

IERG5050 AI Foundation Models, Systems and Applications Spring 2025

Transformers Part I: Basic Architecture

Prof. Wing C. Lau

wclau@ie.cuhk.edu.hk

<http://www.ie.cuhk.edu.hk/wclau>

Acknowledgements

Many of the slides in this lecture are adapted from the sources below. Copyrights belong to the original authors.

- Stanford CS25: Transformer United V4, Spring 2024, <https://web.stanford.edu/class/cs25/>
Instructors: Div Garg, Steven Feng, Seonghee Lee, Emily Bunnapradist ;
Faculty Advisor: Prof. Chris Manning,
Overview Slides https://docs.google.com/presentation/d/1oXPs3LXtIVIsVbwTyGjAWj_aWvak9c1uNC4uhkS6glk/edit?usp=sharing
- Stanford CS336: Language Modeling from Scratch, Spring 2024
○ by Profs. Tatsunori Hashimoto, Percy Liang, <https://stanford-cs336.github.io/spring2024/>
- Stanford CS229S: Systems for Machine Learning, Fall 2023
by Profs. Azalia Mirhoseini, Simran Arora, <https://cs229s.stanford.edu/fall2023/>
- Stanford CS224N: Natural Language Processing with Deep Learning, Winter 2021
by Prof. Chris Manning, <https://web.stanford.edu/class/archive/cs/cs224n/cs224n.1214/>
- Stanford CS231n: Deep Learning for Computer Vision, Spring 2023
by Prof. Fei-fei Li, <https://cs231n.stanford.edu/slides/2023/>
- CMU 11-667: Large Language Models: Methods and Applications, Fall 2024
by Profs. Chenyan Xiong and Daphne Ippolito, <https://cmu-llms.org>
- CMU 11-711: Advanced Natural Language Processing (ANLP), Spring 2024
by Prof. Graham Neubig, <https://phontron.com/class/anlp2024/lectures/>
- UPenn CIS7000: Large Language Models, Fall 2024
by Prof. Mayur Naik, <https://llm-class.github.io/schedule.html>
- Princeton COS597G: Understanding Large Language Models, Fall 2022
by Prof. Danqi Chen, <https://www.cs.princeton.edu/courses/archive/fall22/cos597G/>
- UWaterloo CS886: Recent Advances on Foundation Models, Winter 2024
by Prof. Wenhua Chen, <https://cs.uwaterloo.ca/~wenhuche/teaching/cs886/>
- UMD CMSC848K: Multimodal Foundation Models, Fall 2024
by Prof. Jia-Bin Huang, <https://jbbhuang0604.github.io/teaching/CMSC848K/>

Natural Language Processing (NLP) & Language Modeling

- NLP (natural language processing) tasks
 - Translation, question answering, recommendations, sentence completion, etc
- Language model
 - Model the probability of a sequence of tokens in a text
- Examples
 - I was eating ___.
 - an apple (0.02)
 - a banana (0.01)
 - popcorns (0.0001)
 - I was in a ___. I was eating popcorns
 - house (0.01)
 - mansion (0.001)
 - movie theater (0.2)

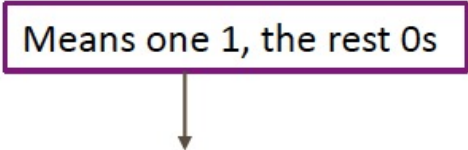
$$p(x) = p(x_1, \dots, x_T) = \prod_{t=1}^T p(x_t | x_{<t})$$

Representing Words as Discrete Symbols

In traditional NLP, we regard words as discrete symbols:

hotel, conference, motel – a **localist** representation

Means one 1, the rest 0s



Such symbols for words can be represented by **one-hot** vectors:

motel = [0 0 0 0 0 0 0 0 0 0 1 0 0 0 0]

hotel = [0 0 0 0 0 0 0 1 0 0 0 0 0 0 0]

Vector dimension = number of words in vocabulary (e.g., 500,000+)

Problem with Words as Discrete Symbols

Example: in web search, if a user searches for “Seattle motel”, we would like to match documents containing “Seattle hotel”

But:

motel = [0 0 0 0 0 0 0 0 0 1 0 0 0 0]

hotel = [0 0 0 0 0 0 0 1 0 0 0 0 0 0]

These two vectors are orthogonal

There is no natural notion of **similarity** for one-hot vectors!

Solution:

- Could try to rely on WordNet’s list of synonyms to get similarity?
 - But it is well-known to fail badly: incompleteness, etc.
- **Instead: learn to encode similarity in the vectors themselves**

Representing Words by their Context



- **Distributional semantics: A word's meaning is given by the words that frequently appear close-by**
 - *"You shall know a word by the company it keeps"* (J. R. Firth 1957: 11)
 - One of the most successful ideas of modern statistical NLP!
- When a word w appears in a text, its **context** is the set of words that appear nearby (within a fixed-size window).
- We use the many contexts of w to build up a representation of w

*...government debt problems turning into **banking** crises as happened in 2009...*
*...saying that Europe needs unified **banking** regulation to replace the hodgepodge...*
*...India has just given its **banking** system a shot in the arm...*

↑ ↑
These **context words** will represent **banking**

Word Vectors (aka Word Embeddings)

We will build a dense vector for each word, chosen so that it is similar to vectors of words that appear in similar contexts, measuring similarity as the vector **dot** (scalar) **product**

$$\begin{matrix} \textit{banking} = \\ \\ \\ \\ \\ \\ \\ \end{matrix} \begin{pmatrix} 0.286 \\ 0.792 \\ -0.177 \\ -0.107 \\ 0.109 \\ -0.542 \\ 0.349 \\ 0.271 \end{pmatrix} \qquad \begin{matrix} \textit{monetary} = \\ \\ \\ \\ \\ \\ \\ \end{matrix} \begin{pmatrix} 0.413 \\ 0.582 \\ -0.007 \\ 0.247 \\ 0.216 \\ -0.718 \\ 0.147 \\ 0.051 \end{pmatrix}$$


Note: **word vectors** are also called **(word) embeddings** or **(neural) word representations**
They are a **distributed** representation

Word2vec: How to learn the Word Embedding

Word2vec is a framework for learning word vectors
(Mikolov et al. 2013)

Idea:

- We have a large corpus (“body”) of text: a long list of words
- Every word in a fixed vocabulary is represented by a **vector**
- Go through each position t in the text, which has a center word c and context (“outside”) words o
- Use the **similarity of the word vectors** for c and o to **calculate the probability** of o given c (or vice versa)
- **Keep adjusting the word vectors** to maximize this probability



