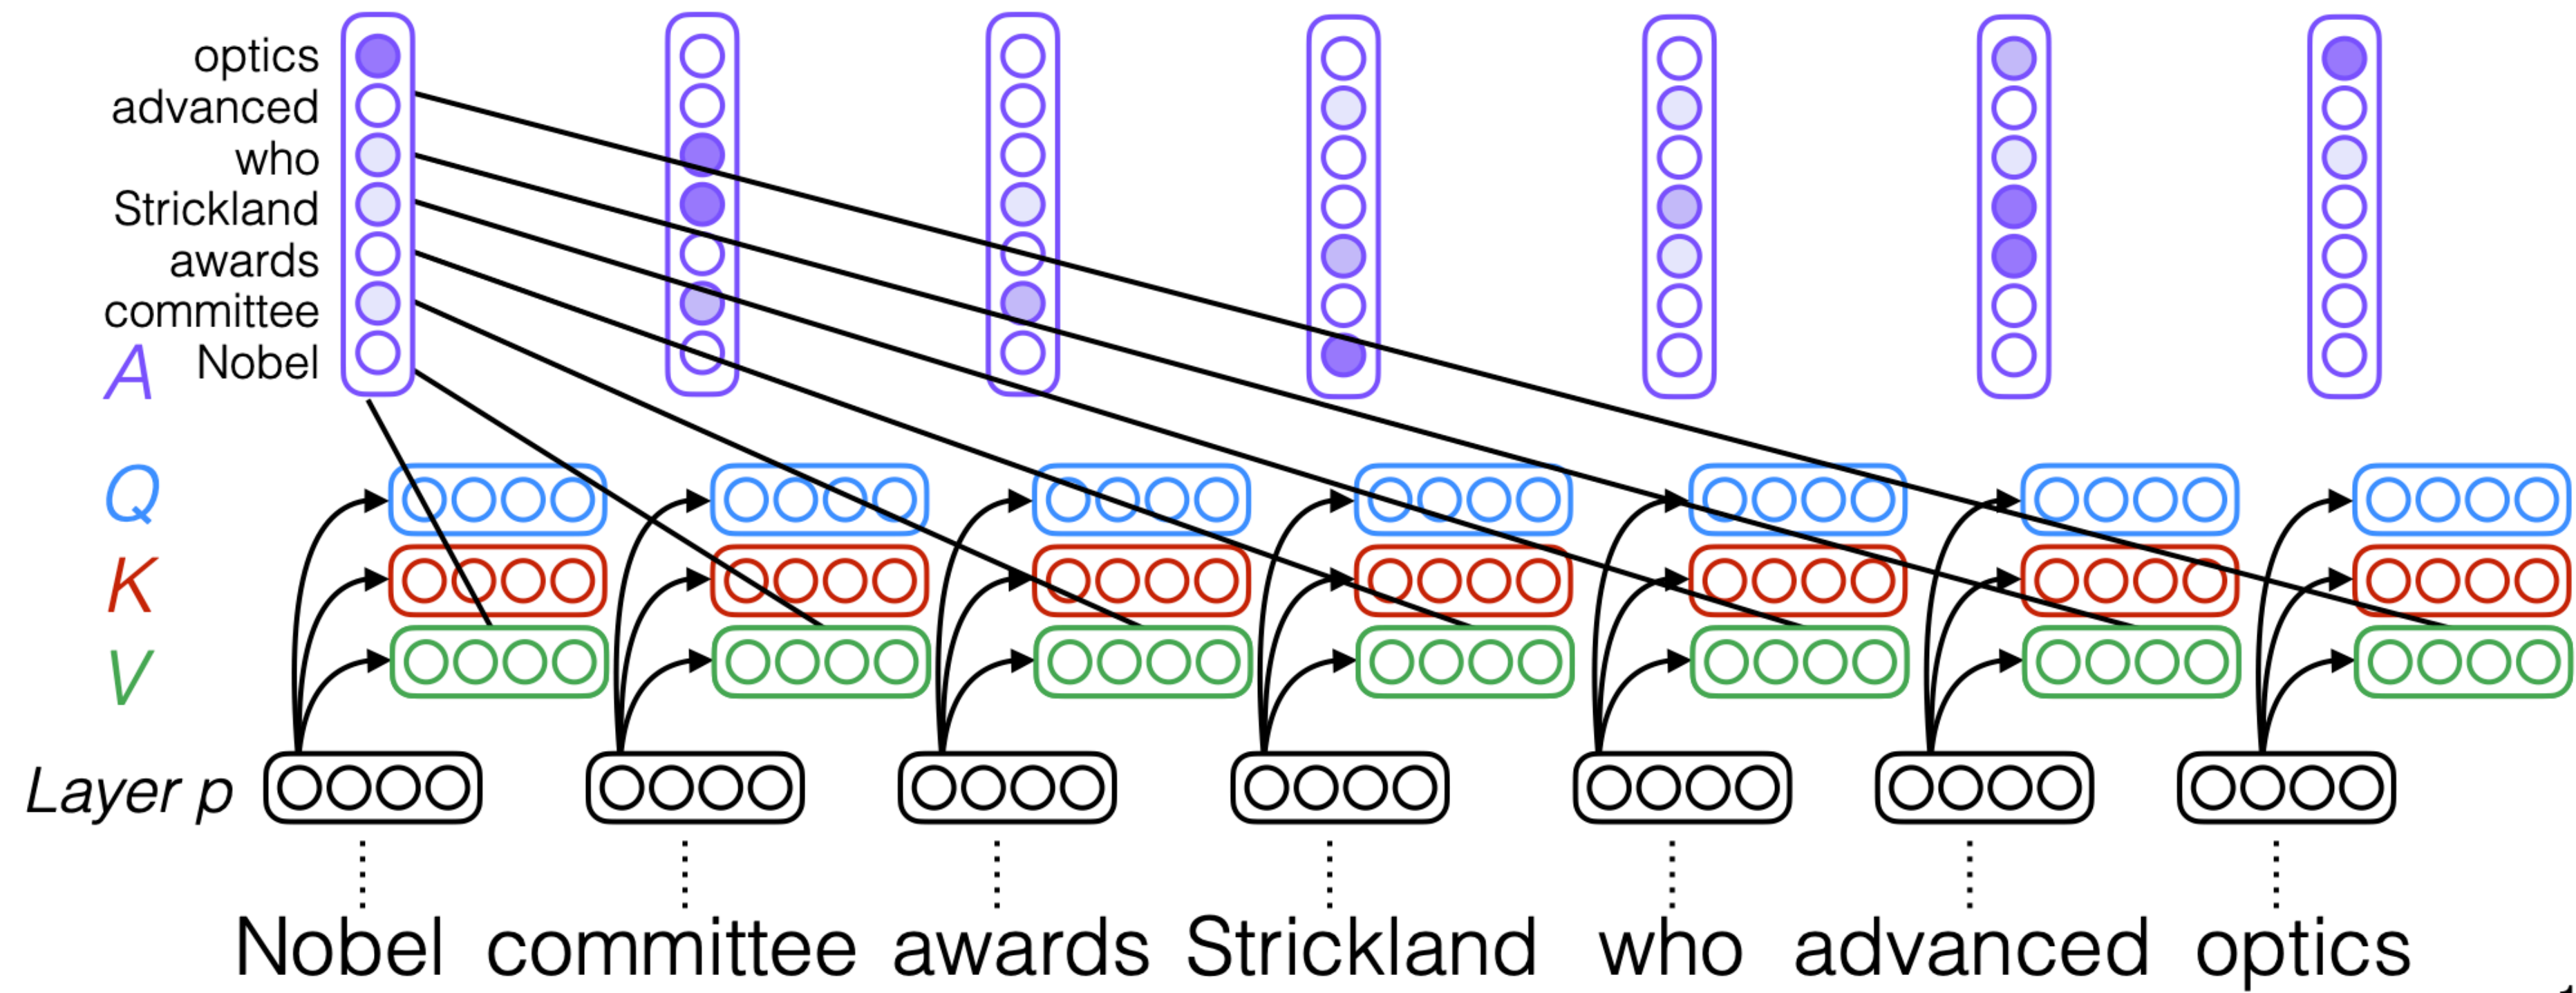
Self-Attention



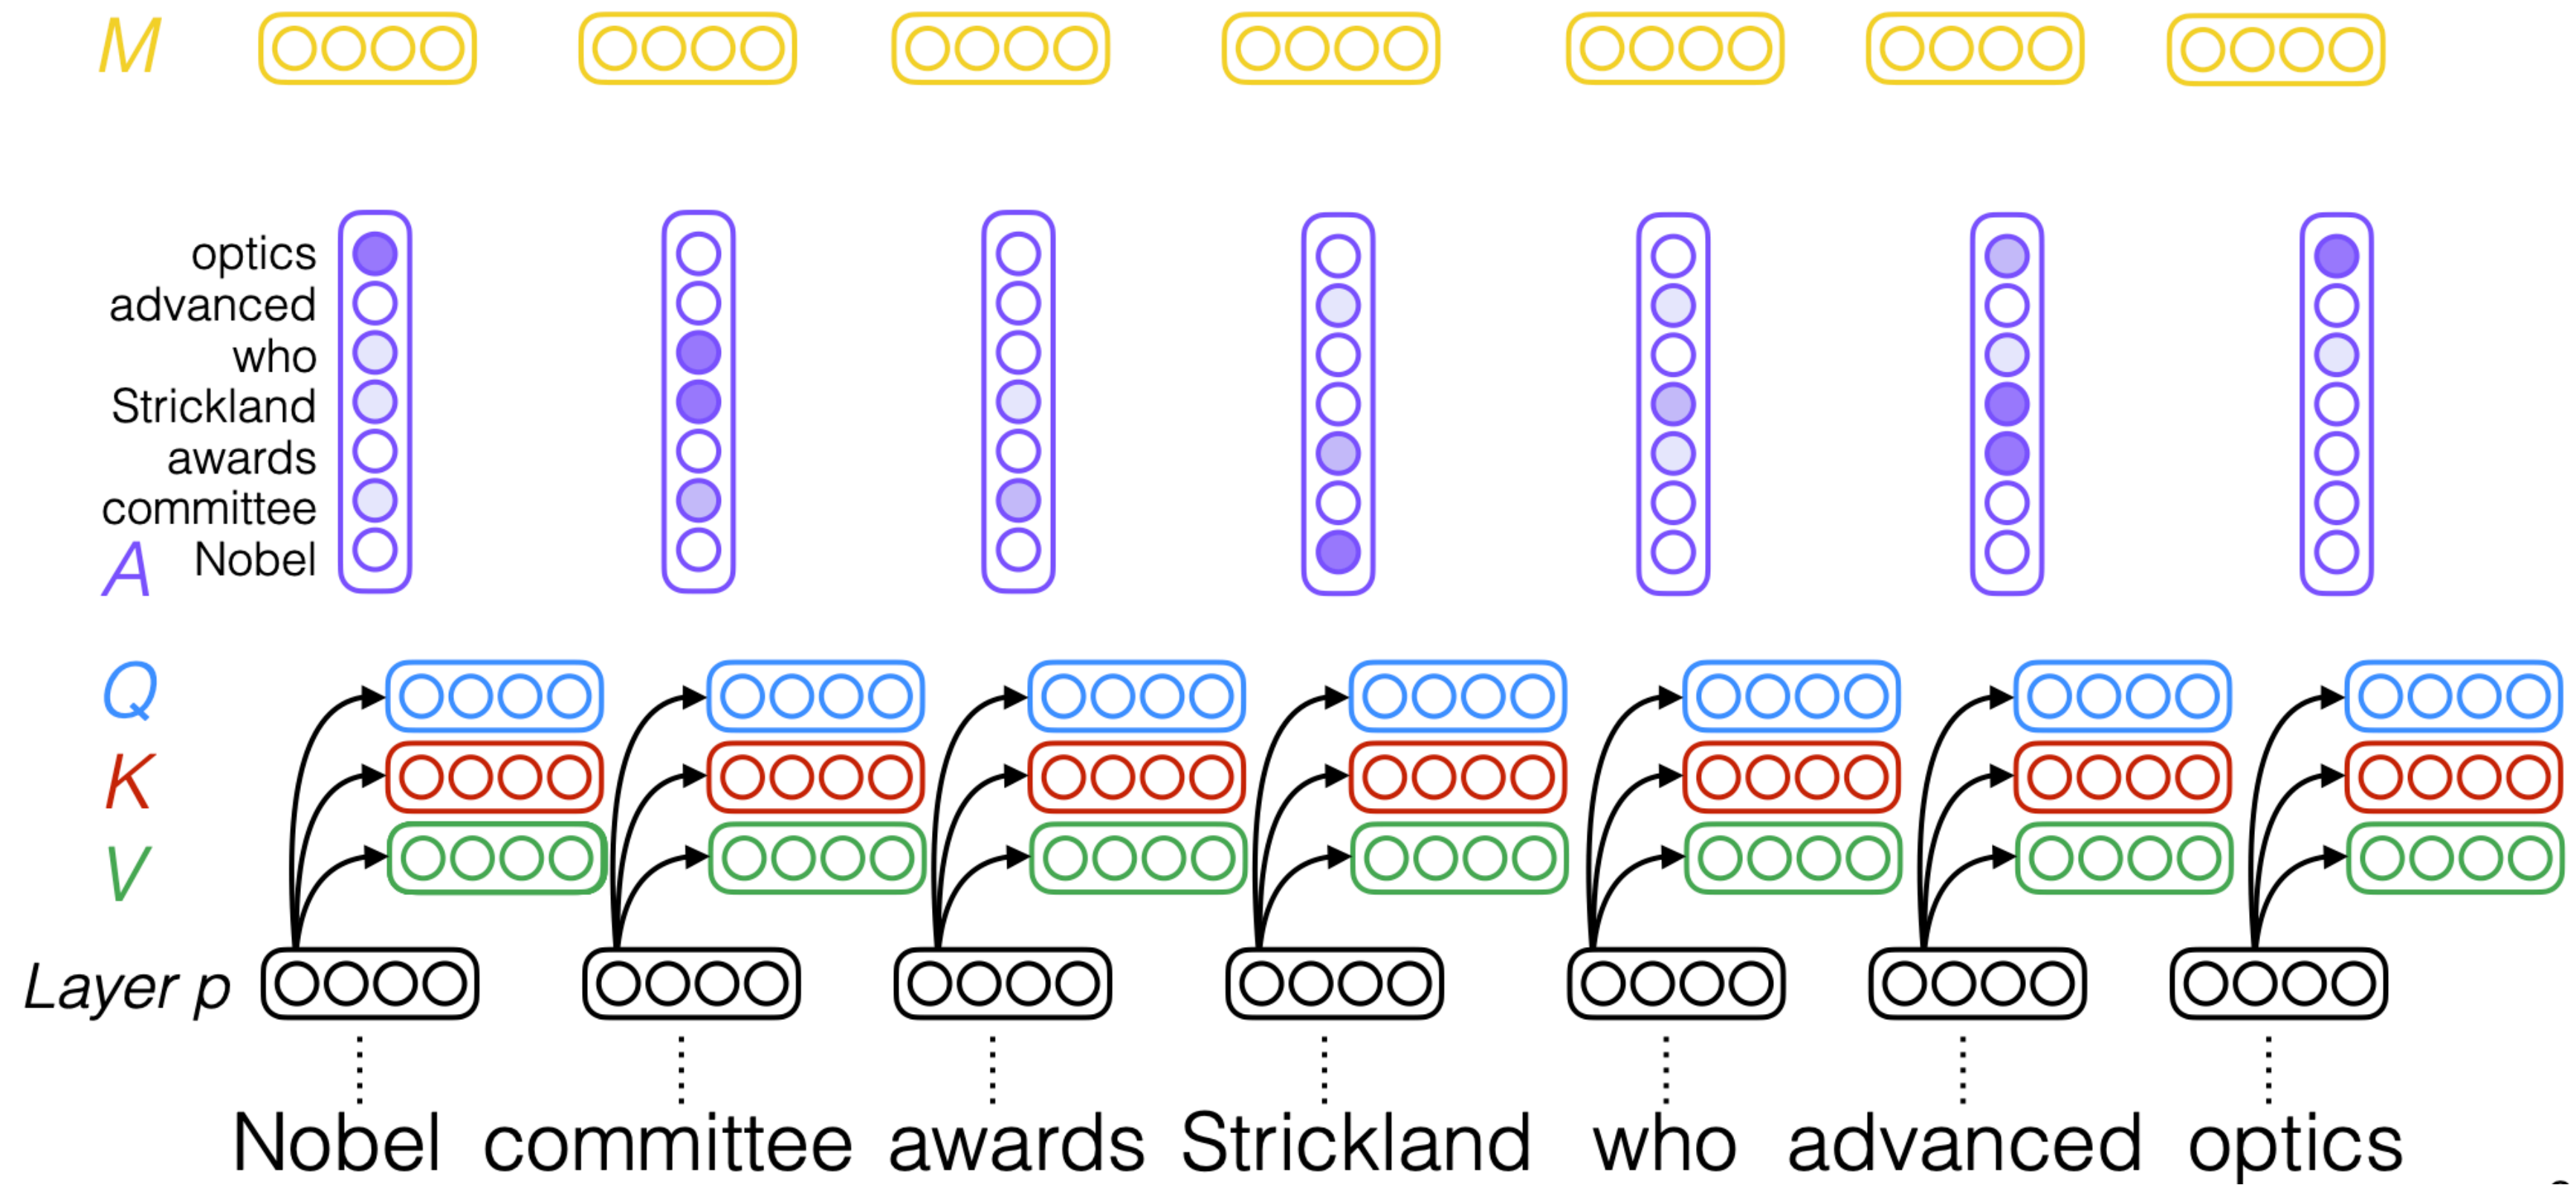
Self-Attention



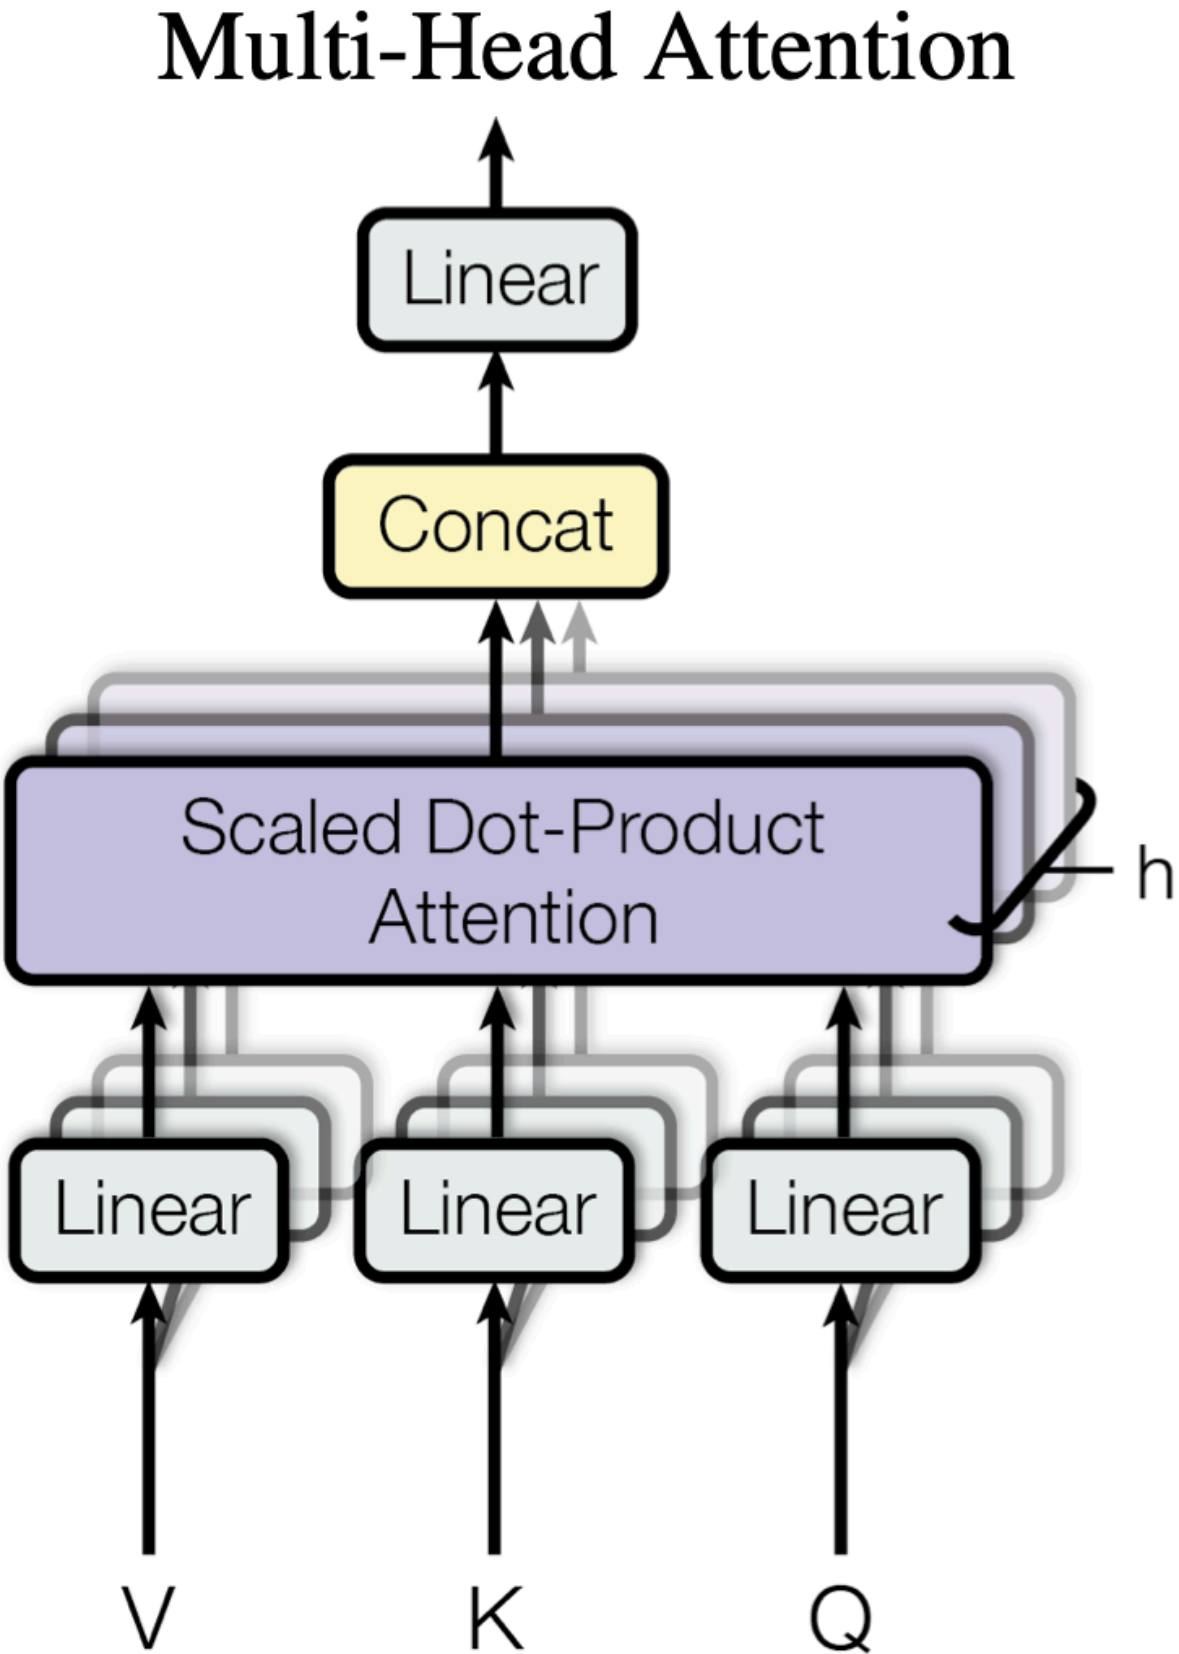
Self-Attention



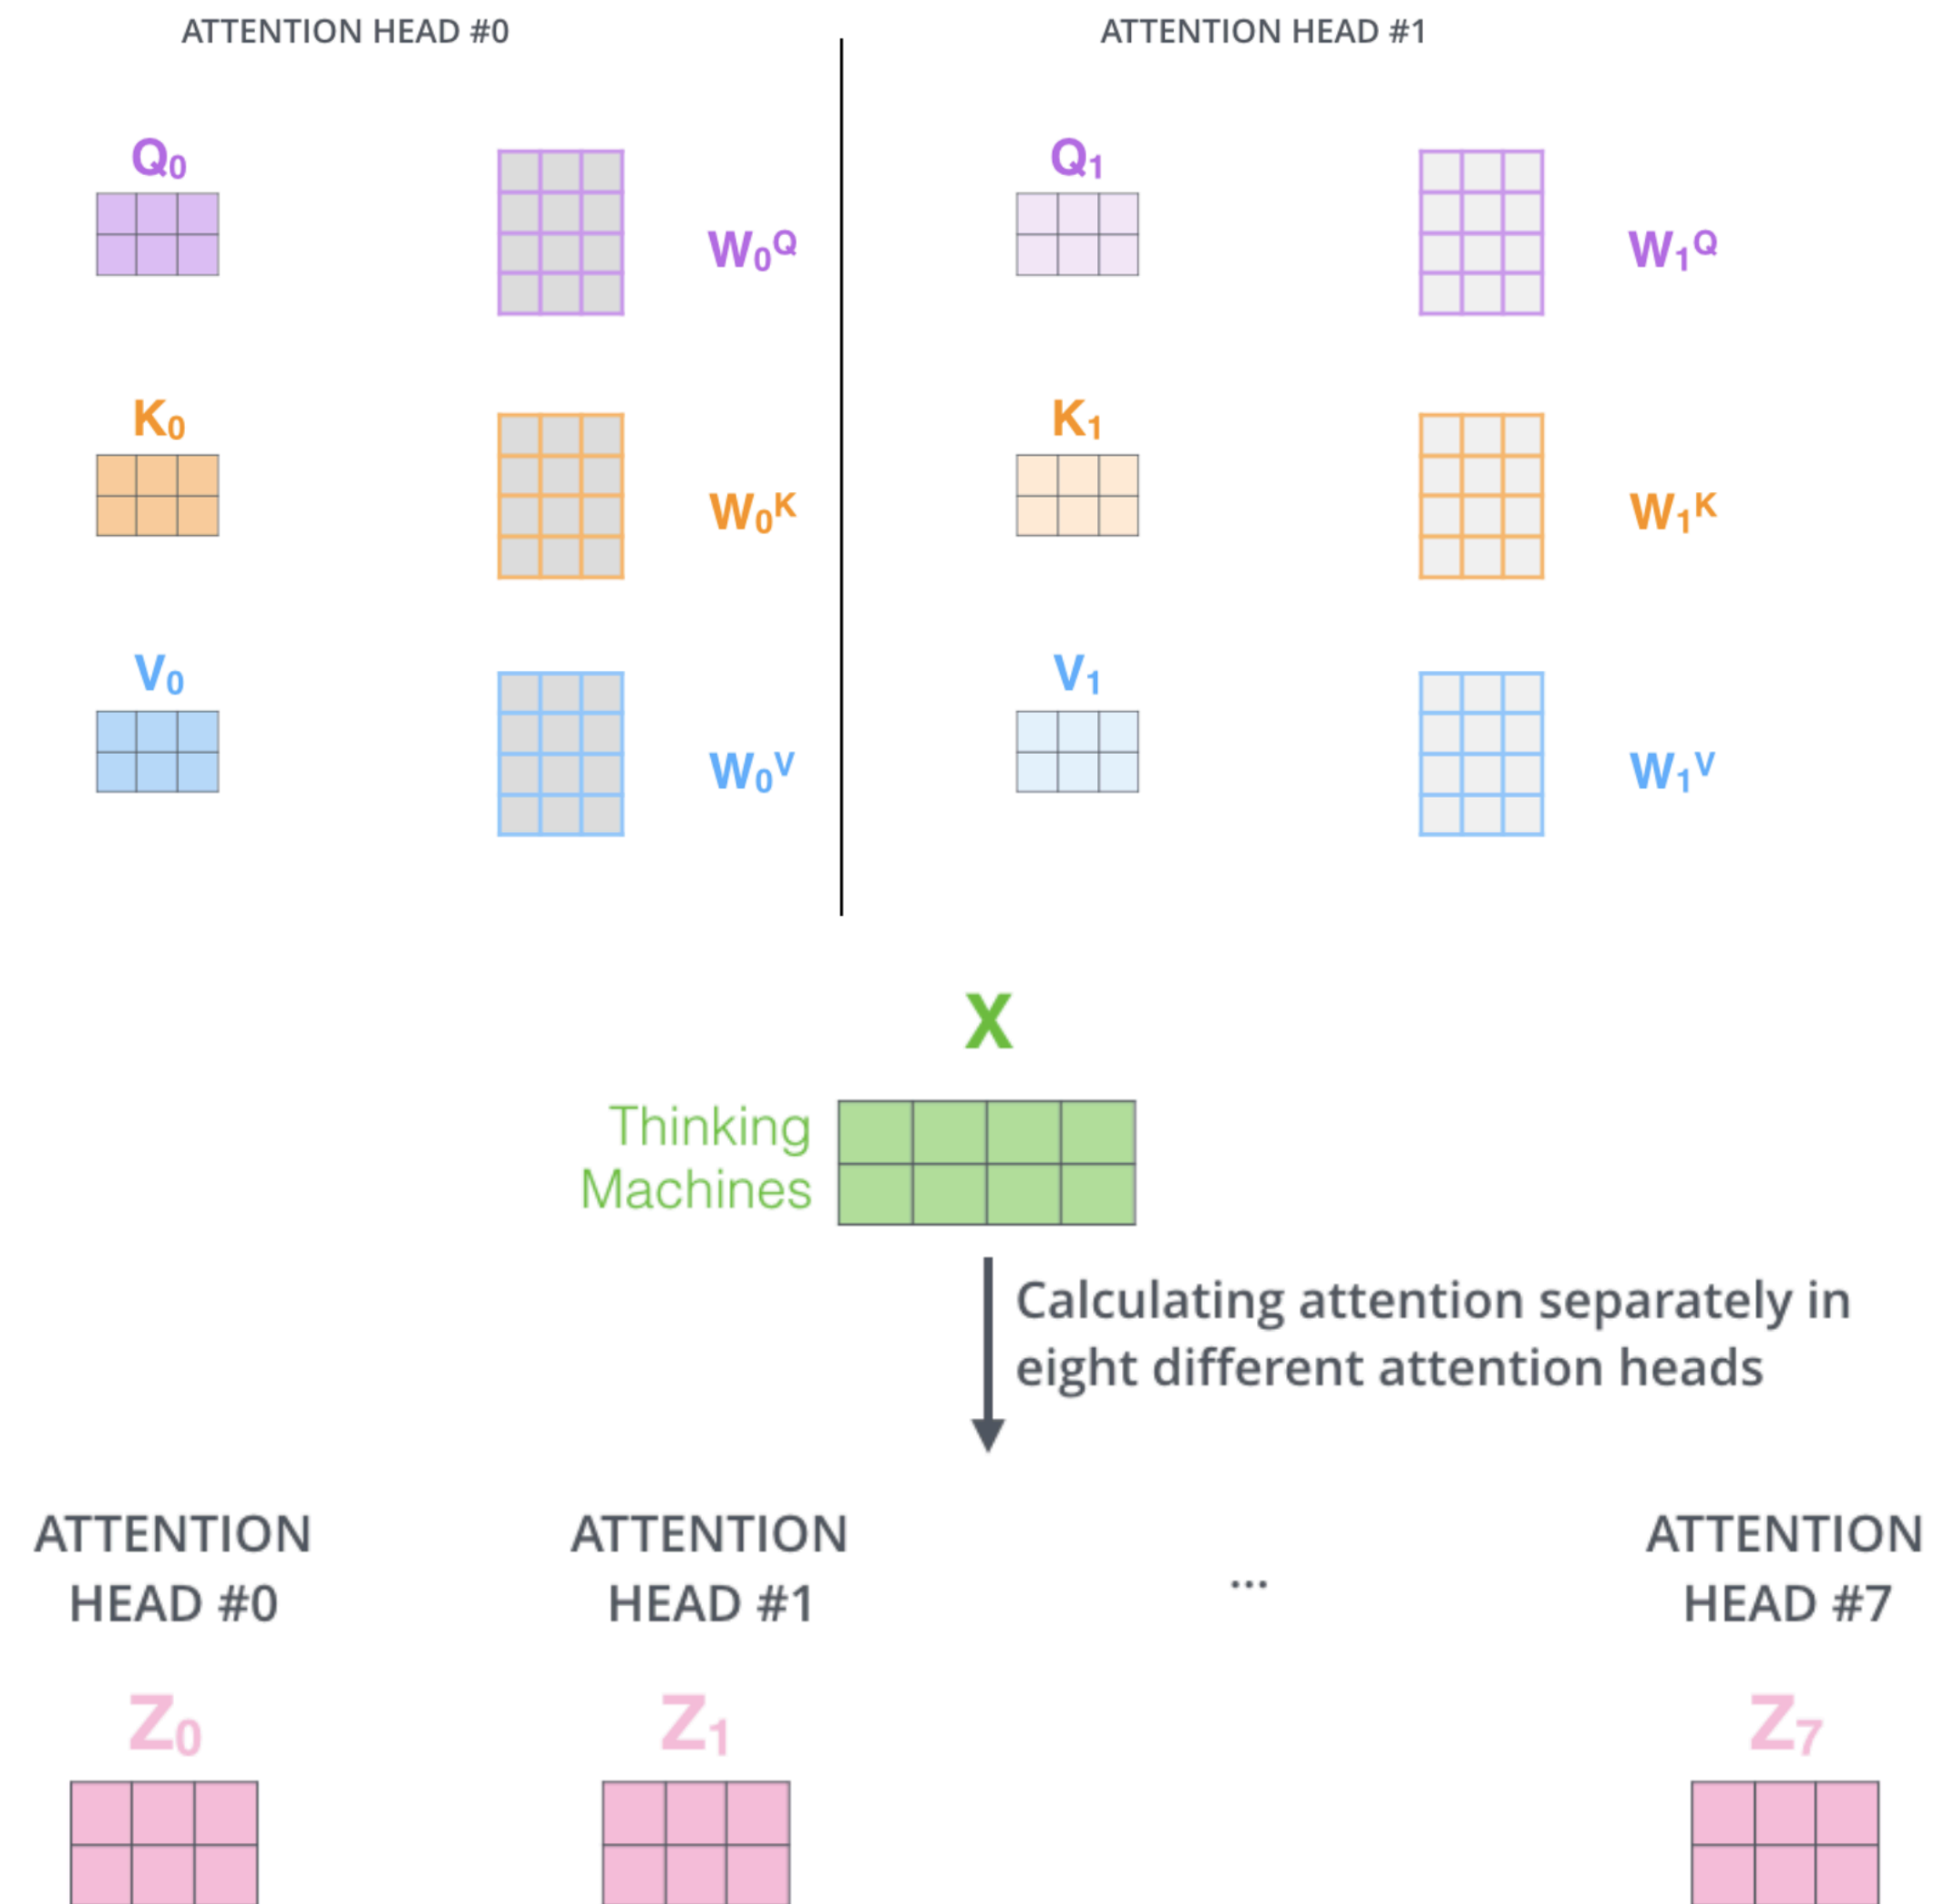
