Modern DL: From Embeddings to Transformers

Dr. Riad Sonbol

The Core Problem in AI

"How can computers understand language, images, and meaning?"

- Computers only understand numbers.
- So AI needs a way to convert:
  - words
  - images
  - sounds

into numerical representations.

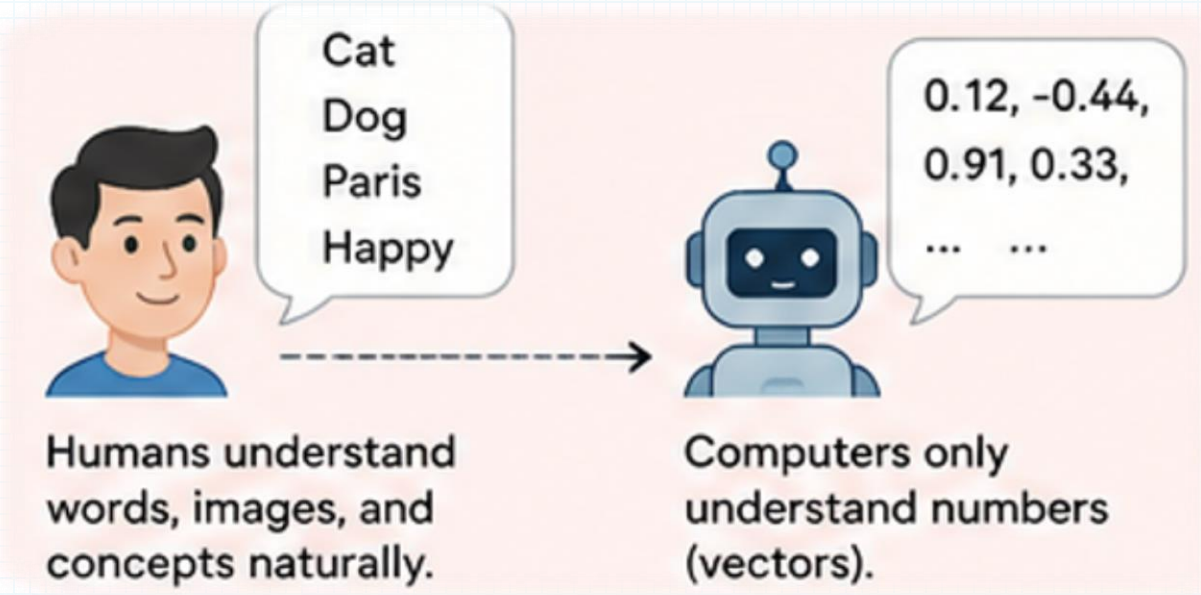

Embeddings: The Foundation of Modern AI

- Embeddings are numerical representations that capture meaning.

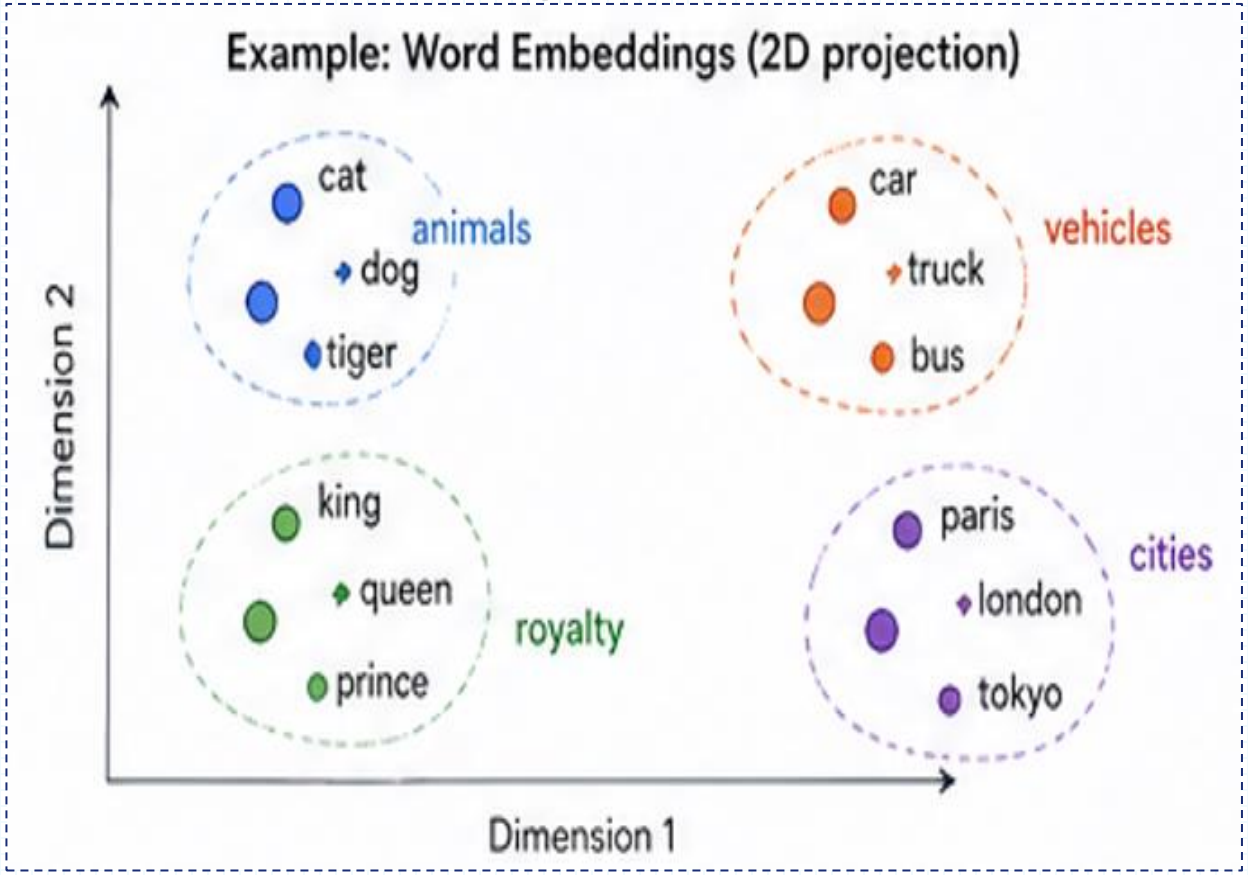

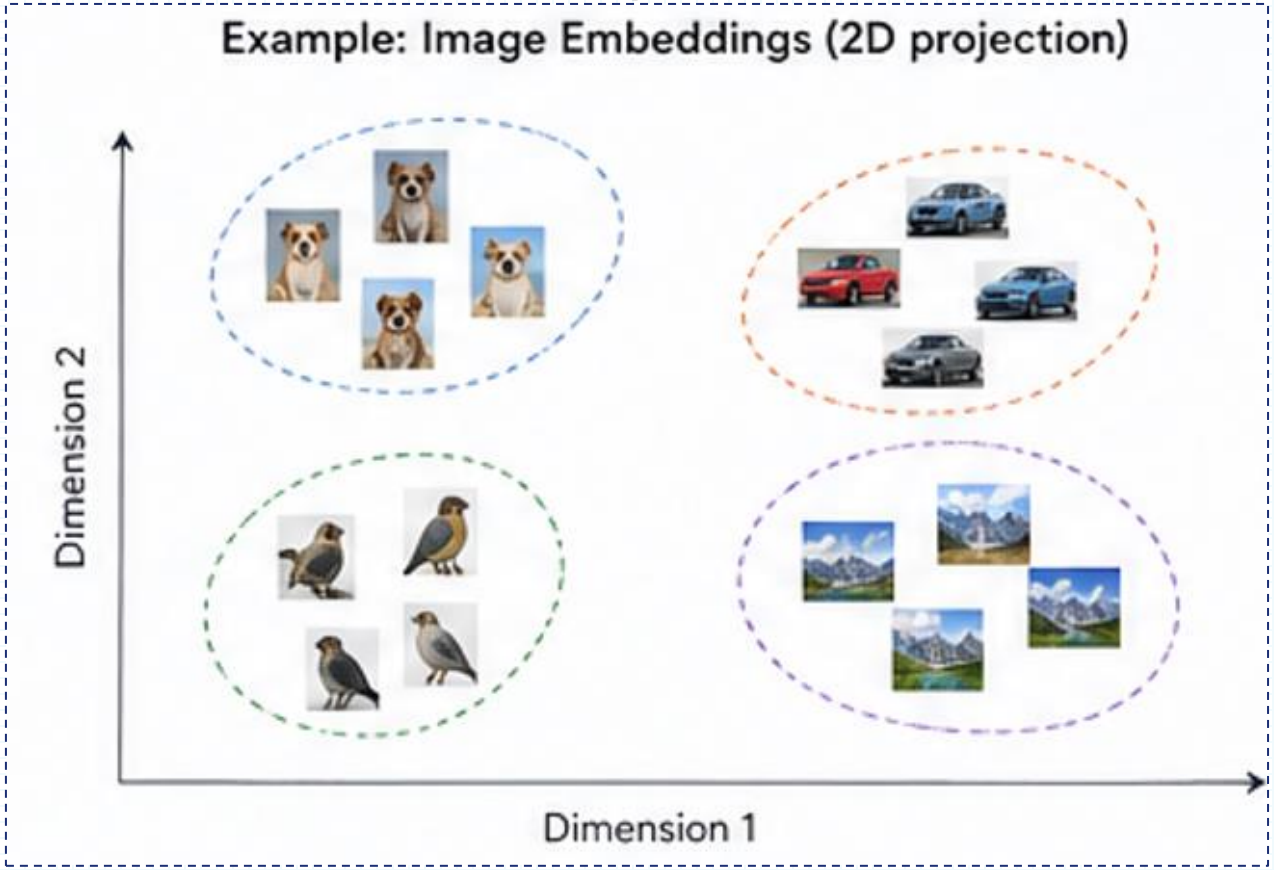

- Words with similar meanings are represented by nearby vectors.

The “Good” Embedding?

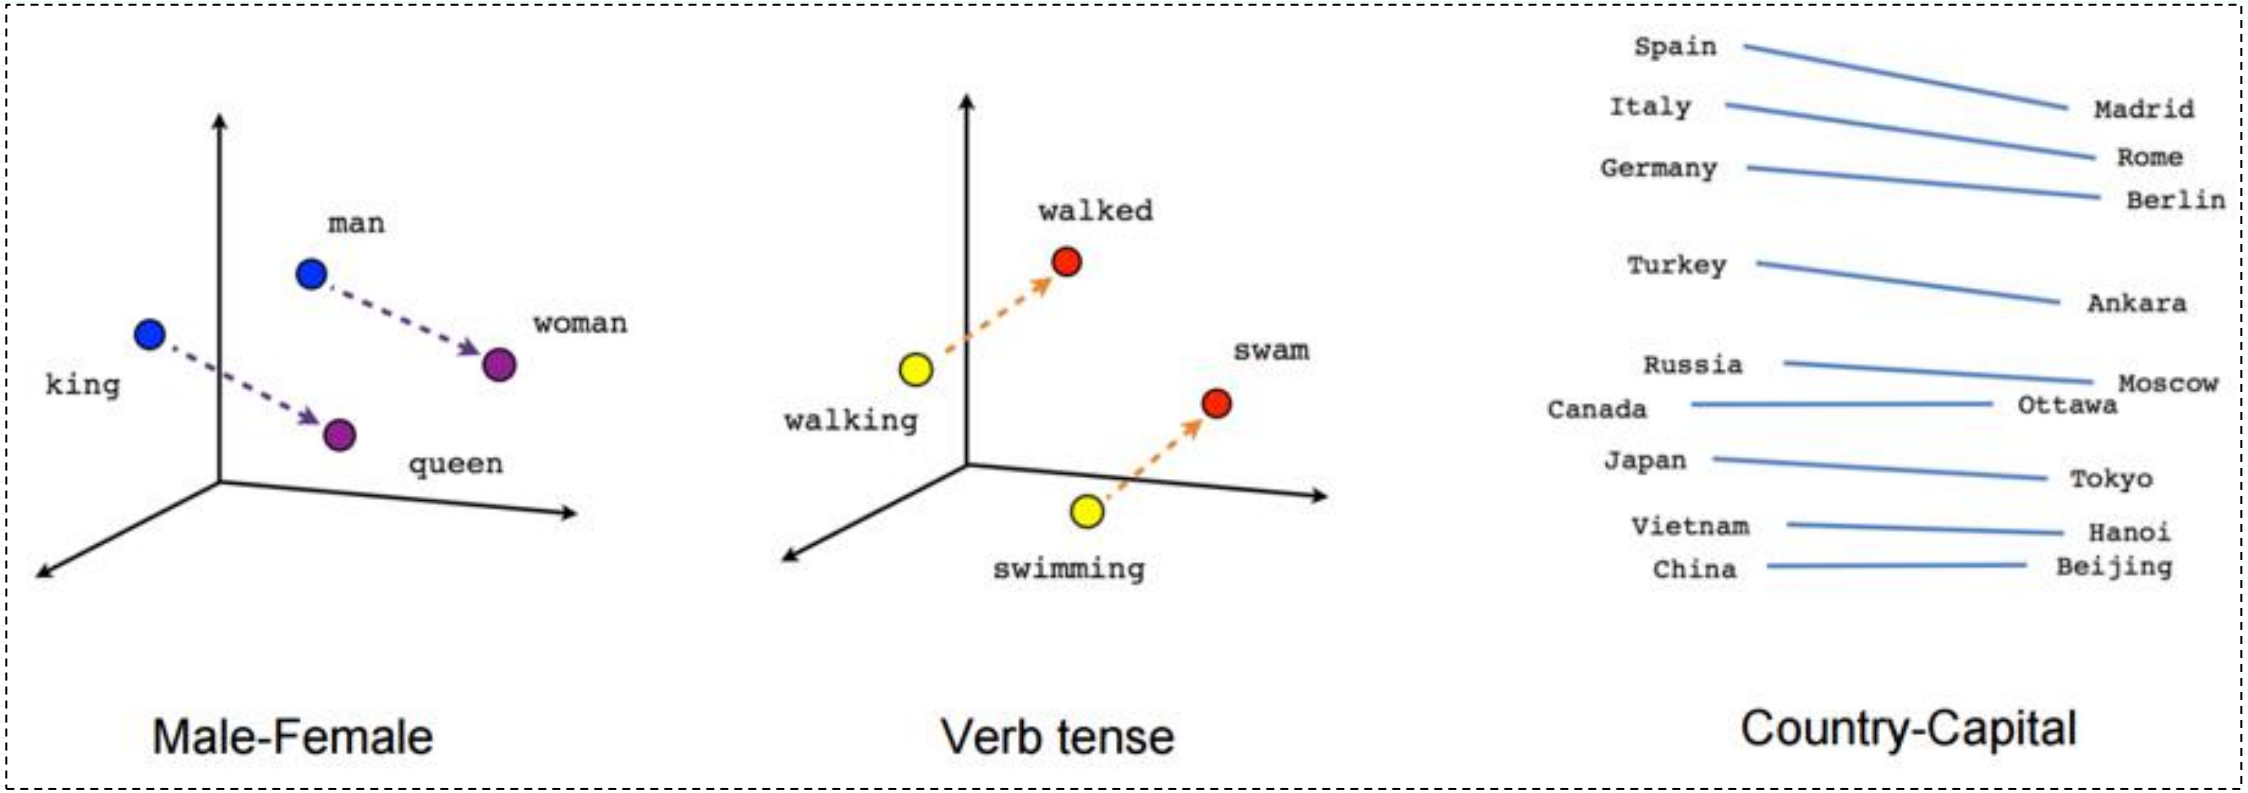

- vector[Queen] ≈ vector[King] - vector[Man] + vector[Woman]
- vector[Paris] ≈ vector[France] - vector[Italy] + vector[Rome]
  - This can be interpreted as “France is to Paris as Italy is to Rome”.

Embeddings are used for many types of data

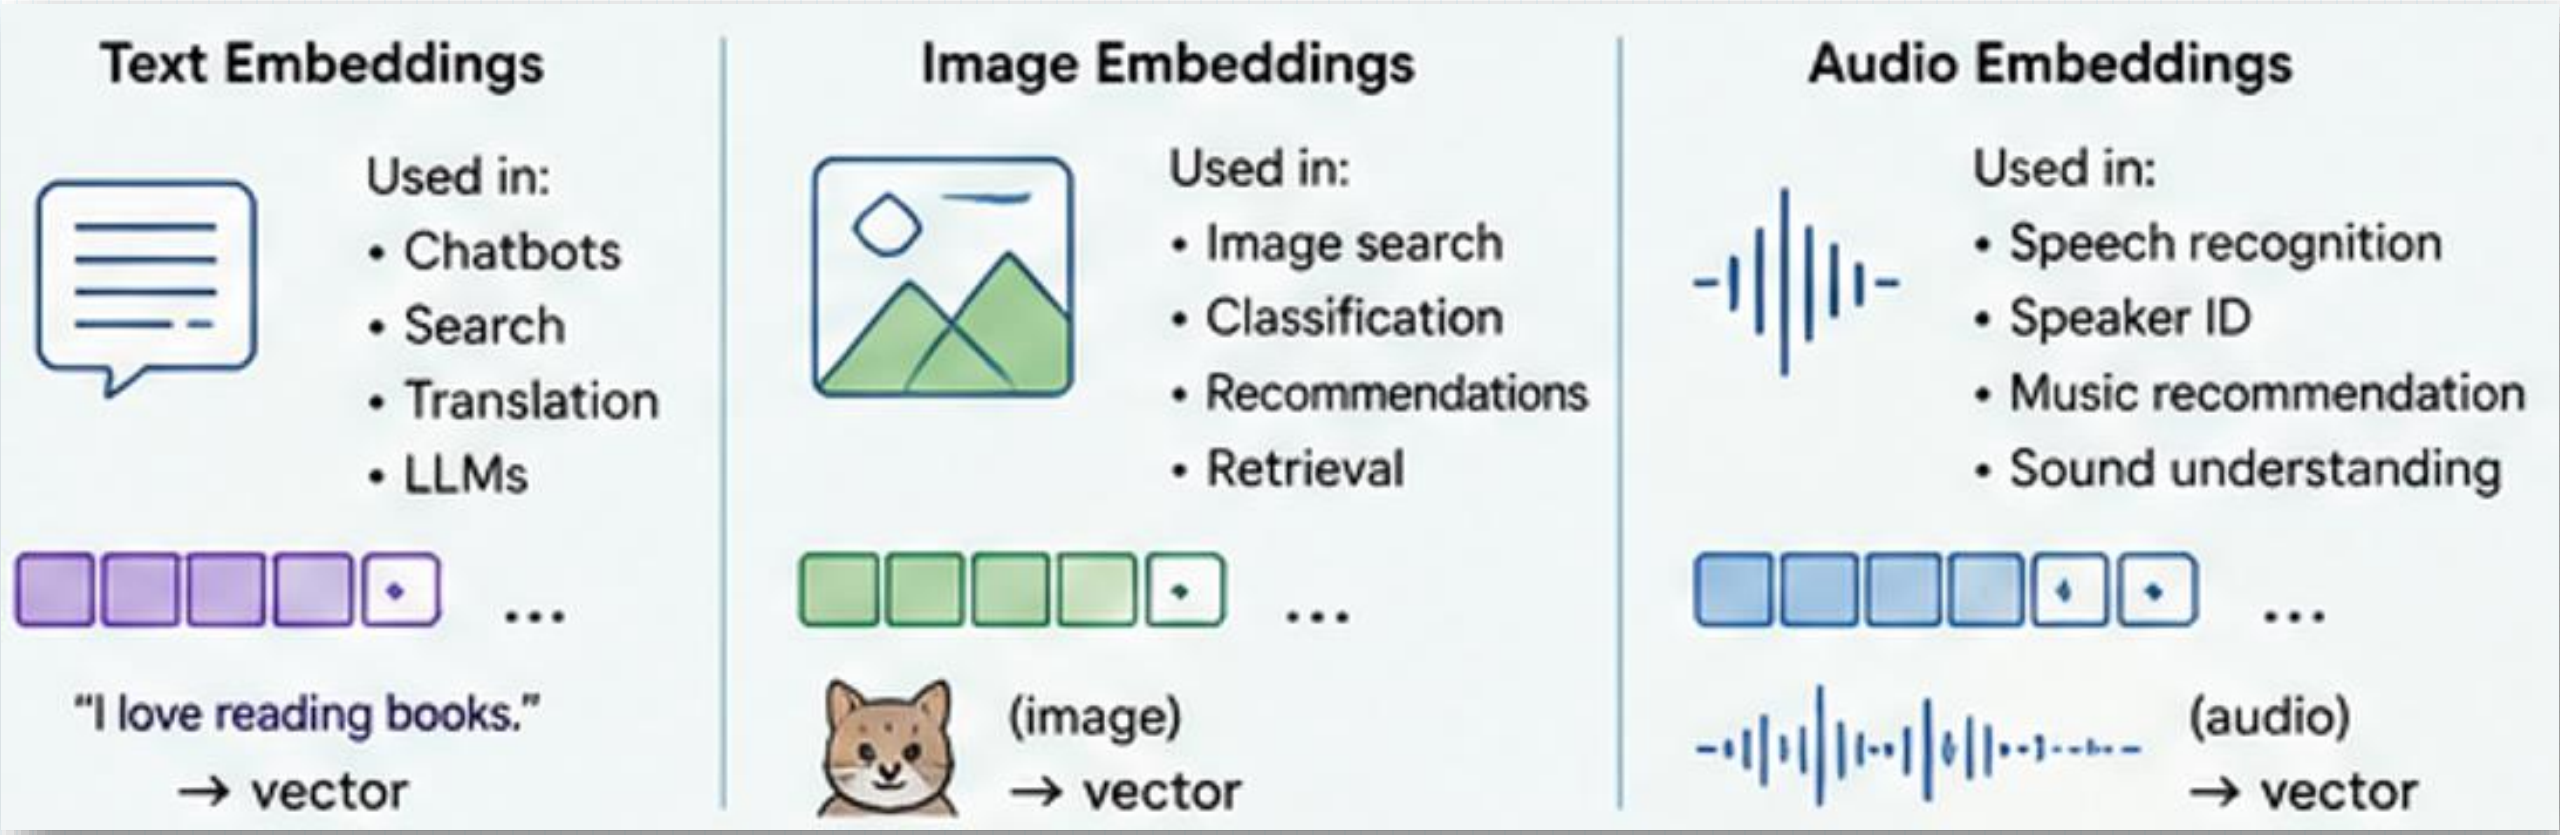

How to build these Embedding Vectors

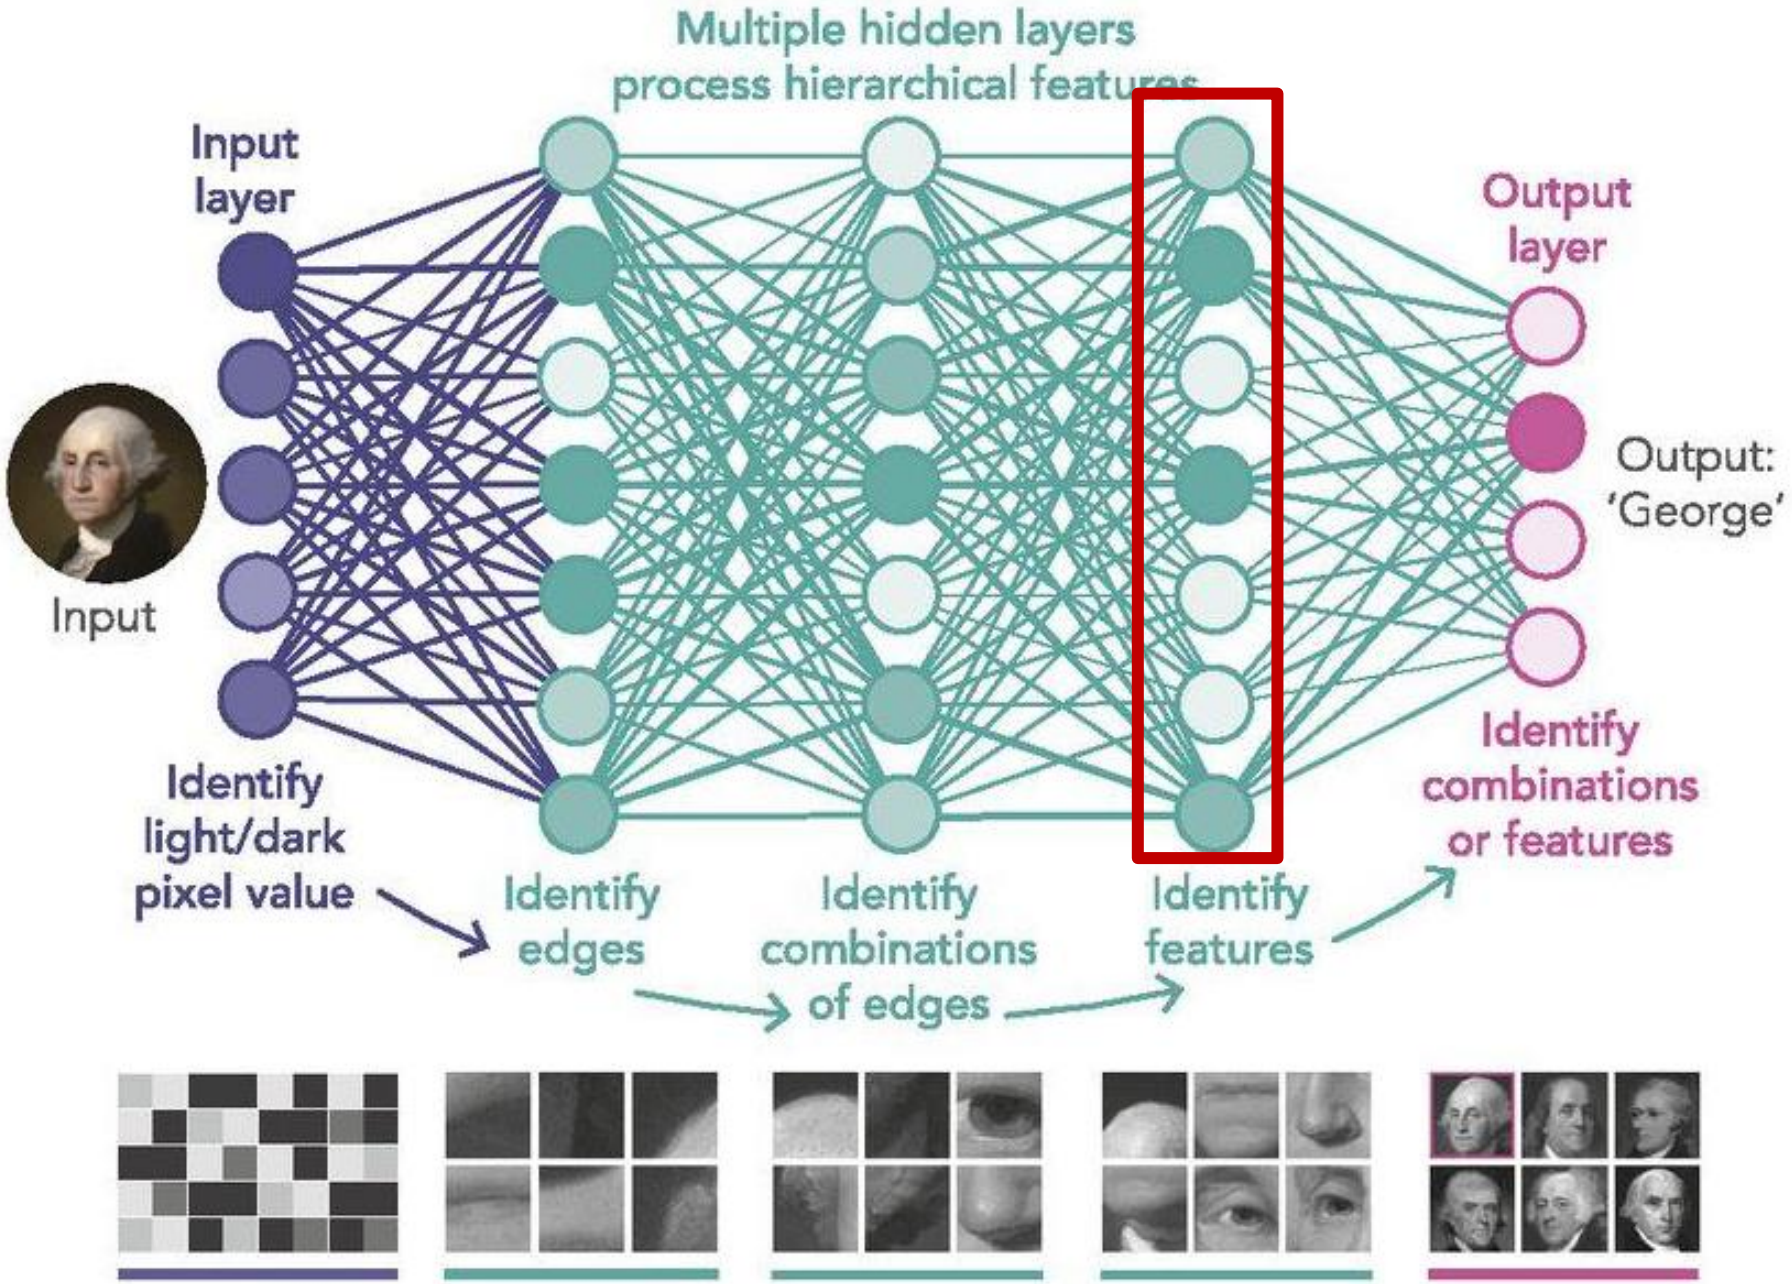
Multiple hidden layers
process hierarchical features
Input
layer
Output
layer
Input
Output:
'George'
Identify
light/dark
pixel value
Identify
edges
Identify
combinations
of edges
Identify
features
Identify
combinations
or features

Example: CNN

Convolution

Max Pooling

A new image

Convolution

Max Pooling

A new image

Flatten

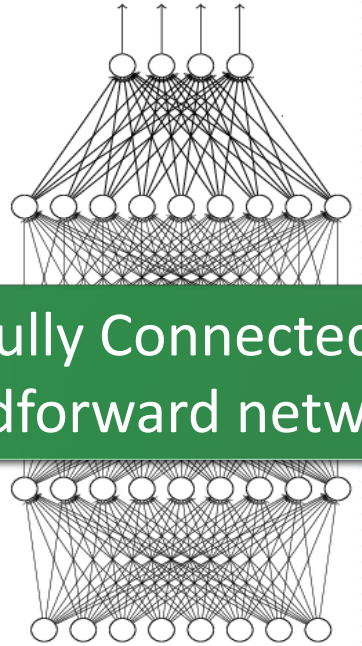

Fully Connected Feedforward network

cat dog ……

BUT…

The embedding depends entirely on the task the network is trained on.