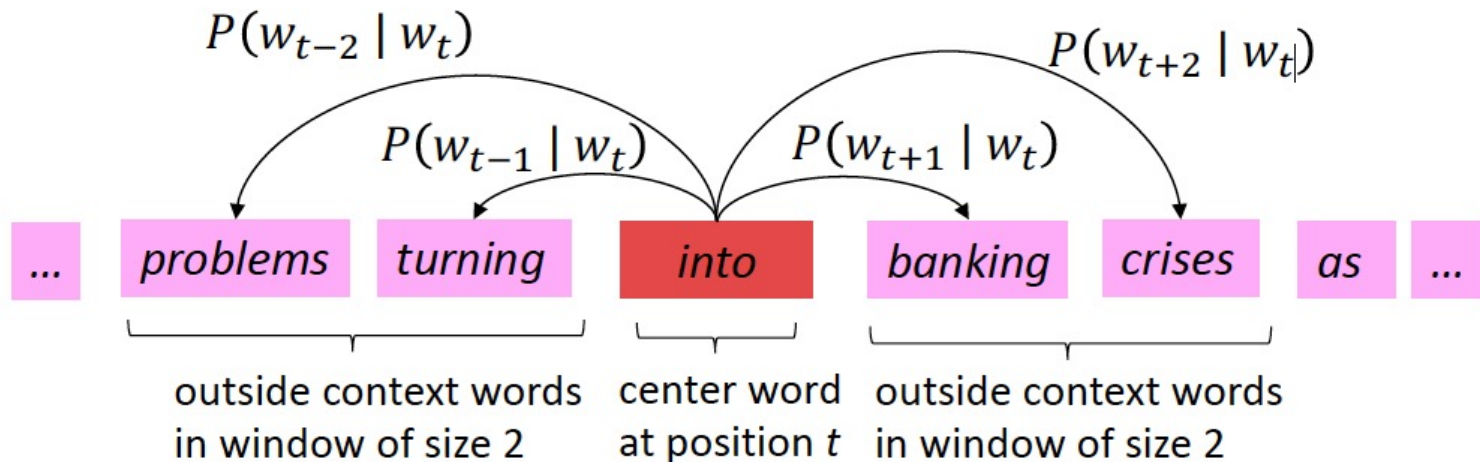
Skip-gram model
(Mikolov et al. 2013)

Source: <https://web.stanford.edu/class/archive/cs/cs224n/cs224n.1244/slides/cs224n-2024-lecture01-wordvecs1-public.pdf>

For details, refer to: https://web.stanford.edu/class/archive/cs/cs224n/cs224n.1244/readings/cs224n_winter2023_lecture1_notes_draft.pdf

Word2vec Overview

Example windows and process for computing $P(w_{t+j} | w_t)$

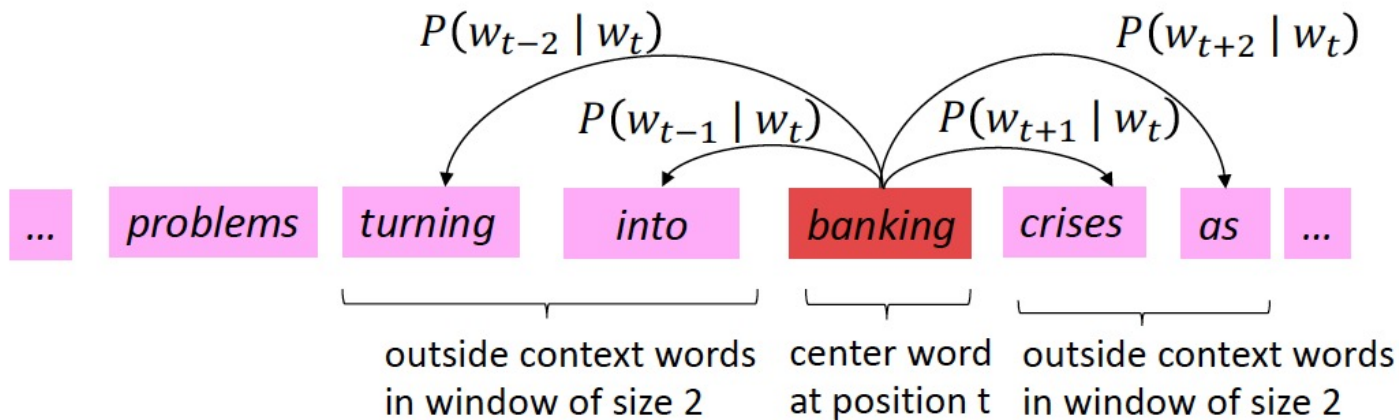


Source: <https://web.stanford.edu/class/archive/cs/cs224n/cs224n.1244/slides/cs224n-2024-lecture01-wordvecs1-public.pdf>

For details, refer to: https://web.stanford.edu/class/archive/cs/cs224n/cs224n.1244/readings/cs224n_winter2023_lecture1_notes_draft.pdf

Word2vec Overview

Example windows and process for computing $P(w_{t+j} | w_t)$



Source: <https://web.stanford.edu/class/archive/cs/cs224n/cs224n.1244/slides/cs224n-2024-lecture01-wordvecs1-public.pdf>

For details, refer to: https://web.stanford.edu/class/archive/cs/cs224n/cs224n.1244/readings/cs224n_winter2023_lecture1_notes_draft.pdf

Word2vec: Objective Function

For each position $t = 1, \dots, T$, predict context words within a window of fixed size m , given center word w_t . Data likelihood:

$$\text{Likelihood} = L(\theta) = \prod_{t=1}^T \prod_{\substack{-m \leq j \leq m \\ j \neq 0}} P(w_{t+j} | w_t; \theta)$$

θ is all variables to be optimized

sometimes called a *cost* or *loss* function

The **objective function** $J(\theta)$ is the (**average**) negative log likelihood:

$$J(\theta) = -\frac{1}{T} \log L(\theta) = -\frac{1}{T} \sum_{t=1}^T \sum_{\substack{-m \leq j \leq m \\ j \neq 0}} \log P(w_{t+j} | w_t; \theta)$$

Minimizing objective function \Leftrightarrow Maximizing predictive accuracy

Source: <https://web.stanford.edu/class/archive/cs/cs224n/cs224n.1244/slides/cs224n-2024-lecture01-wordvecs1-public.pdf>

For details, refer to: https://web.stanford.edu/class/archive/cs/cs224n/cs224n.1244/readings/cs224n_winter2023_lecture1_notes_draft.pdf

Word2vec: Objective Function

- We want to minimize the objective function:

$$J(\theta) = -\frac{1}{T} \sum_{t=1}^T \sum_{\substack{-m \leq j \leq m \\ j \neq 0}} \log P(w_{t+j} | w_t; \theta)$$

- **Question:** How to calculate $P(w_{t+j} | w_t; \theta)$?

- **Answer:** We will use two vectors per word w :

- v_w when w is a center word
- u_w when w is a context word

} These word vectors are subparts of the big vector of all parameters θ

- Then for a center word c and a context word o :

$$P(o|c) = \frac{\exp(u_o^T v_c)}{\sum_{w \in V} \exp(u_w^T v_c)}$$

Word2vec: Prediction Function = Prob[o|c]

② Exponentiation makes anything positive

$$P(o|c) = \frac{\exp(u_o^T v_c)}{\sum_{w \in V} \exp(u_w^T v_c)}$$

① Dot product compares similarity of o and c .

$$u^T v = u \cdot v = \sum_{i=1}^n u_i v_i$$

Larger dot product = larger probability

③ Normalize over entire vocabulary to give probability distribution

• This is an example of the **softmax function** $\mathbb{R}^n \rightarrow (0,1)^n$

Open region

$$\text{softmax}(x_i) = \frac{\exp(x_i)}{\sum_{j=1}^n \exp(x_j)} = p_i$$

• The softmax function maps arbitrary values x_i to a probability distribution p_i

- “max” because amplifies probability of largest x_i
- “soft” because still assigns some probability to smaller x_i
- Frequently used in Deep Learning

But sort of a weird name because it returns a distribution!