Self-Attention



Multi-Head Attention




Multi-Head Self-Attention




Multi-Head Self-Attention

1) Concatenate all the attention heads




2) Multiply with a weight matrix W^O that was trained jointly with the model


\times



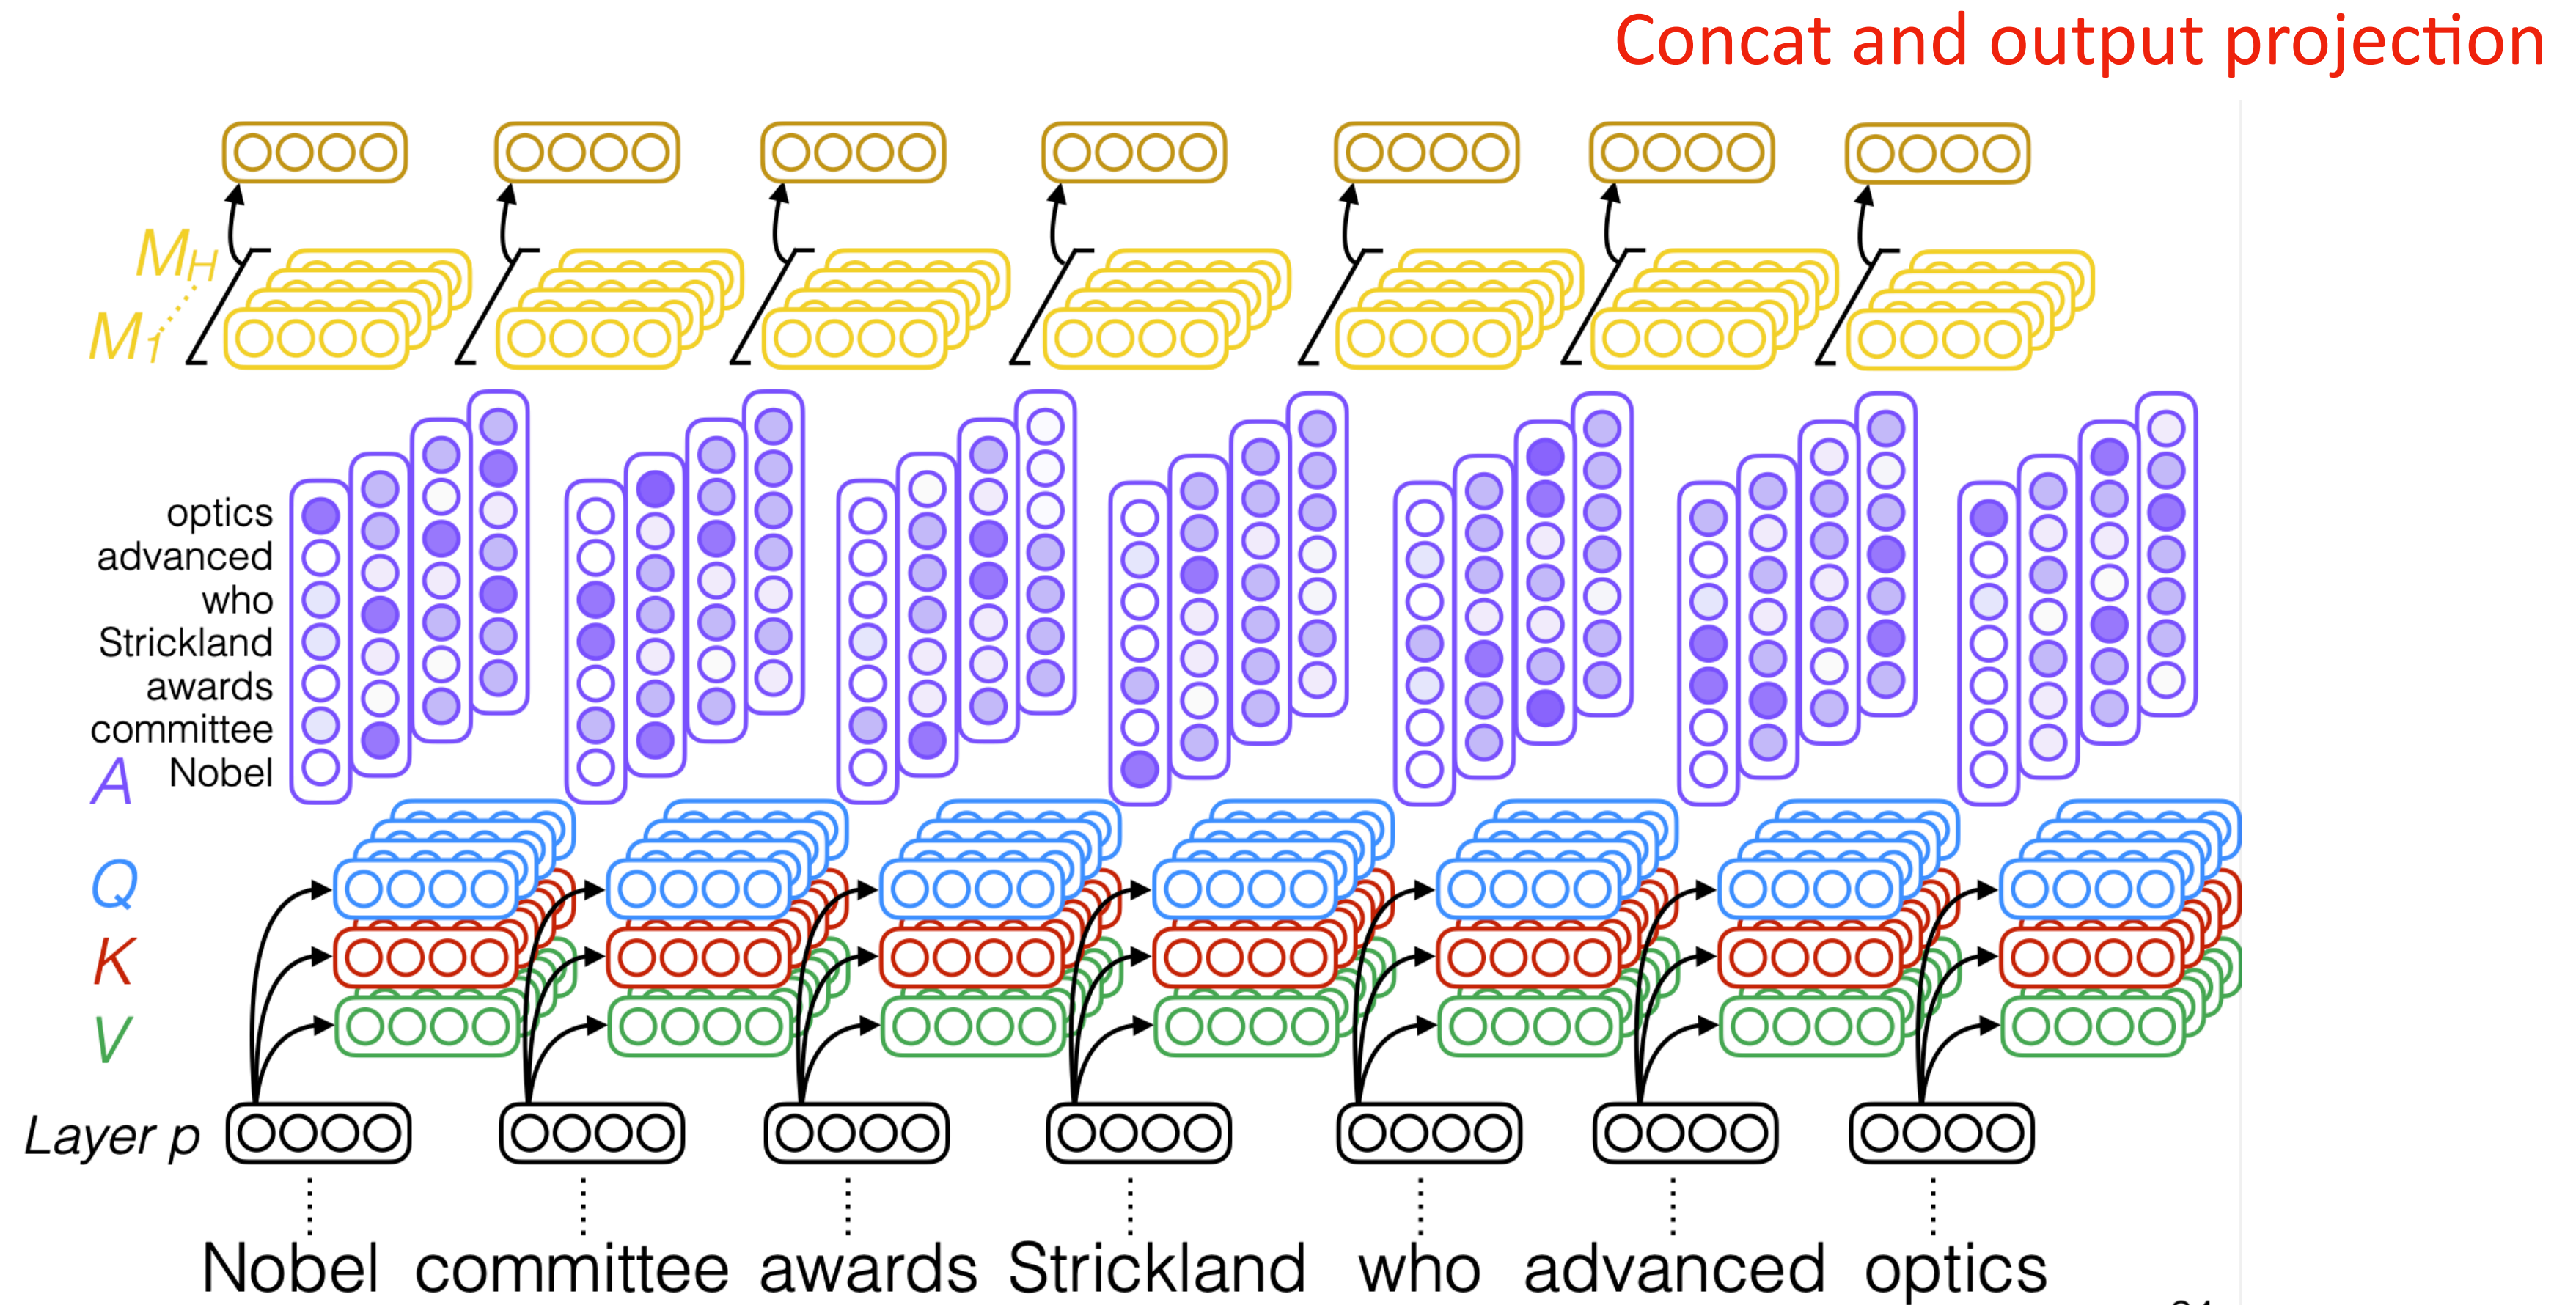
3) The result would be the Z matrix that captures information from all the attention heads. We can send this forward to the FFNN