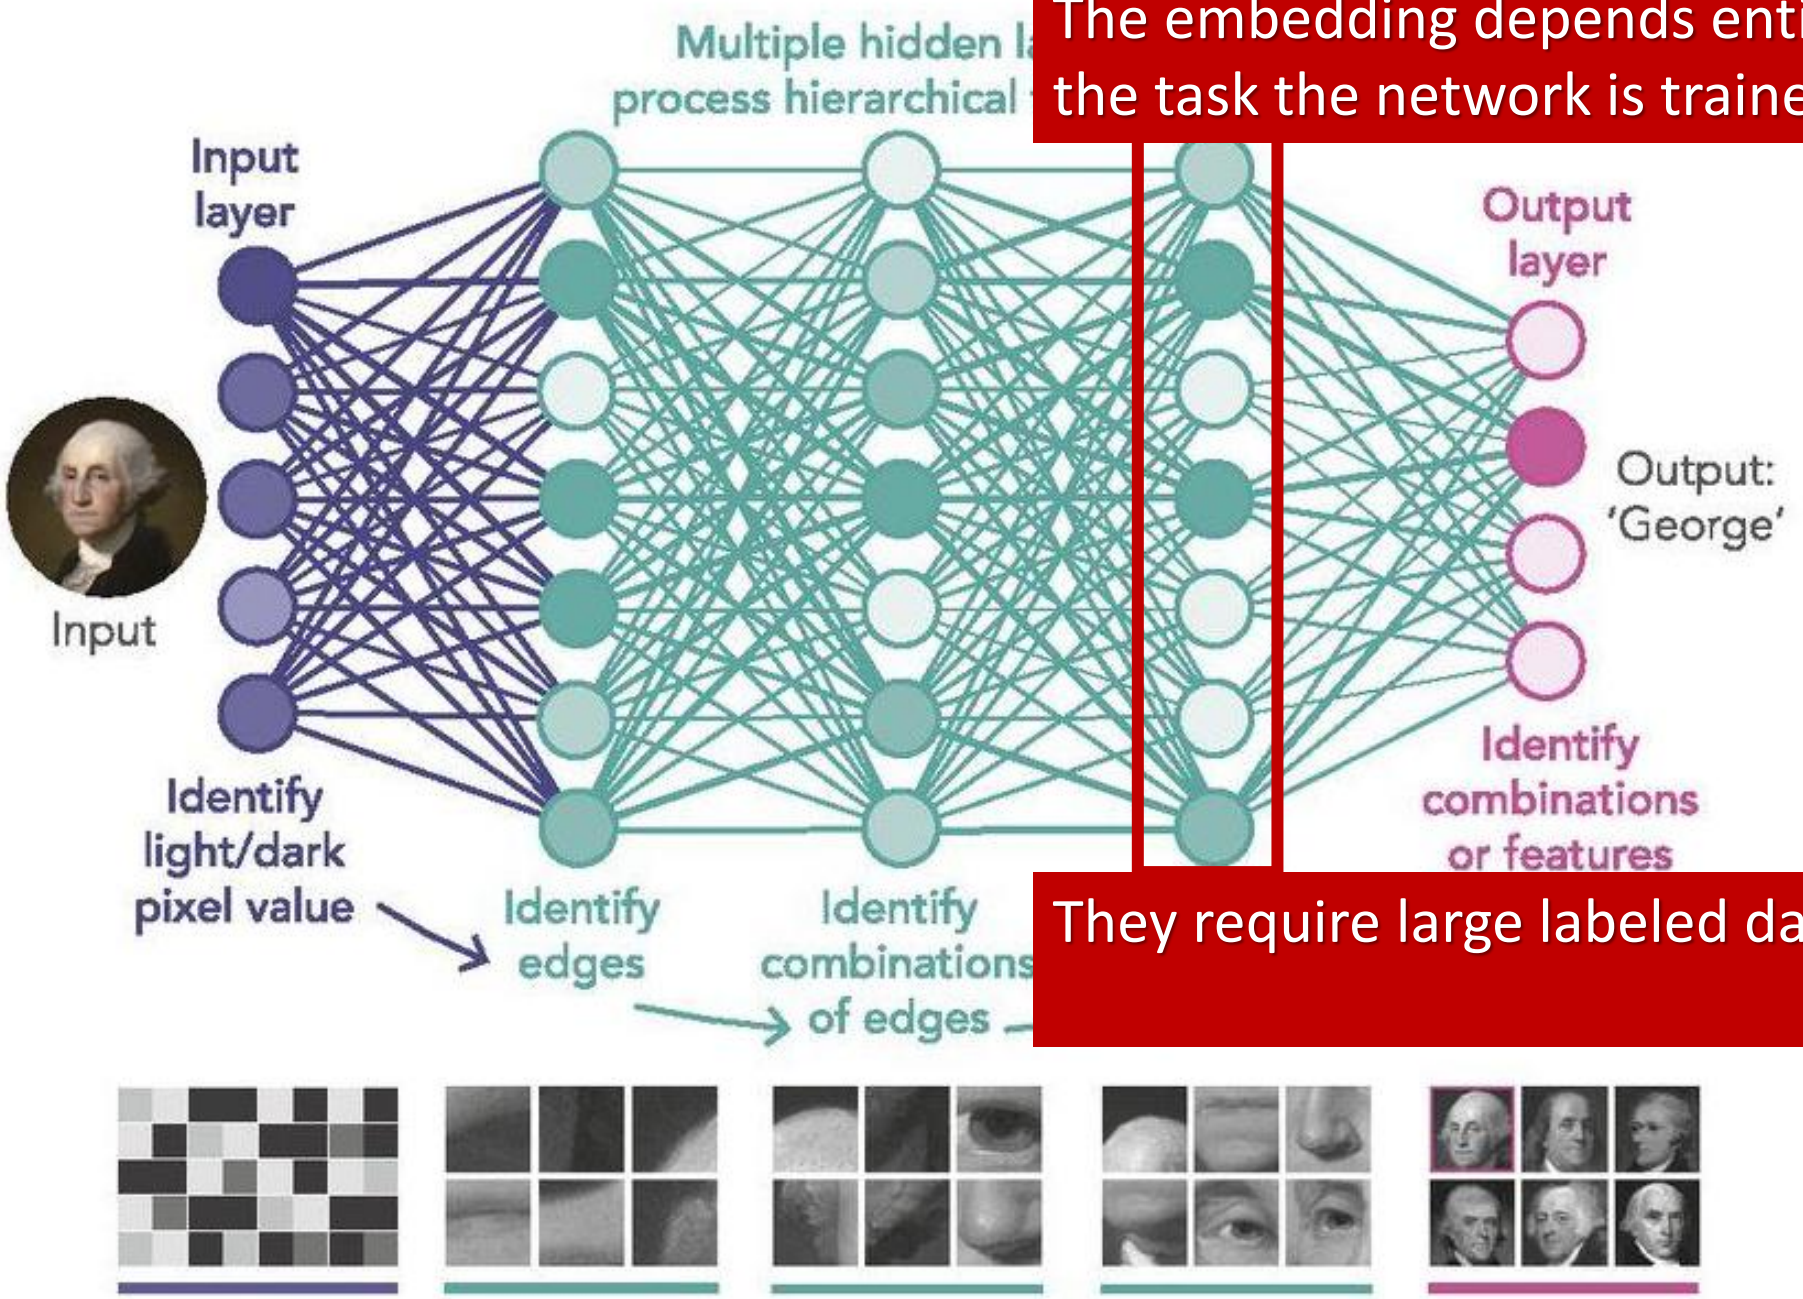

They require large labeled datasets!

Solution: Self-supervised learning

- This idea—**“given a word, predict the surrounding words”**

she had to pay using ___ instead of any digital or card-based method.

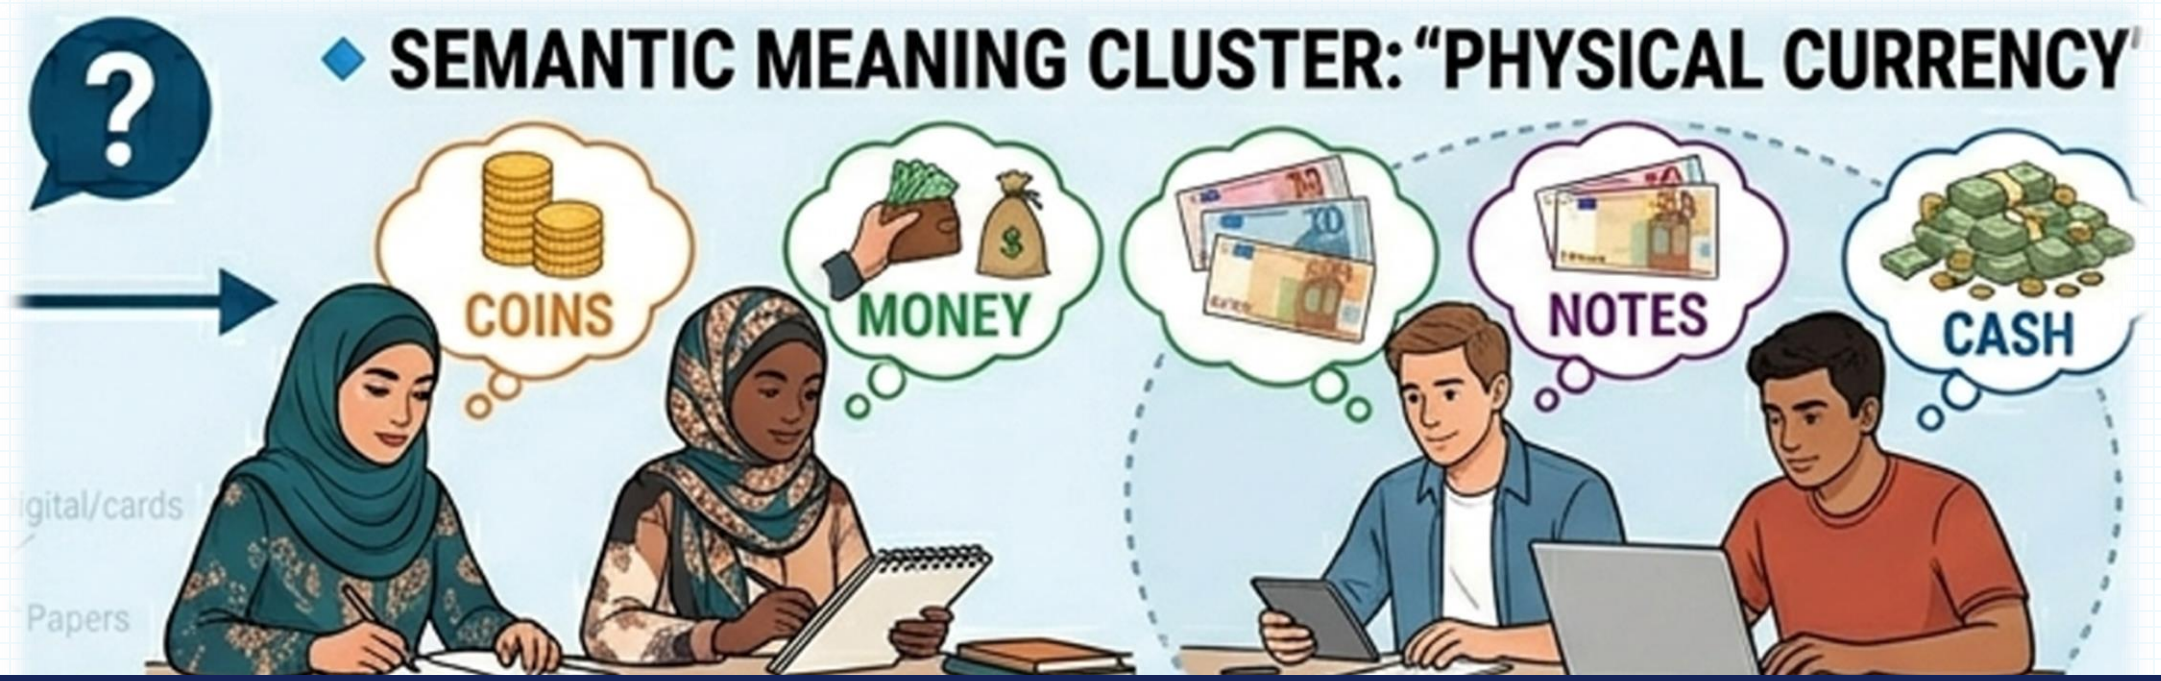

The model is not learning exact word replacement. It is learning a region of meaning.

word2vec

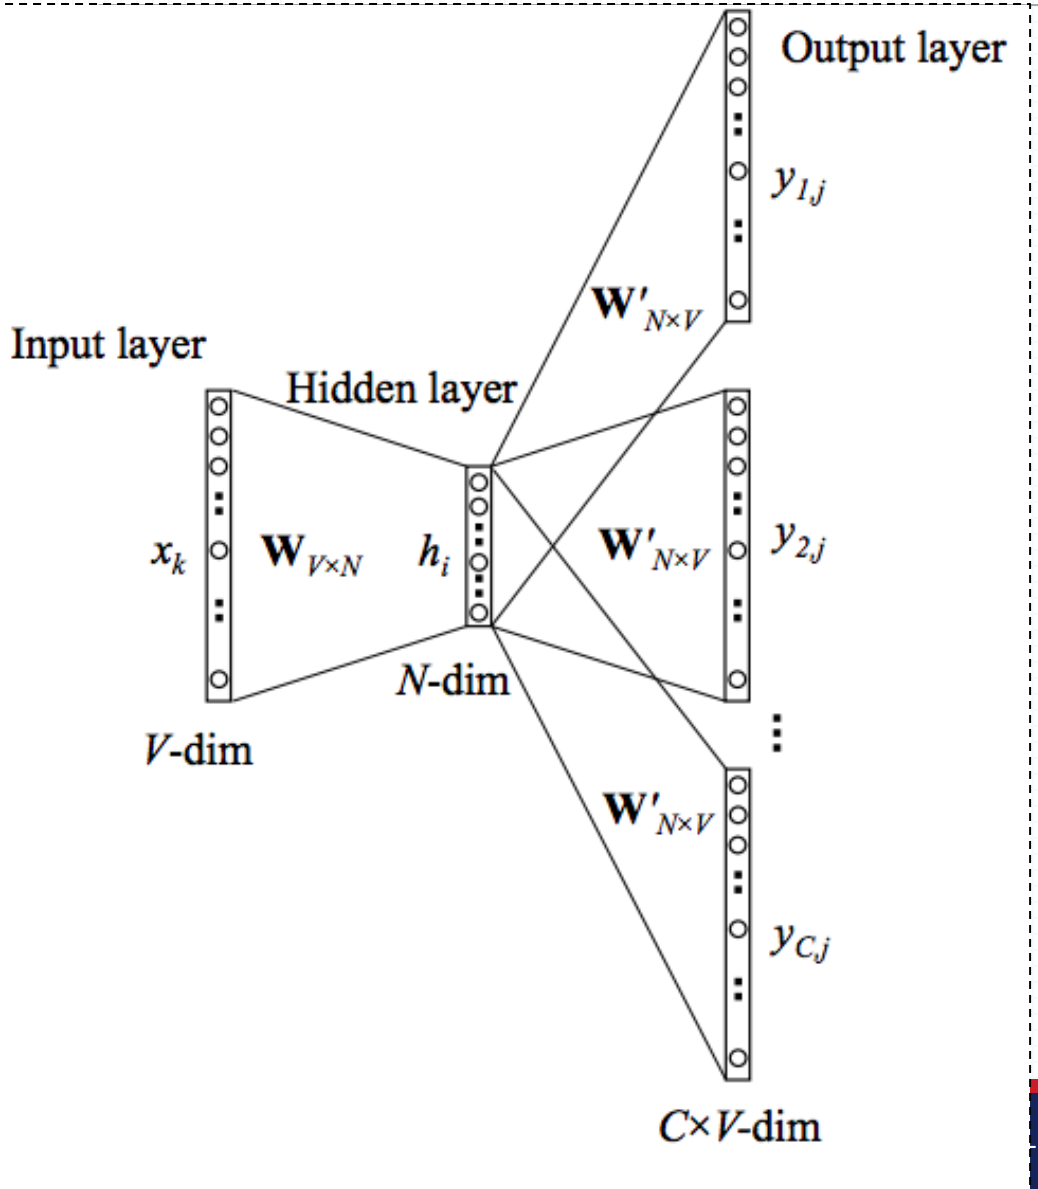
Output layer
Input layer
Hidden layer
$y_{1,j}$
$\mathbf{W}'_{N\times V}$
x_k
$\mathbf{W}_{V\times N}$
h_i
$\mathbf{W}'_{N\times V}$
$y_{2,j}$
N-dim
V-dim
$\mathbf{W}'_{N\times V}$
$y_{C,j}$
$C\times V$-dim

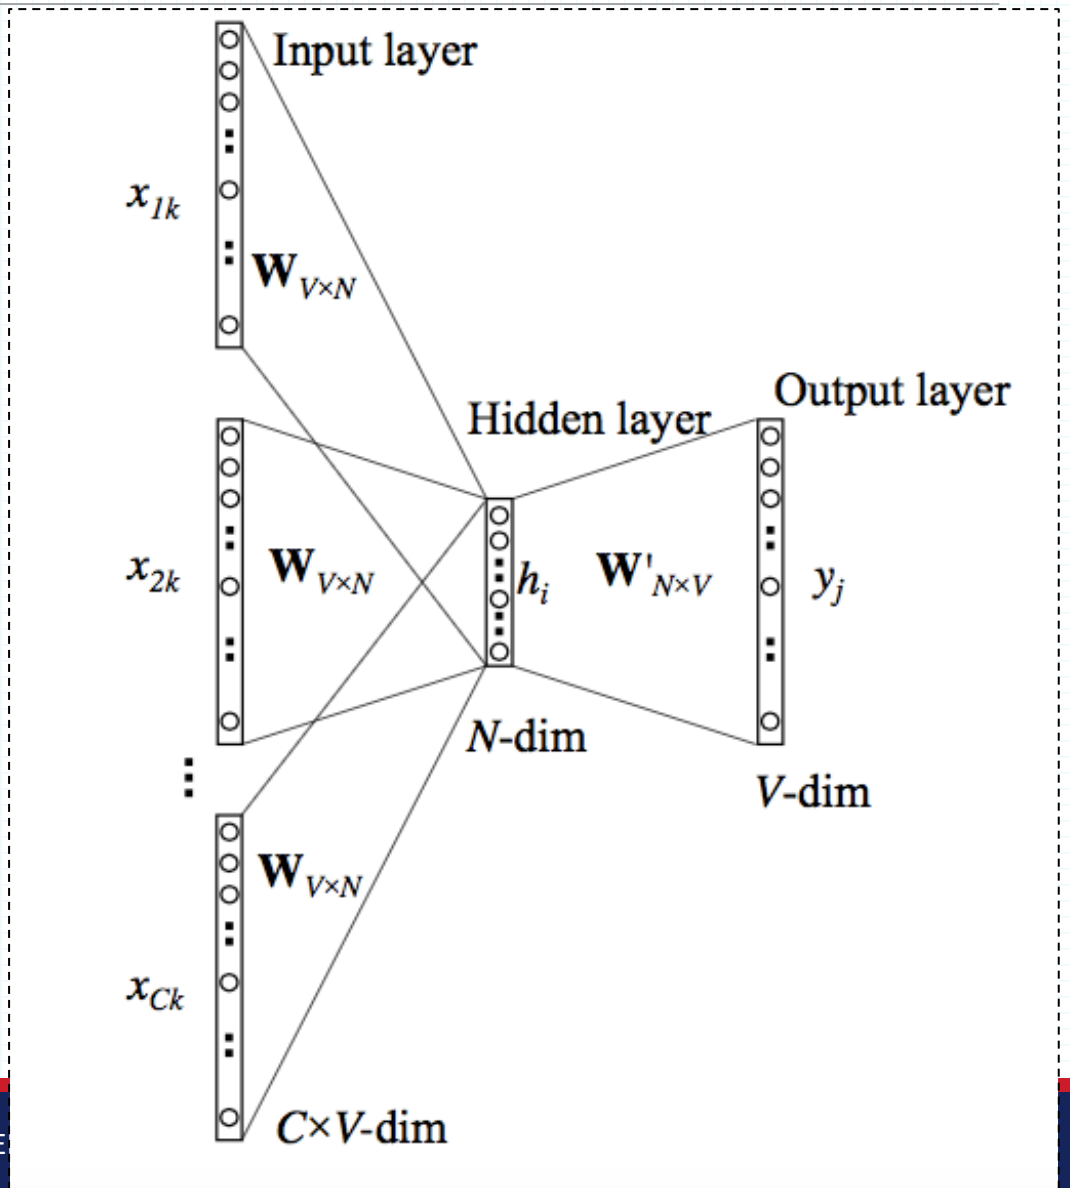
Input layer
x_{1k}
$\mathbf{W}_{V\times N}$
Hidden layer
Output layer
x_{2k}
$\mathbf{W}_{V\times N}$
h_i
$\mathbf{W}'_{N\times V}$
y_j
N-dim
V-dim
$\mathbf{W}_{V\times N}$
x_{Ck}
$C\times V$-dim