Source: <https://web.stanford.edu/class/archive/cs/cs224n/cs224n.1244/slides/cs224n-2024-lecture01-wordvecs1-public.pdf>


For details, refer to: https://web.stanford.edu/class/archive/cs/cs224n/cs224n.1244/readings/cs224n_winter2023_lecture1_notes_draft.pdf

Word2vec: To Train the Model

To train a model, we gradually adjust parameters to minimize a loss

- Recall: θ represents **all** the model parameters, in one long vector
- In our case, with d -dimensional vectors and V -many words, we have \rightarrow
- Remember: every word has two vectors

$$\theta = \begin{bmatrix} v_{aardvark} \\ v_a \\ \vdots \\ v_{zebra} \\ u_{aardvark} \\ u_a \\ \vdots \\ u_{zebra} \end{bmatrix} \in \mathbb{R}^{2dV}$$



- We optimize these parameters by walking down the gradient (see right figure)
- We compute **all** vector gradients!

Source: <https://web.stanford.edu/class/archive/cs/cs224n/cs224n.1244/slides/cs224n-2024-lecture01-wordvecs1-public.pdf>

For details, refer to: https://web.stanford.edu/class/archive/cs/cs224n/cs224n.1244/readings/cs224n_winter2023_lecture1_notes_draft.pdf

The problematic $P[o|c]$

$$P(o|c) = \frac{\exp(u_o^T v_c)}{\sum_{i=1}^N \exp(u_i^T v_c)}$$

The denominator is difficult to evaluate as it involves the embedding of **ALL** words in the Universe (Vocabulary)


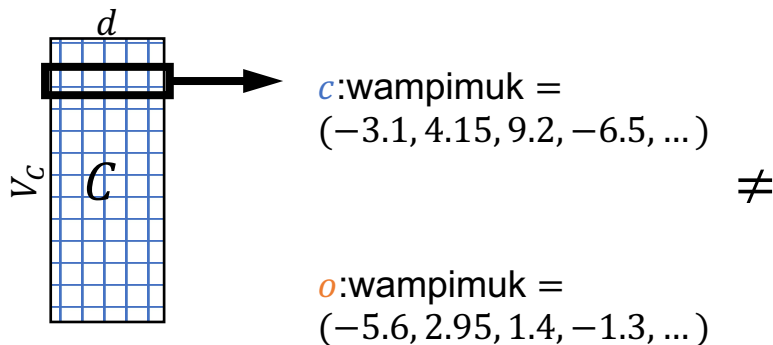
=> Use a trick to circumvent the problem via Negative Sampling

What is Word2vec ?

- ▶ `word2vec` is **not** a single algorithm
- ▶ It is a **software package** for representing words as vectors, containing:
 - ▶ Two distinct models
 - ▶ CBoW
 - ▶ **Skip-Gram** (SG)
 - ▶ Various training methods
 - ▶ **Negative Sampling** (NS)
 - ▶ Hierarchical Softmax
 - ▶ A rich preprocessing pipeline
 - ▶ Dynamic Context Windows
 - ▶ Subsampling
 - ▶ Deleting Rare Words
- ▶ We will focus on the **Skip-Grams with Negative Sampling (SGNS)** approach !

SGNS starts with SAME basic setting

- SGNS finds a vector v_c for each word c in our vocabulary V_C
- Each such vector has d latent dimensions (e.g. $d = 100$)
- Effectively, it learns a matrix C whose rows represent the center vectors v_c
- **Key point:** it also learns a similar auxiliary matrix O of outside context vectors
- In fact, each word has two embeddings: v_c and u_o



“word2vec Explained...”
Goldberg & Levy, arXiv
2014

Skip-Grams with Negative Sampling (SGNS)

You first observe (**actually sample**) the following sentence from the training corpus:

Marco saw a furry little **wampimuk** hiding in the tree.

“word2vec Explained...”
Goldberg & Levy, arXiv
2014

Skip-Grams with Negative Sampling (SGNS)

Marco saw a furry little wampimuk hiding in the tree.

center word

wampimuk

wampimuk

wampimuk

wampimuk

...

outside context word

furry

little

hiding

in

...



D (observed data)

“word2vec Explained...” Goldberg & Levy, arXiv 2014

Skip-Grams with Negative Sampling (SGNS)

Maximize $\prod_i \sigma(\vec{c} \cdot \vec{o}_i)$

- o_i was **observed** with c

where $\sigma(z) = 1 / [1 + \exp(-z)]$

<u>center word</u>	<u>outside context</u>
--------------------	------------------------

wampimuk	furry
wampimuk	little
wampimuk	hiding
wampimuk	in

AND Minimize $\prod_i \sigma(\vec{c} \cdot \vec{o}_i')$

\equiv Maximize $\prod_i [1 - \sigma(\vec{c} \cdot \vec{o}_i')]$

- o_i' was **NOT observed** with c , they are from the set of Negative Samples D' **randomly generated** by the algorithm.

<u>center word</u>	<u>NOT outside context</u>
--------------------	----------------------------

wampimuk	Australia
wampimuk	cyber
wampimuk	the
wampimuk	1985

Take Log and the optimization problem becomes “similar” to the training of a binary logistic-regression classifier:

$$\arg \max_{\theta} \sum_{(w,c) \in D} \log \frac{1}{1 + e^{-v_c \cdot v_w}} + \sum_{(w,c) \in D'} \log \left(1 - \frac{1}{1 + e^{-v_c \cdot v_w}} \right) = \arg \max_{\theta} \sum_{(w,c) \in D} \log \frac{1}{1 + e^{-v_c \cdot v_w}} + \sum_{(w,c) \in D'} \log \left(\frac{1}{1 + e^{v_c \cdot v_w}} \right)$$

Summary: How to learn Word2vec Embeddings via SGNS

For a vocabulary of size V : Start with V random 300-dimensional vectors as initial embeddings

Train a logistic regression classifier to distinguish words that co-occur in corpus from those that don't

Pairs of words that co-occur are positive examples

Pairs of words that don't co-occur are negative examples

Train the classifier to distinguish these by slowly adjusting all the embeddings to improve the classifier performance


Throw away the classifier code and keep the embeddings.