Multi-head Self-Attention



Multi-head Self-Attention

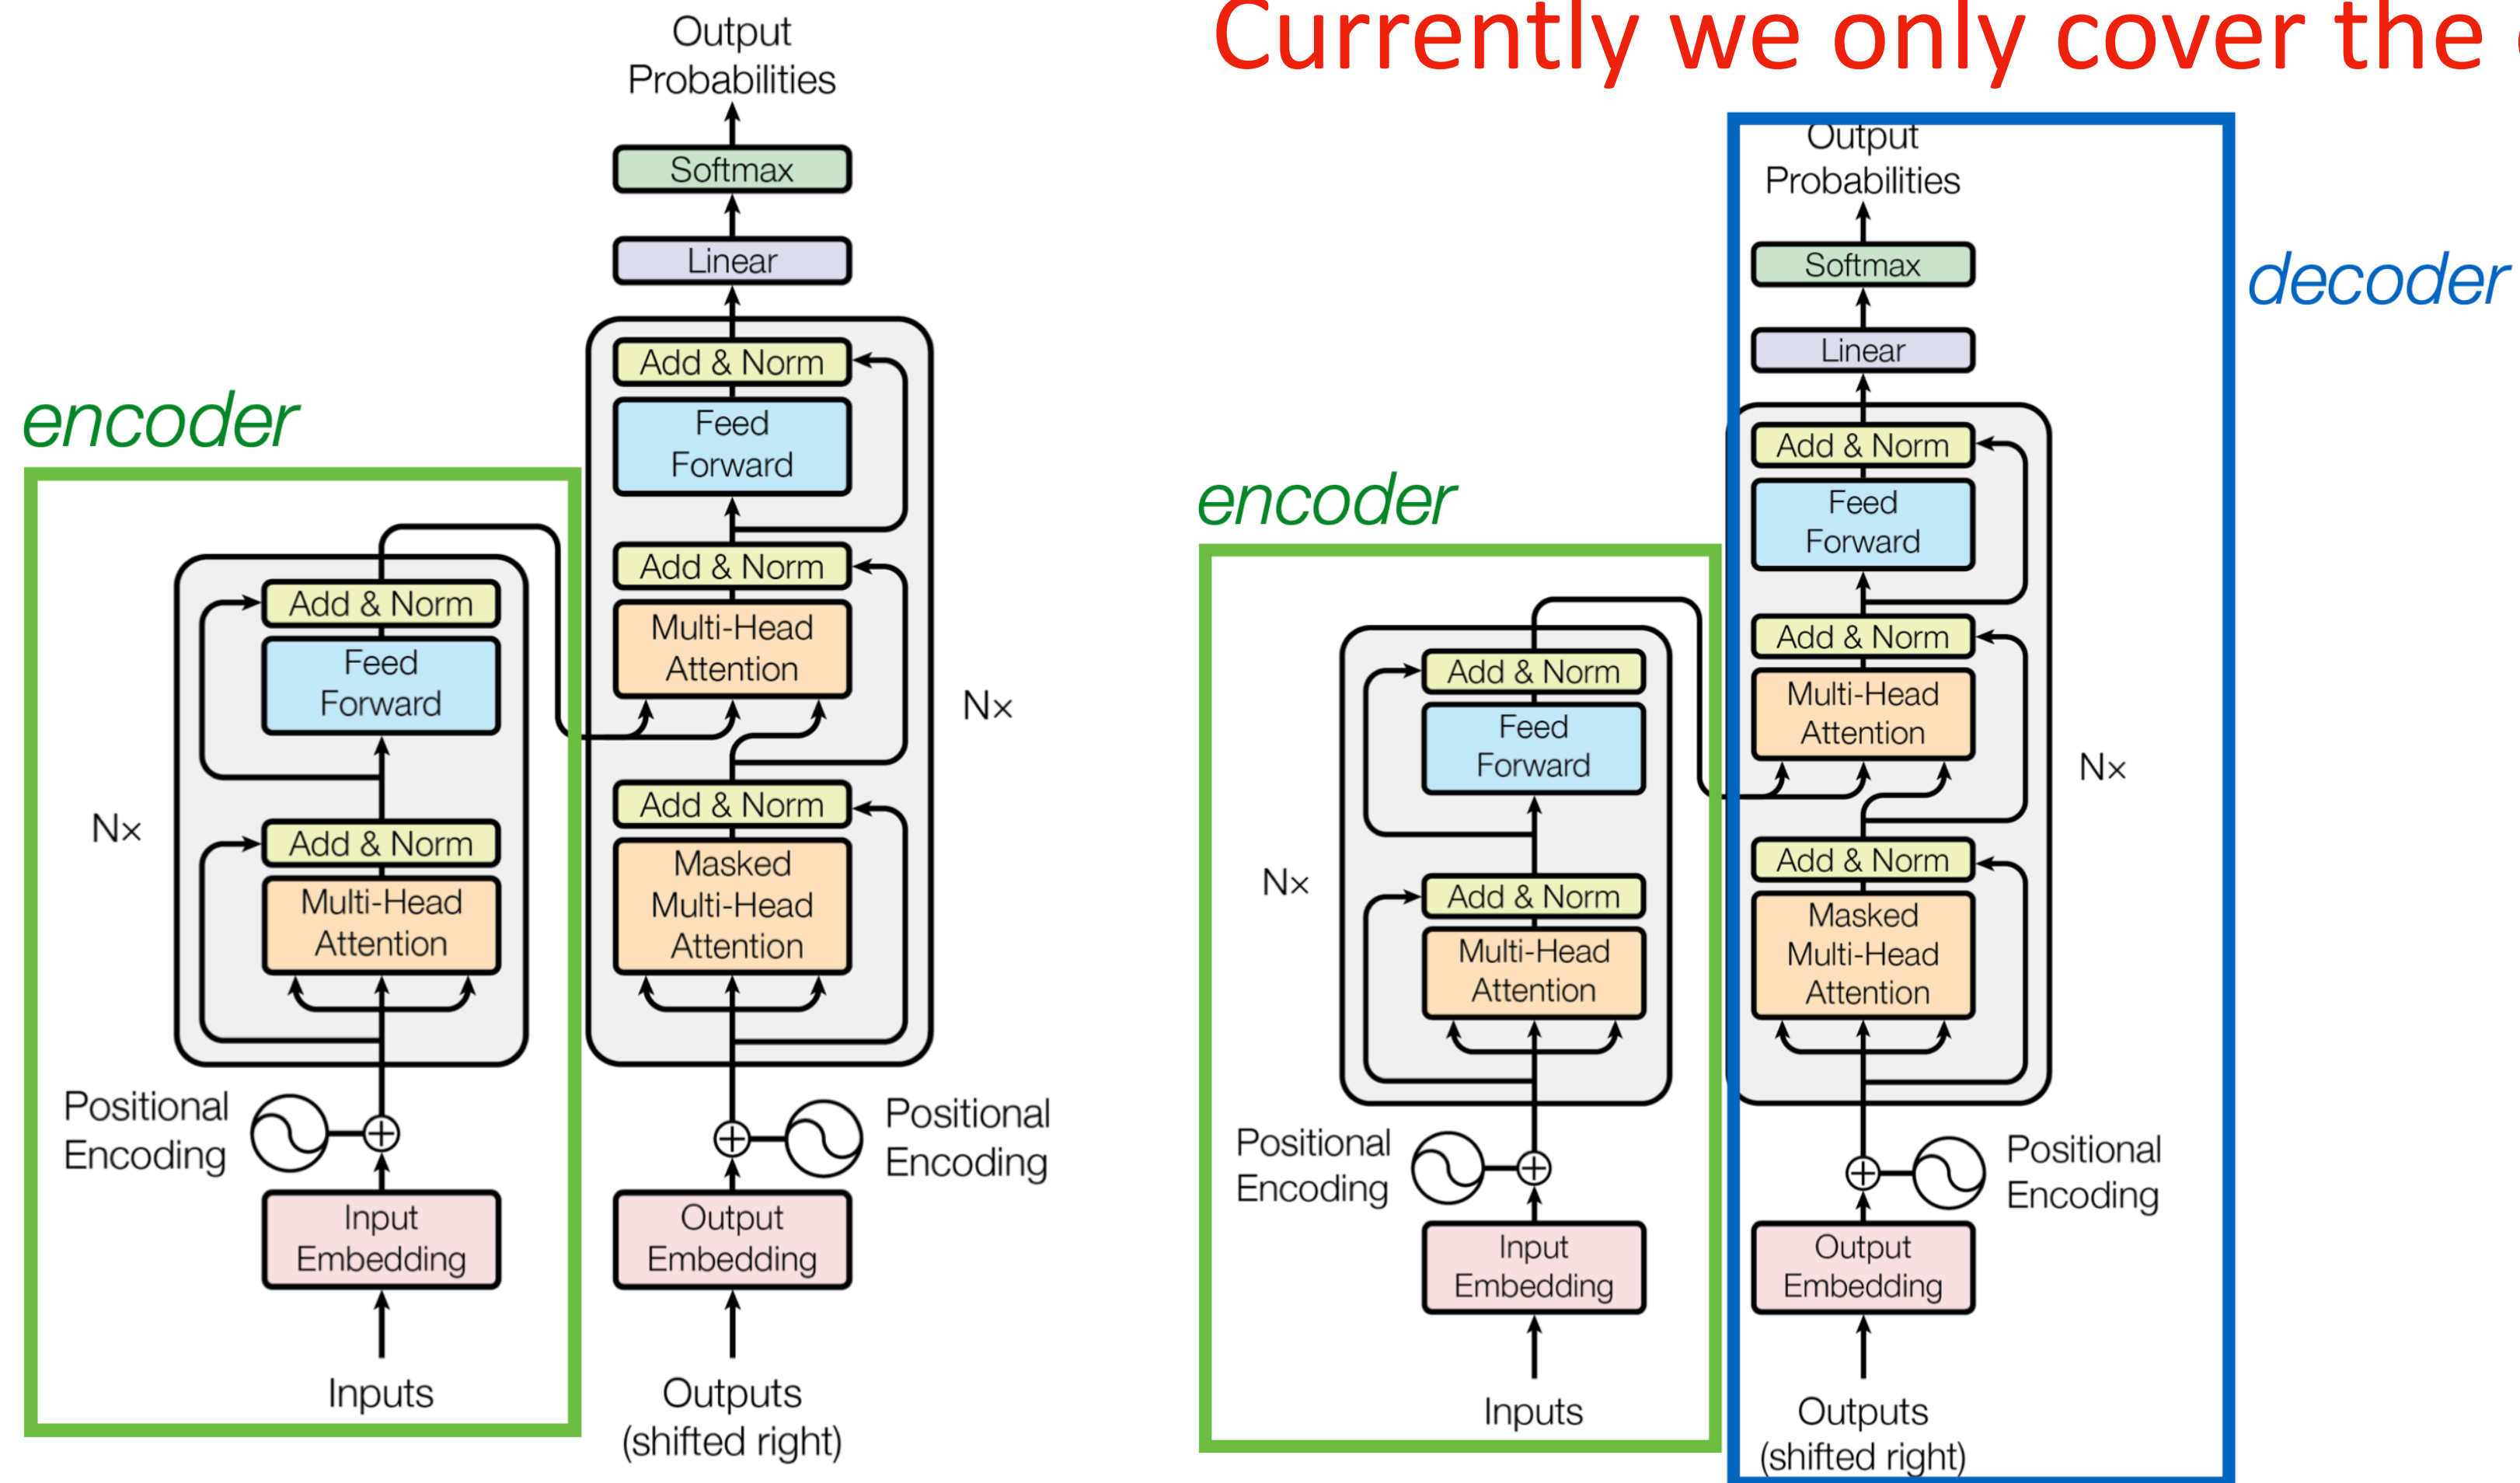


Multi-head Self-Attention + FFN



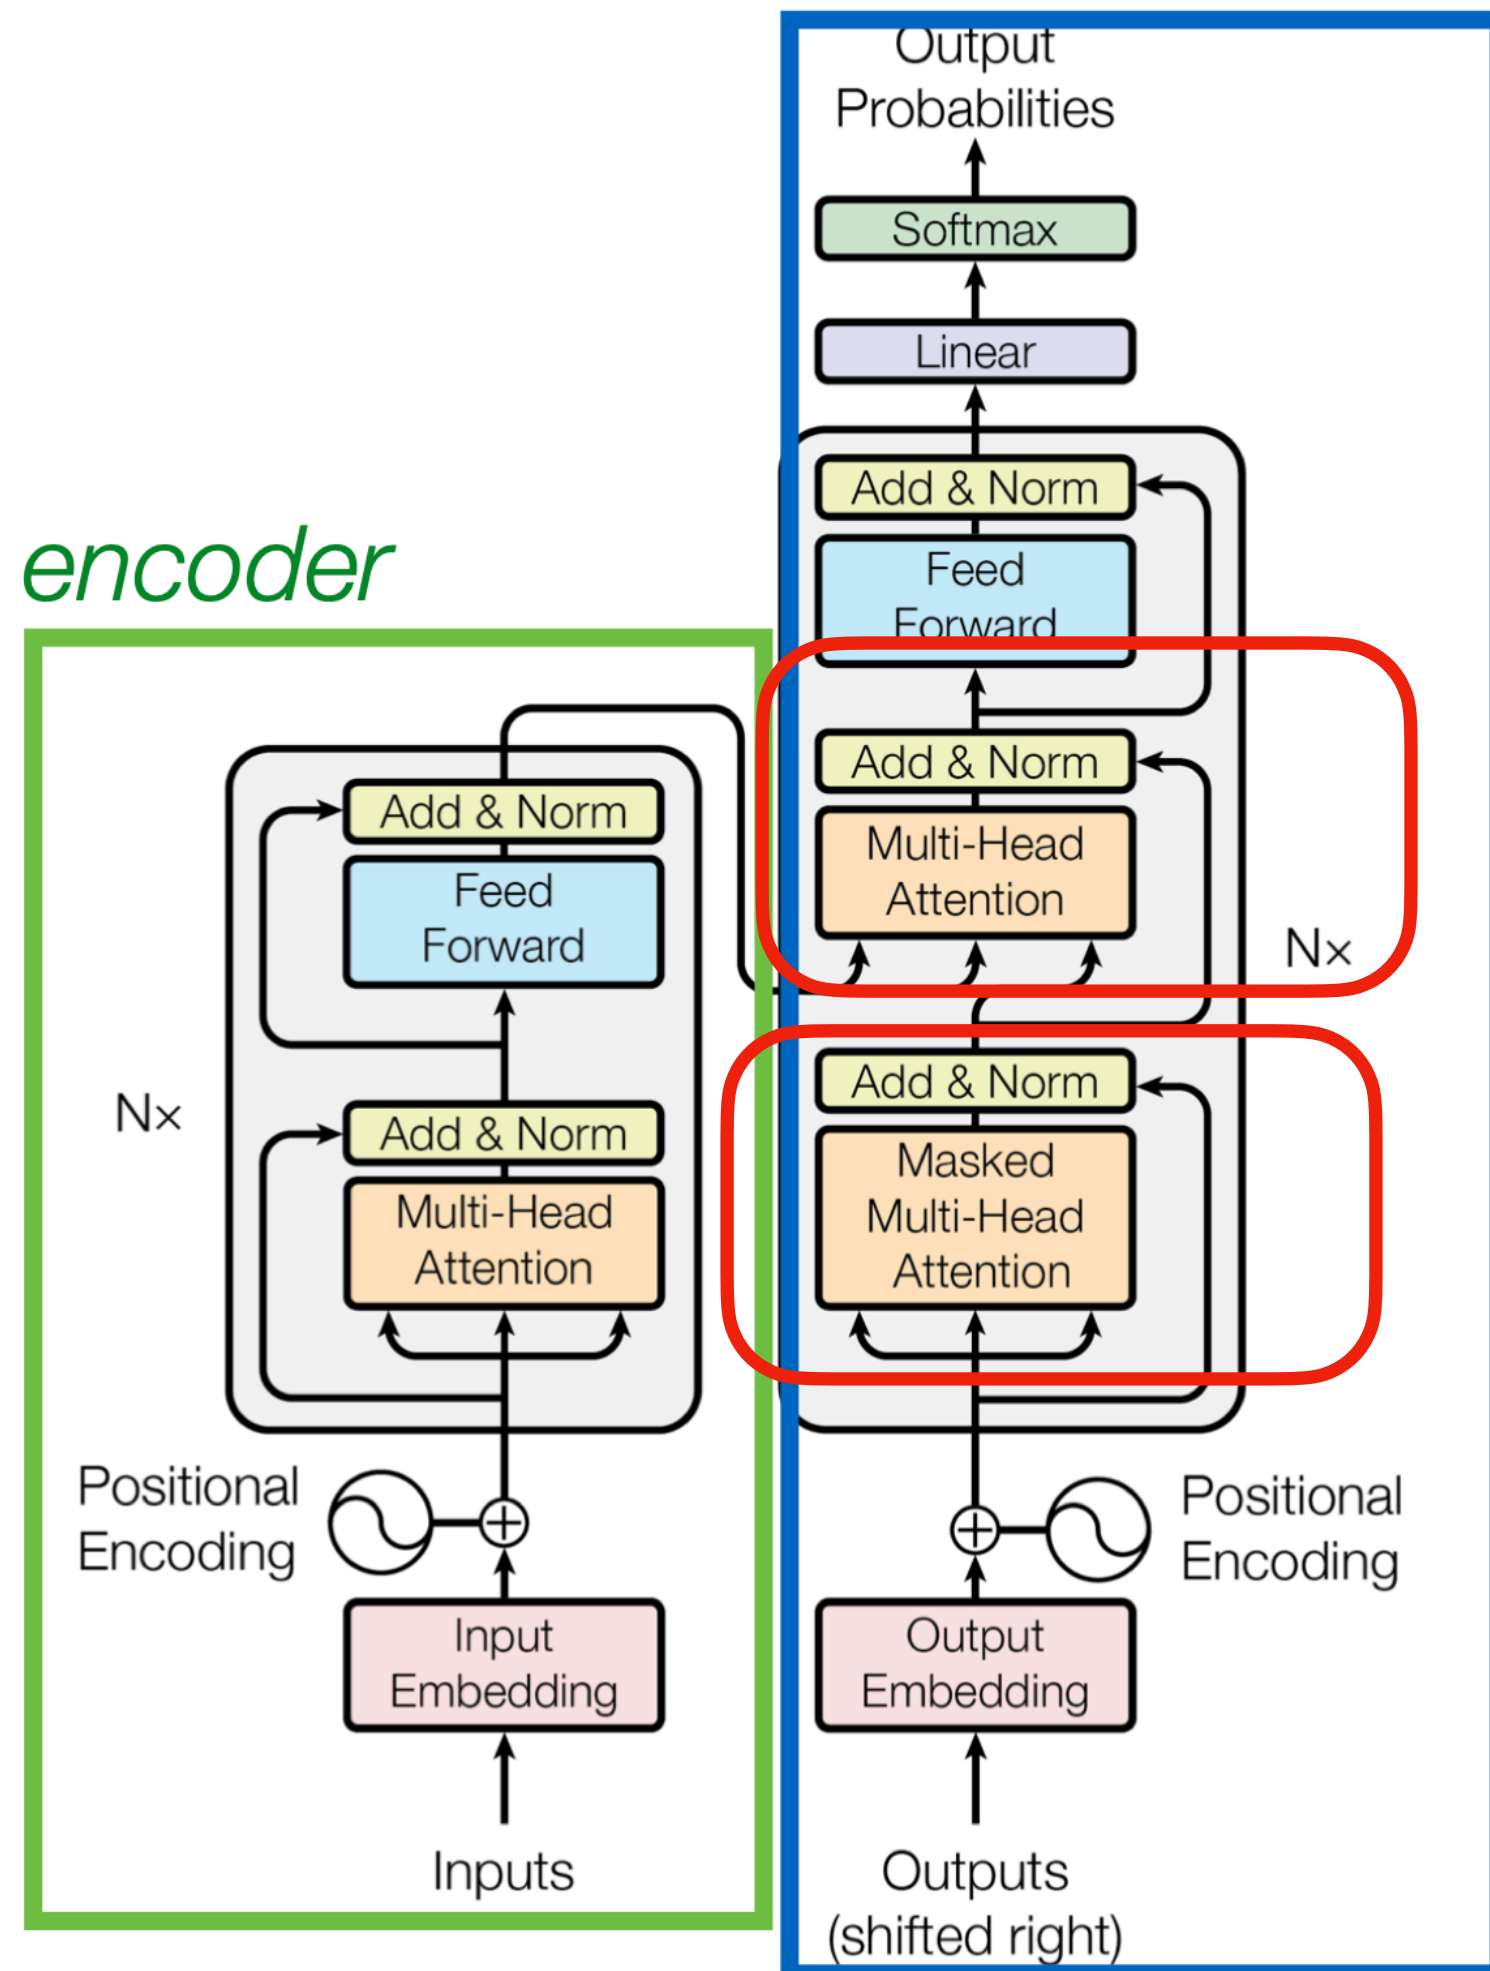
Transformer Encoder

Currently we only cover the encoder side



This encoder-decoder arch is originally proposed as a seq2seq arch, for classification tasks, often only encoder is used. And language models often only have a decoder