Embedding Vector

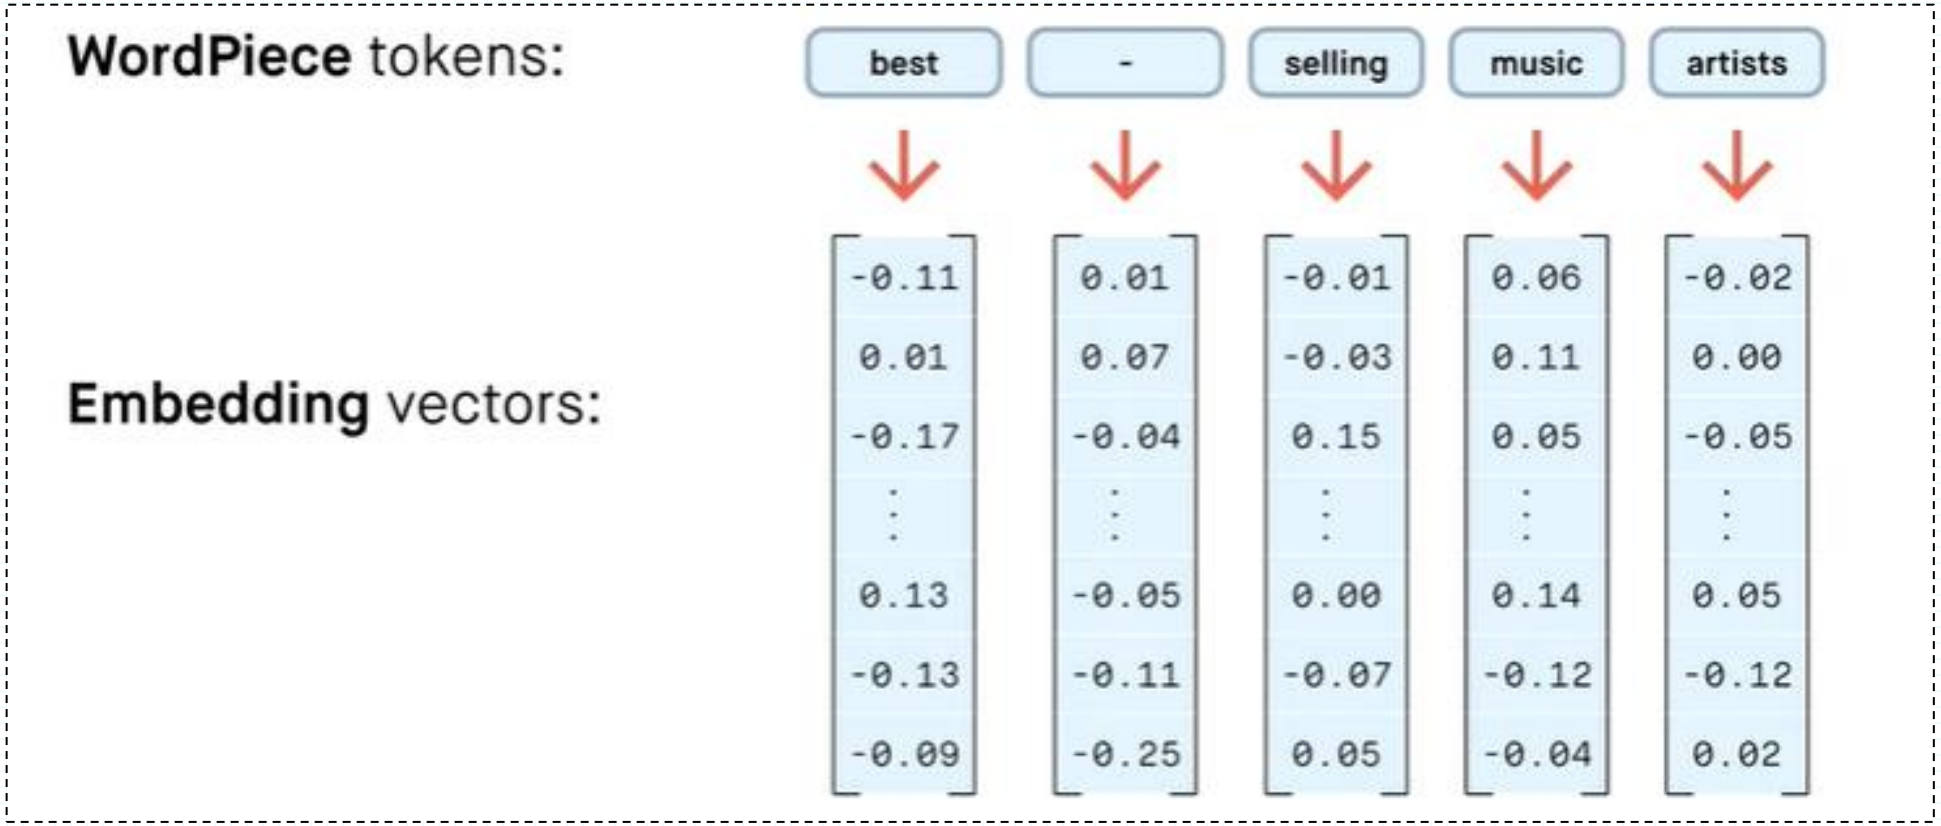

But… Contextual Embedding

Solution.. Attention

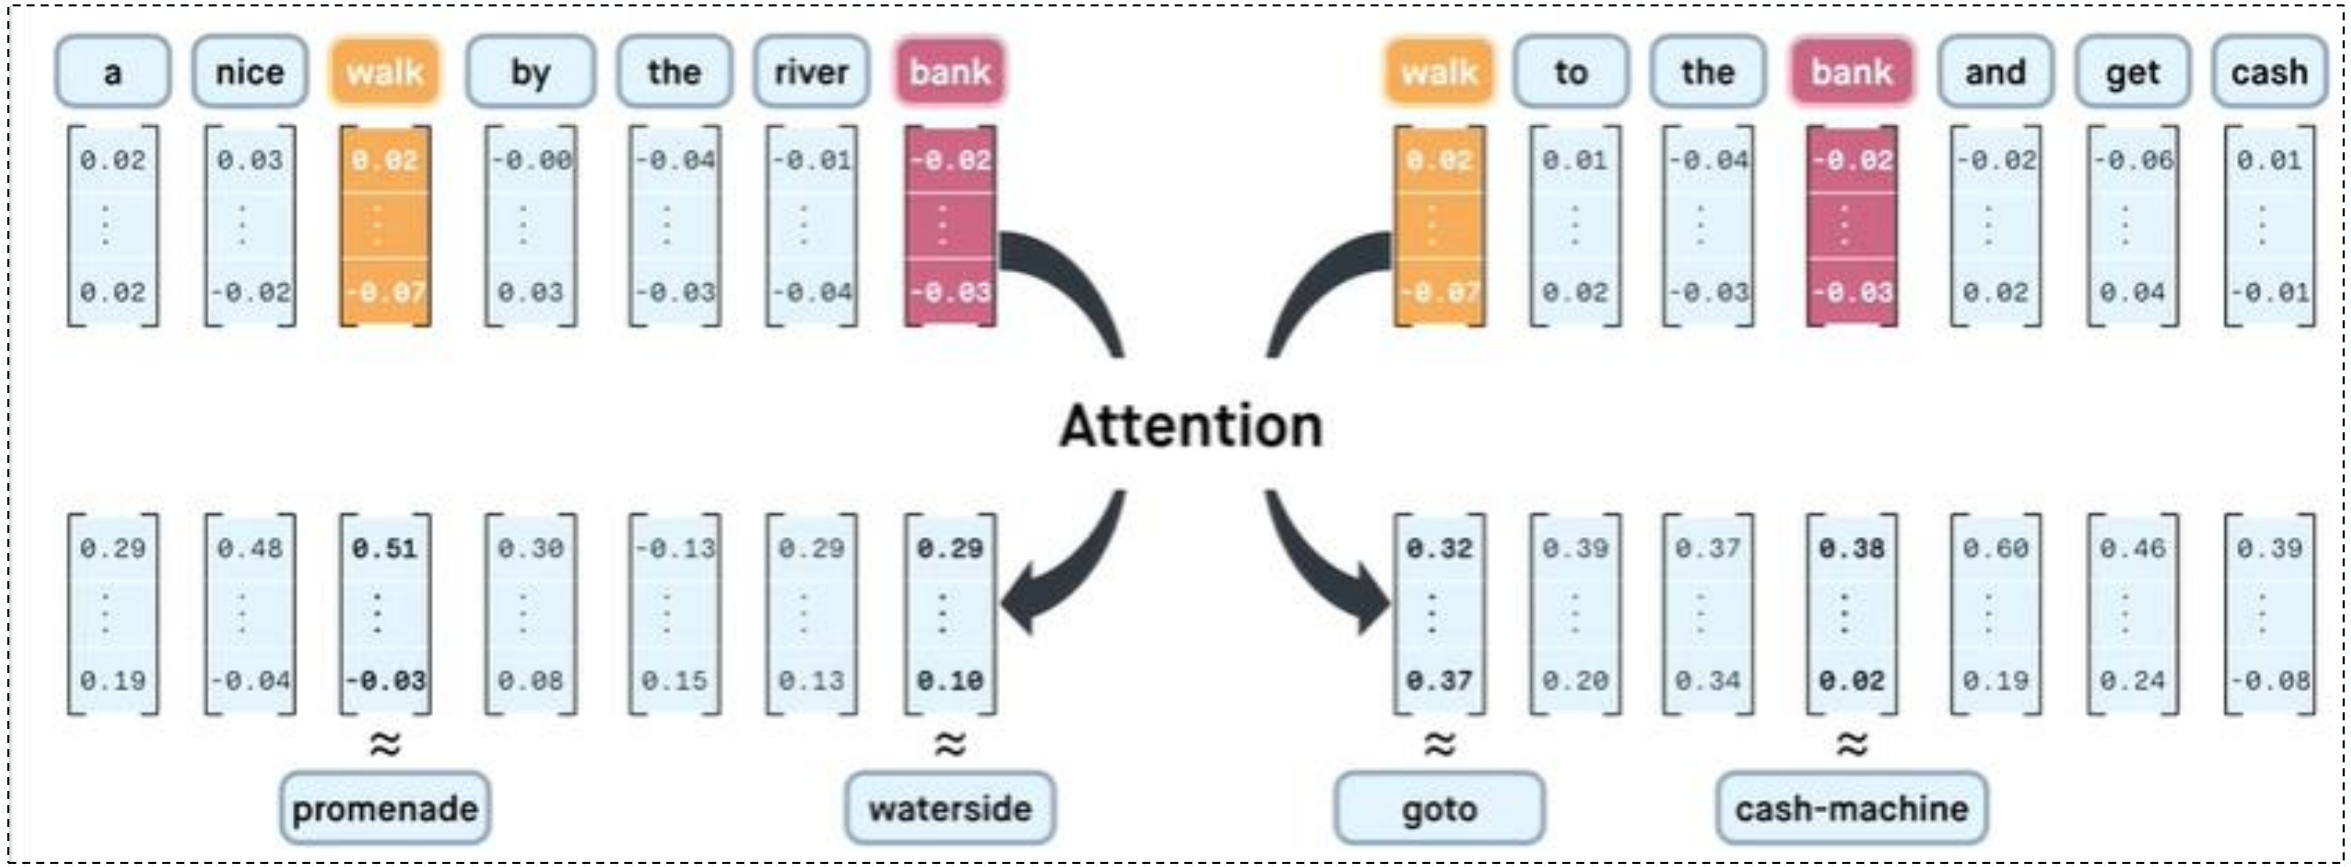

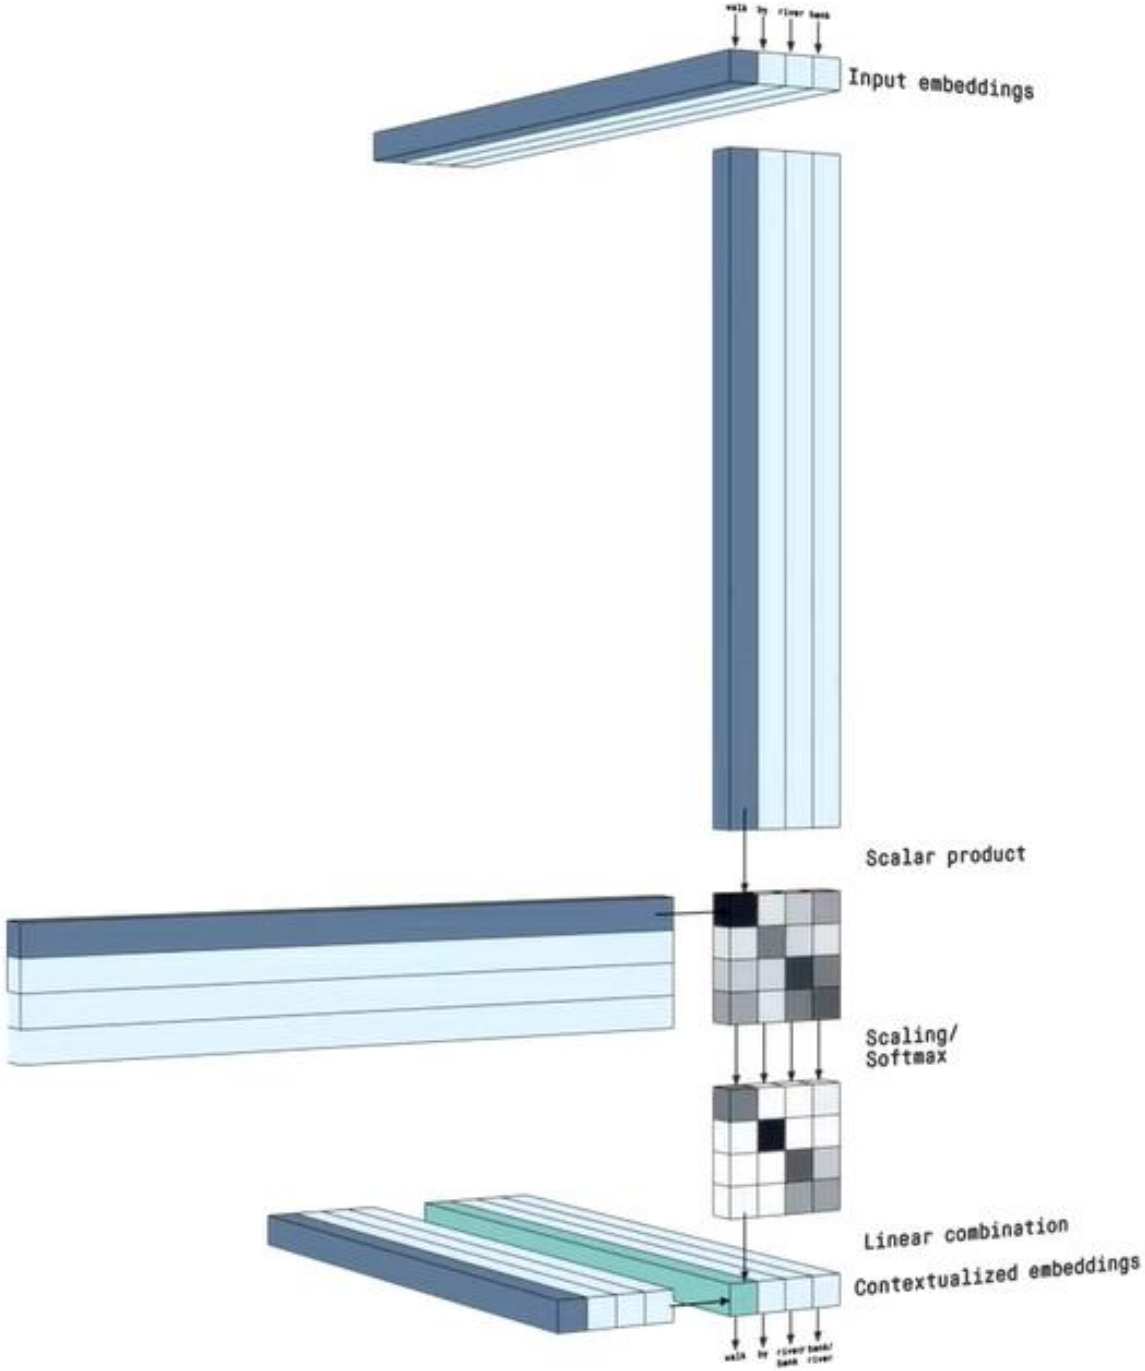
Input embeddings
Scalar product
Scaling/
Softmax
Linear combination
Contextualized embeddings

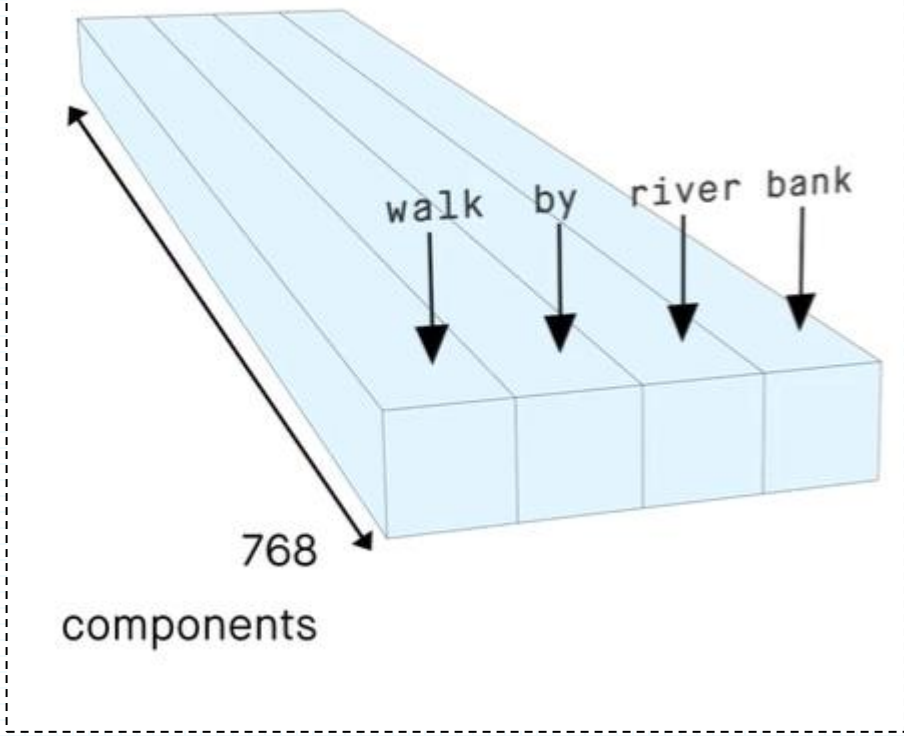
walk
by
river bank
768
components

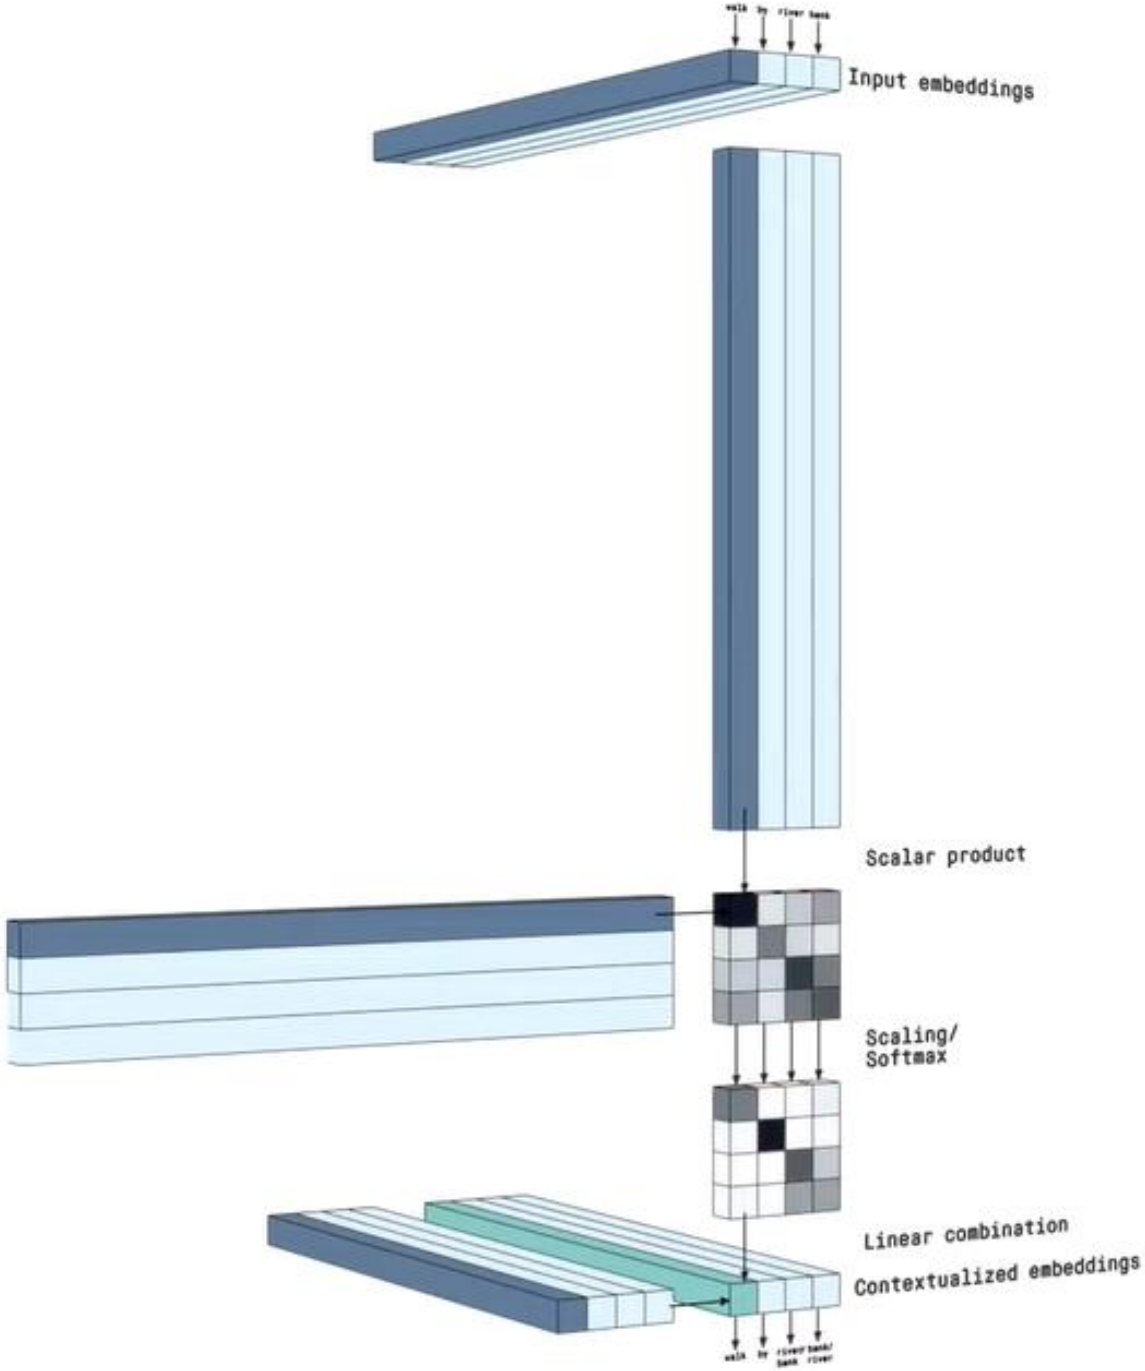
Input embeddings
Scalar product
Scaling/
Softmax
Linear combination
Contextualized embeddings

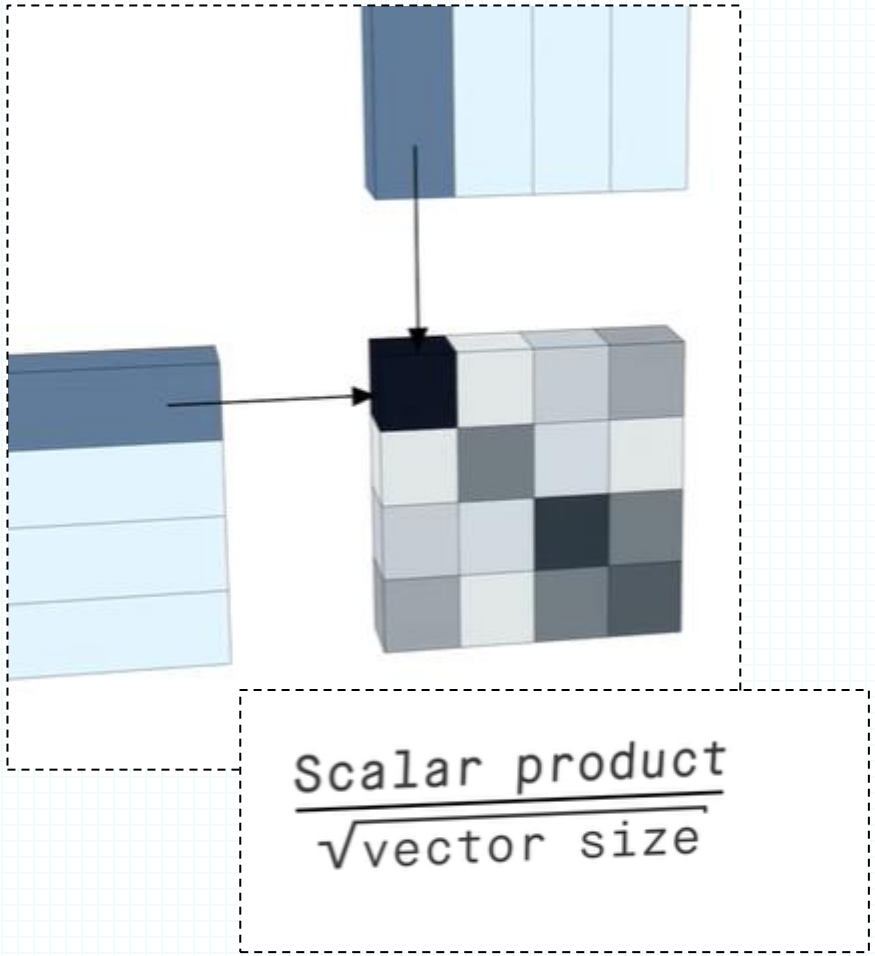
Scalar product
√vector size

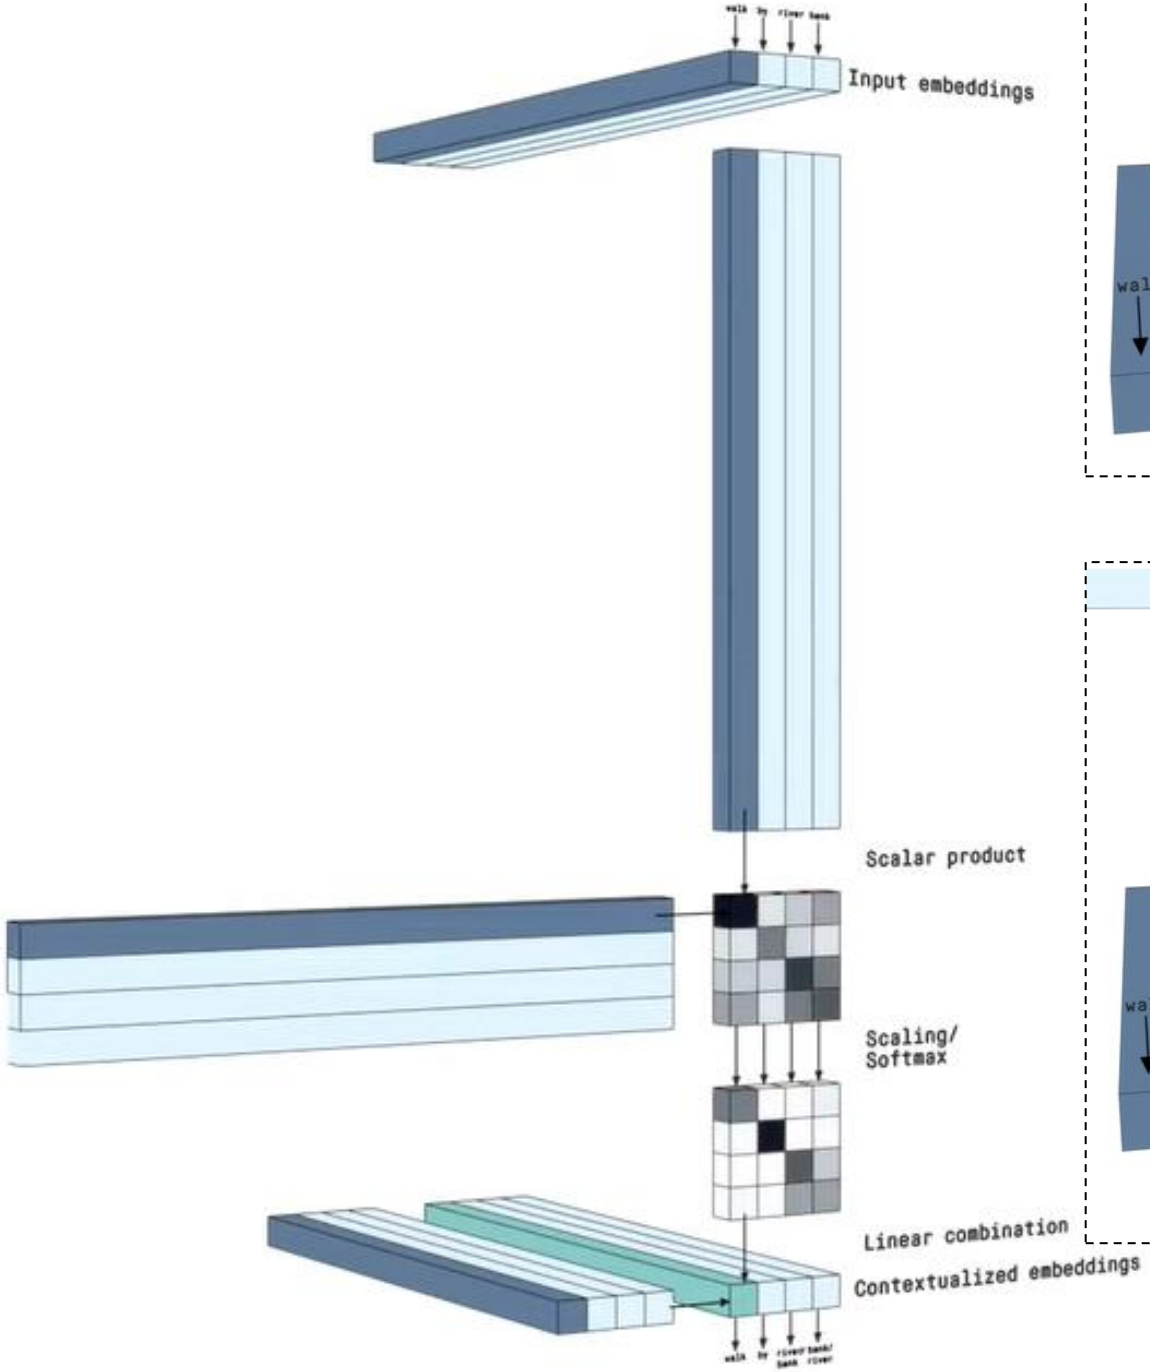
Input embeddings
Scalar product
Scaling/
Softmax
Linear combination
Contextualized embeddings

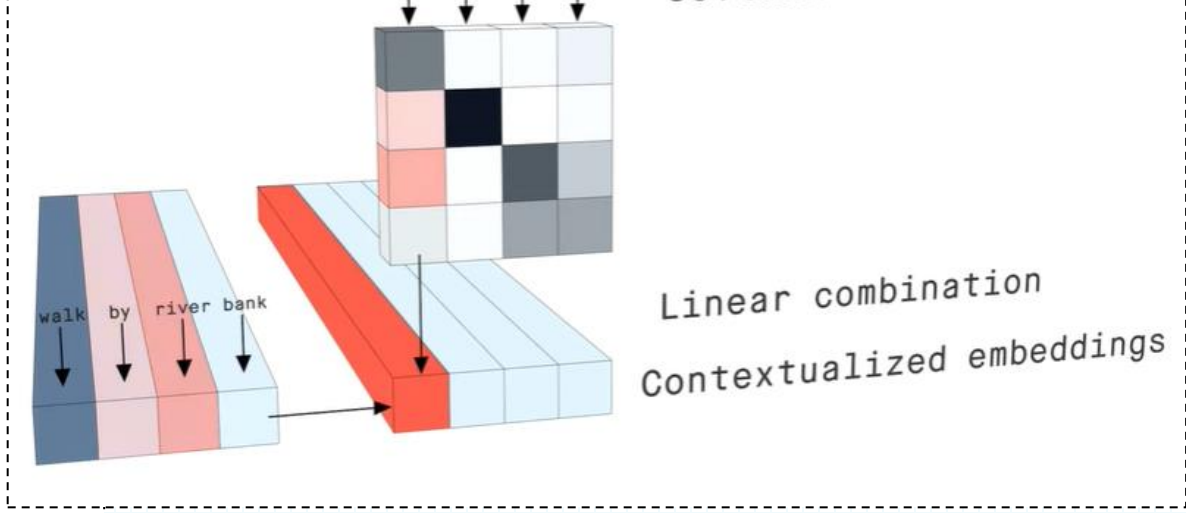
walk
by
river bank
Linear combination
Contextualized embeddings

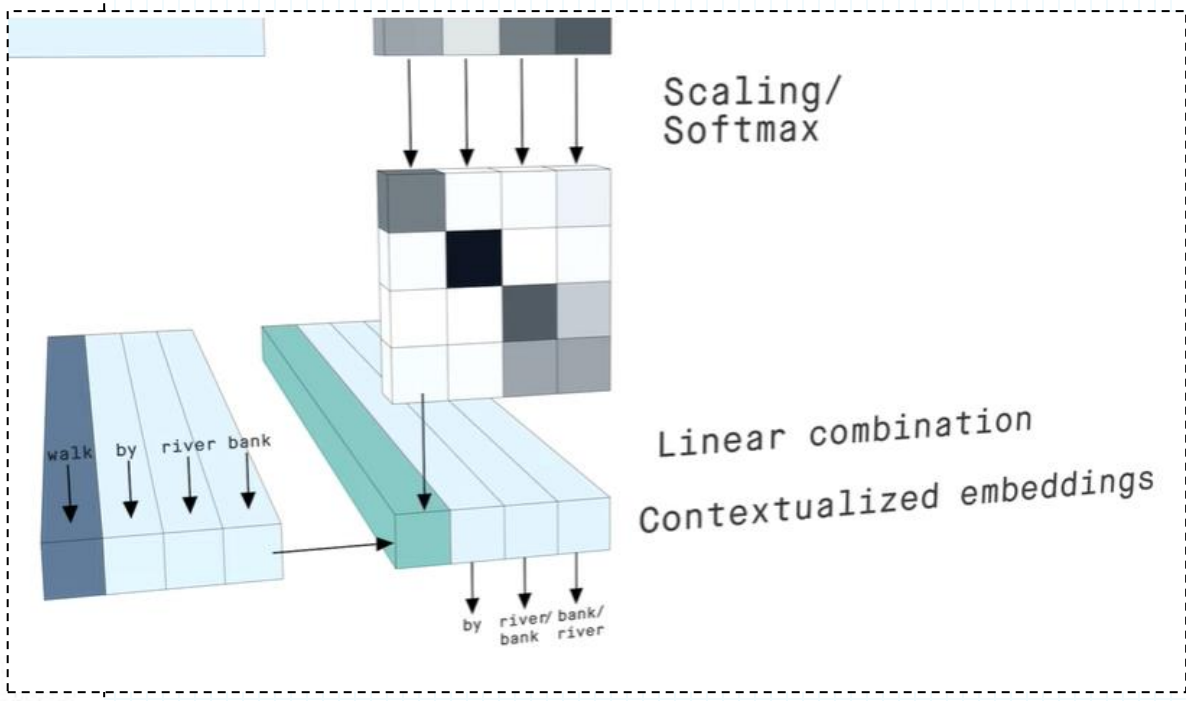
Scaling/
Softmax
walk
by
river bank
Linear combination
Contextualized embeddings
by
river/
bank
bank/
river

Attention

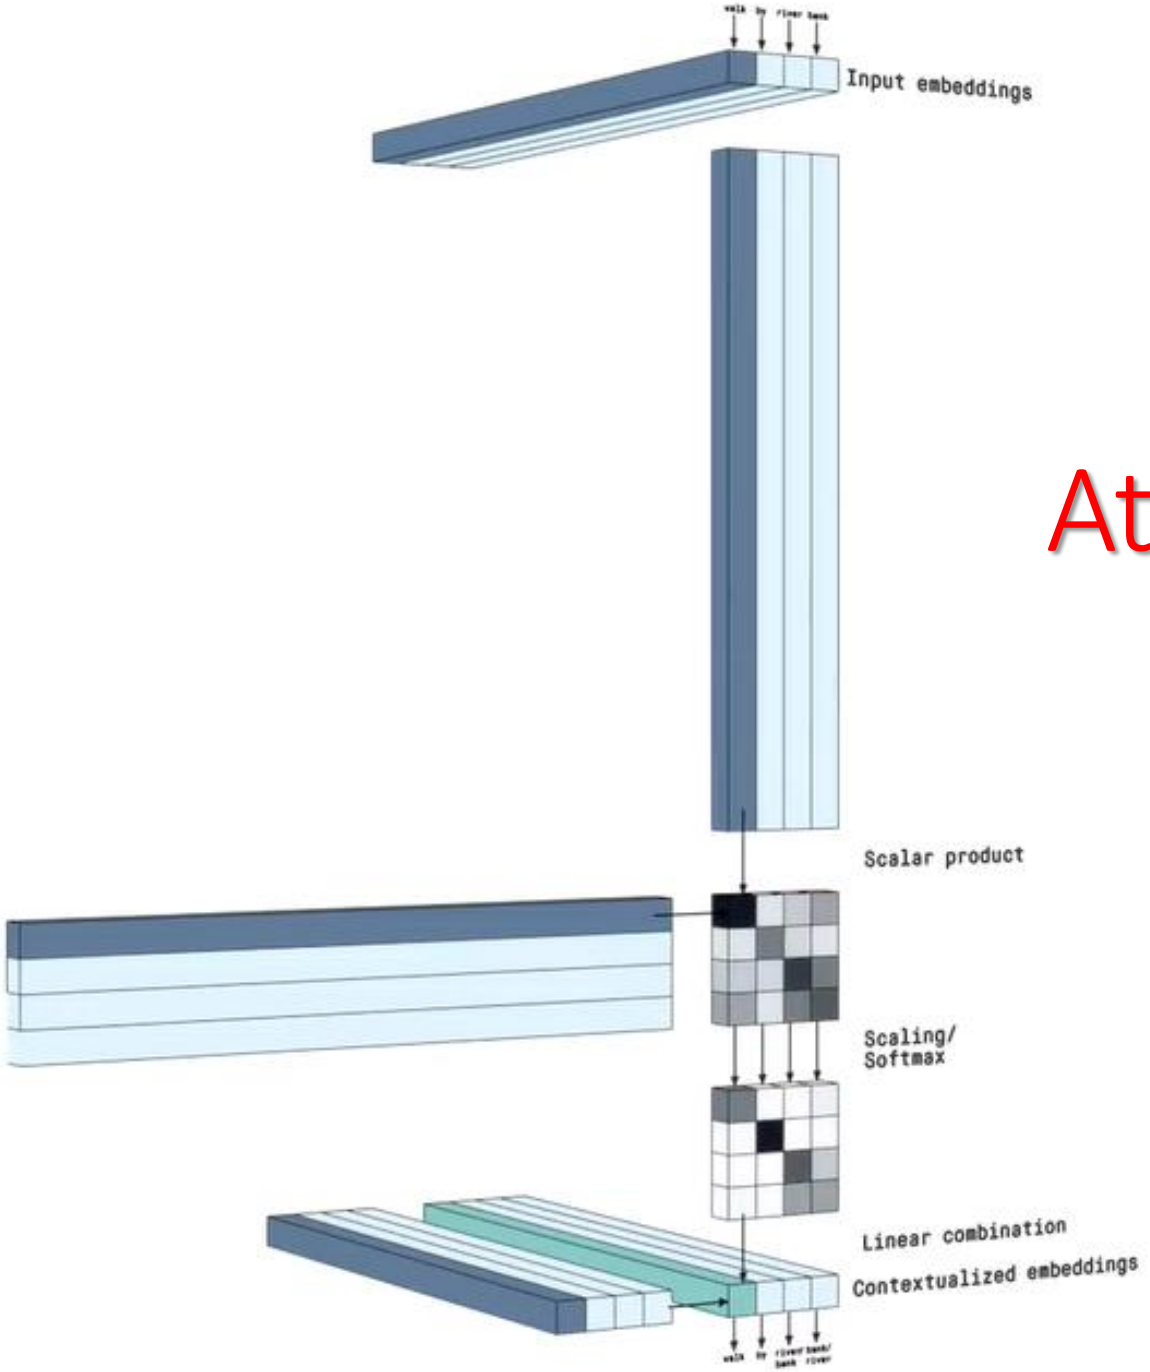
Input embeddings
Scalar product
Scaling/
Softmax
Linear combination
Contextualized embeddings