NLP Tasks as Sequence to Sequence Modeling

NLP Tasks as Sequence to Sequence Modeling

► Example Scenarios

- Text \rightarrow Text (e.g. Q/A, translation, text summarization)
- Image \rightarrow Text (e.g. image captioning)



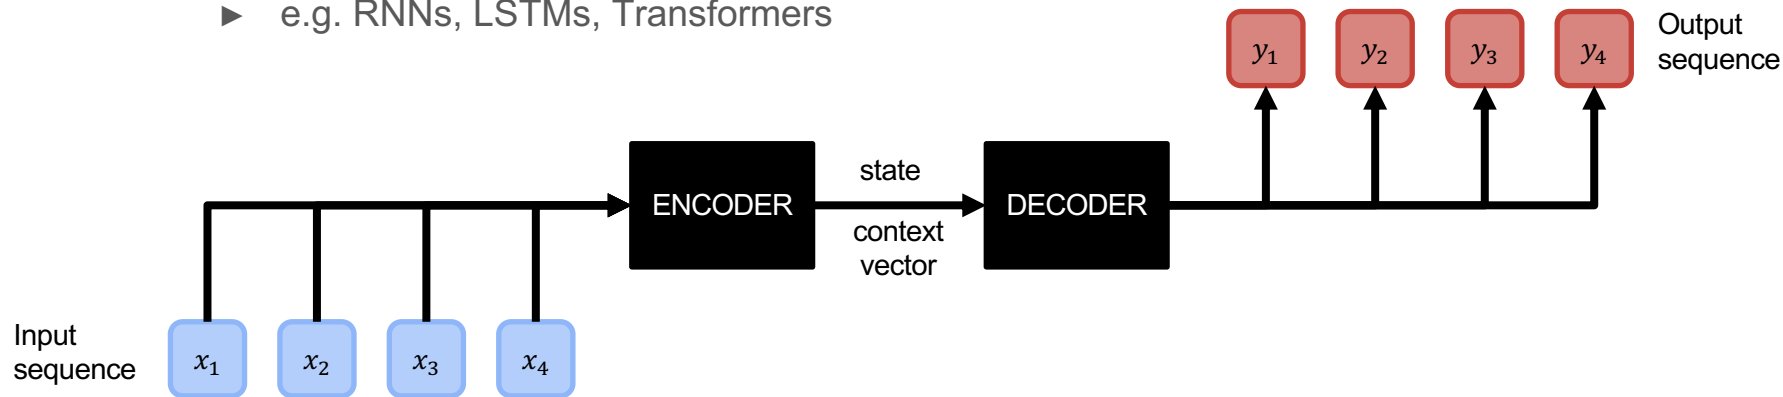
NLP Tasks as Sequence to Sequence Modeling

▶ Example Scenarios

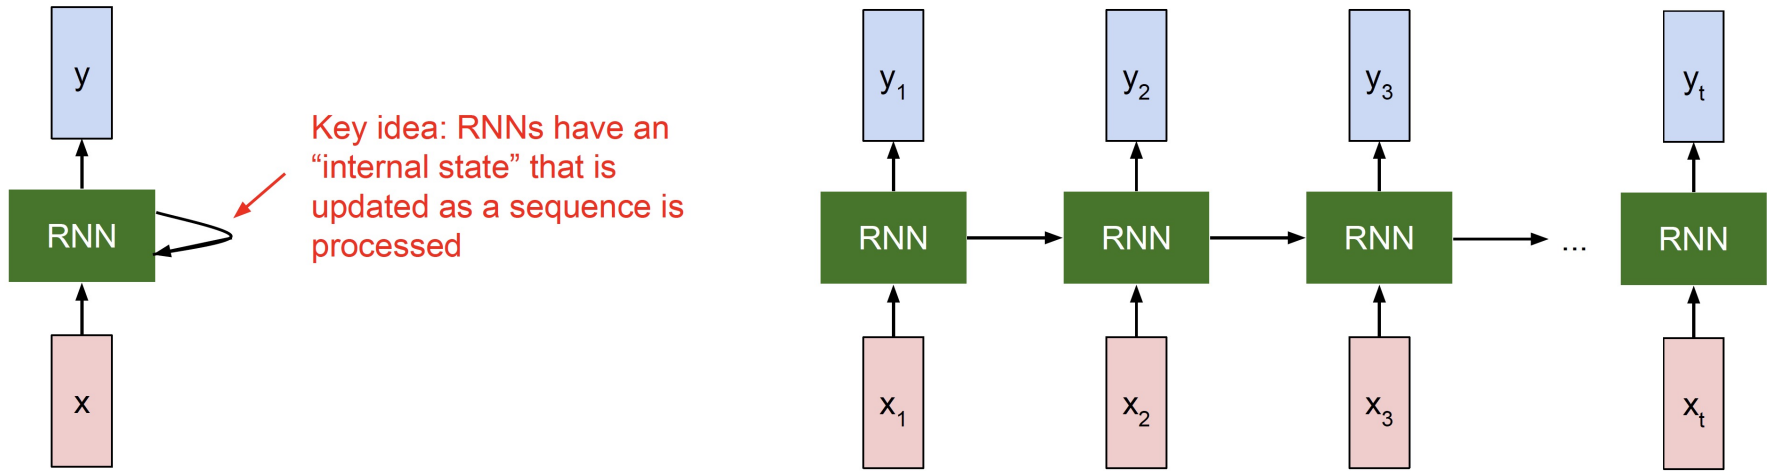
- ▶ Text → Text (e.g. Q/A, translation, text summarization)
- ▶ Image → Text (e.g. image captioning)

▶ How? Usually Encoder-Decoder models

- ▶ e.g. RNNs, LSTMs, Transformers



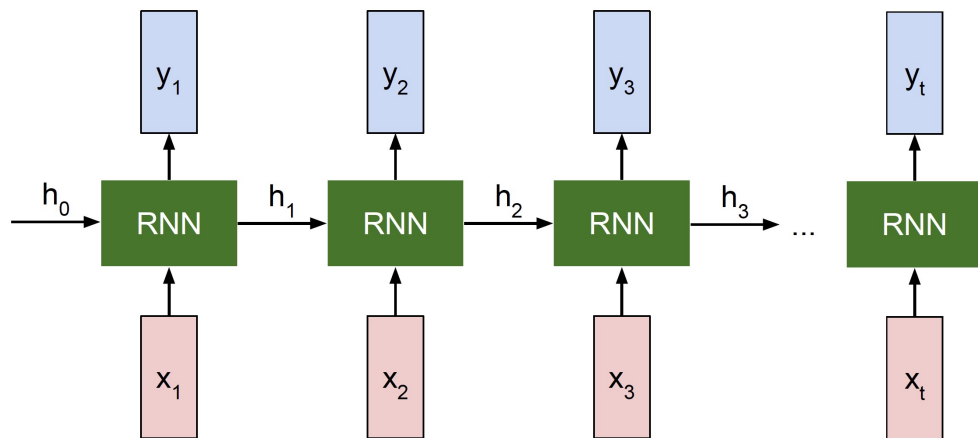
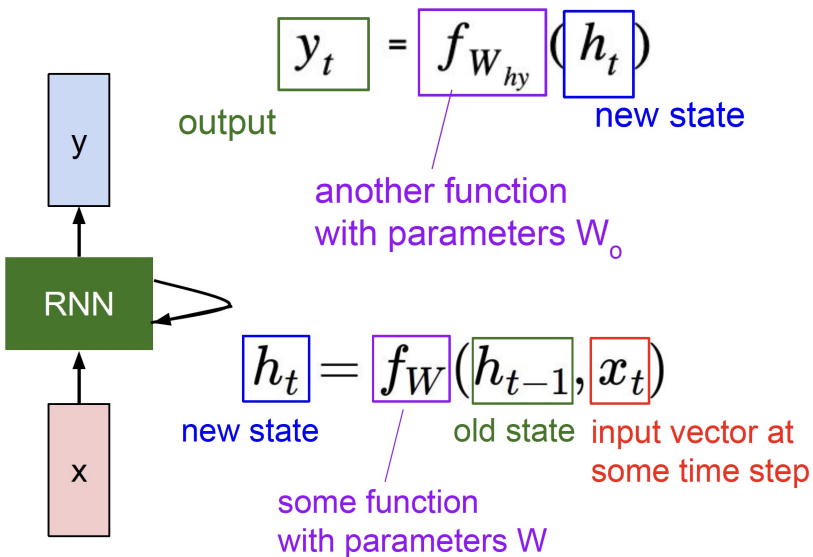
Recurrent Neural Networks (RNN) – a Seq2Seq NN model