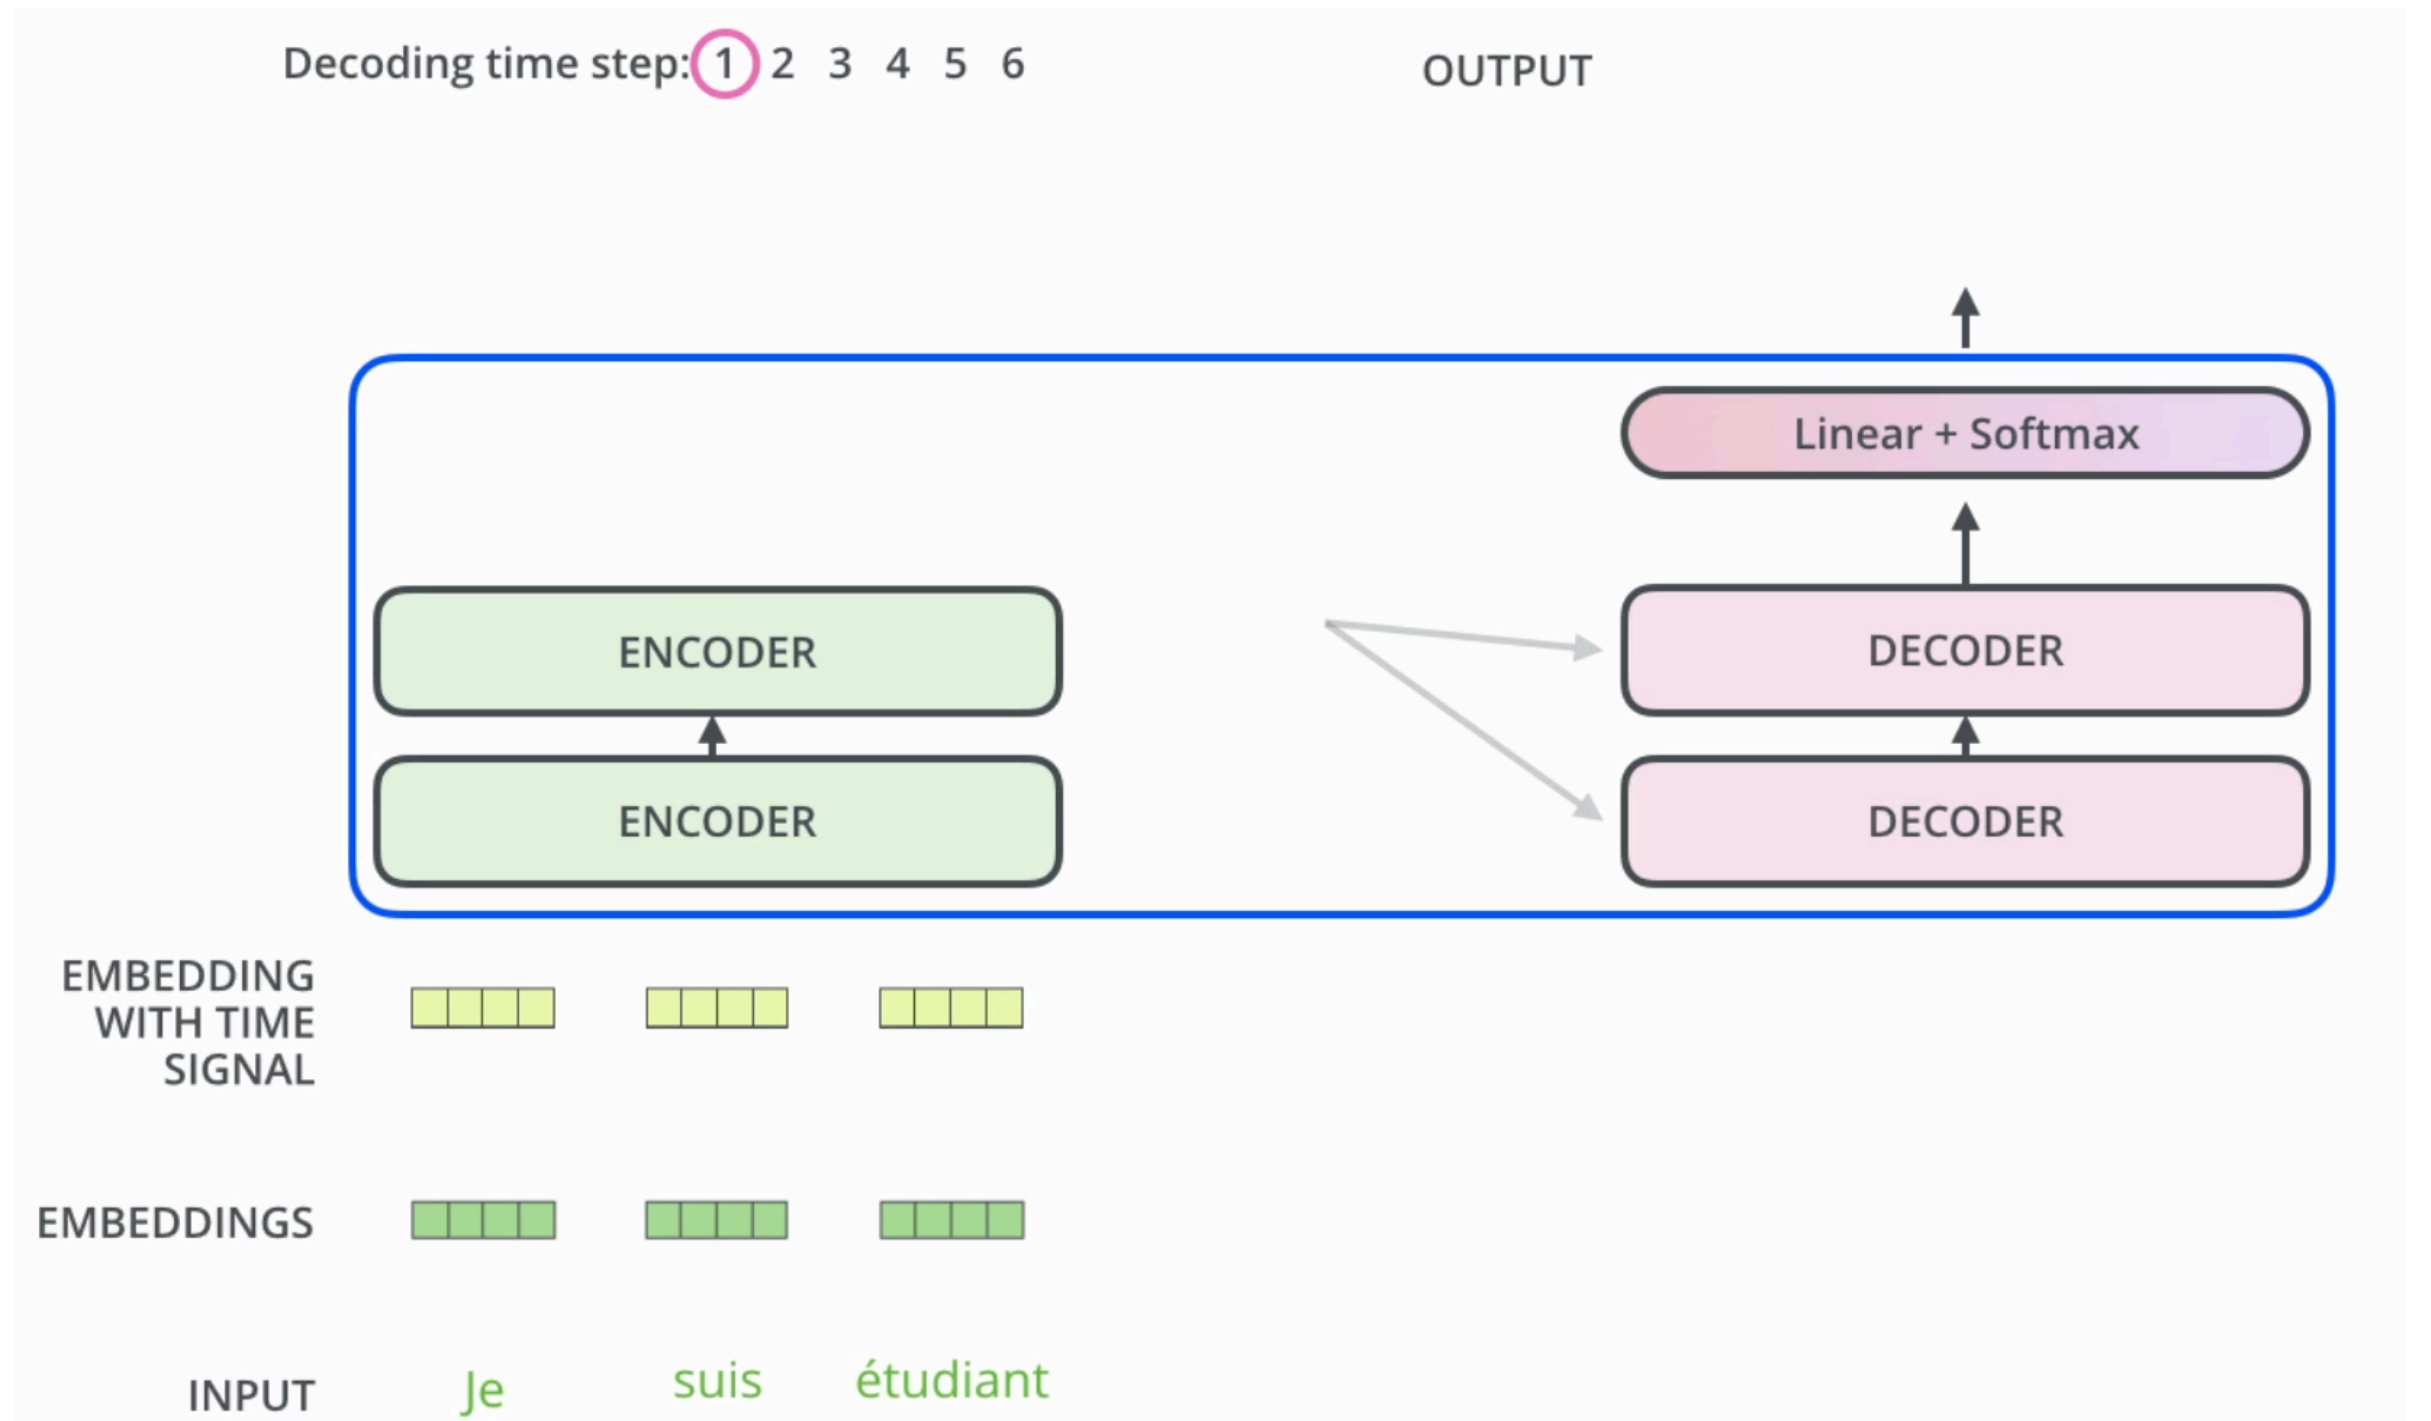
Transformer Decoder in Seq2Seq



decoder

Cross-attention