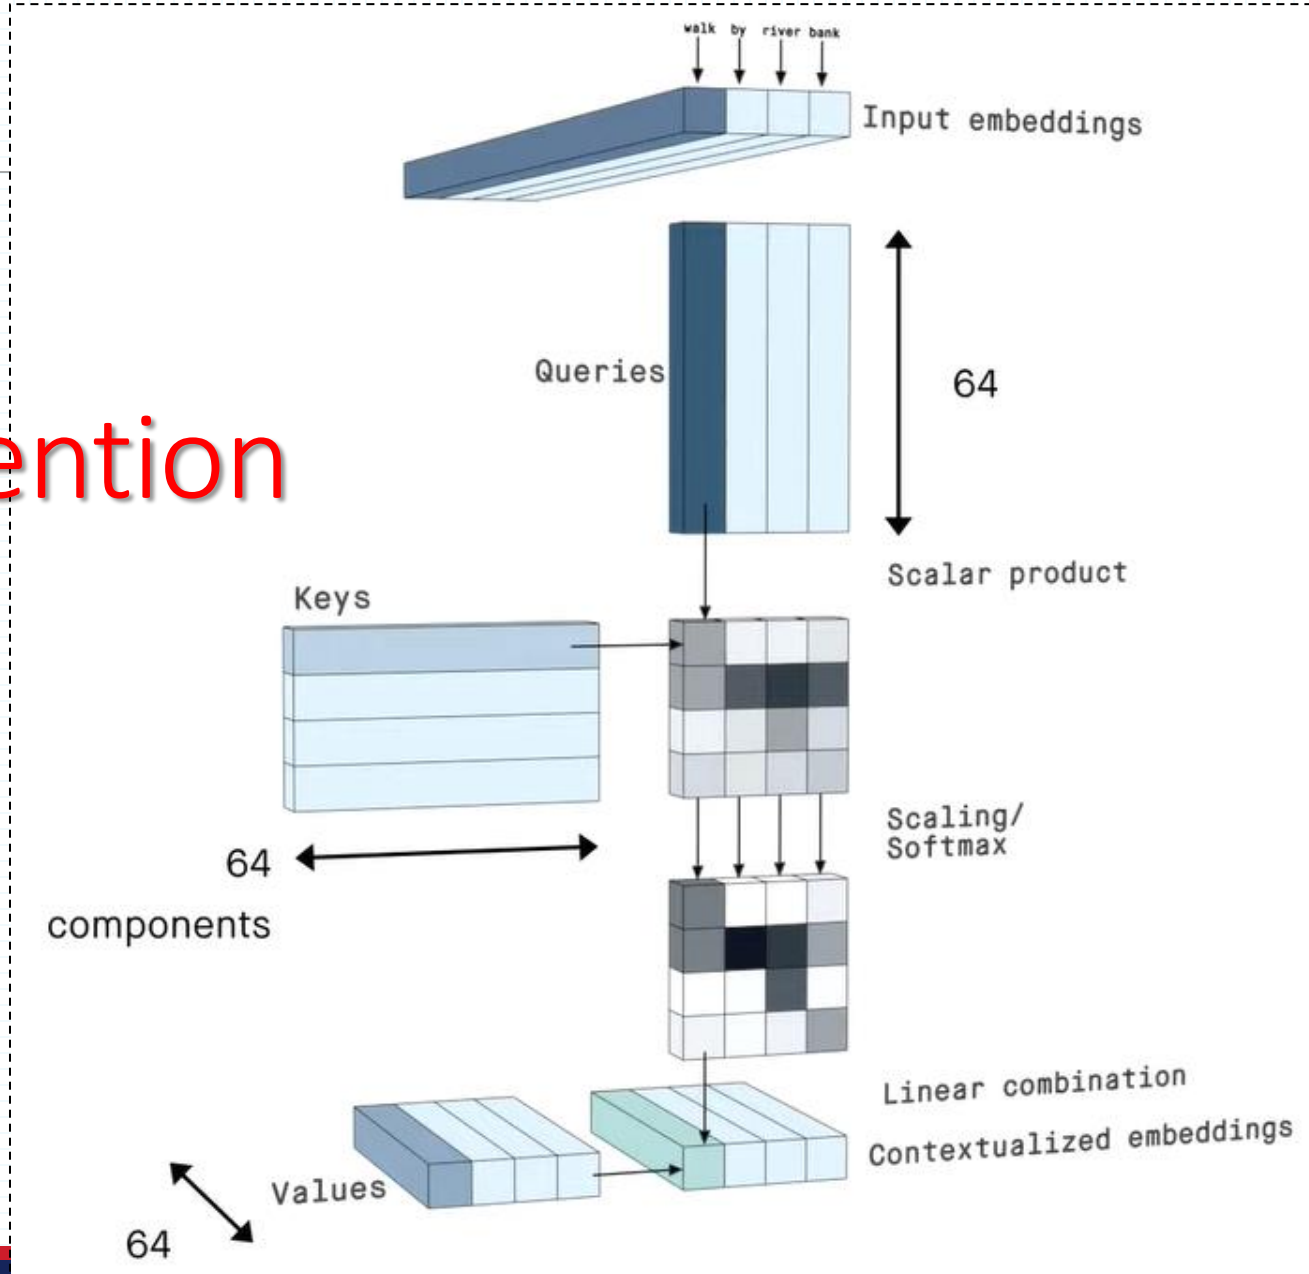
walk by river bank
Input embeddings
Queries
64
Scalar product
Keys
64
components
Scaling/
Softmax
Linear combination
Contextualized embeddings
Values
64

Multi-head attention

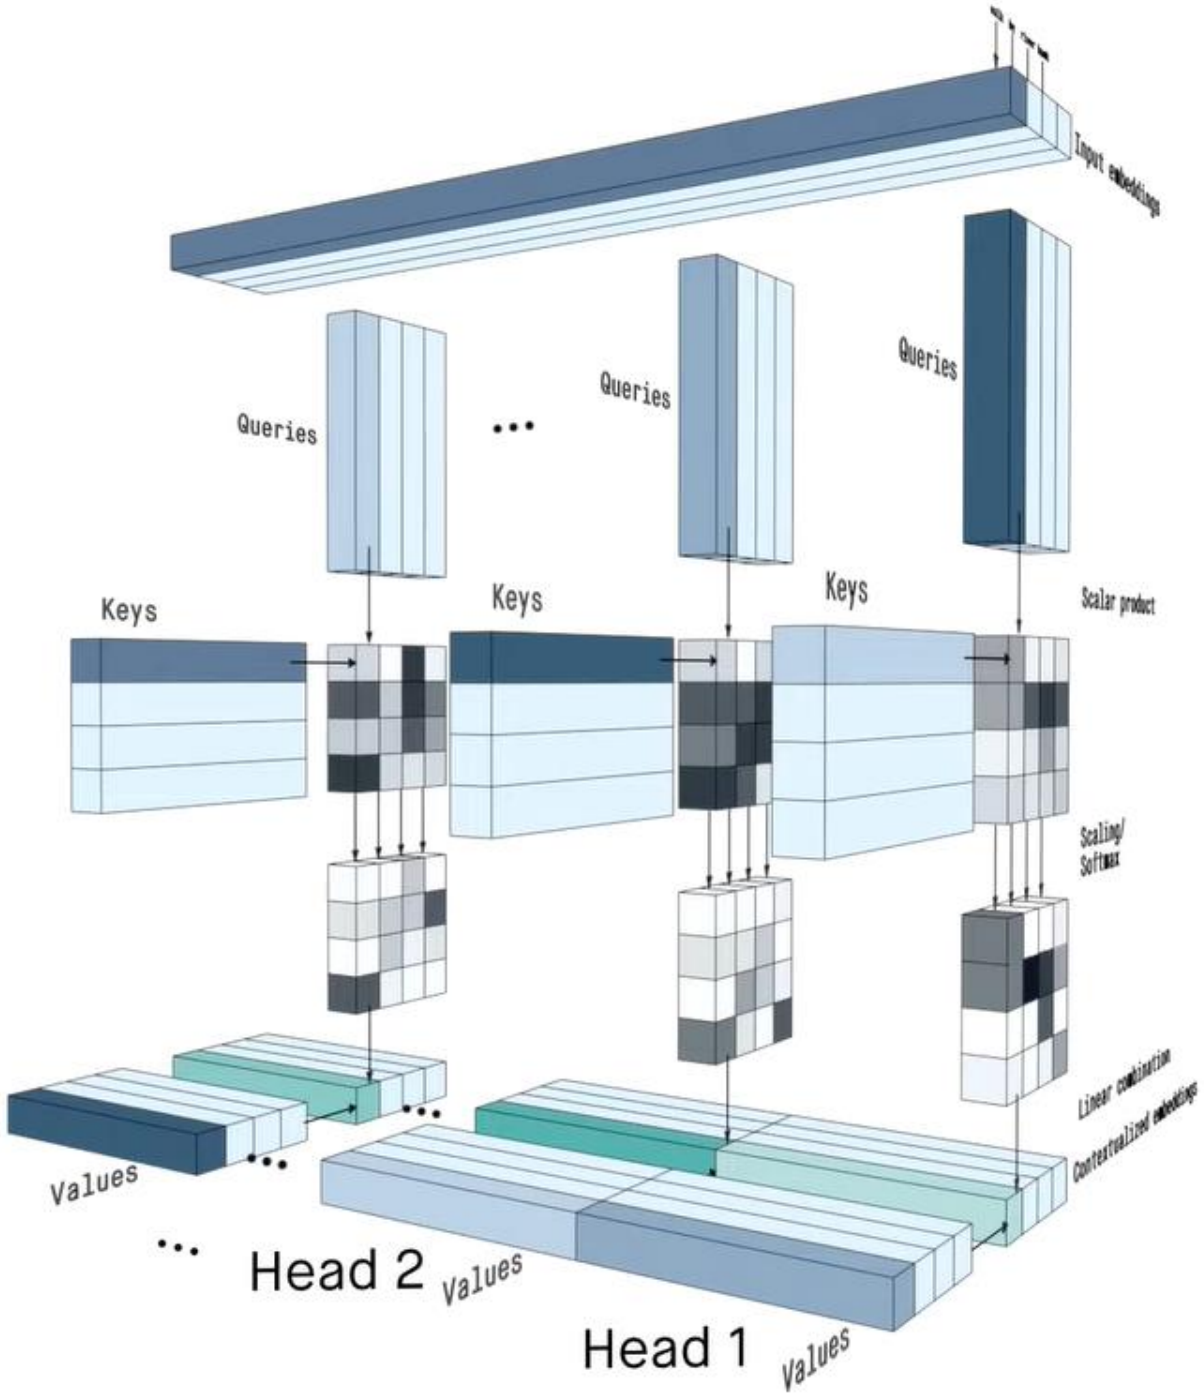

Enrich the input!

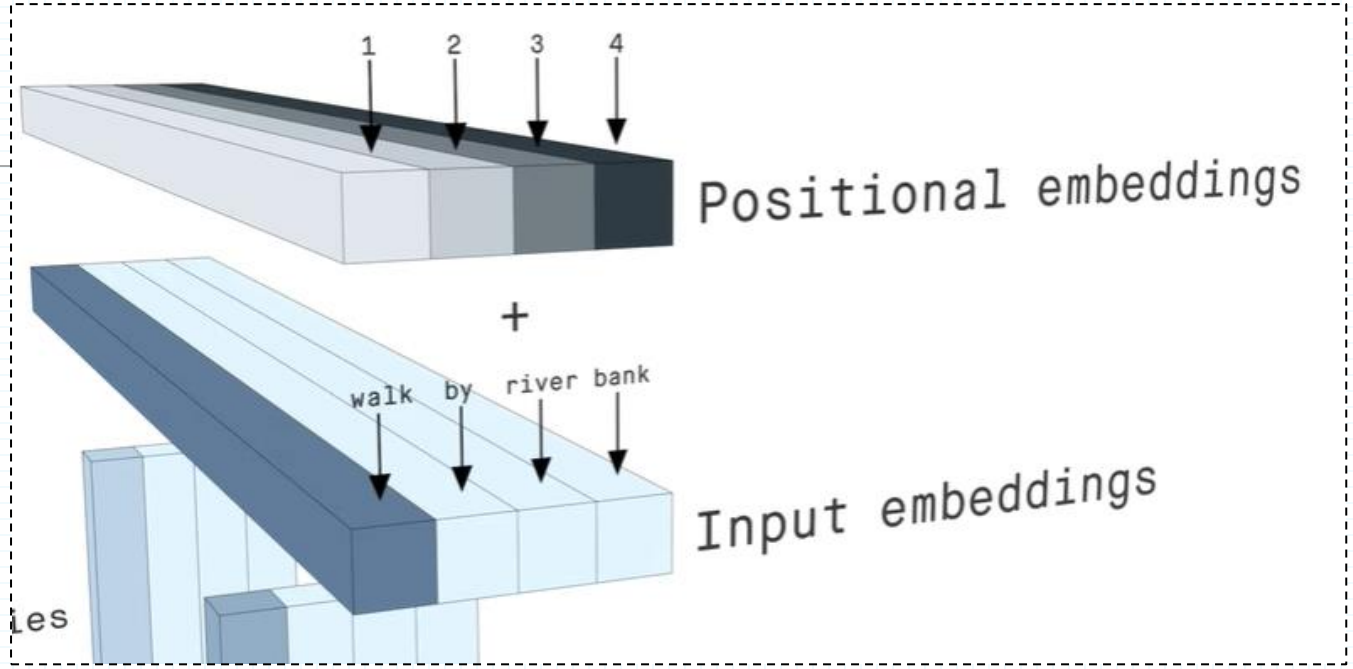

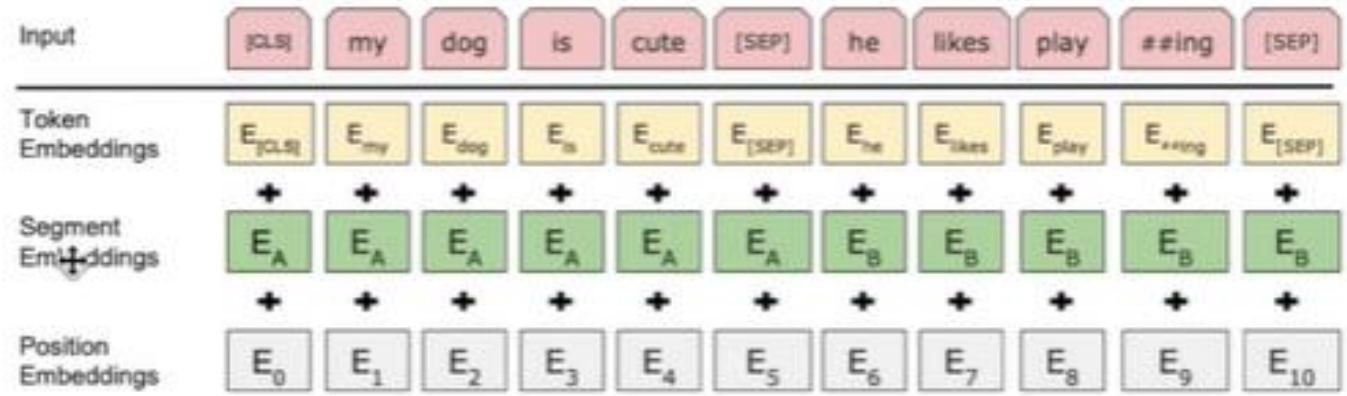

Figure 2: BERT input representation. The input embeddings are the sum of the token embeddings, the segmentation embeddings and the position embeddings.

Transformers

Transformers

- Transformers are a type of neural network architecture that transforms or changes an input sequence into an output sequence.
- They do this by learning context and tracking relationships between sequence components.
- And break the problem into two parts:
  - An encoder (e.g., Bert)
  - A decoder (e.g., GPT)