An RNN "Unrolled" along the Time axis

Recurrent Neural Networks (RNN)


We can process a sequence of vectors \mathbf{x} by applying a **recurrence formula** at every time step:



Notice: the same function and the same set of parameters are used at every time step.

(Vanilla) Recurrent Neural Networks (RNN)

The state consists of a single “*hidden*” vector \mathbf{h} :



$$h_t = f_W(h_{t-1}, x_t)$$

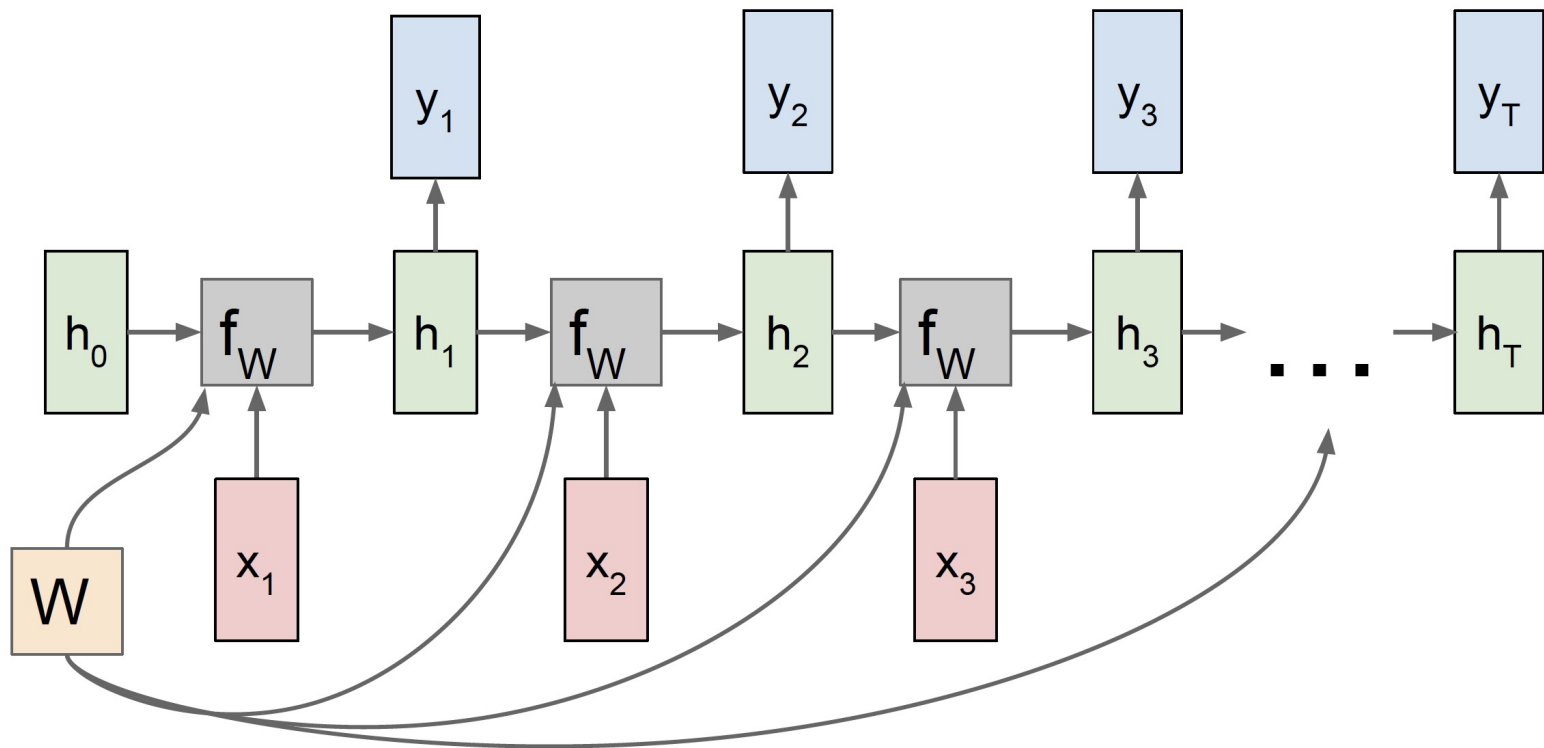


$$h_t = \tanh(W_{hh}h_{t-1} + W_{xh}x_t)$$

$$y_t = W_{hy}h_t$$

Sometimes called a “Vanilla RNN” or an “Elman RNN” after Prof. Jeffrey Elman

Computational Graph for an RNN



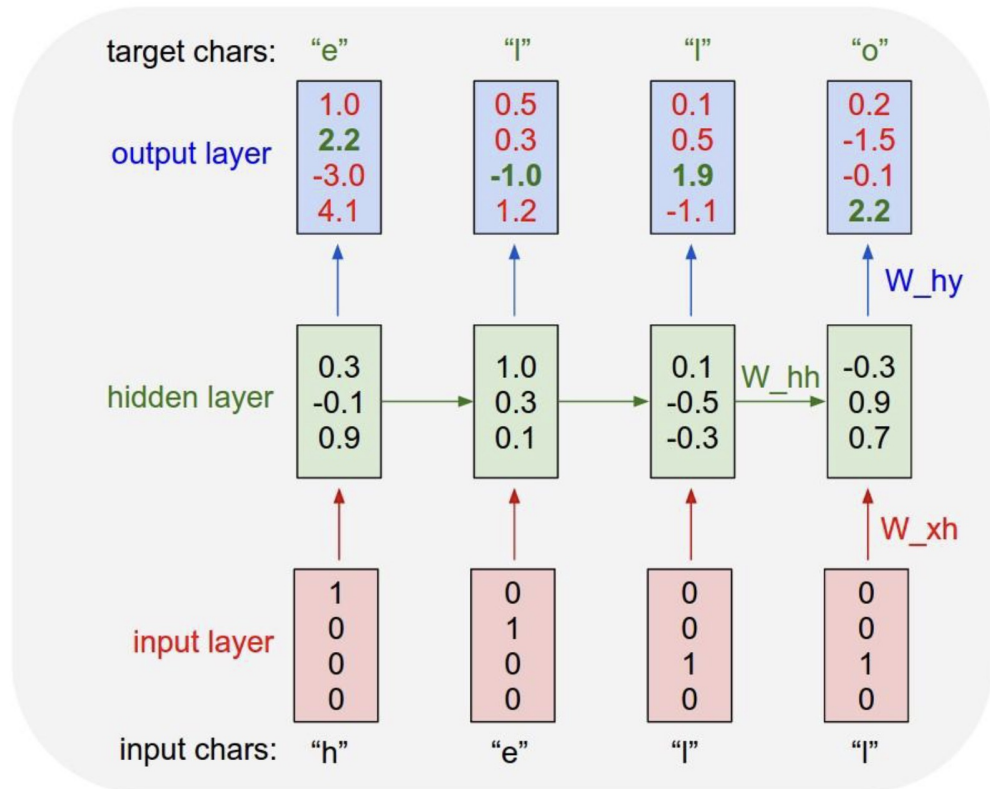
Note the reusing of the **SAME** weight matrix f_W at every time-step !