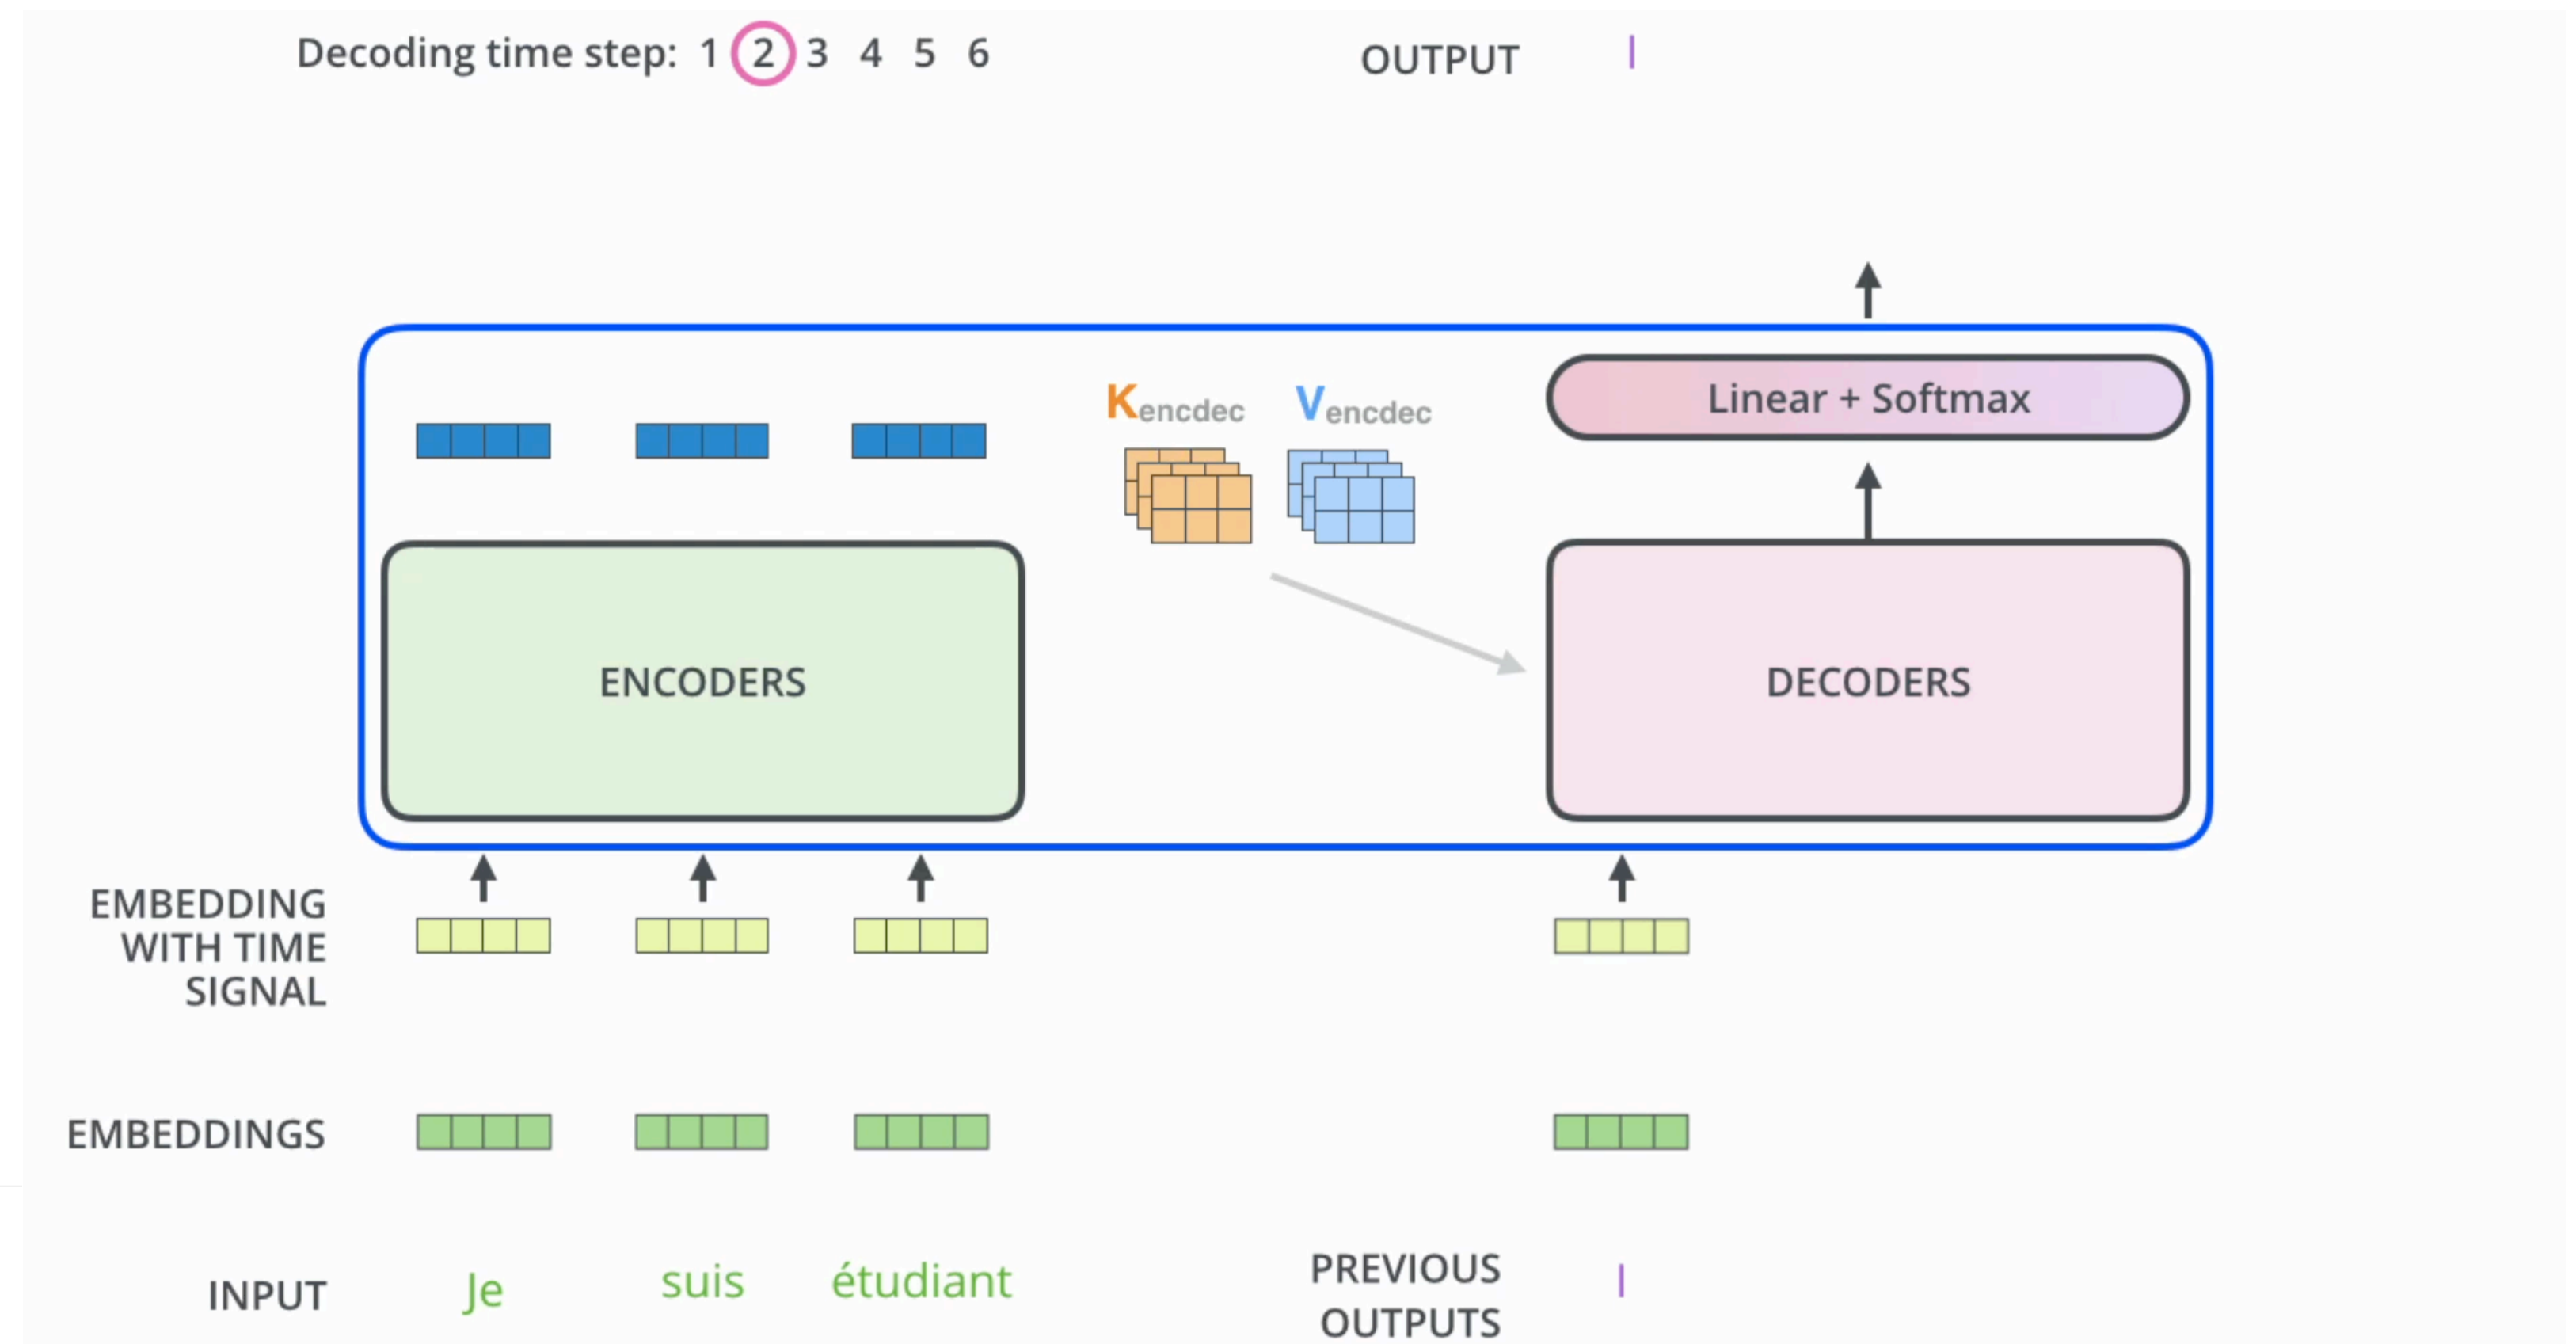
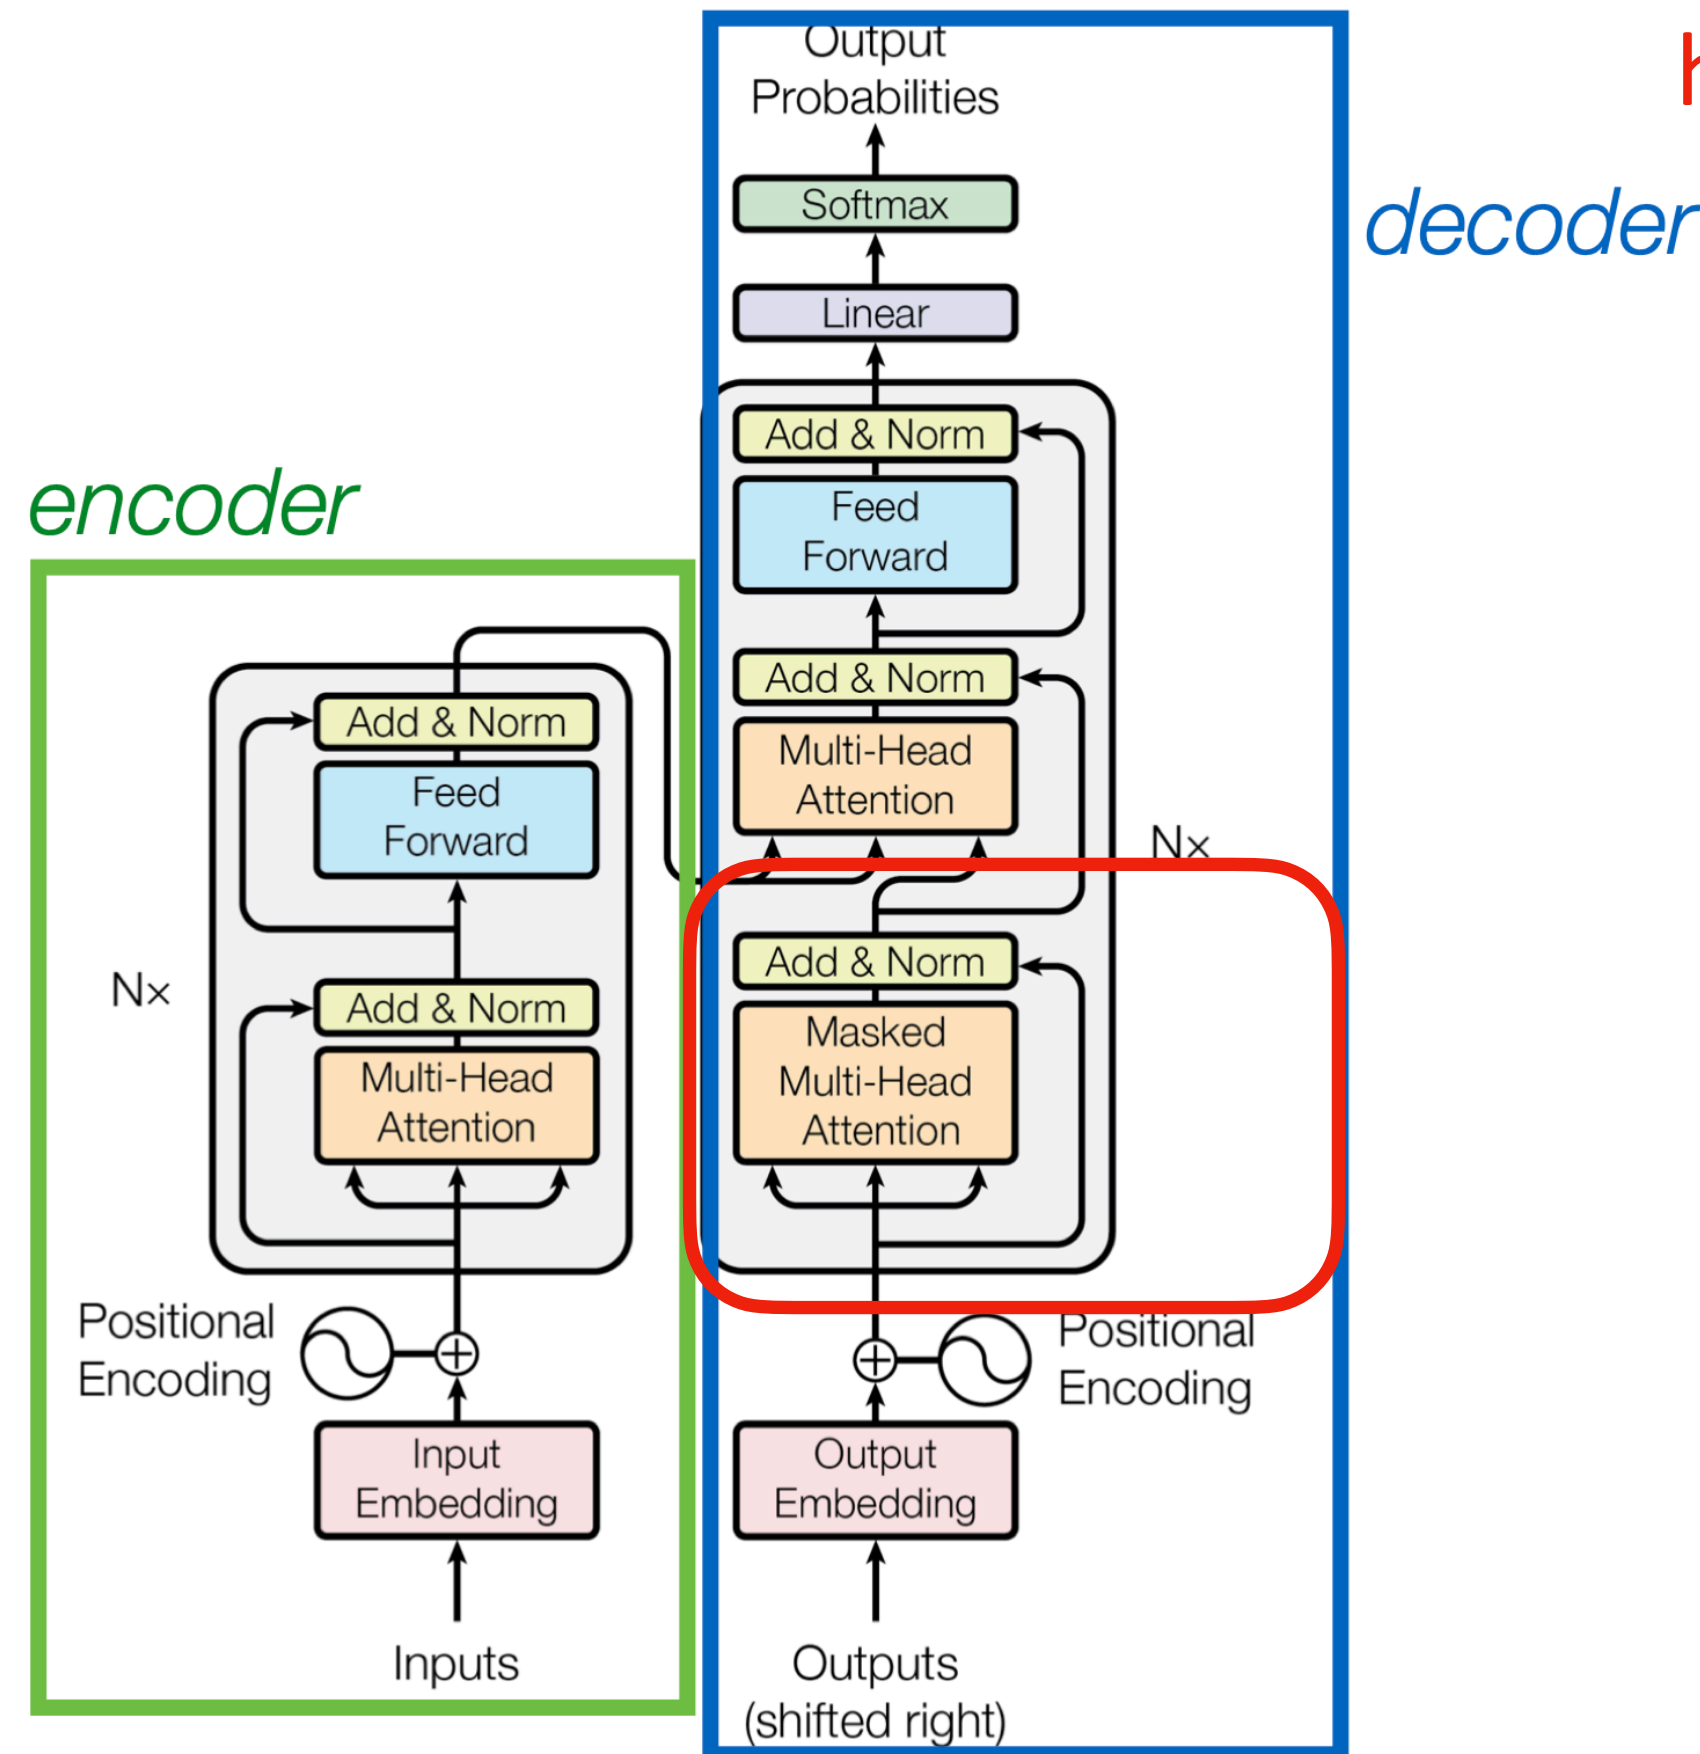
Self-attention