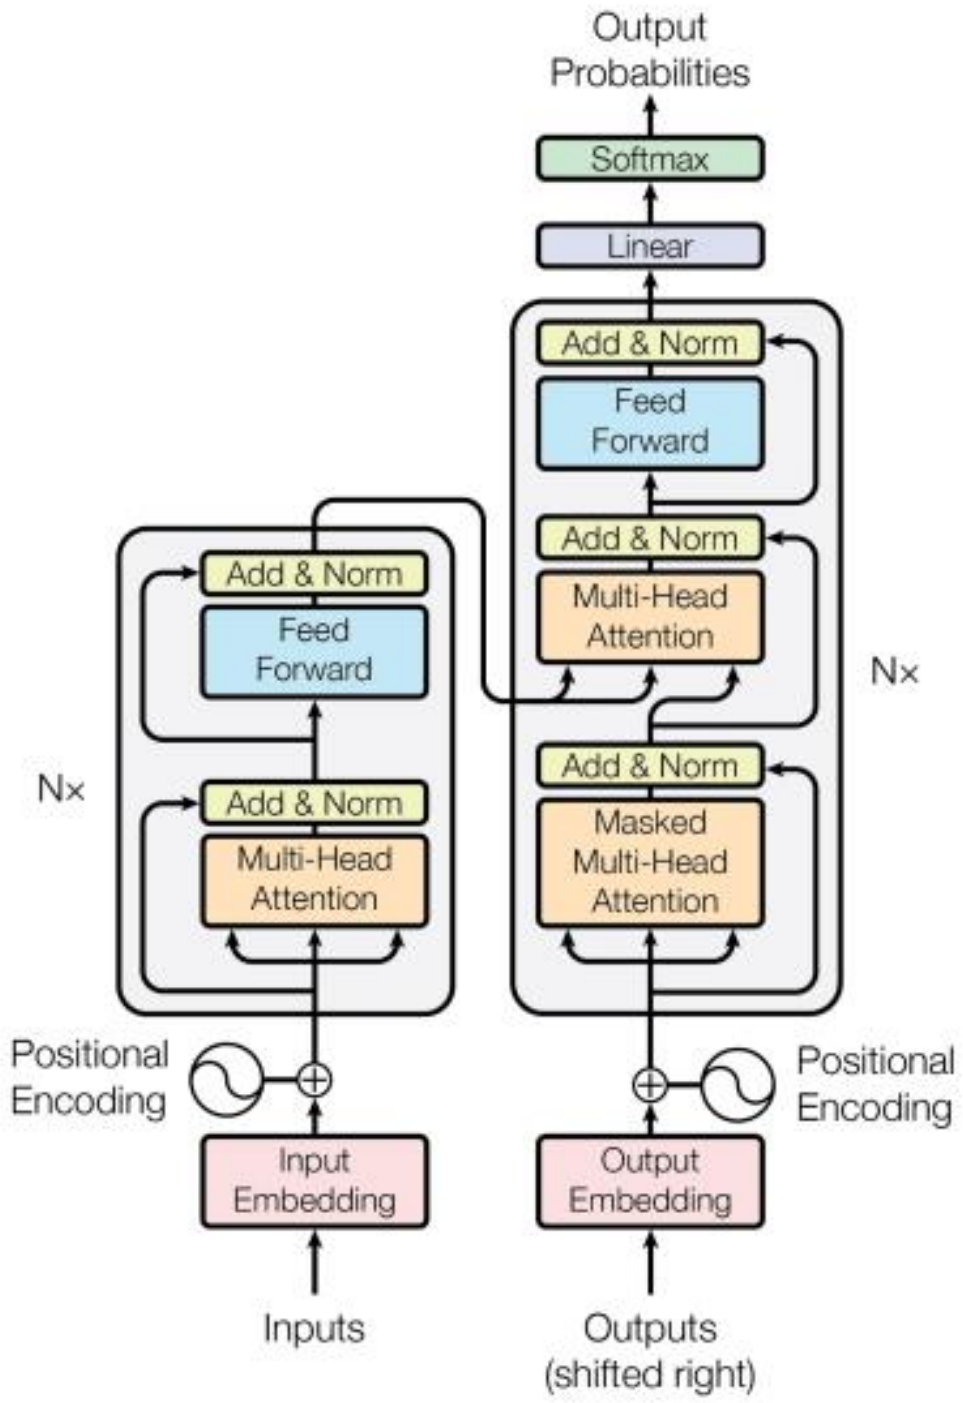

BERT
Oct 2018

Representation

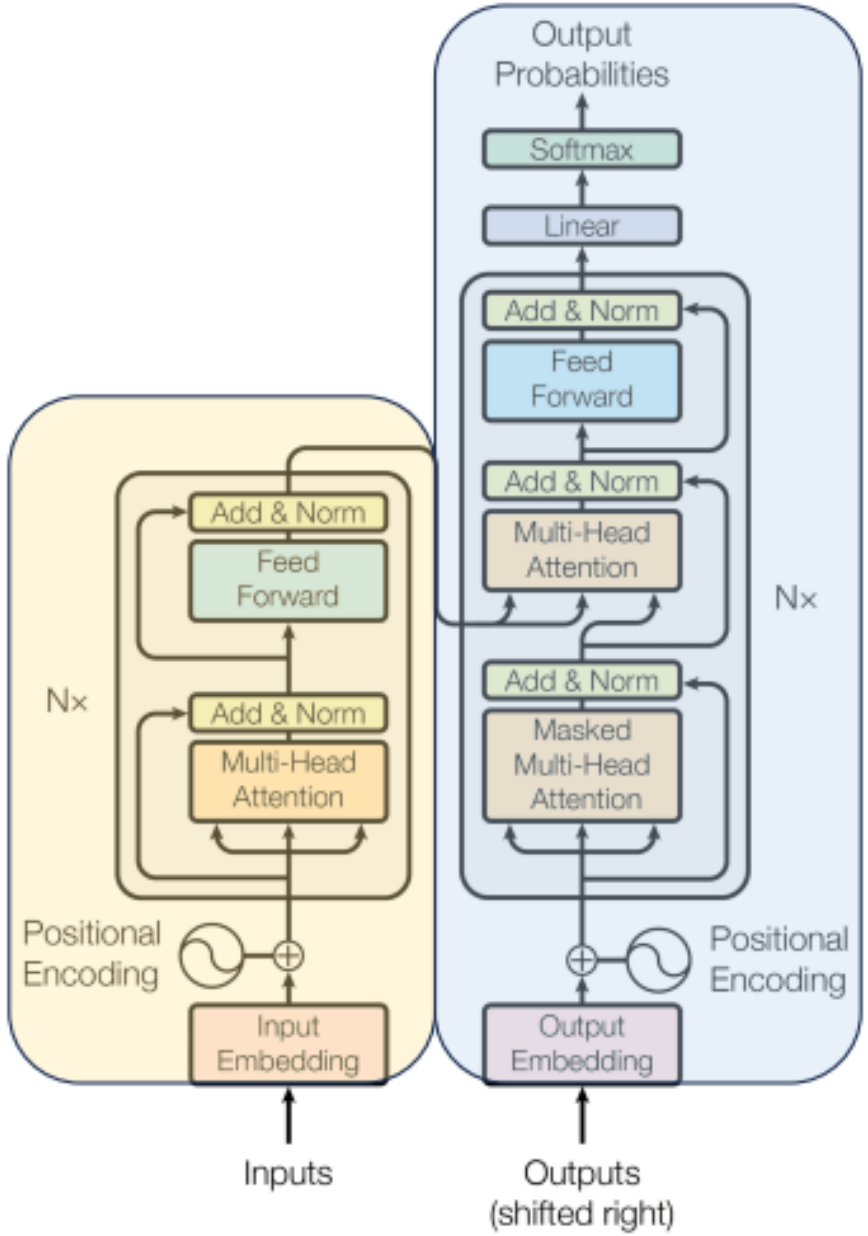

GPT
Jun 2018

Generation

Transformers

- Example: Machine Translation

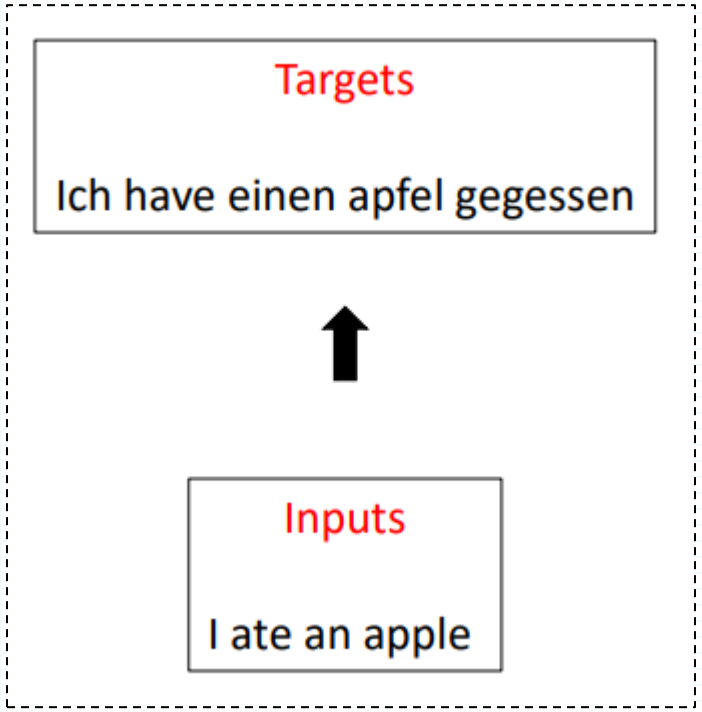

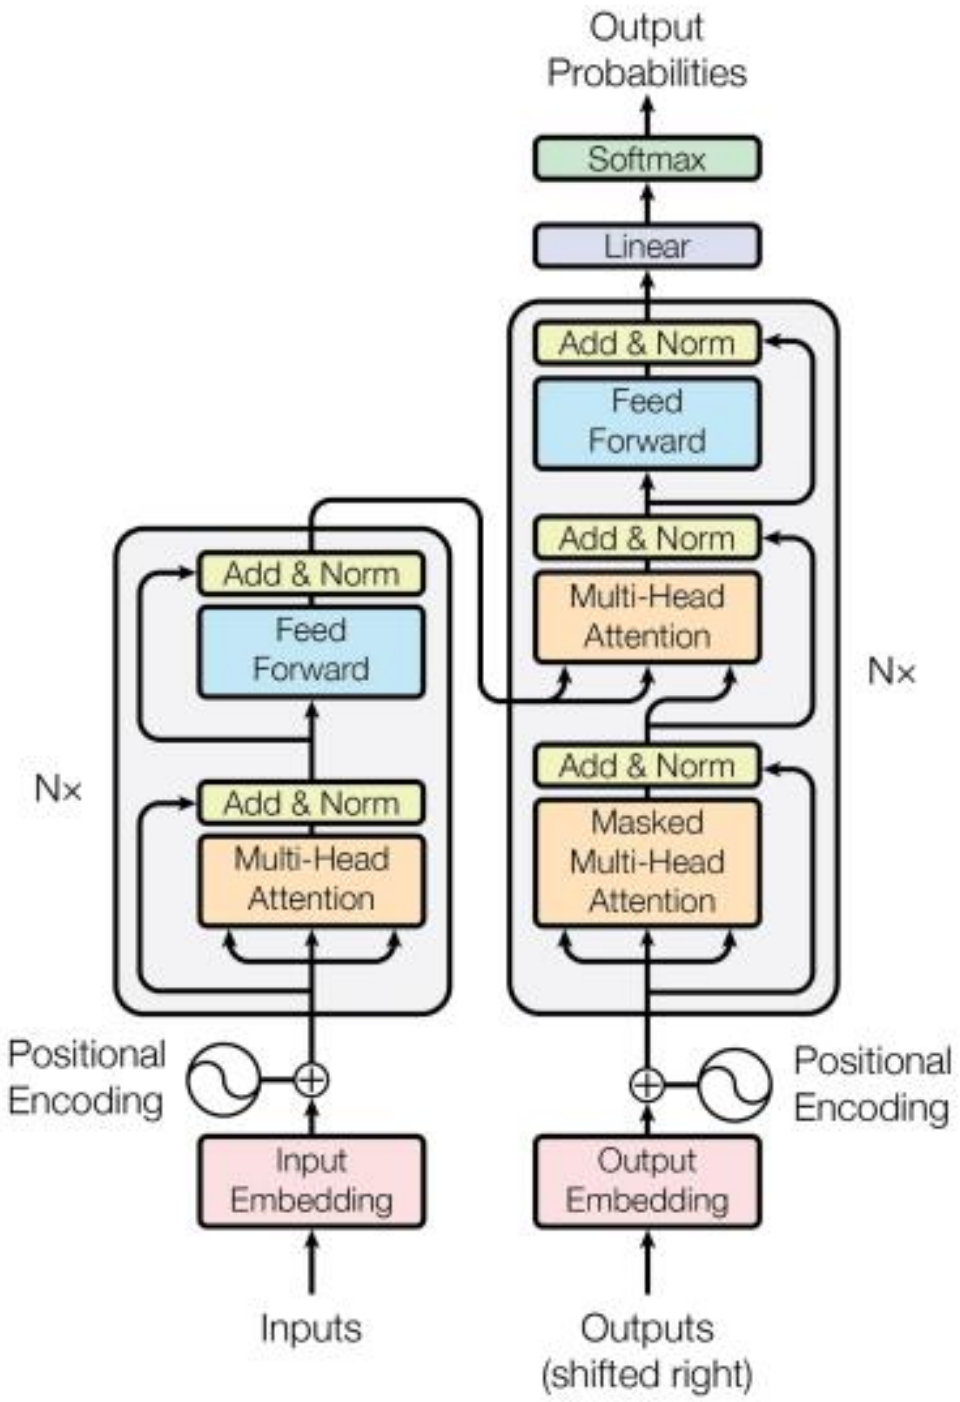

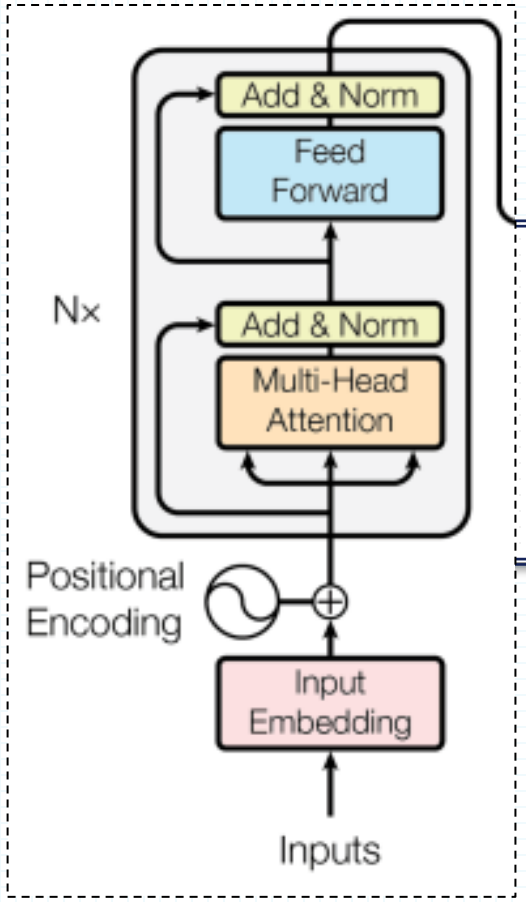

Encoder Part

Transformers

- Processing Input

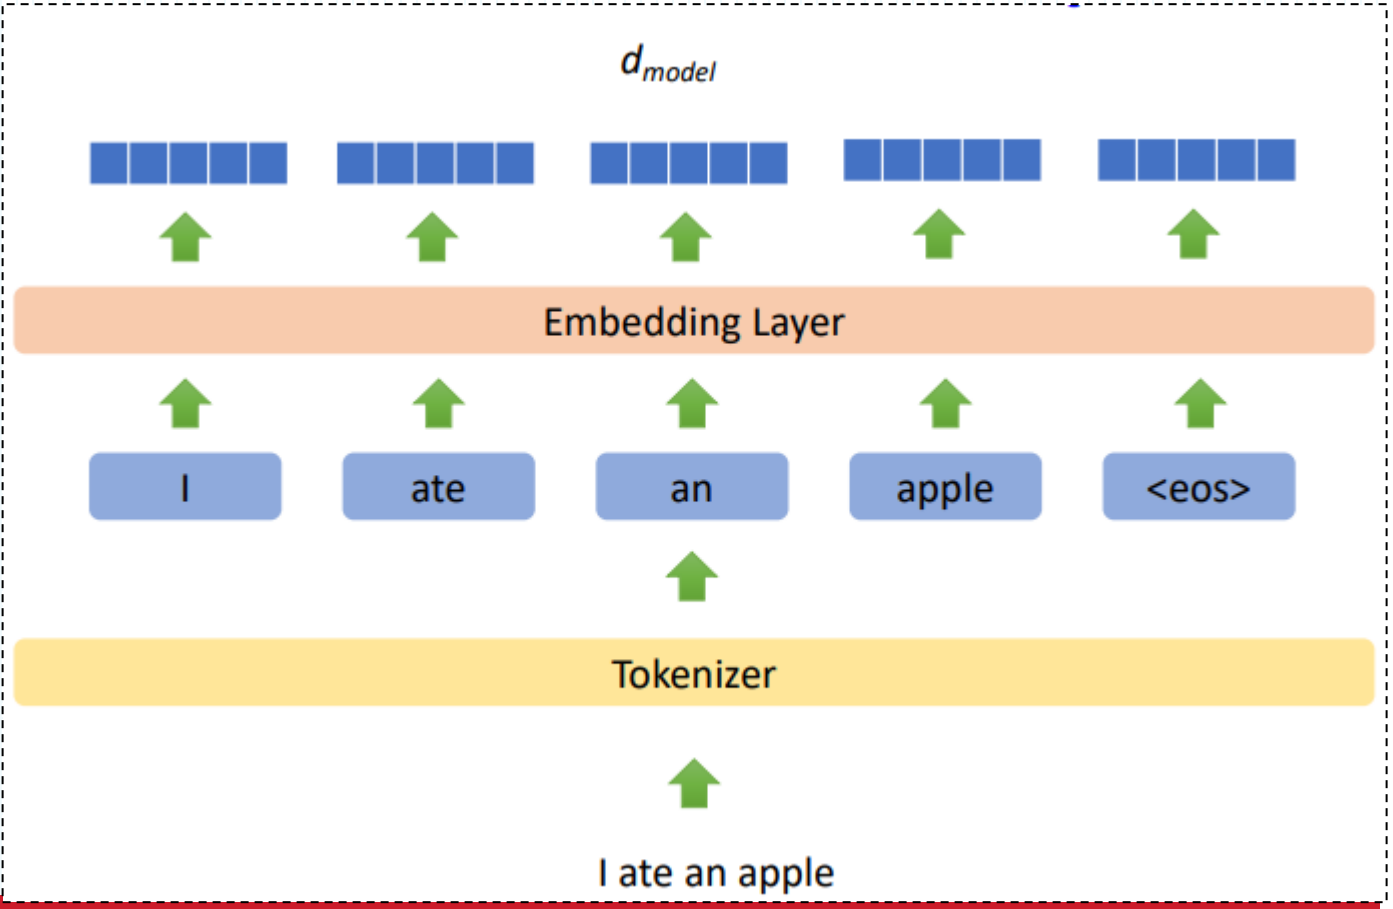

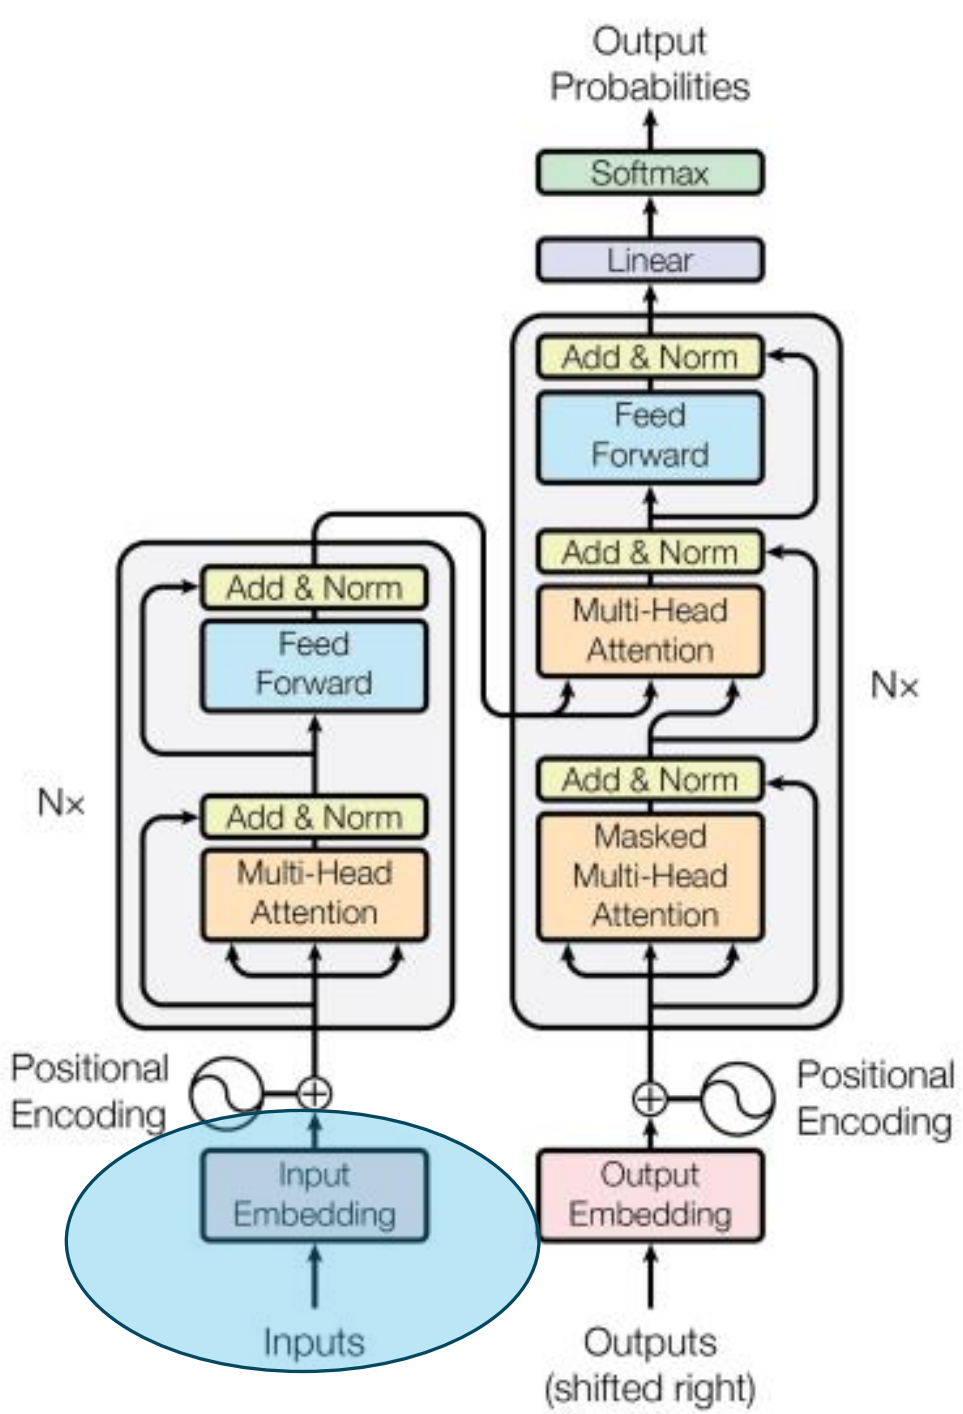

Transformers

- Learn to add context

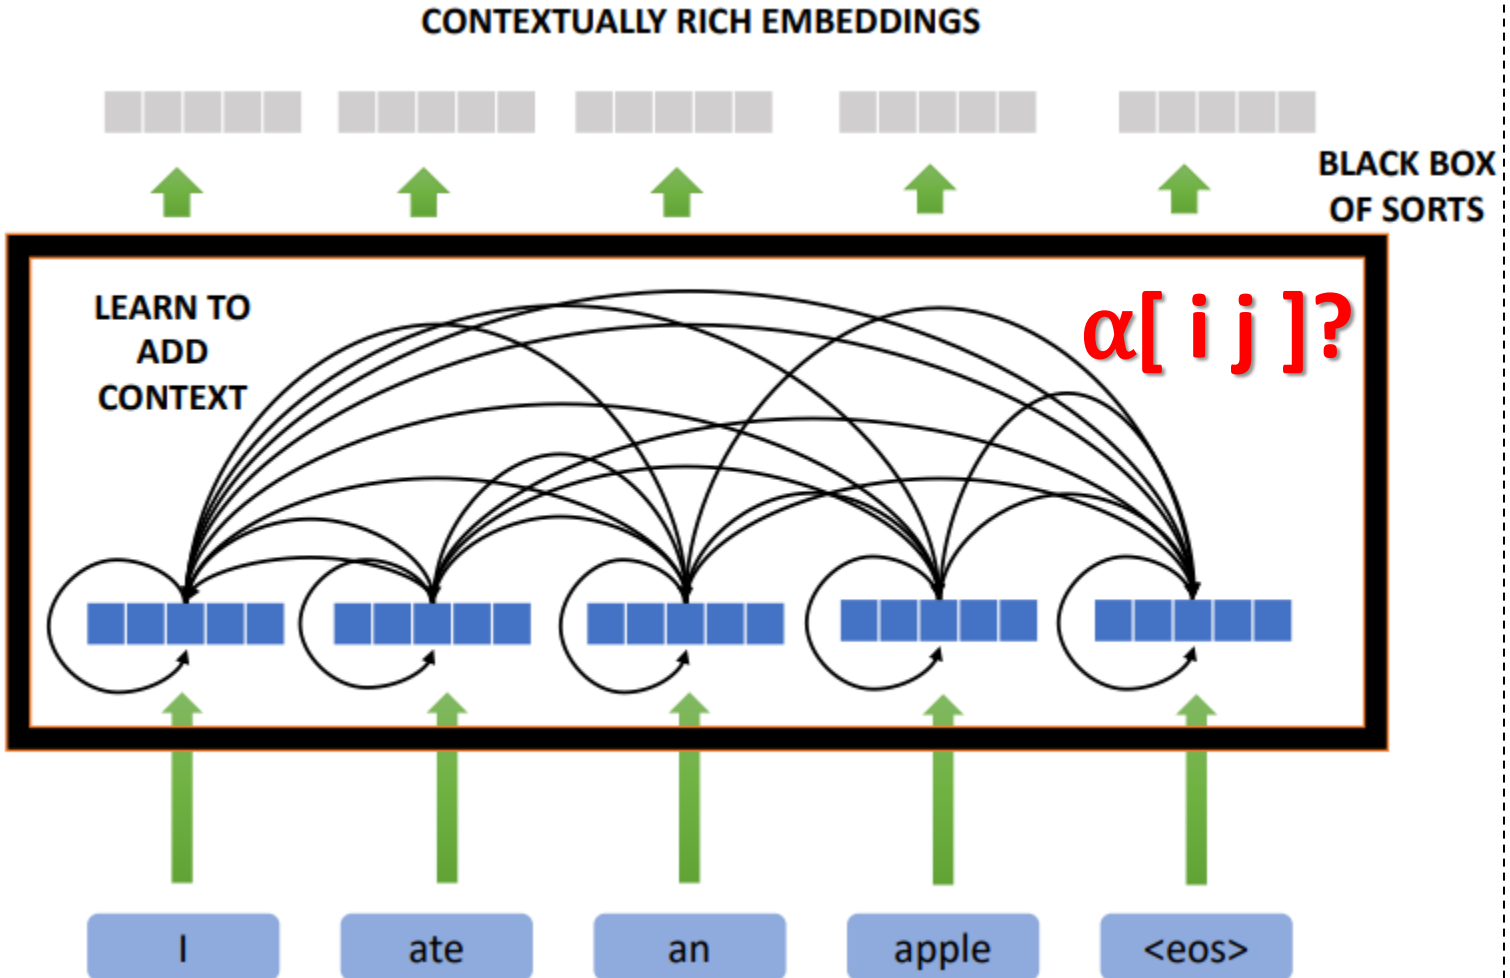

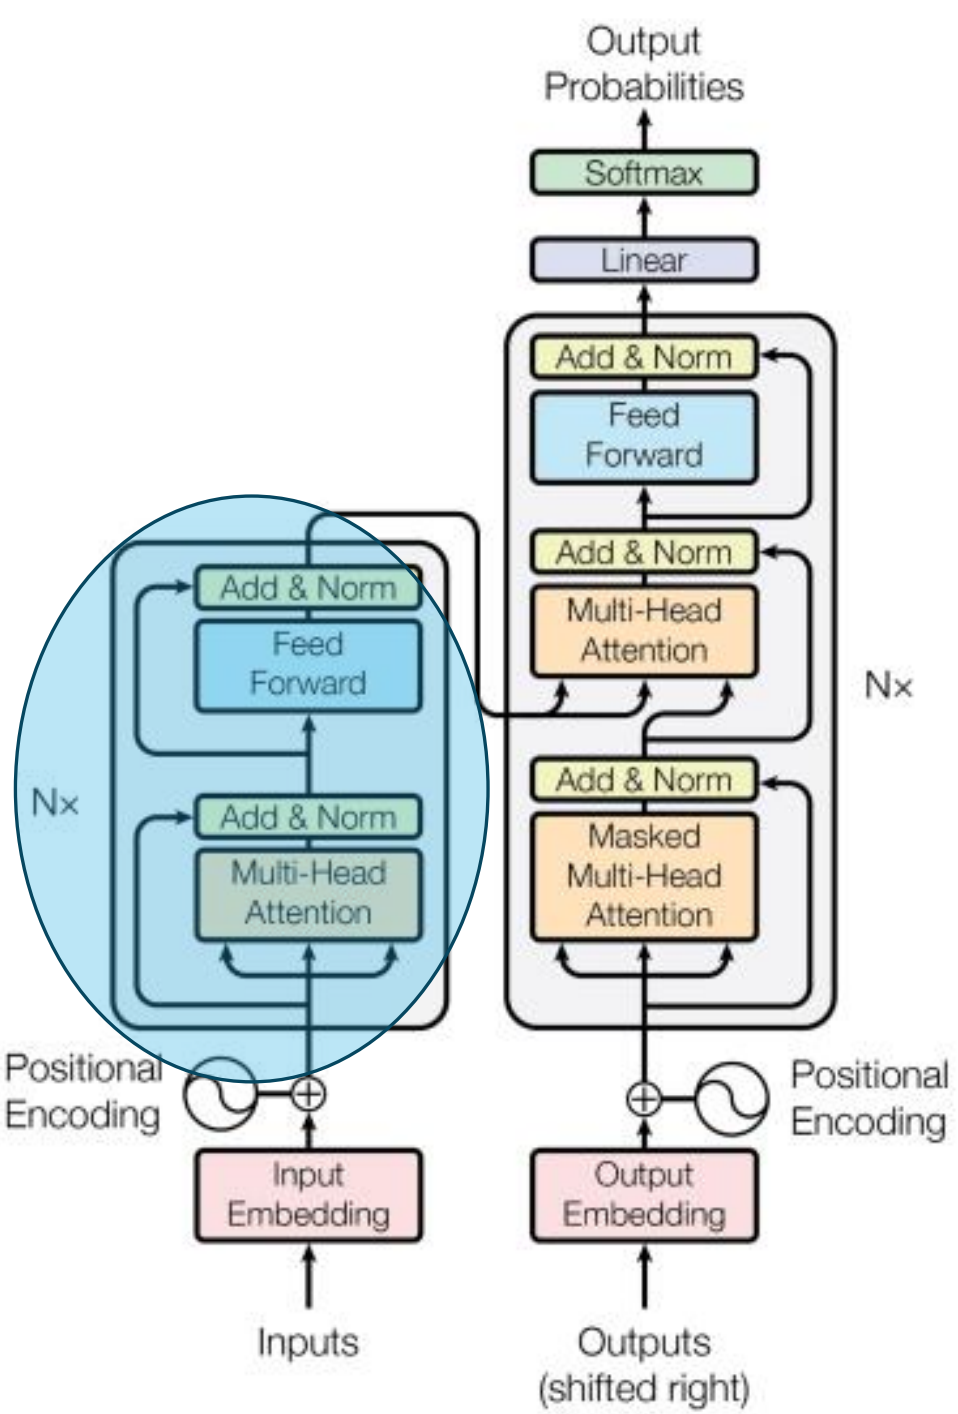

Transformers

- Learn to add context

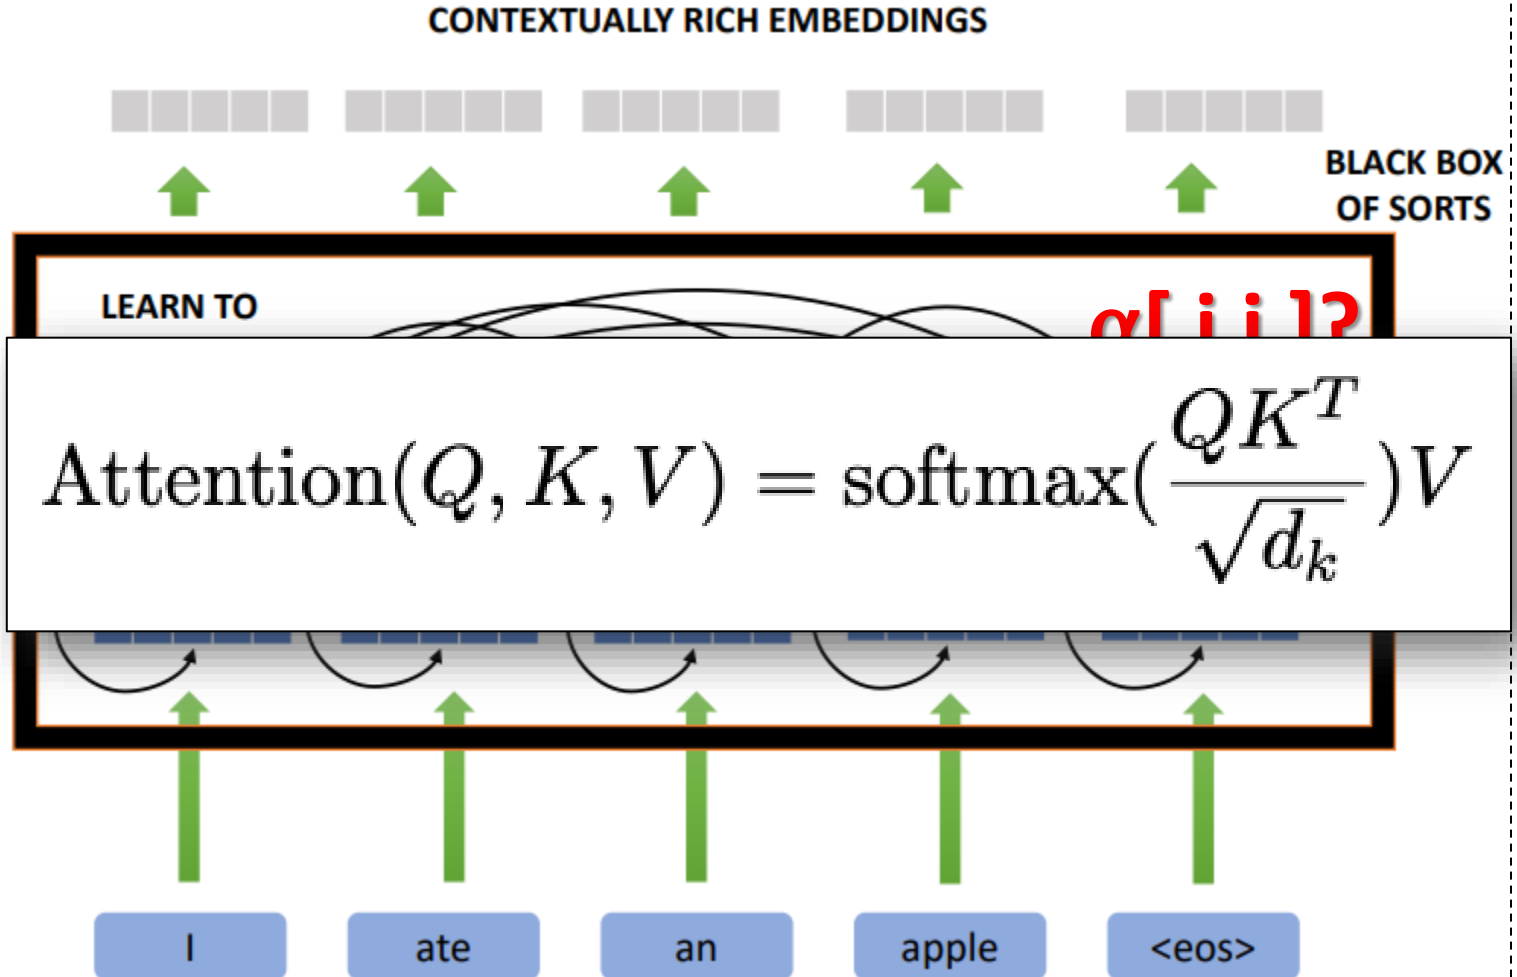

$$\text{Attention}(Q, K, V) = \text{softmax}(\frac{QK^T}{\sqrt{d_k}})V$$

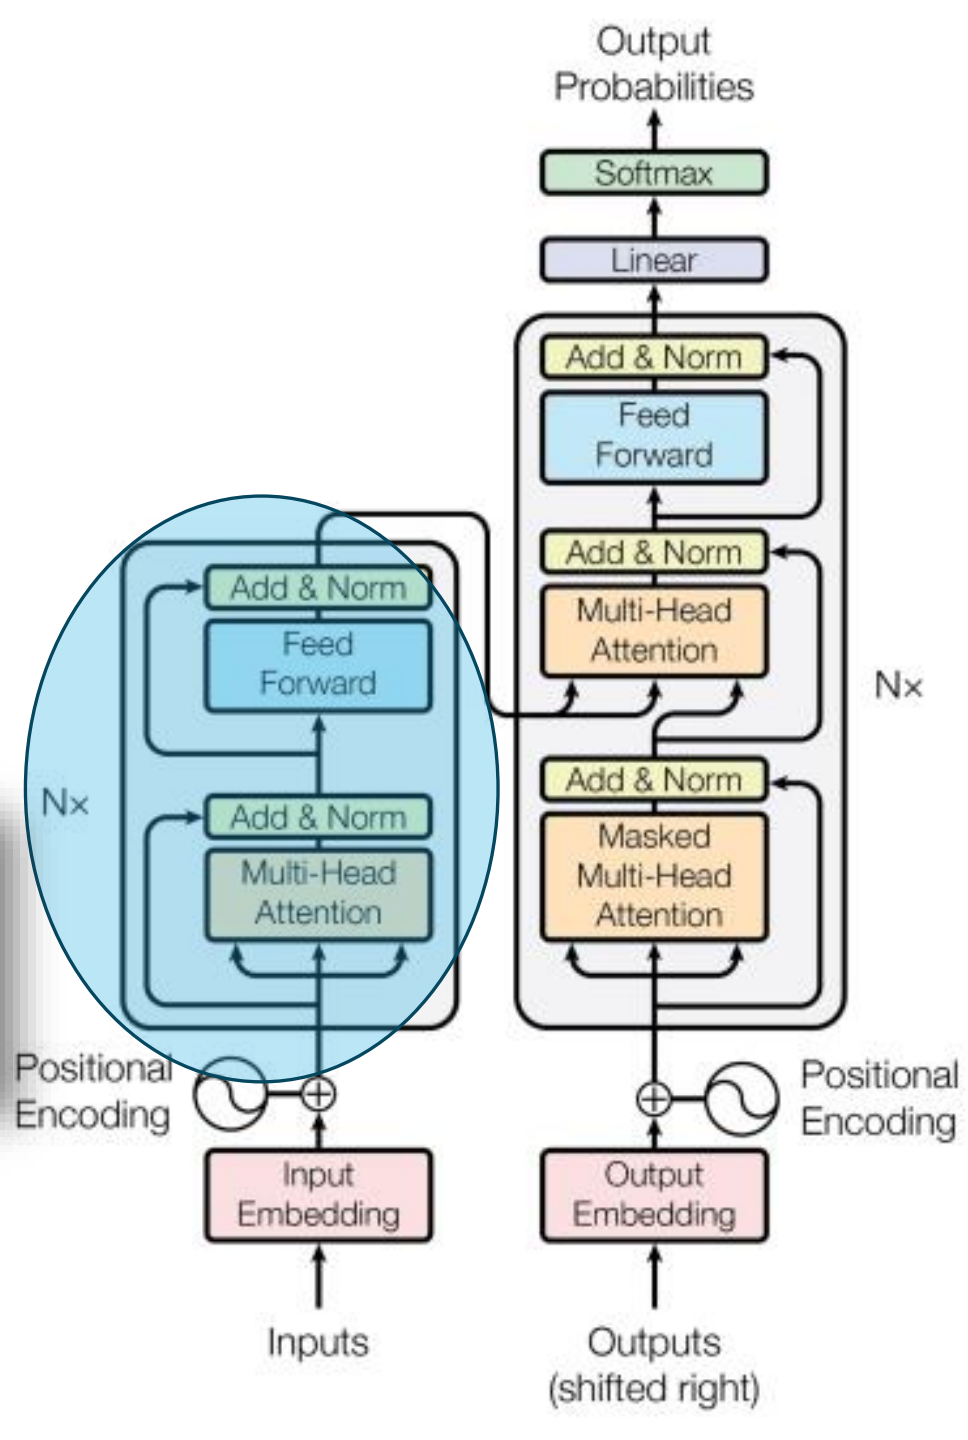

Attention Concepts

- In the attention mechanism, we can draw an analogy to Information Retrieval (IR).

{Query: "Order details of order_104"}

OR

{Query: "Order details of order_106"}

A **query** (an embedding vector representing what we want to find more about — e.g., a specific token or position in a sequence)

A set of **keys** (representing the indexed or stored information — i.e., all other tokens in the context)

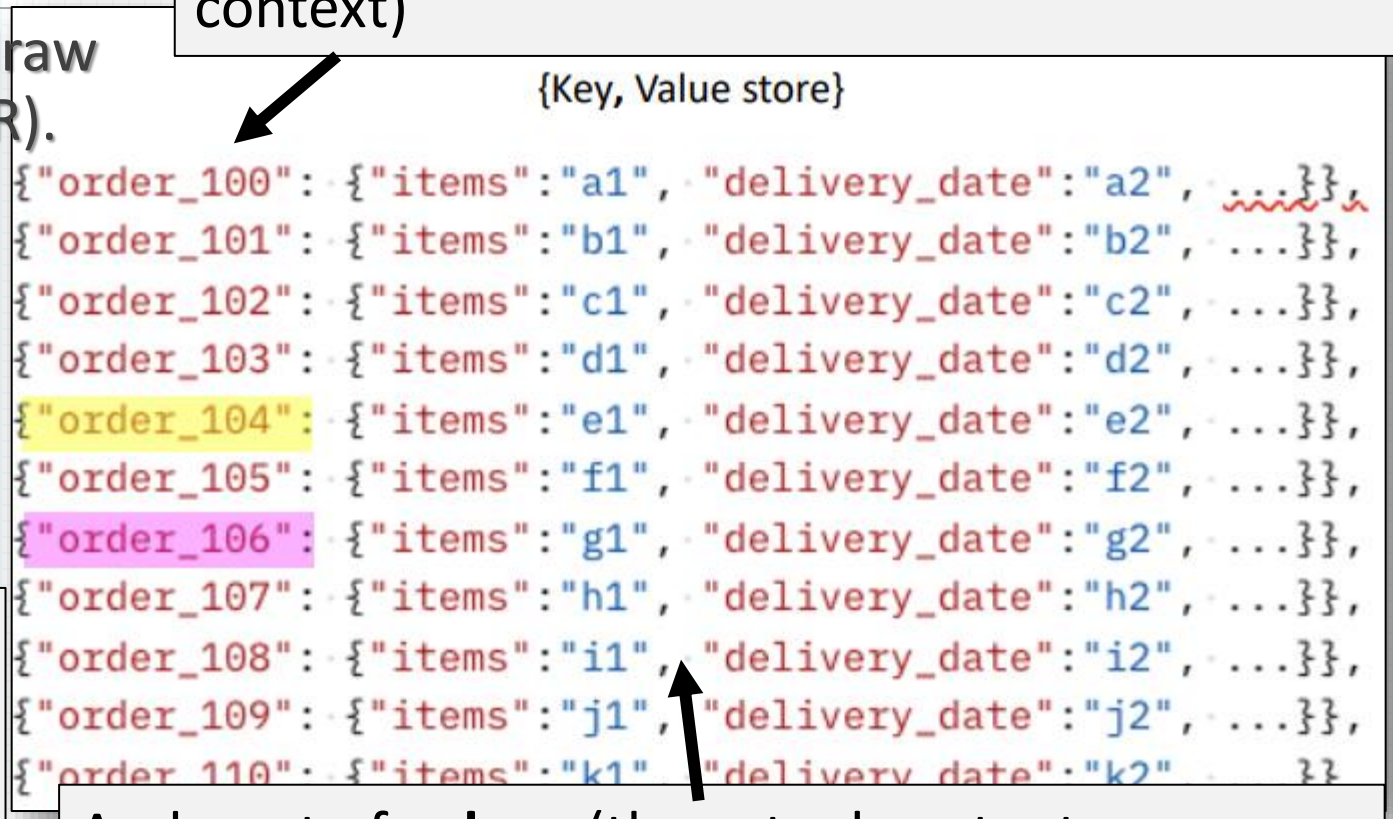

And a set of **values** (the actual content or information associated with each key).