Example: Character-level Language Model

Vocabulary:
[h,e,l,o]

Example training
sequence:
“hello”

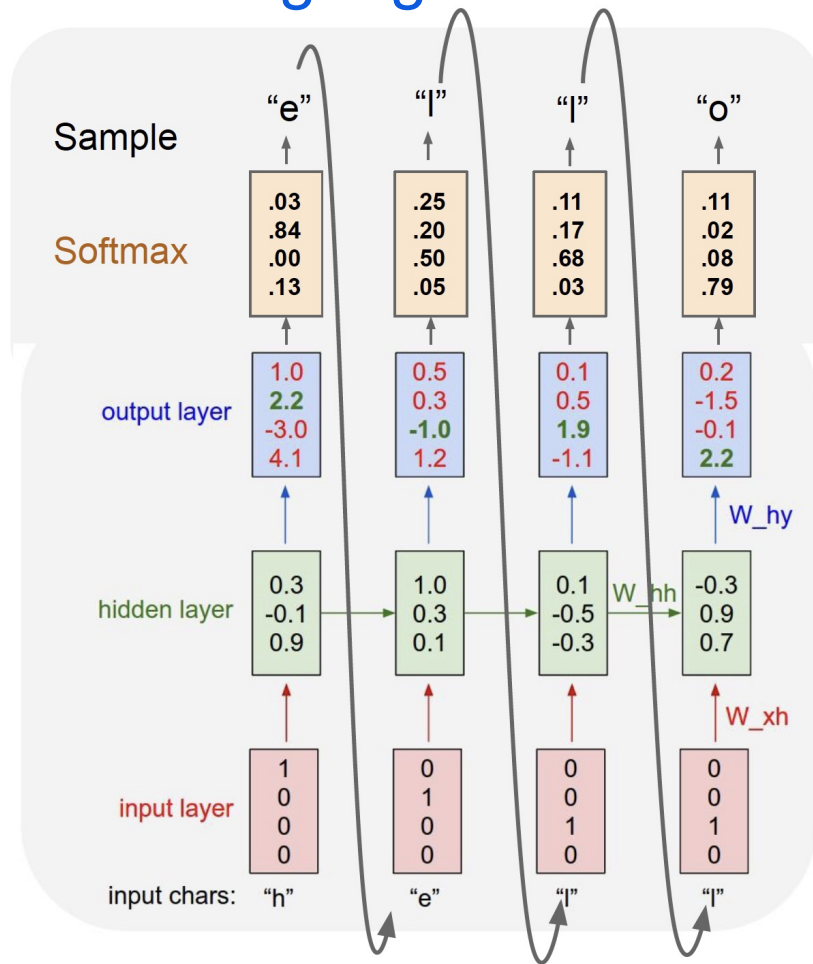
$$h_t = \tanh(W_{hh}h_{t-1} + W_{xh}x_t)$$



Example: Character-level Language Model

Vocabulary:
[h,e,l,o]

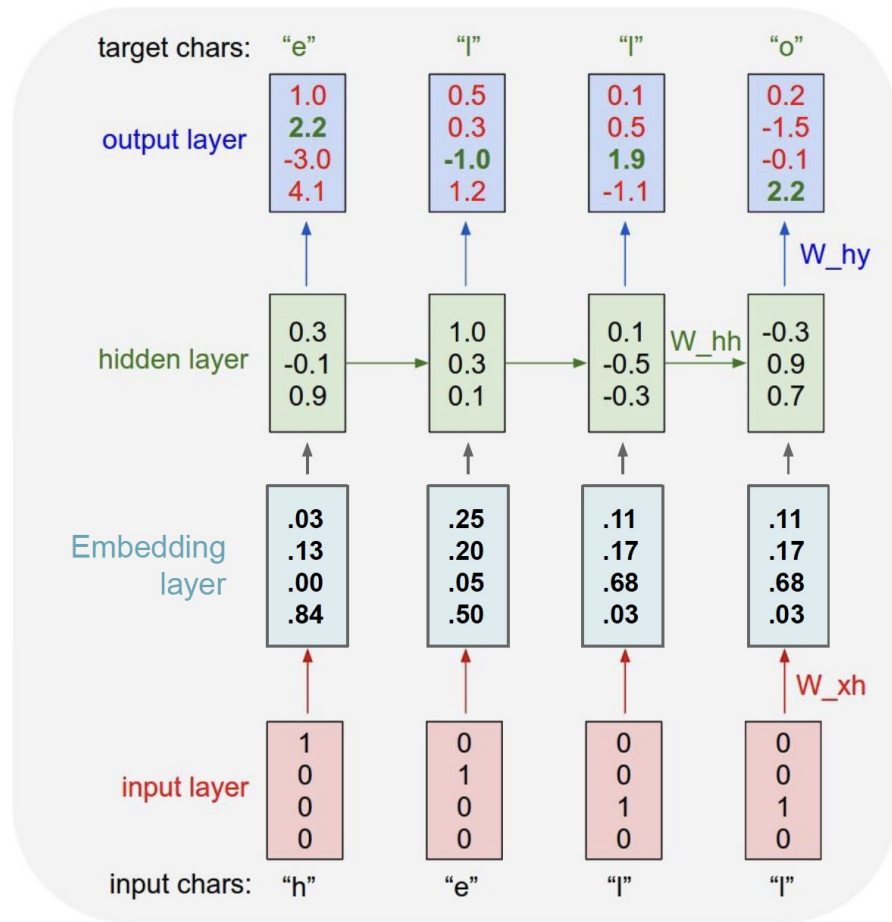
At test-time sample
characters one at a time, feed
back to model



Example: Character-level Language Model

$$\begin{bmatrix} w_{11} & w_{12} & w_{13} & w_{14} \\ w_{21} & w_{22} & w_{23} & w_{24} \\ w_{31} & w_{32} & w_{33} & w_{34} \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} w_{11} \\ w_{21} \\ w_{31} \end{bmatrix}$$

Matrix multiply with a one-hot vector just extracts a column from the weight matrix. We often put a separate **embedding** layer between input and hidden layers.