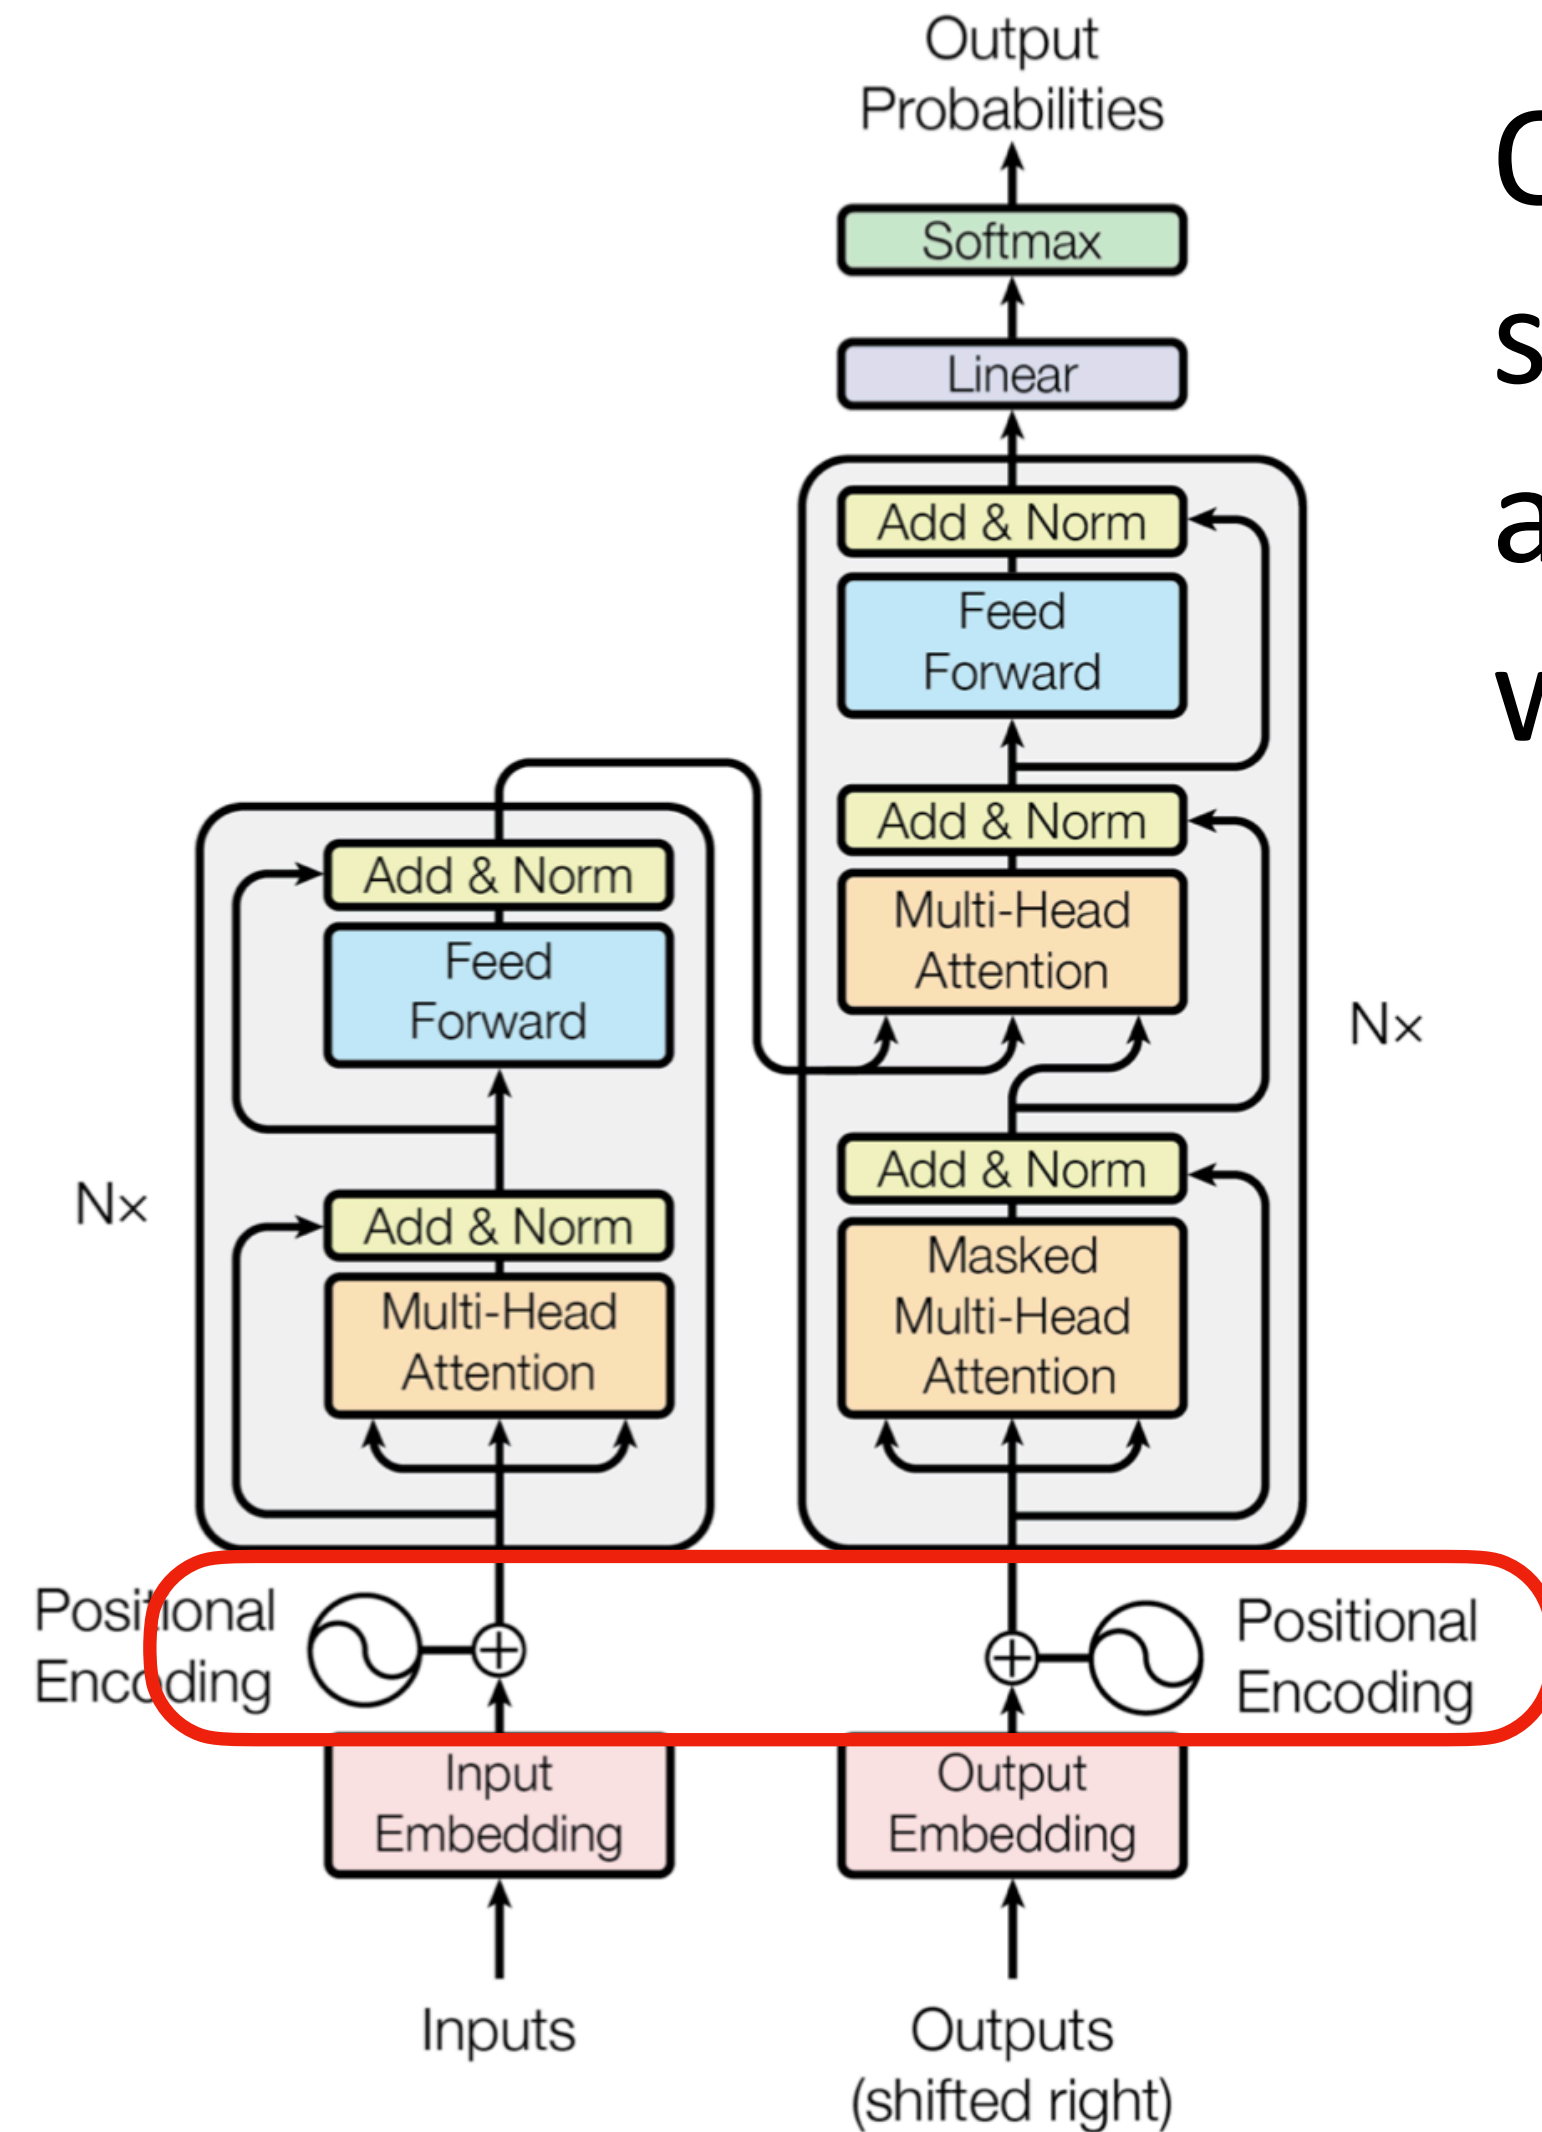
Cross-attention uses the output of encoder as input

Masked Attention

Typical attention attends to the entire sequence, while masked attention only attends to the ones on the left because future words have not been generated




Position Embeddings



Question: If we shuffle the order of words in the sequence, will that change the attention output and feed forward output of the corresponding word?