Attention Concepts

- In the attention mechanis
  an analogy to Information

A set of **keys** (representing the ind

Attention allows the model to dynamically retrieve relevant information (values) from the context, based on how similar other tokens (keys) are to the current focus (query).

The attention process works as follows:

1. We compute the similarity between the query and each key (usually using a dot product).
2. This gives us a set of attention scores, indicating how much focus we should give to each position in the sequence.
3. We then apply a softmax to normalize the scores into a probability distribution.
4. Finally, we compute a weighted sum of the values based on these scores to produce the attention output.

Attention

- More formal!

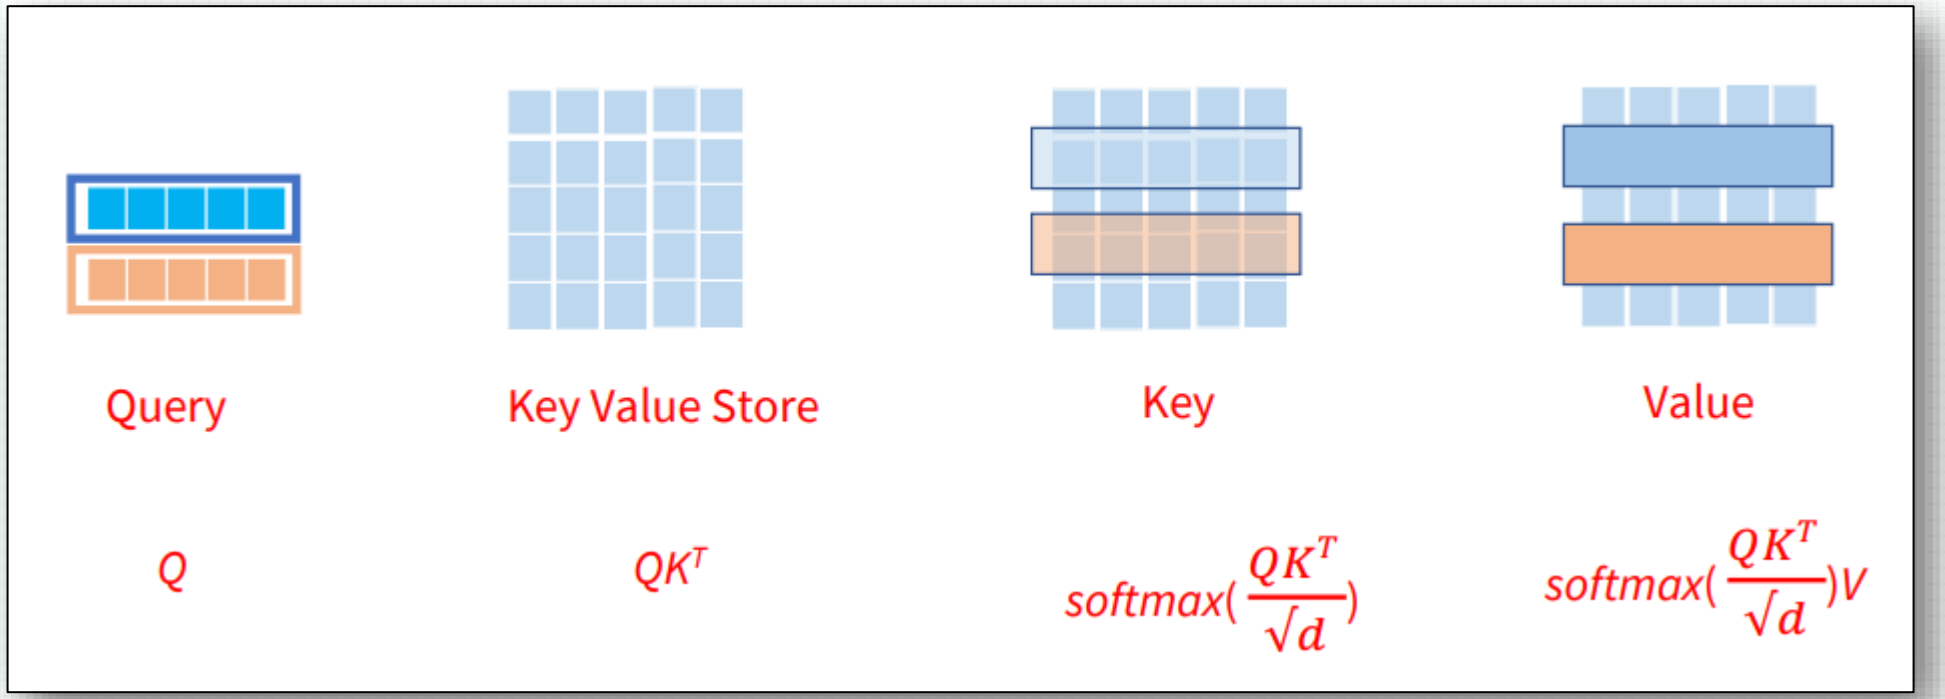

Attention

To calculate attention weights for input I1, you would use query q1, and all keys

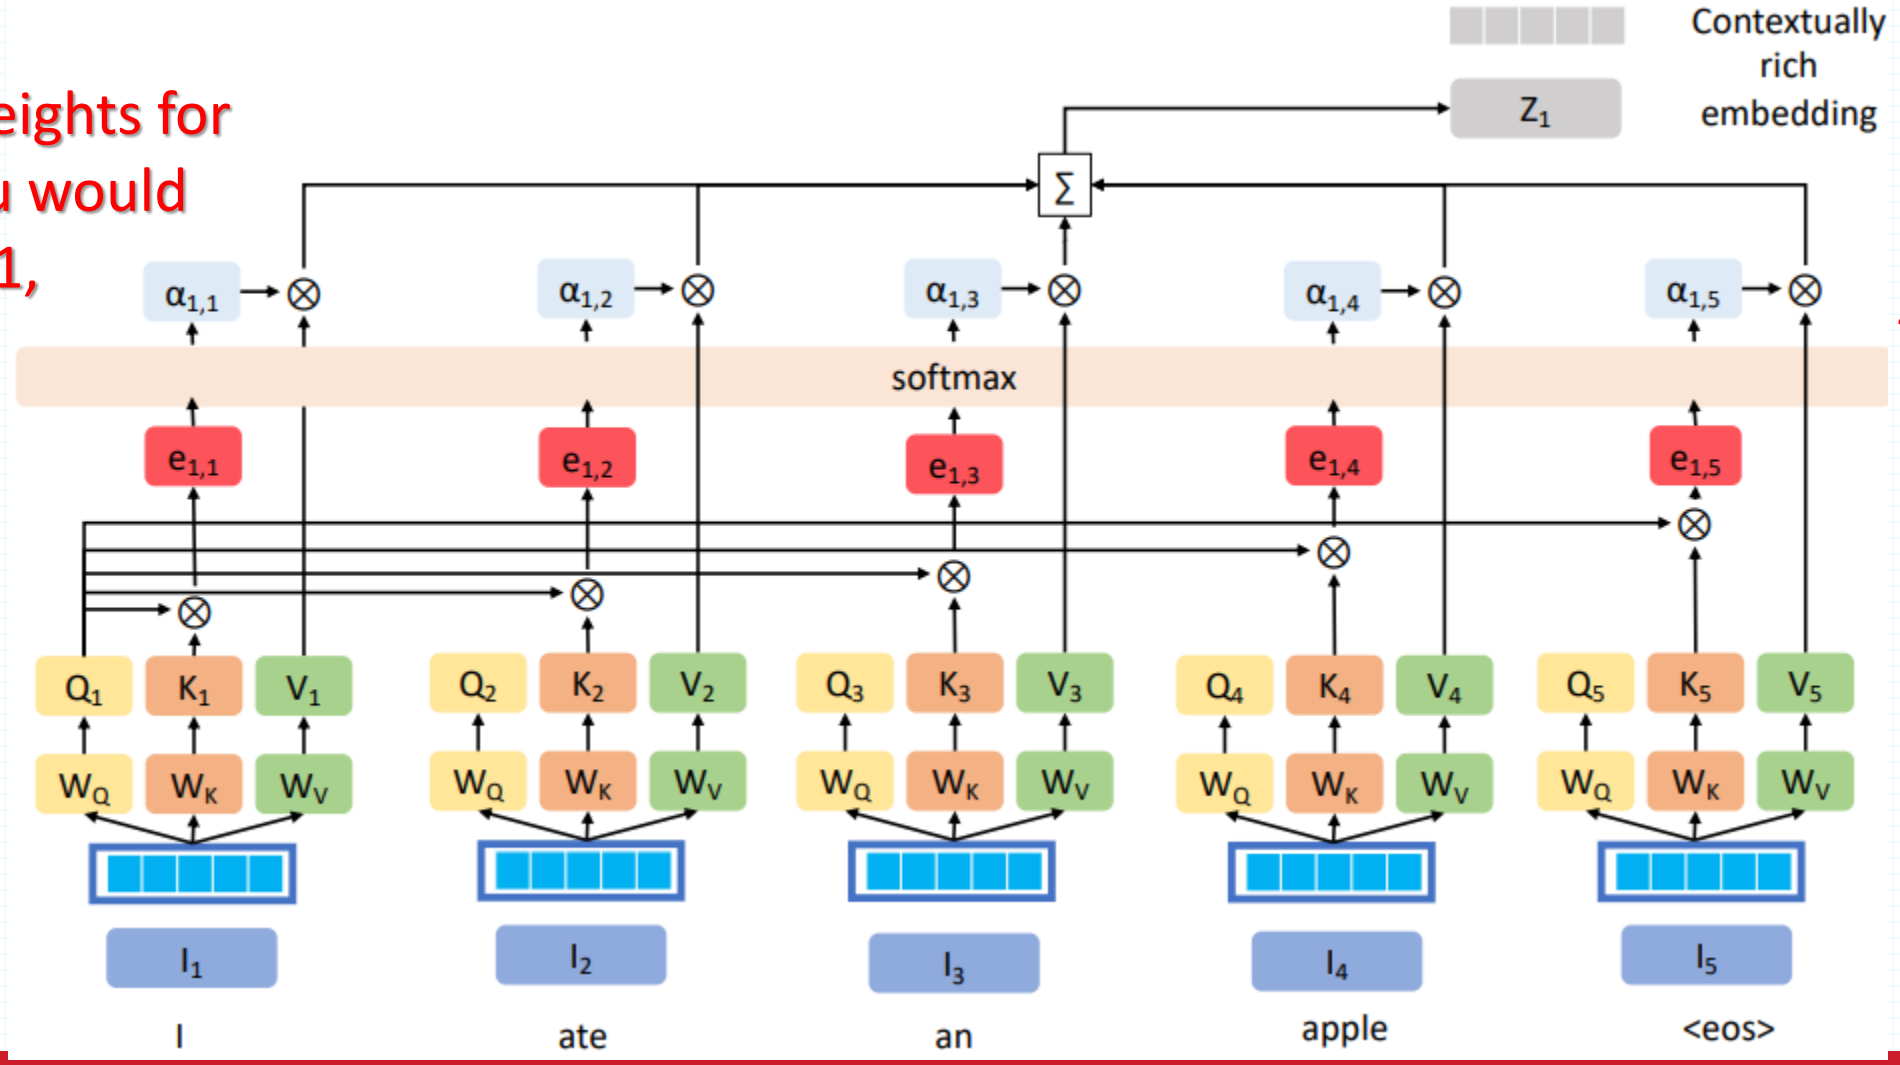

Of I1

We do the same for I2, I3, etc

Self Attention

Query Inputs = Key Inputs = Value Inputs

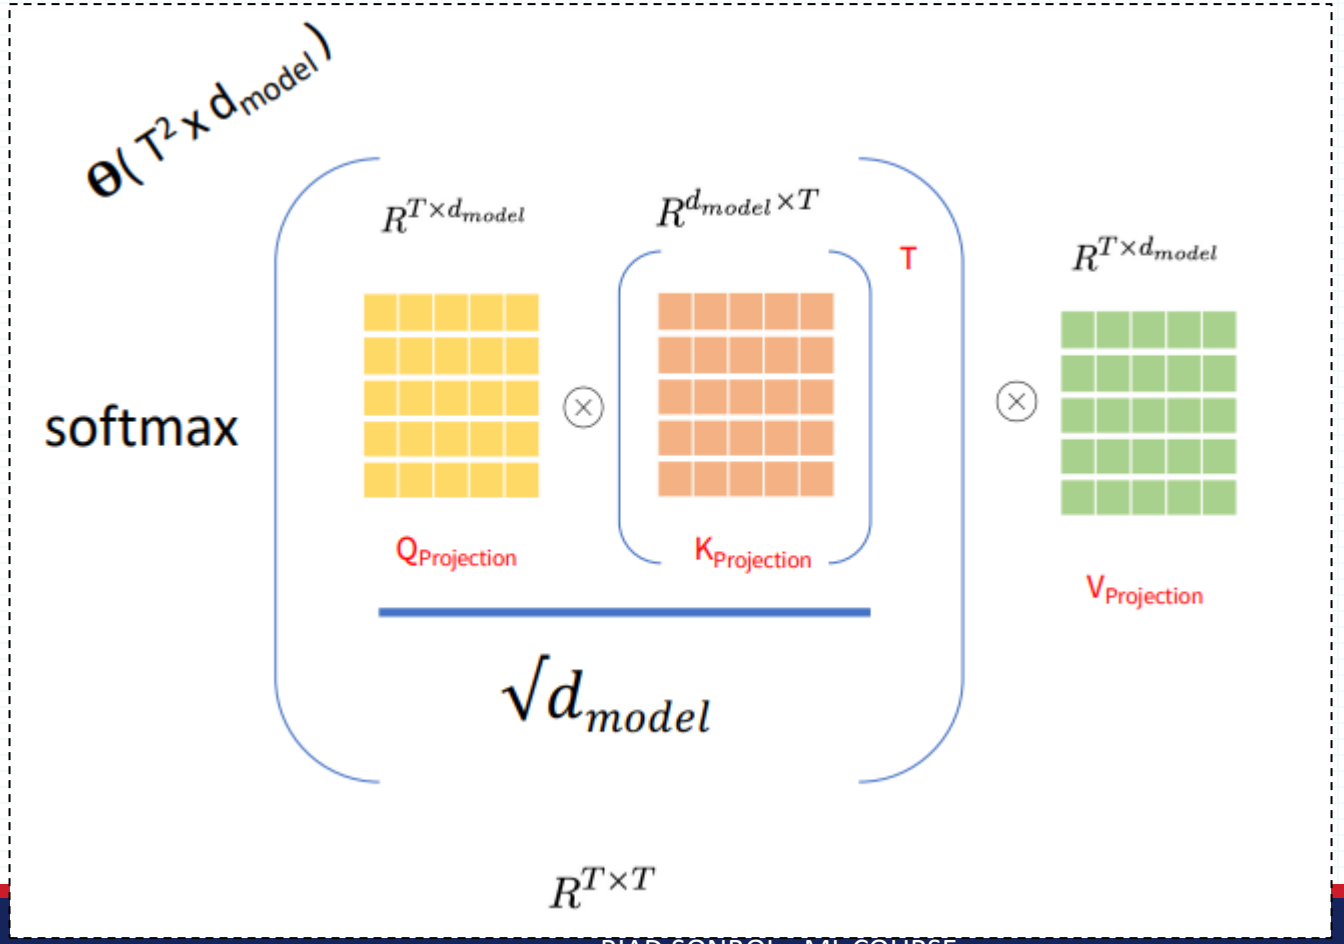

Positional Encoding

- Injects sequence order information into transformer models.
- Added to token embeddings to help the model understand word positions.
- Needed because transformers process tokens in parallel and lack inherent order sensitivity.
- Can be fixed (e.g., sinusoidal – sine, cosine) or learned during training.