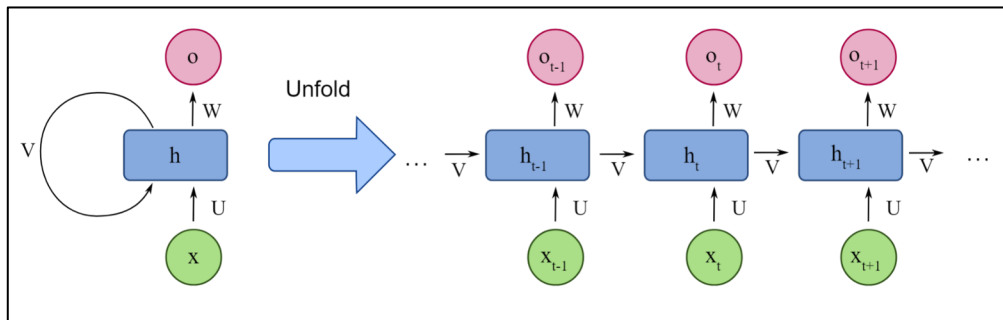


Weaknesses of early NN-based NLP approaches

- ▶ Short context length
- ▶ “Linear” reasoning - no attention mechanism to focus on other parts
- ▶ Earlier approaches (e.g. word2vec) do not adapt based on context.

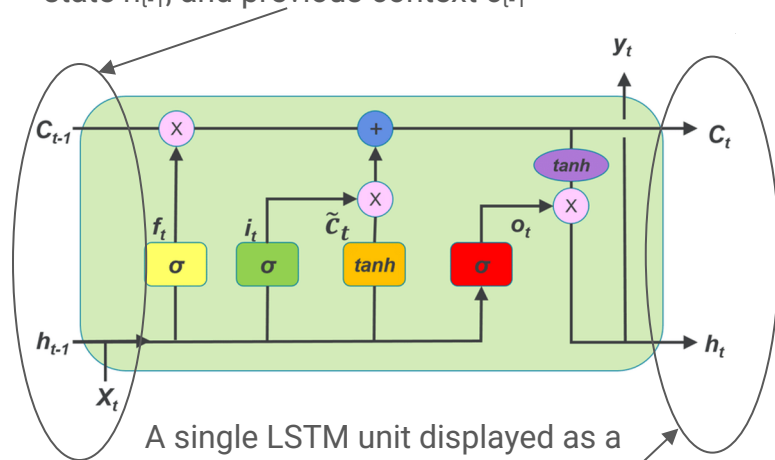
Seq2Seq Models w/ Neural Nets: the Pre-Transformer Era

- Recurrent Neural Networks (RNNs)
- Long Short-Term Memory Networks (LSTMs)
- Capture dependencies between input tokens
- Gates control the flow of information



A simple RNN shown unrolled in time. Network layers are recalculated for each time step, while weights U , V and W are shared across all time steps.

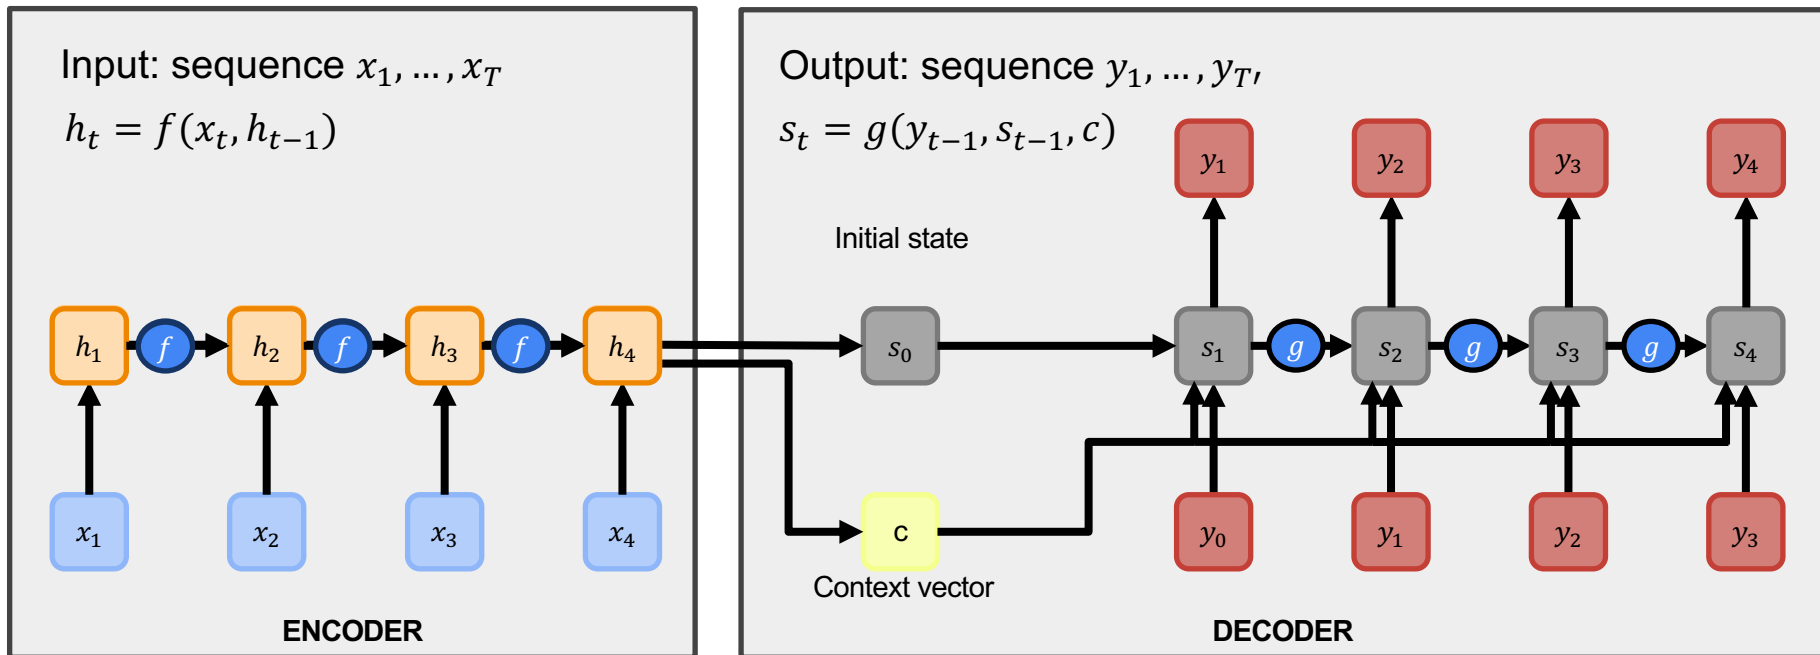
The inputs to each unit consists of the current input x_t , previous hidden state h_{t-1} , and previous context c_{t-1}



The outputs are a new hidden state h_t and an updated context c_t .