Position embeddings are added to each word embedding, otherwise our model is unaware of the position of a word

Positional Encoding



Transformer Positional Encoding

$$PE_{(pos,2i)} = \sin\left(\frac{pos}{10000^{2i/d_{model}}}\right)$$

$$PE_{(pos,2i+1)} = \cos\left(\frac{pos}{10000^{2i/d_{model}}}\right)$$

Positional encoding is a 512d vector
 i = a particular dimension of this vector
 pos = dimension of the word
 $d_{model} = 512$

Complexity

| Layer Type | Complexity per Layer | Sequential Operations |
|-----------------------------|--------------------------|-----------------------|
| Self-Attention | $O(n^2 \cdot d)$ | $O(1)$ |
| Recurrent | $O(n \cdot d^2)$ | $O(n)$ |
| Convolutional | $O(k \cdot n \cdot d^2)$ | $O(1)$ |
| Self-Attention (restricted) | $O(r \cdot n \cdot d)$ | $O(1)$ |

n is sequence length, d is embedding dimension.

Restricted self-attention means not attending all words in the sequence, but only a restricted field

Square complexity of sequence length is a major issue for transformers to deal with long sequence

Auto-Encoding Variational Bayes

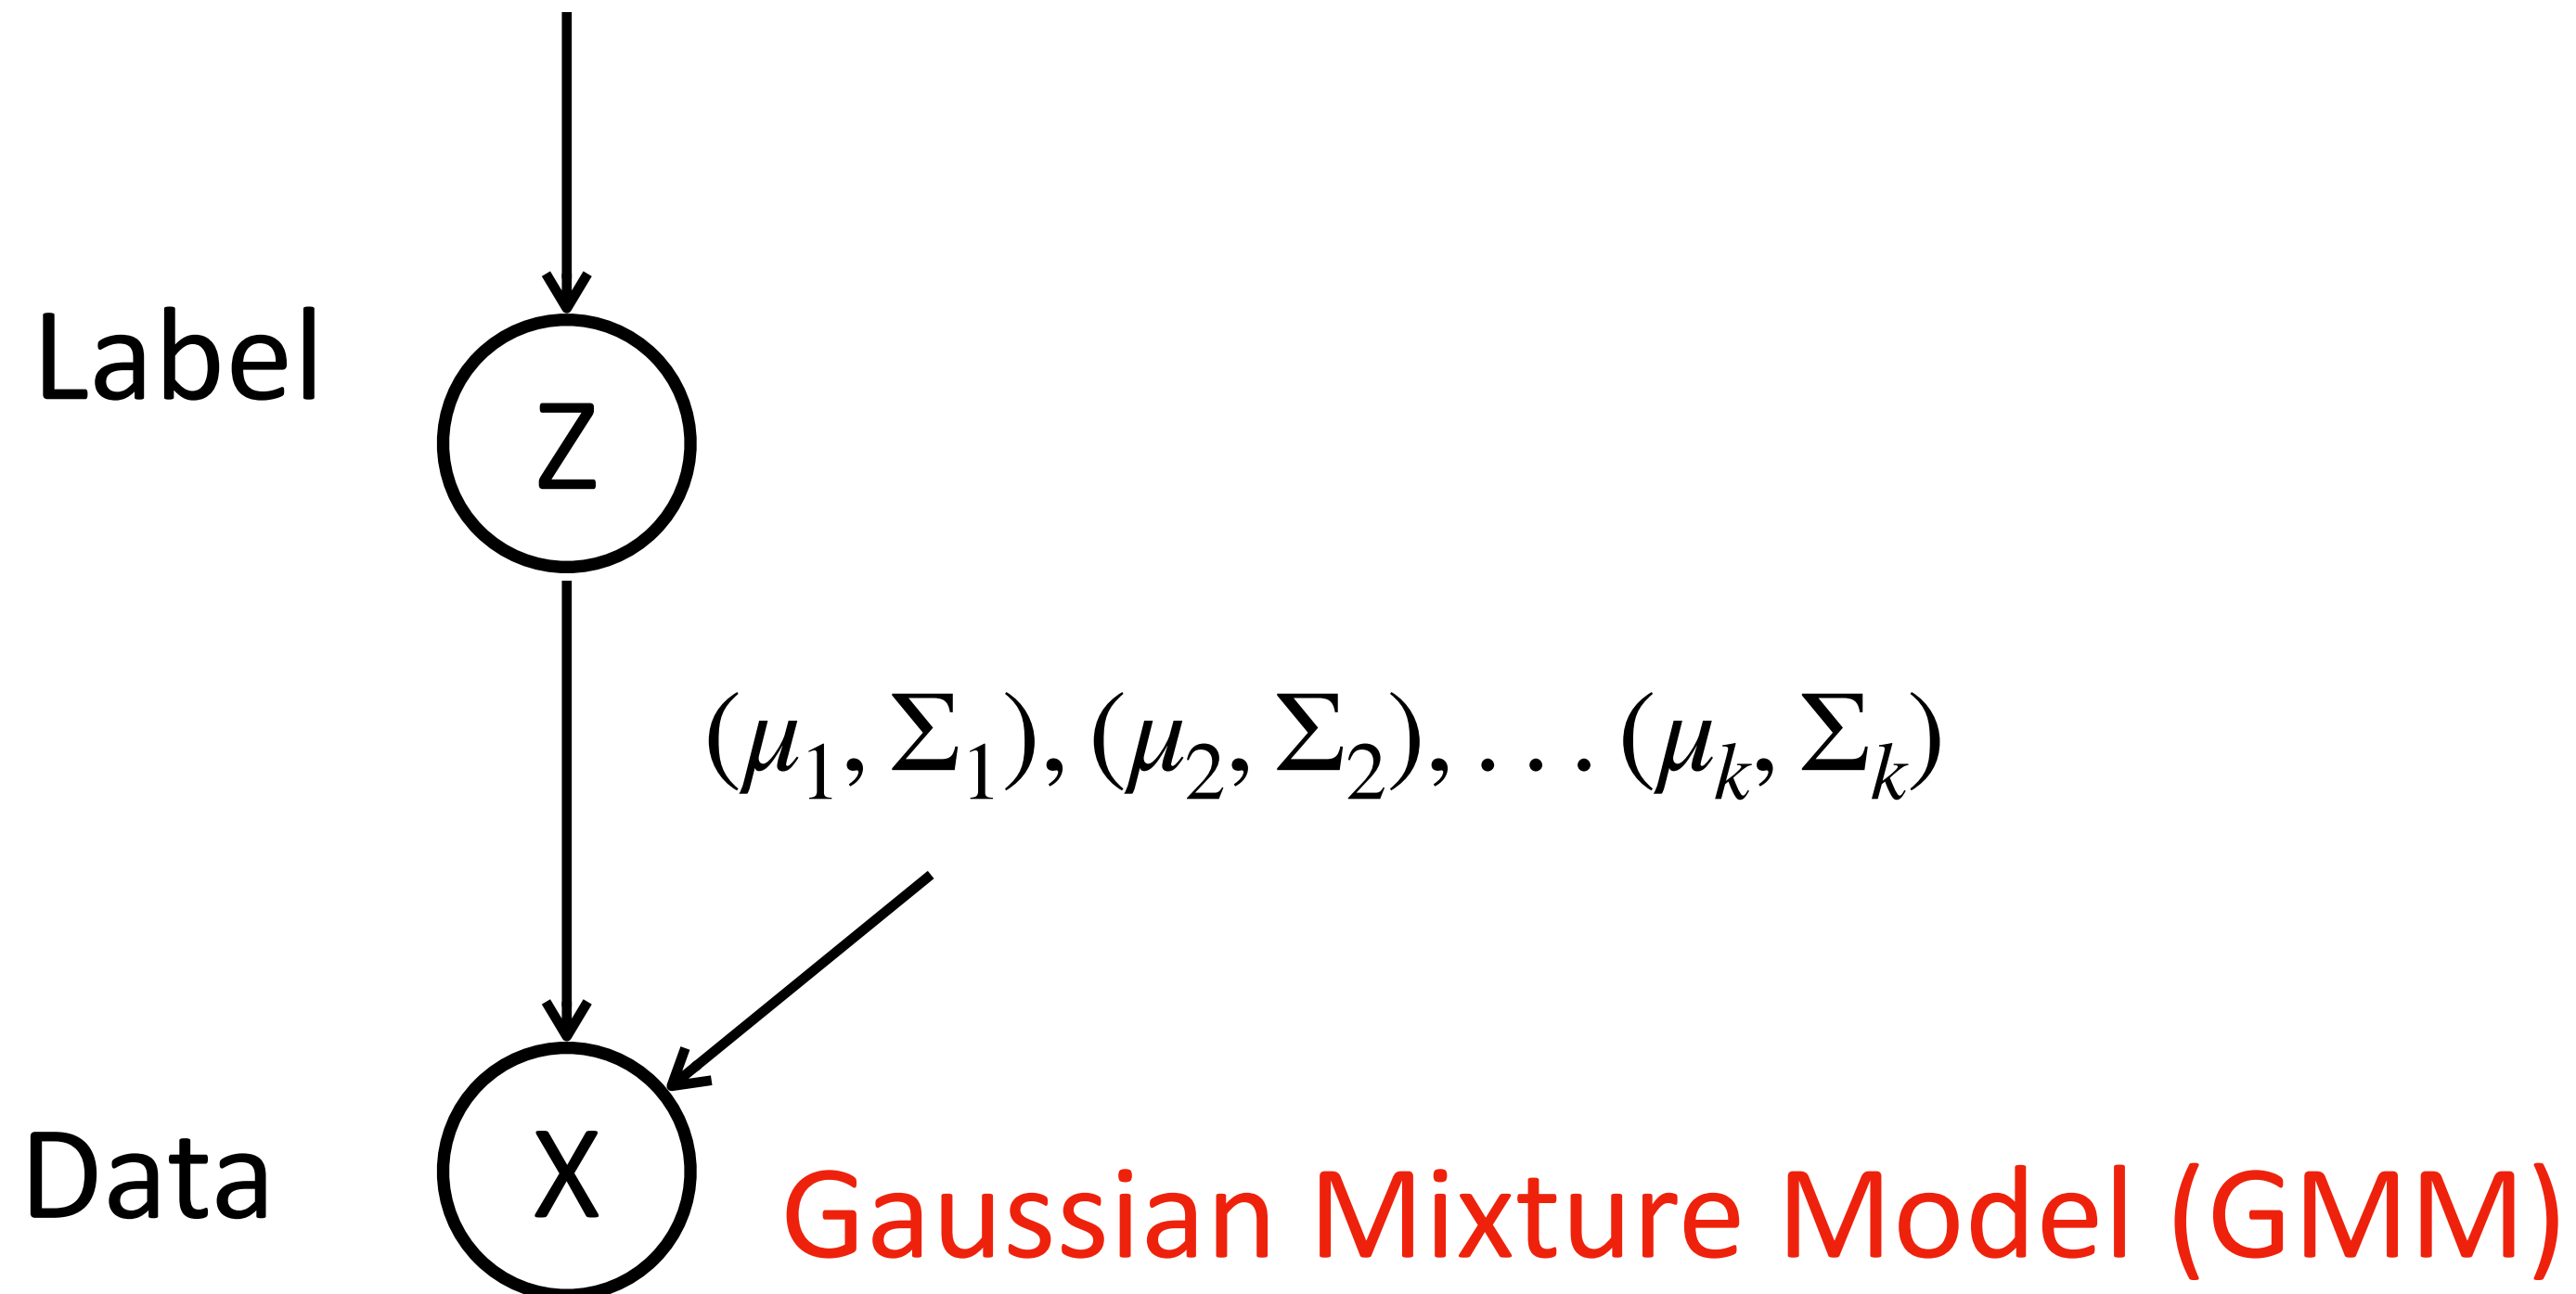
Diederik P. Kingma
Machine Learning Group
Universiteit van Amsterdam
dpkingma@gmail.com

Max Welling
Machine Learning Group
Universiteit van Amsterdam
welling.max@gmail.com

Variational Autoencoders


VAE is a Generative Model

$p(z)$: multinomial, k
classes (e.g. uniform)




The VAE Model

$p(z)$ is a normal distribution in most cases



f is a neural network taking Z as input


Training



How to train the model? Can we do MLE?

Intractable $P(X)$, EM algorithm?

Let's try EM



E-Step: compute $P(z|x)$

$$Q(z) = P(z|x) \propto P(z)P(x|z) \quad \text{This is ok?}$$

M-Step: the ELBO objective

$$\operatorname{argmax}_{\theta} \sum_z Q(z) \log p(x, z; \theta) = \operatorname{argmax}_{\theta} \mathbb{E}_{z \sim Q(z)} \log p(x, z; \theta)$$

In most cases, we cannot do the sum, and cannot easily sample from $Q(z)$ either

Approximate Posterior

We need an easy-to-sample distribution to approximate $P(z|x)$

$q(z|x;\phi)$ to approximate $p(z|x;\theta)$ Why conditioned on x ?

ϕ is the parameter for the approximate function, θ is the generative model parameter

How to train $q(z|x;\phi)$, what would be the loss to find ϕ ?

Recap: ELBO

$$\text{ELBO}(x; Q, \theta) = \sum_z Q(z) \log \frac{p(x, z; \theta)}{Q(z)}$$

What is $\text{argmax}_{Q(z)} \text{ELBO}(x; Q, \theta)$?

ELBO is maximized when $Q(z)$ is equal to $p(z|x)$

Therefore, we can approximate the true posterior by maximizing ELBO:

$$\text{argmax}_{\phi} \sum_z q(z|x; \phi) \log \frac{p(x, z; \theta)}{q(z|x; \phi)}$$

Variational Inference

Training VAEs

E-Step:

$$\operatorname{argmax}_{\phi} \sum_z q(z | x; \phi) \log \frac{p(x, z; \theta)}{q(z | x; \phi)}$$

M-Step:

$$\operatorname{argmax}_{\theta} \sum_z q(z | x; \phi) \log \frac{p(x, z; \theta)}{q(z | x; \phi)}$$

Same objective, different parameters to optimize

Because we use approximate rather than exact posterior, it is also called Variational EM

Training VAEs

E-Step:

$$\operatorname{argmax}_{\phi} \sum_z q(z | x; \phi) \log \frac{p(x, z; \theta)}{q(z | x; \phi)}$$

Can we do gradient descent over ϕ ?

M-Step:

$$\operatorname{argmax}_{\theta} \sum_z q(z | x; \phi) \log \frac{p(x, z; \theta)}{q(z | x; \phi)}$$

We use MC sampling to approximate expectation and use gradient descent to optimize θ