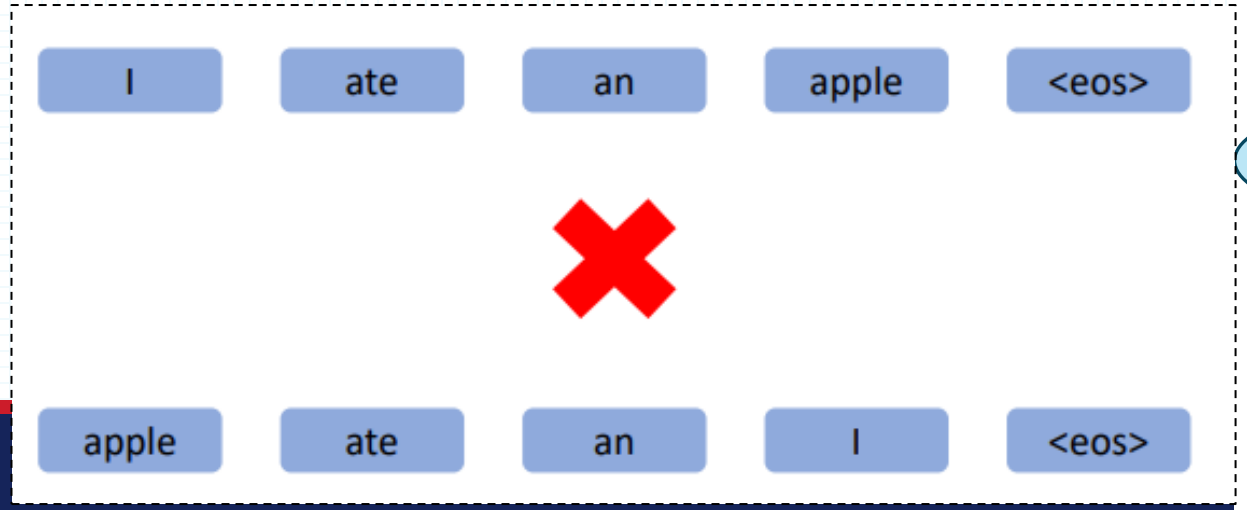

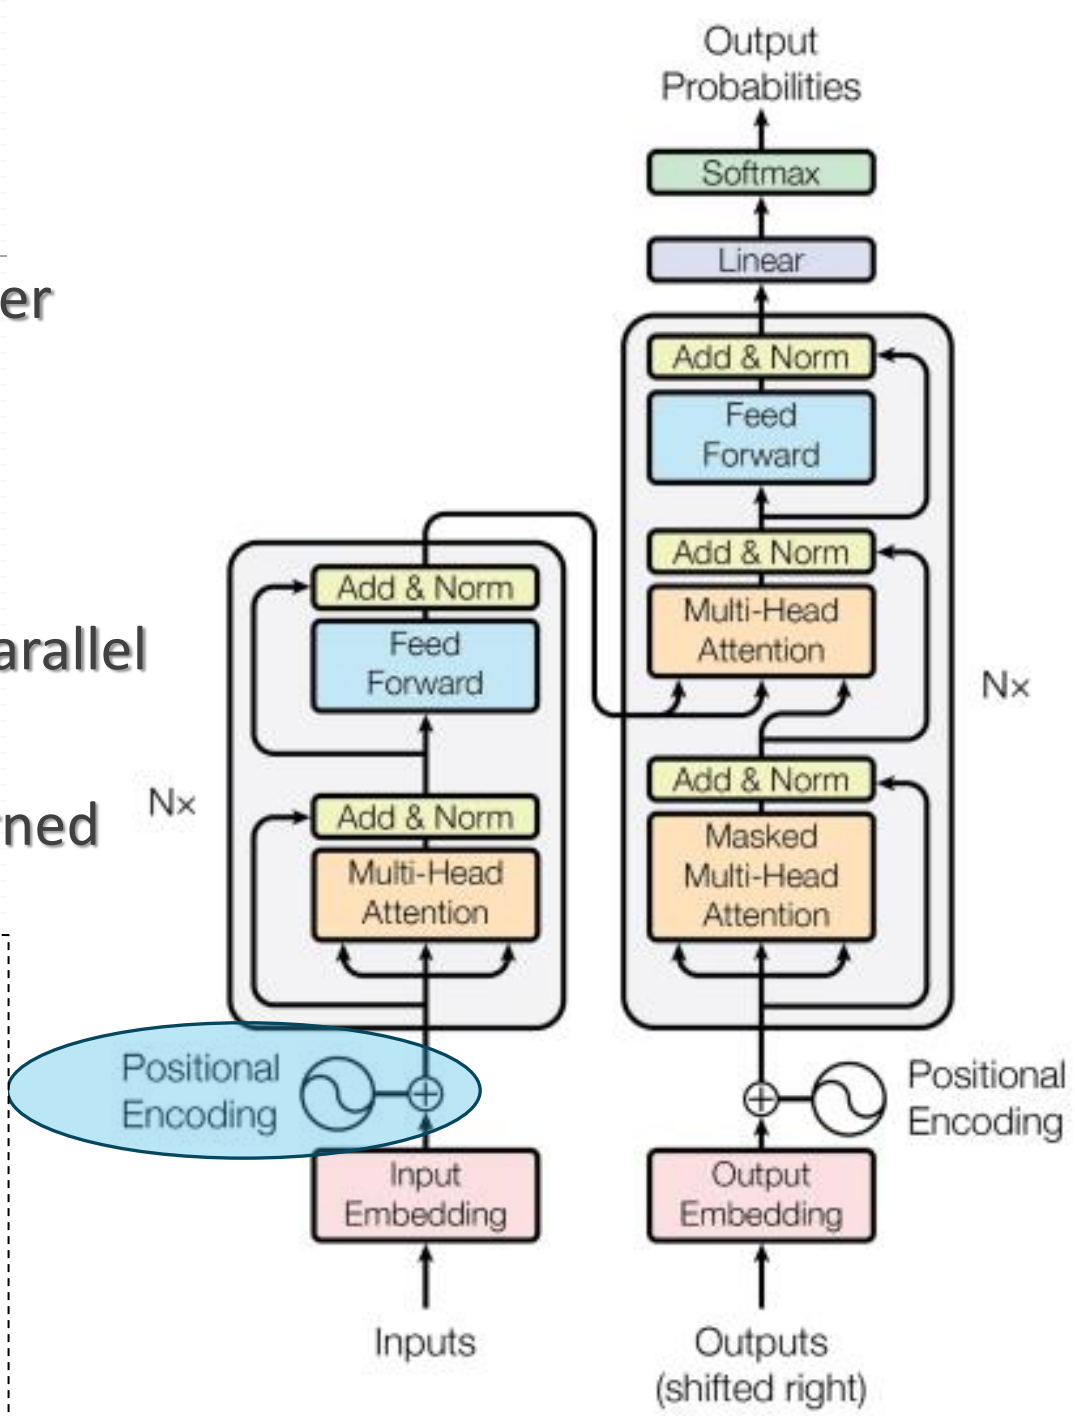

Positional Encoding

pos -> idx of the token in input sentence

i -> i^{th} dimension out of d

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

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

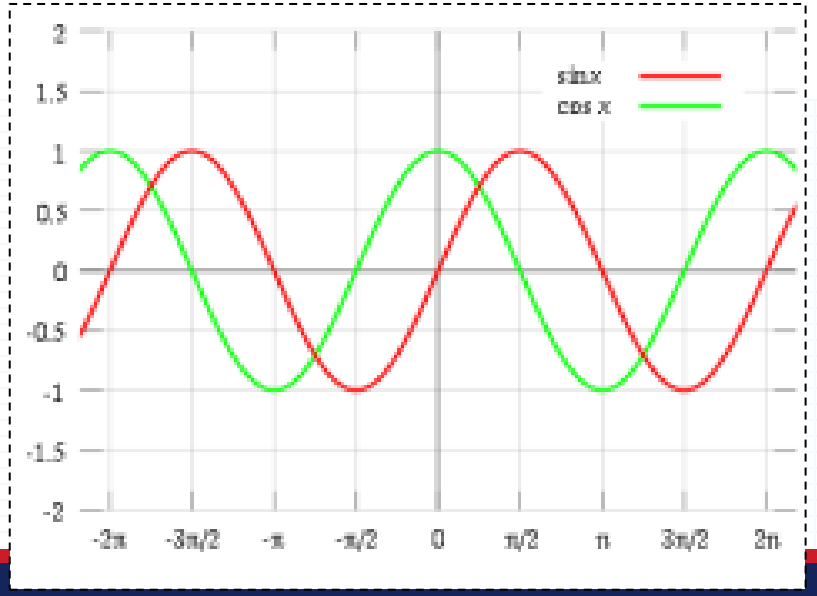

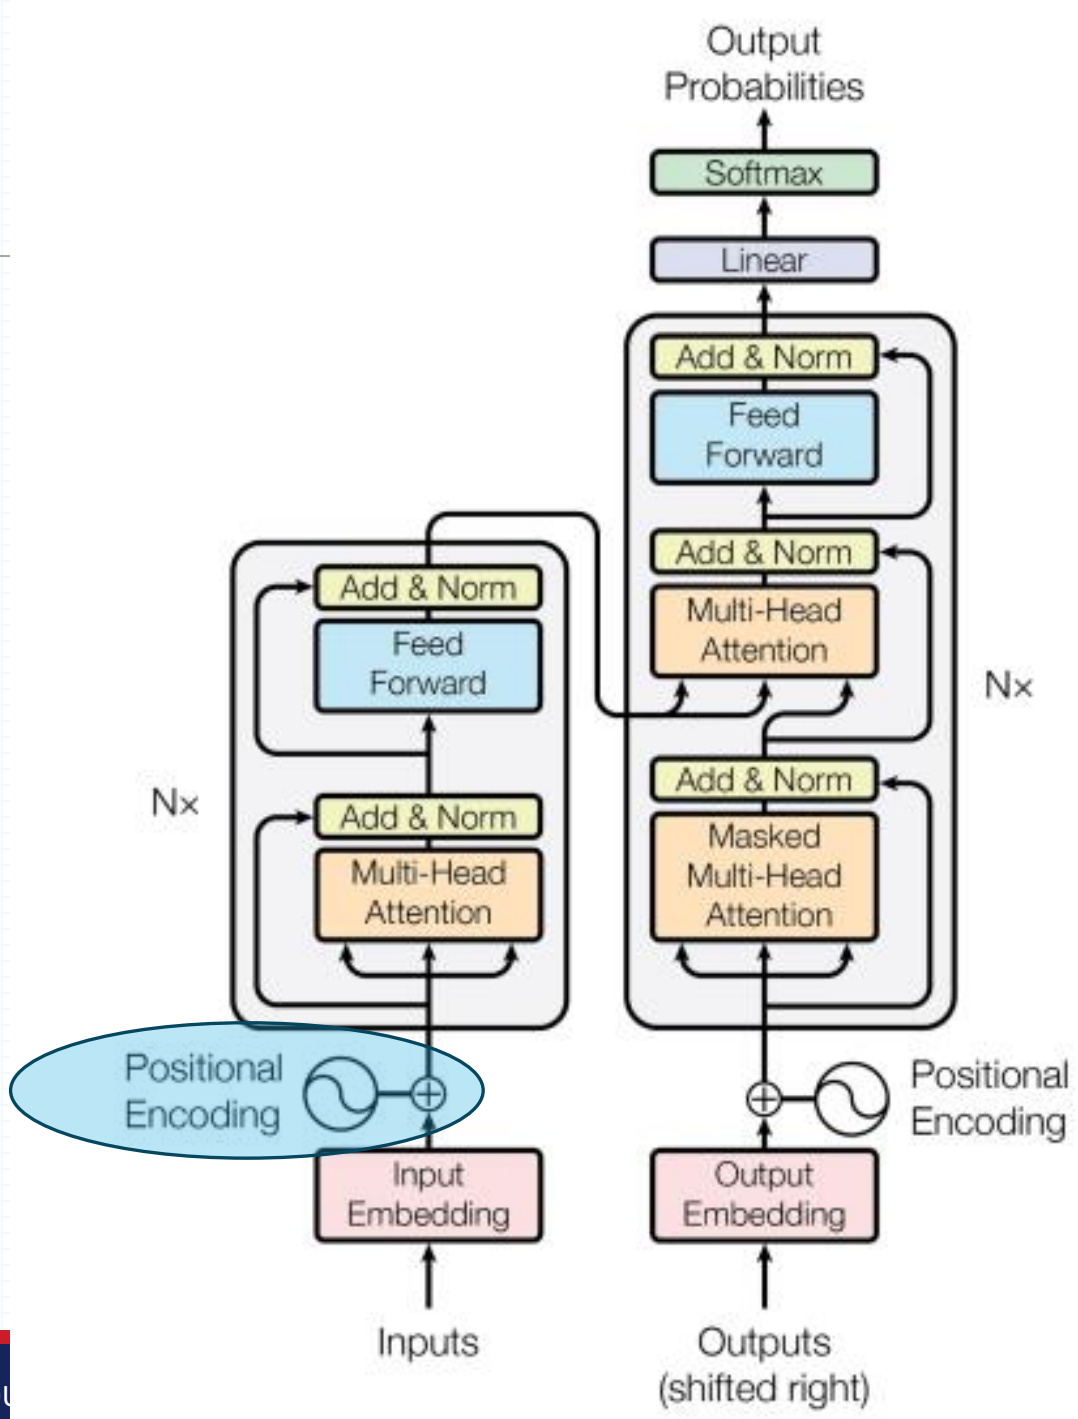

Position Embedding

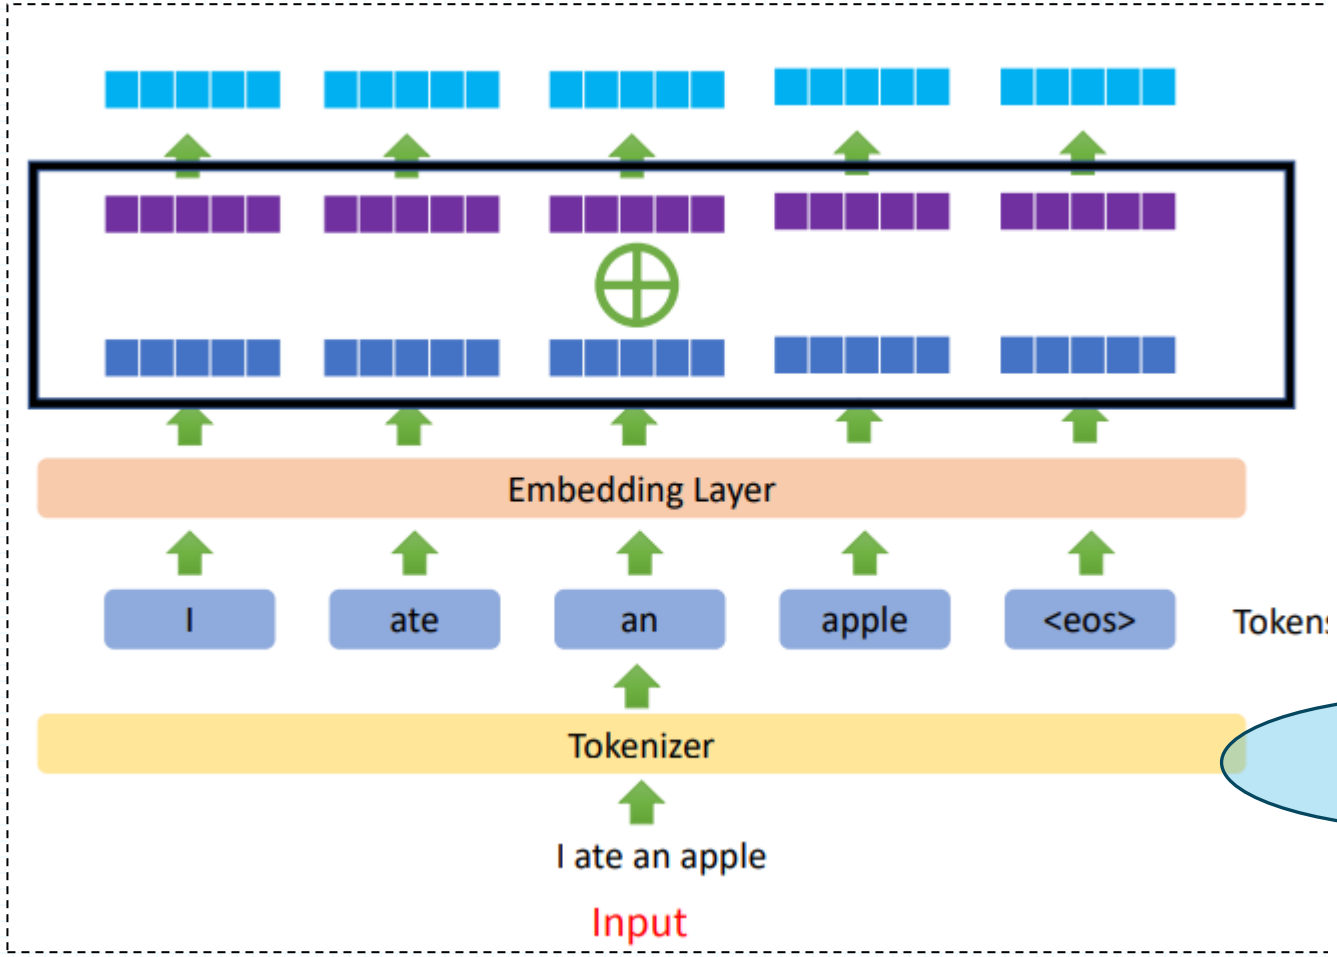

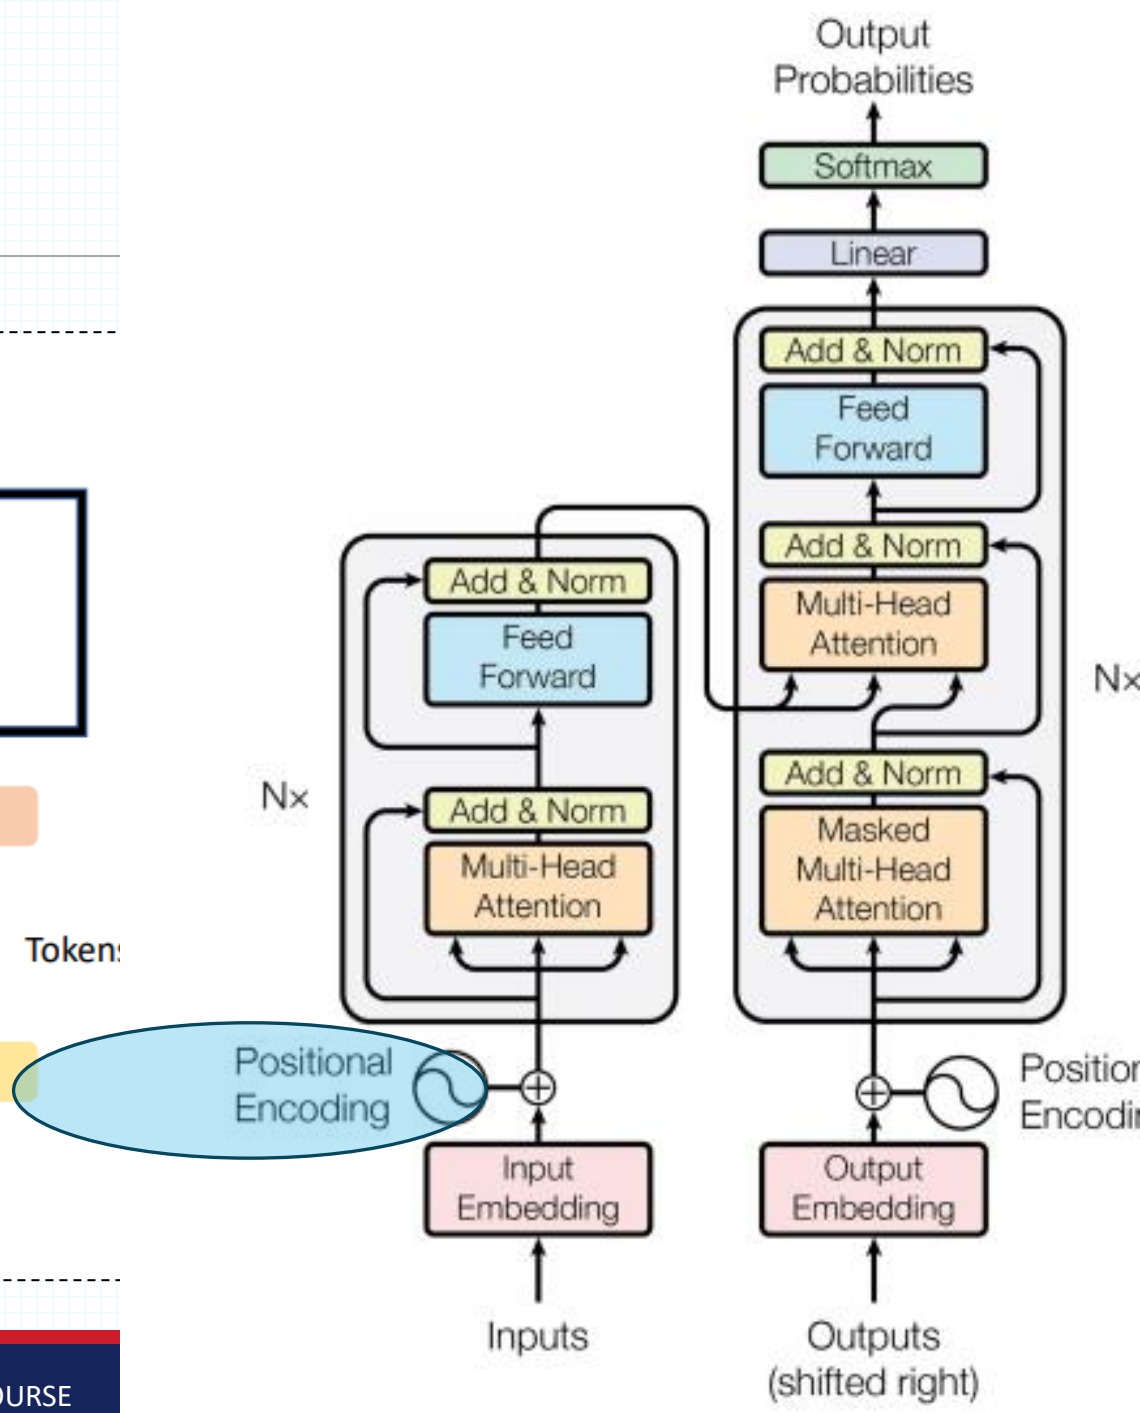

Decoder Part

Transformers

- Transformers are a type of neural network architecture that transforms or changes an input sequence into an output sequence.
- They do this by learning context and tracking relationships between sequence components.
- And break the problem into two parts:
  - An encoder (e.g., Bert)
  - A decoder (e.g., GPT)

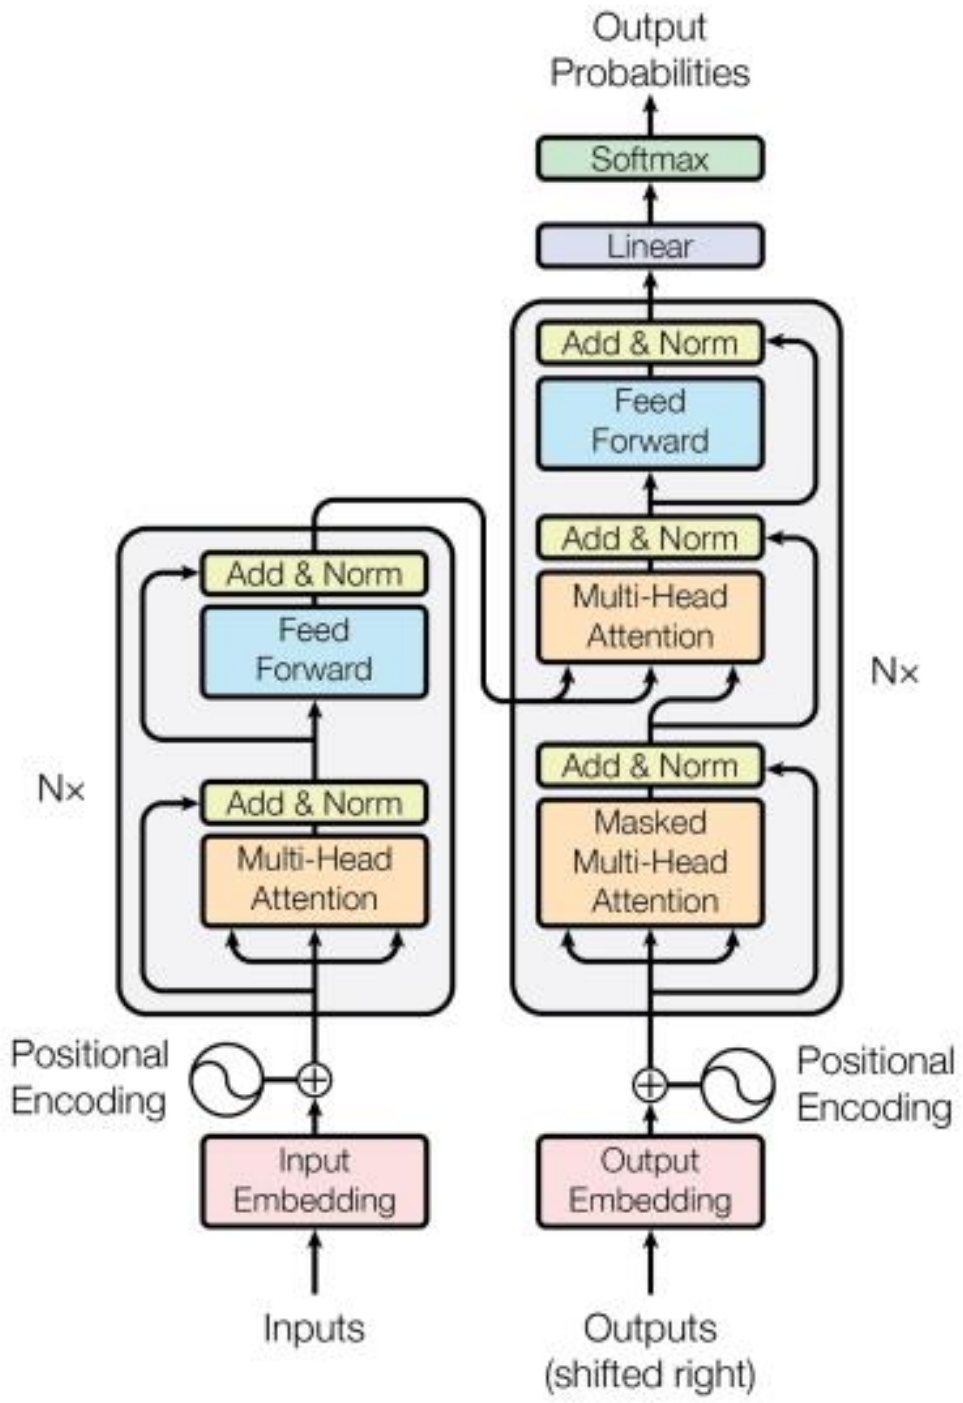

Output Embedding

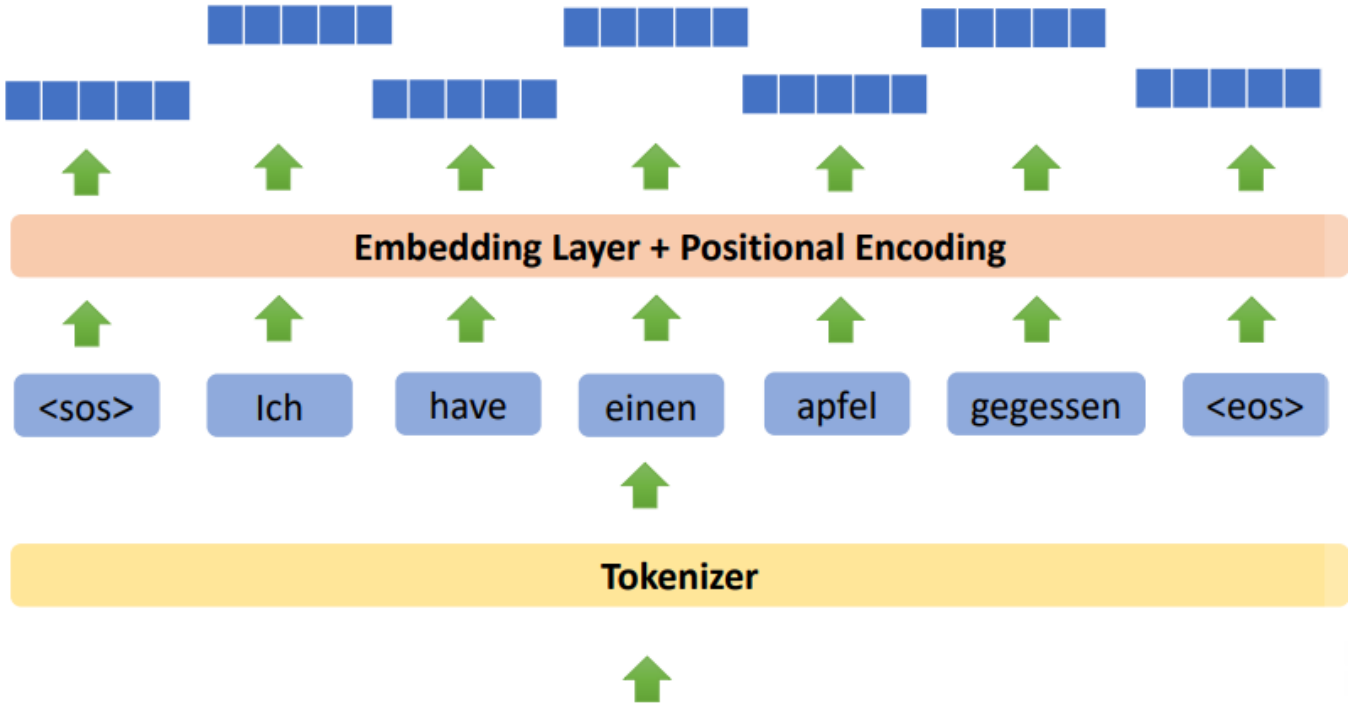

Ich have einen apfel gegessen

Generate Target Emebeddings

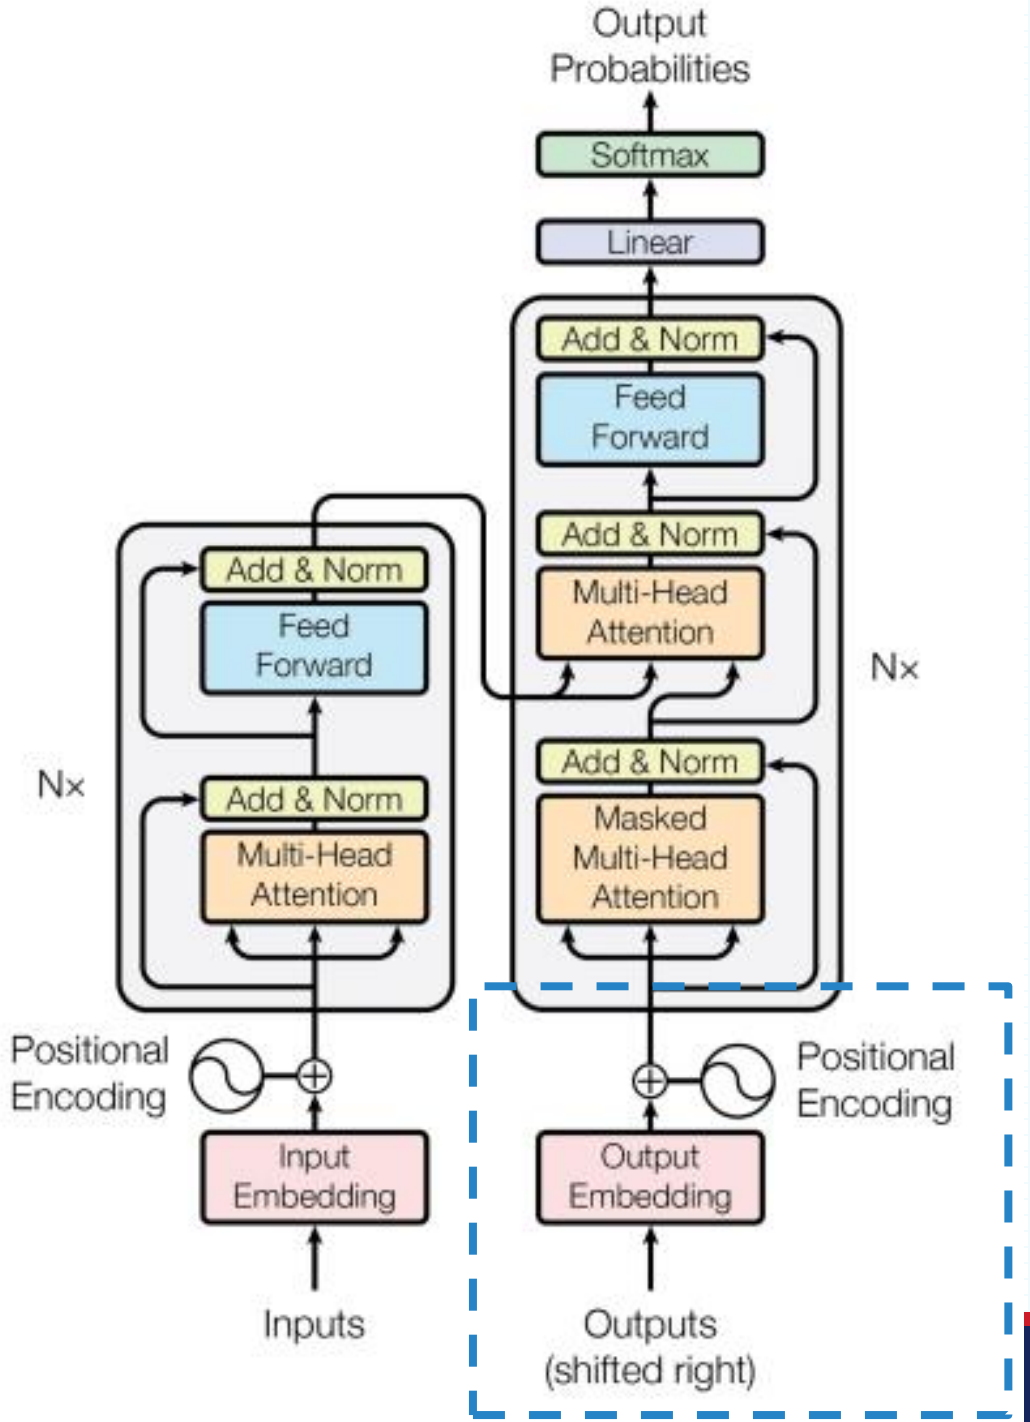

Masked Attention

1	<sos>	Ich	have	einen	apfel	gegessen	<eos>
2	<sos>	Ich	have	einen	apfel	gegessen	<eos>
3	<sos>	Ich	have	einen	apfel	gegessen	<eos>
4	<sos>	Ich	have	einen	apfel	gegessen	<eos>
5	<sos>	Ich	have	einen	apfel	gegessen	<eos>
6	<sos>	Ich	have	einen	apfel	gegessen	<eos>
7	<sos>	Ich	have	einen	apfel	gegessen	<eos>

Mask the available attention values ?

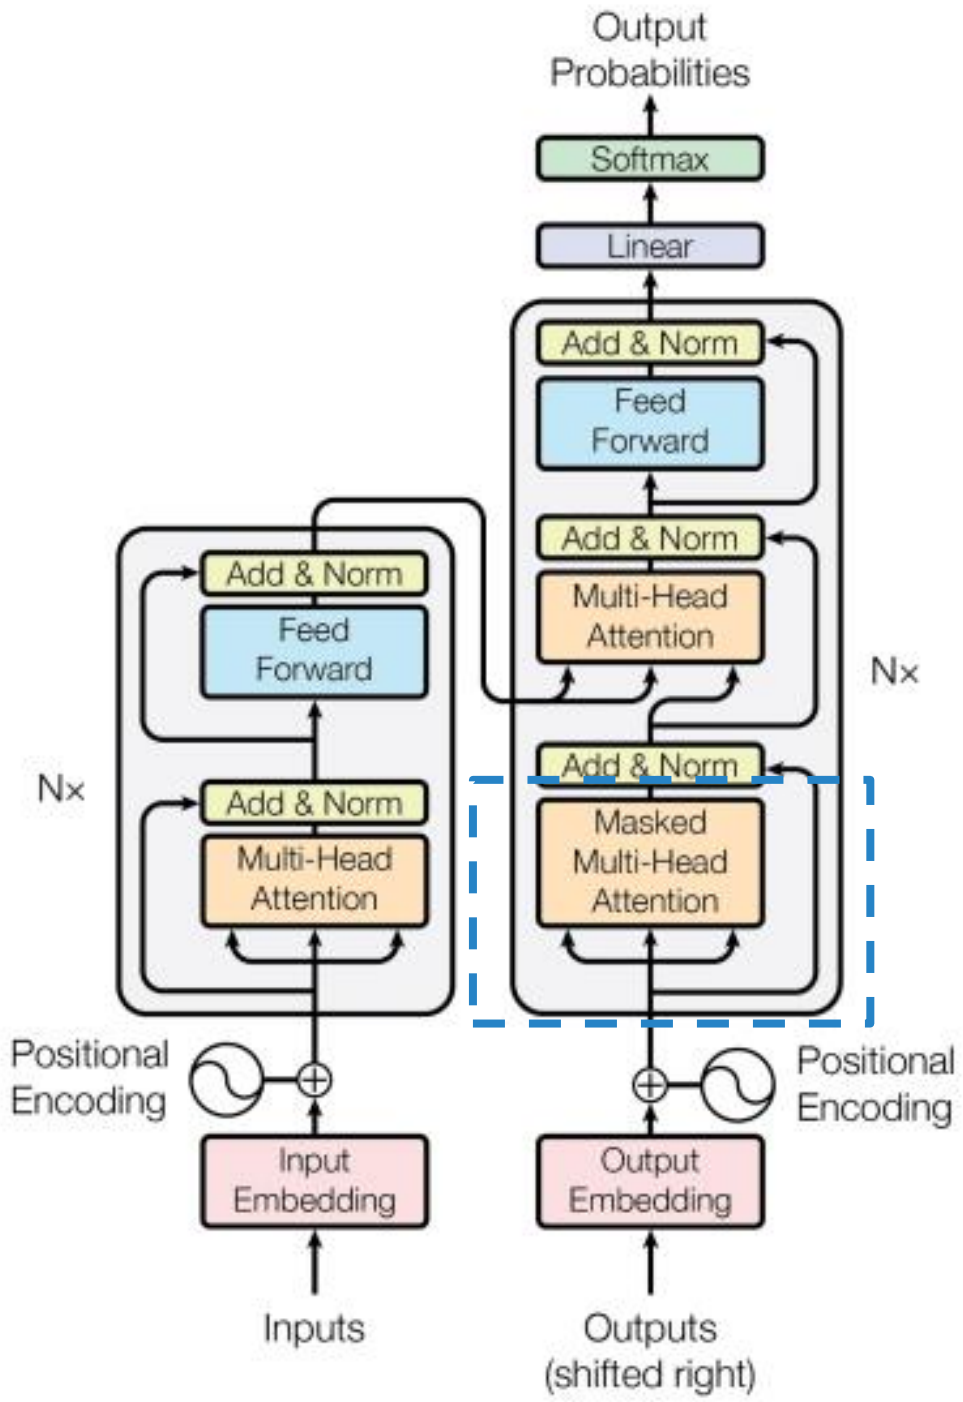

Masked Attention

1	<sos>	- ∞	- ∞	- ∞	- ∞	- ∞	- ∞
2	<sos>	Ich	- ∞	- ∞	- ∞	- ∞	- ∞
3	<sos>	Ich	have	- ∞	- ∞	- ∞	- ∞
4	<sos>	Ich	have	einen	- ∞	- ∞	- ∞
5	<sos>	Ich	have	einen	apfel	- ∞	- ∞
6	<sos>	Ich	have	einen	apfel	gegessen	- ∞
7	<sos>	Ich	have	einen	apfel	gegessen	<eos>

Softmax -> - ∞ -> 0

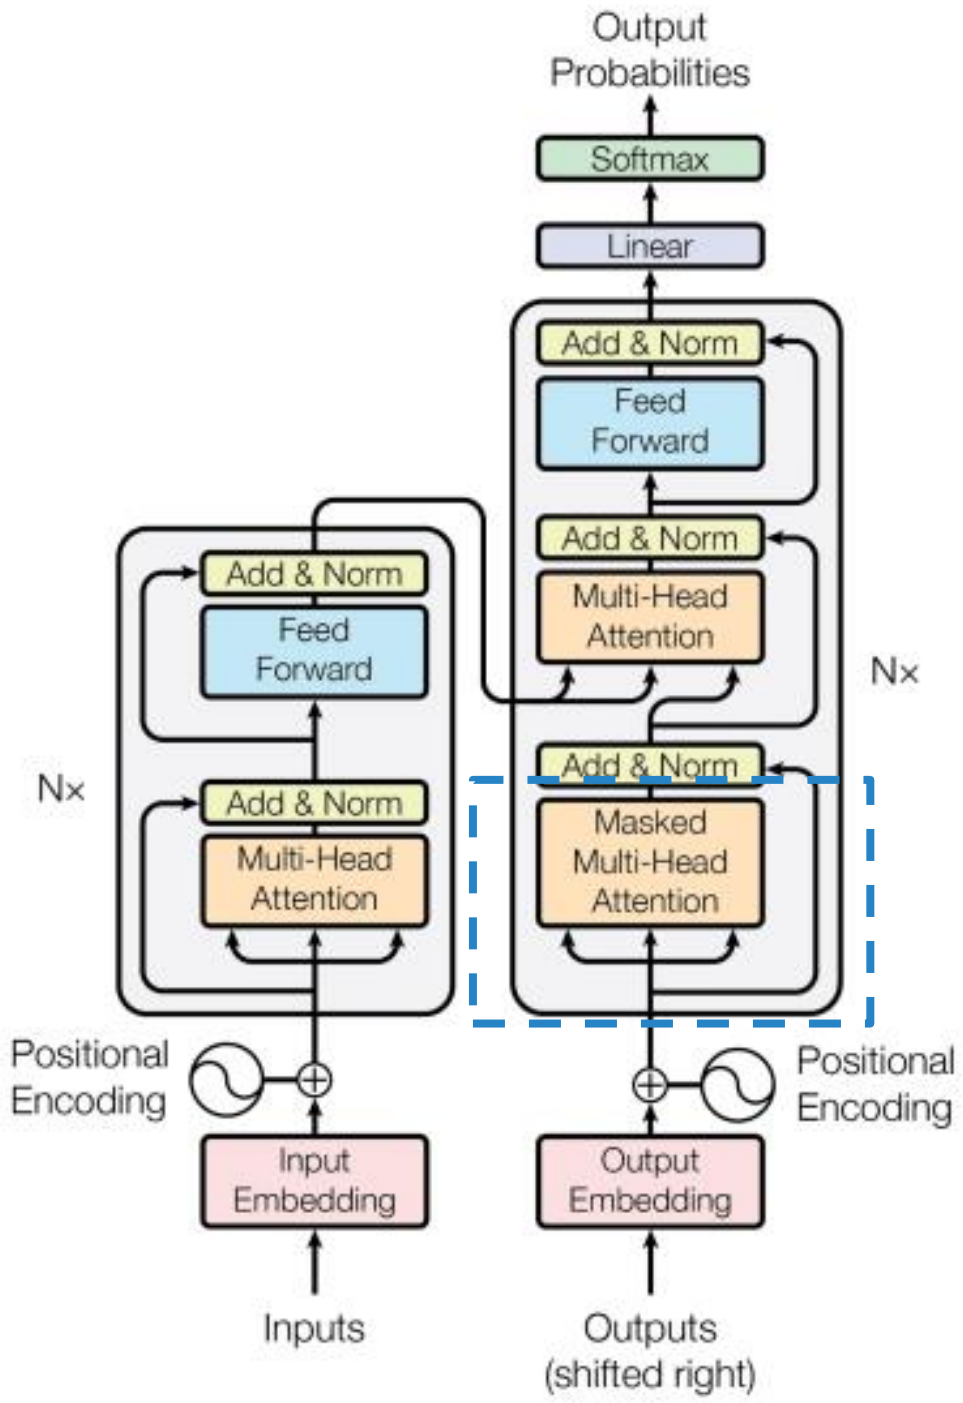

Masked Attention

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

- $Q, K, V \in \mathbb{R}^{T\times d_k}$: represent queries, keys, and values.
- M: a mask matrix with $-\infty$ in positions corresponding to future tokens, ensuring softmax output is zero for those.

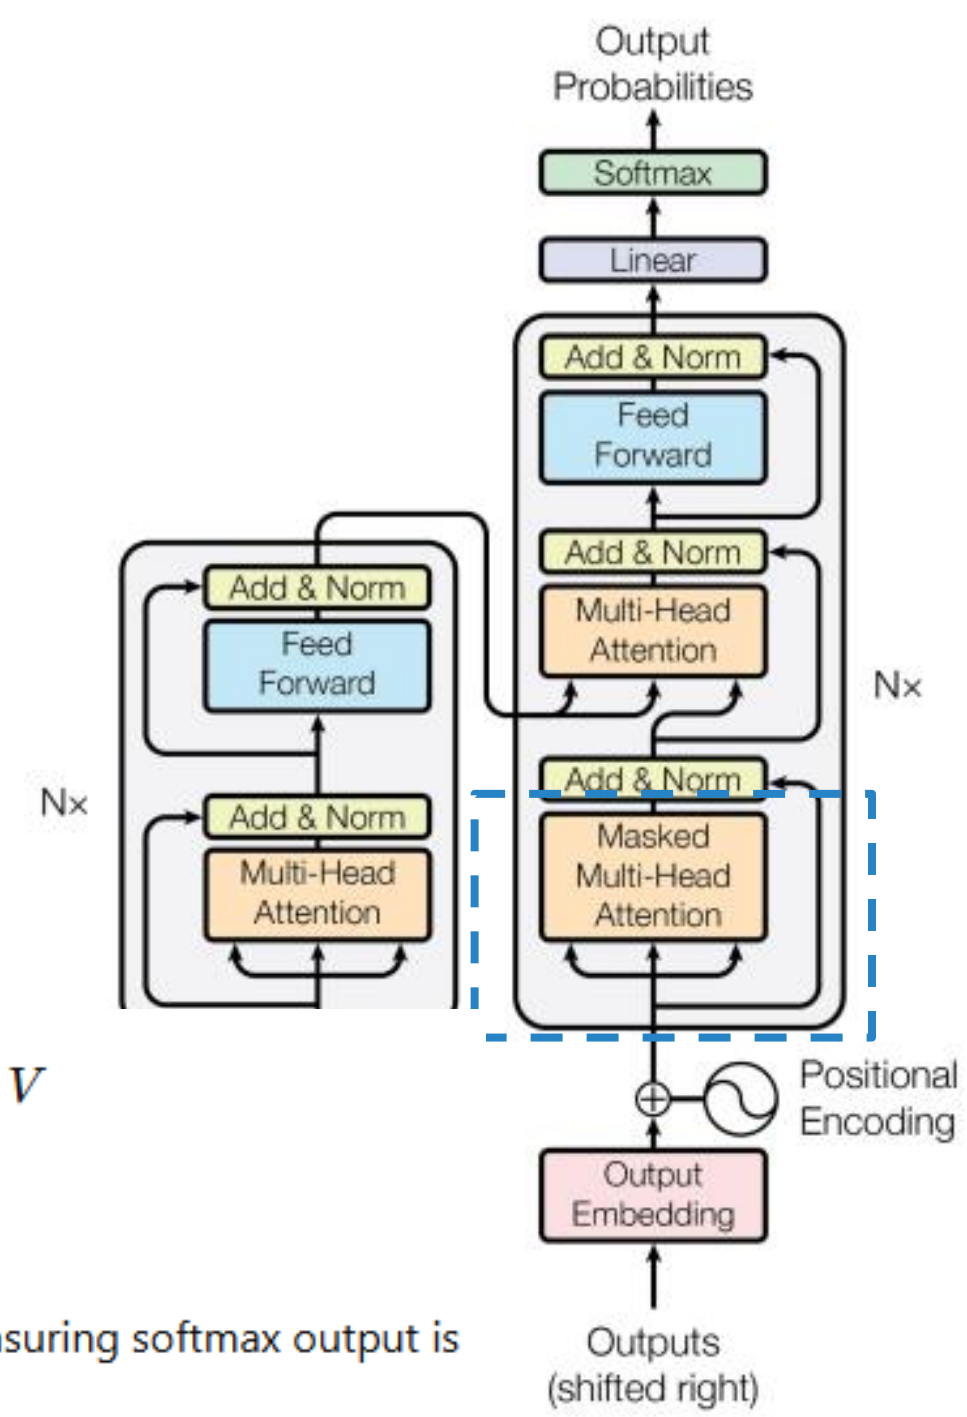

Masked Attention

$R^{T \times T}$ $R^{T \times d_h}$

Masked Attention Values

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

- $Q, K, V \in \mathbb{R}^{T \times d_k}$: represent queries, keys, and values.
- M: a mask matrix with $-\infty$ in positions corresponding to future tokens, ensuring softmax output is zero for those.

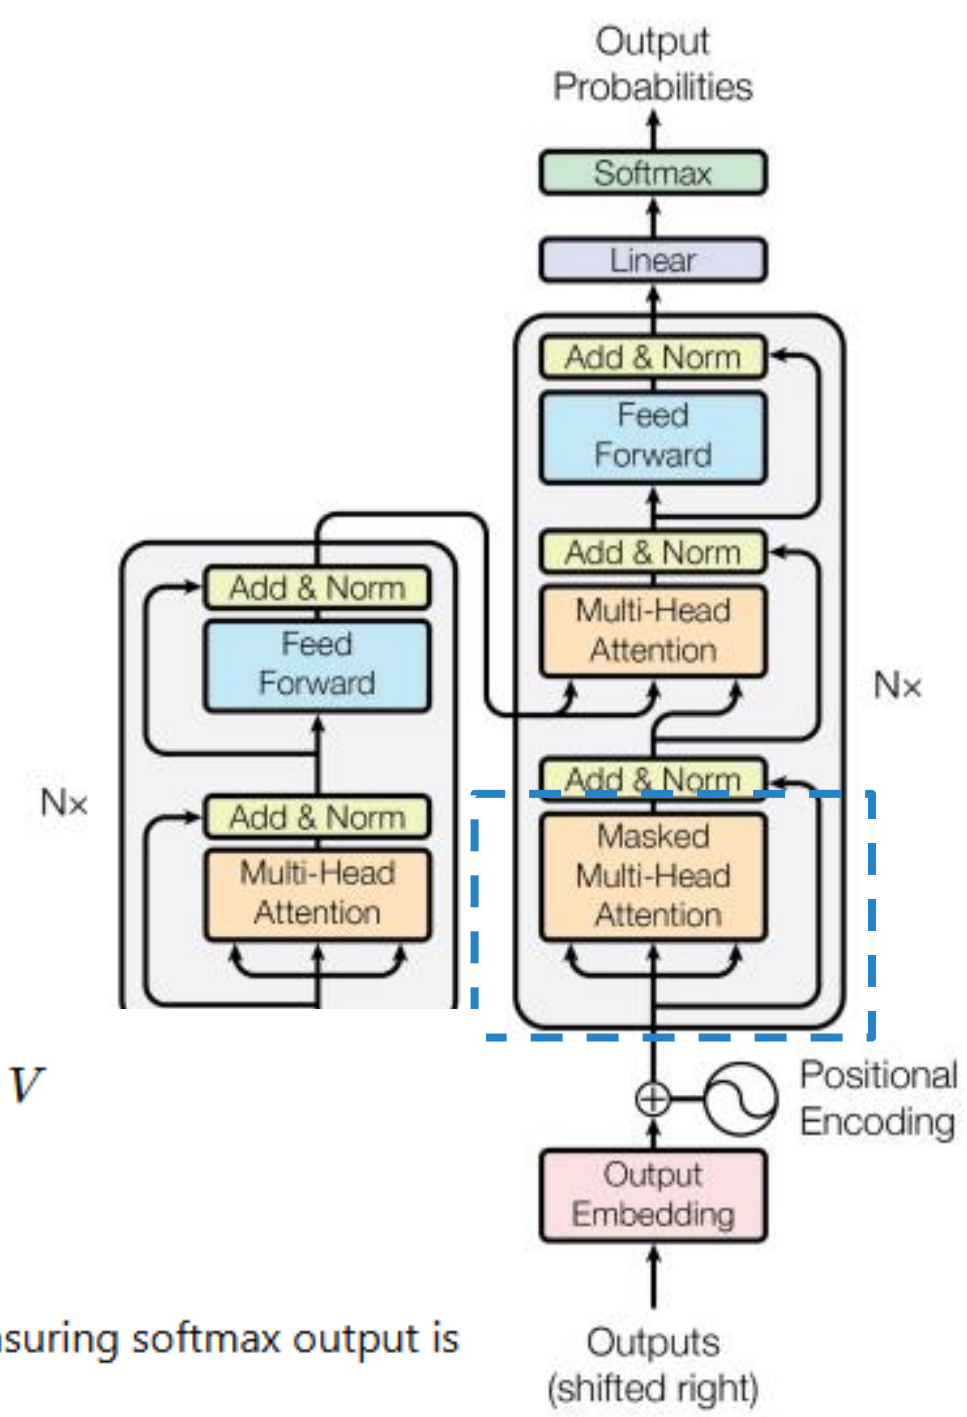

Encoder-Decoder Attention (Cross-Attention)

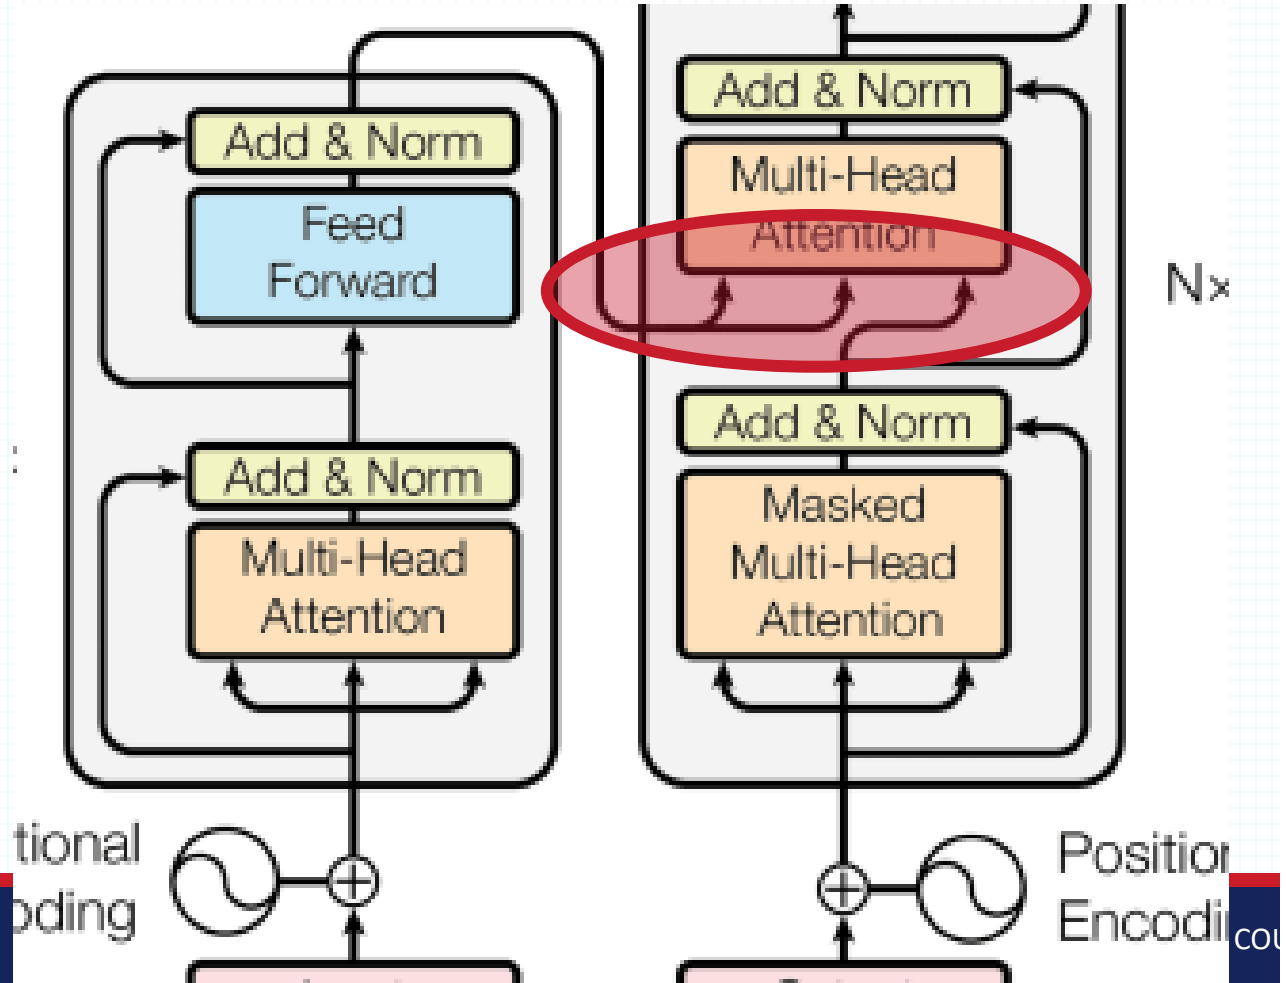

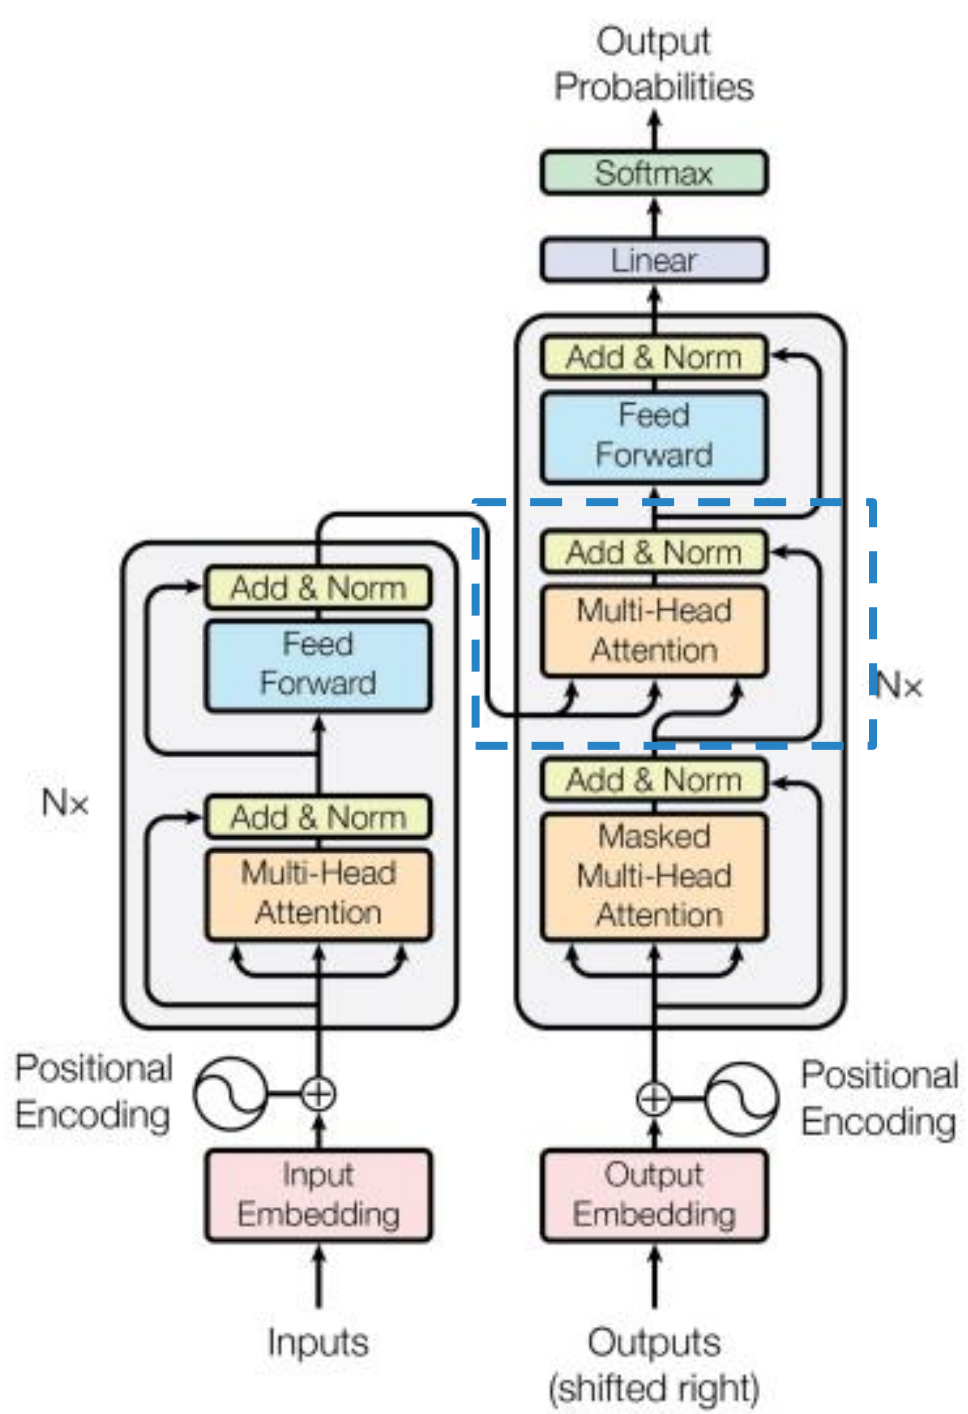

Encoder-Decoder Attention (Cross-Attention)

The outputs of the **Encoder** act as Keys and Values, and the output from the **Decoder's** self-attention acts as Queries.

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

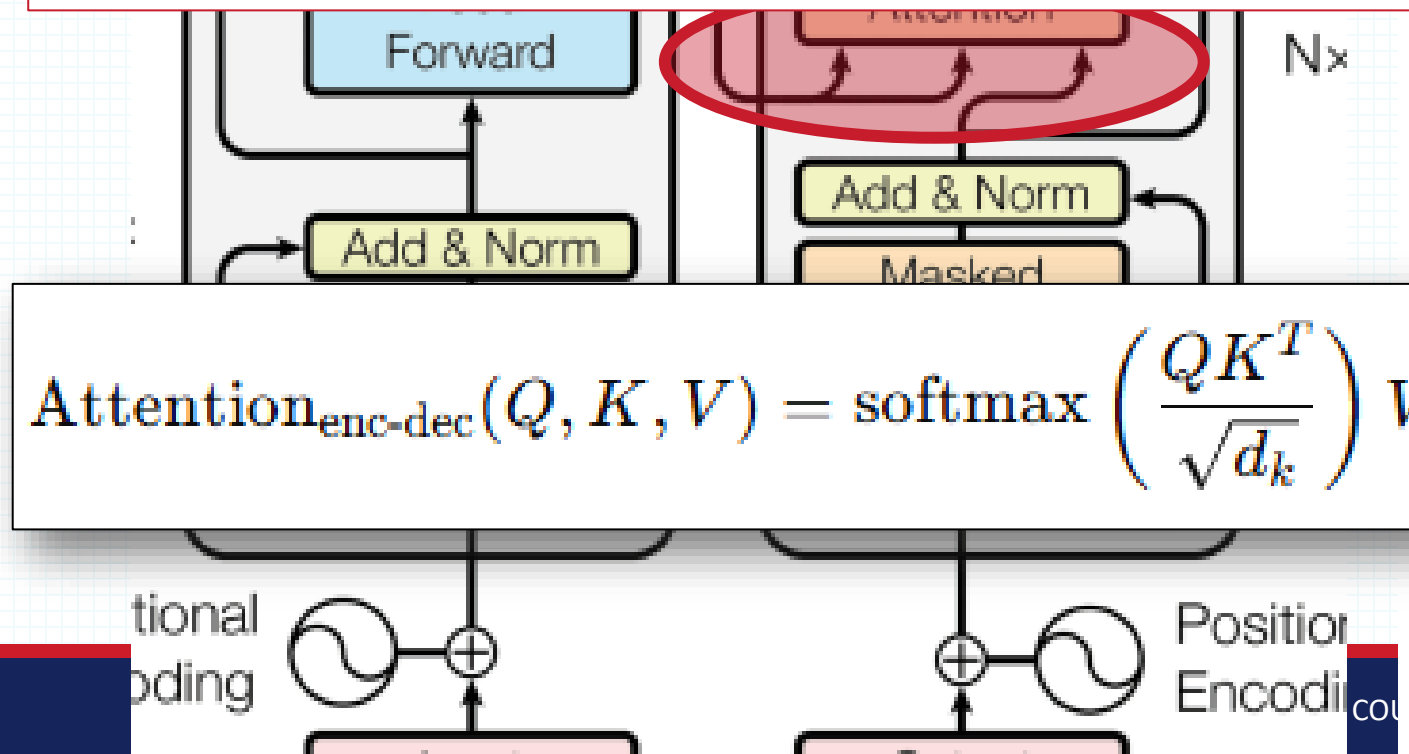

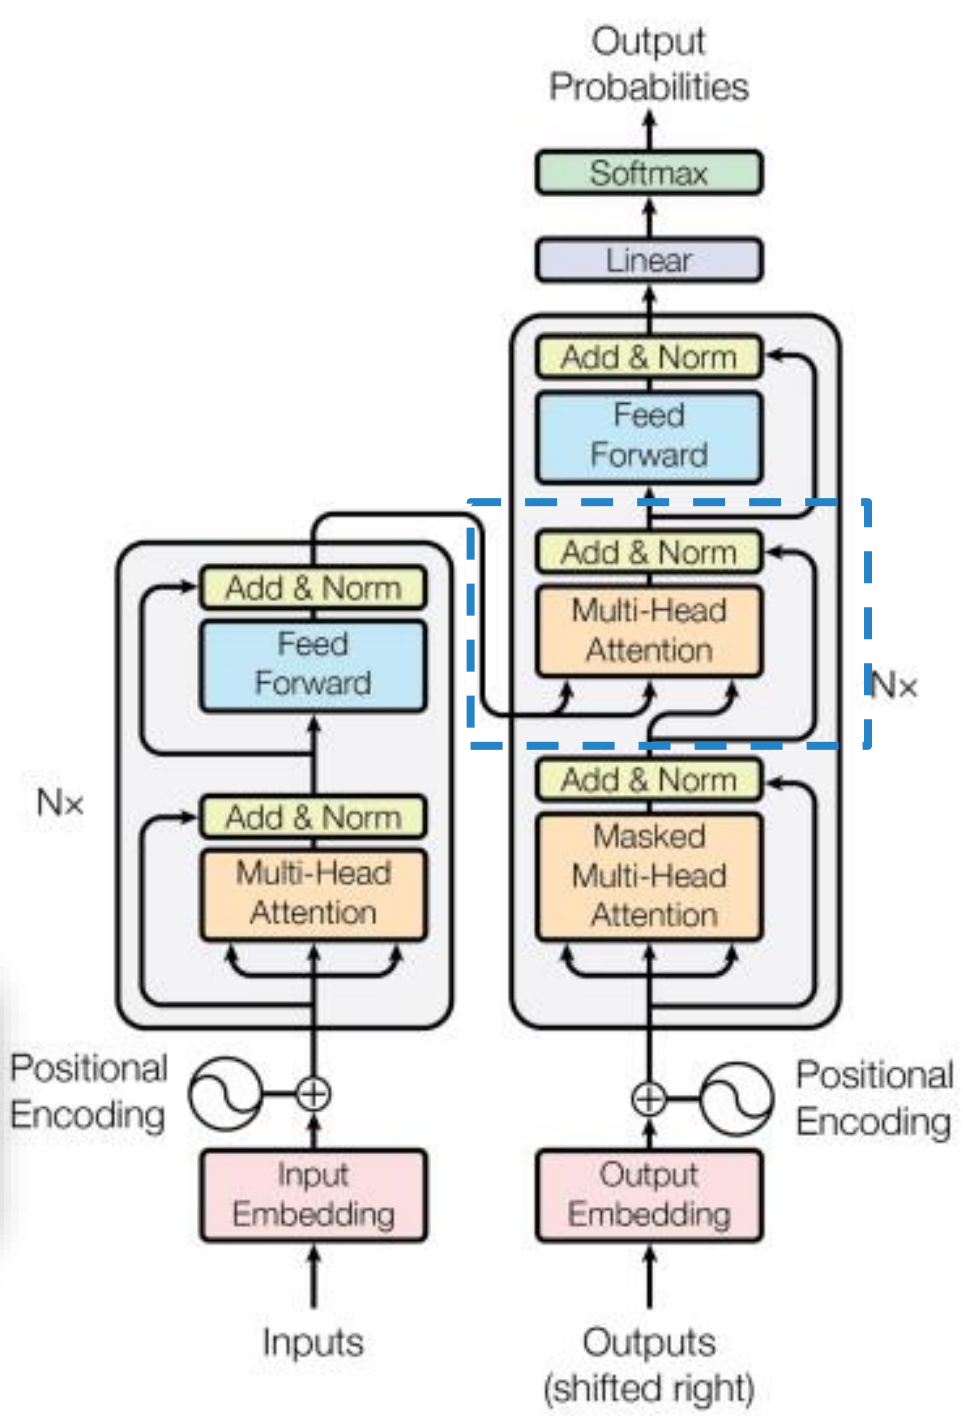

Post-Transformer Evolution and Milestones

Transformers

- a Transformer architecture can be designed in **three different ways** depending on the task and objective:
- 1. **Encoder-Only Transformers**
- **2. Decoder-Only Transformers**
- 3. **Encoder-Decoder Transformers (Seq2Seq)**

BERT
Oct 2018

Representation

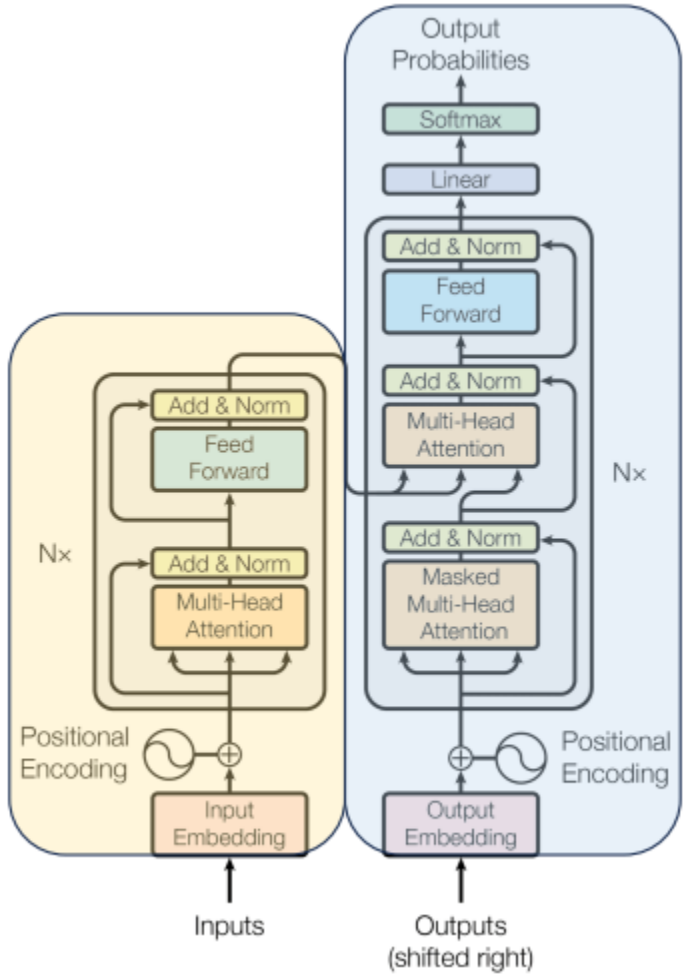

GPT
Jun 2018

Generation

Transformers… The LLM Era

BERT – 2018
DistilBERT – 2019
RoBERTa – 2019
ALBERT – 2019
ELECTRA – 2020
DeBERTa – 2020
…

Representation

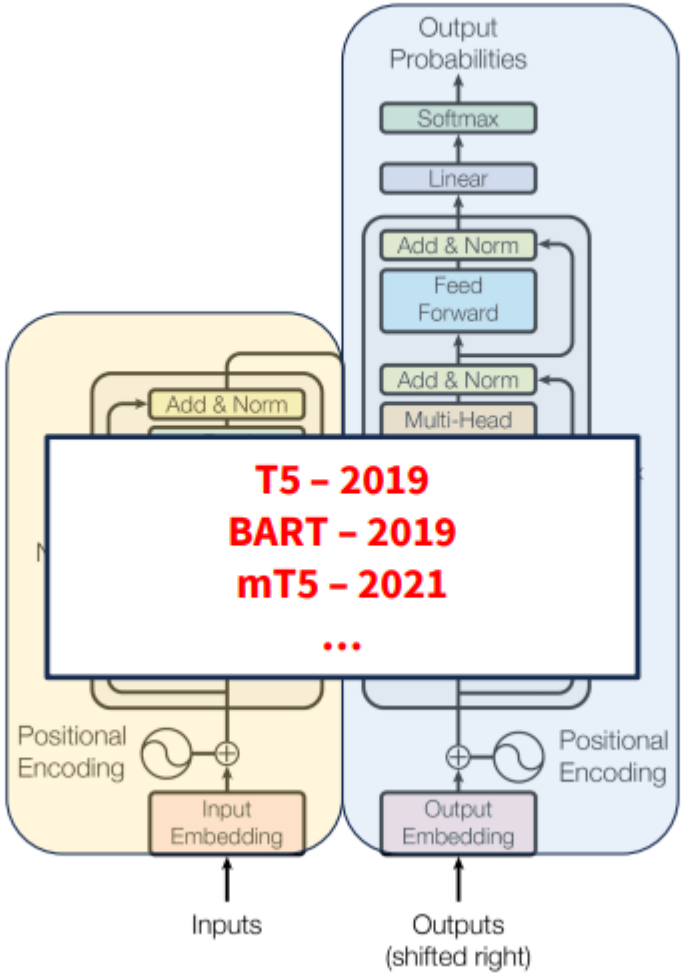

GPT – 2018
GPT-2 – 2019
GPT-3 – 2020
GPT-Neo – 2021
GPT-3.5 (ChatGPT) – 2022
LLaMA – 2023
GPT-4 – 2023
…

Generation