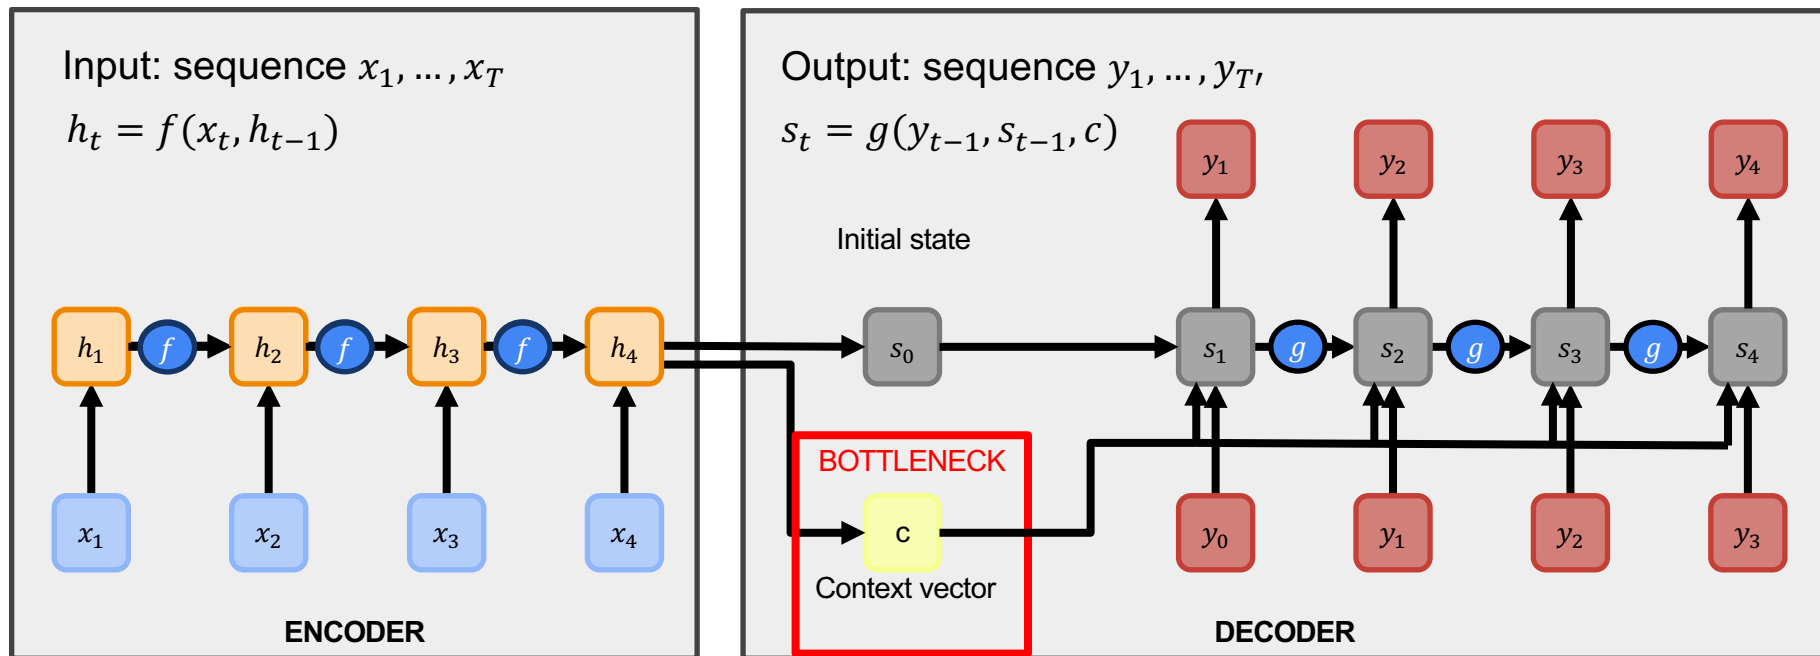
Better Capturing of Long-Range Dependence using LSTM for Seq2Seq Modeling

- ▶ Encoder (LSTM) and decoder (LSTM)
- ▶ Fixed-length context vector

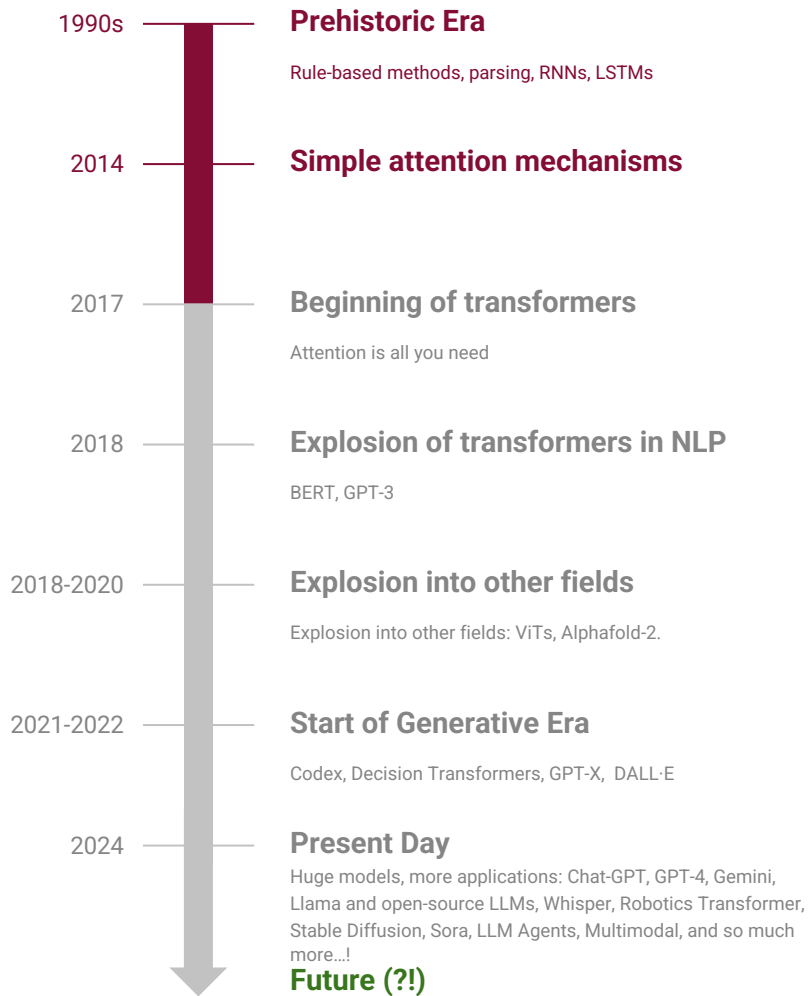


LSTM still suffers from Information Bottleneck

- ▶ Encoder (LSTM) and decoder (LSTM)
- ▶ Fixed-length context vector (**bottleneck**)



Attention Timeline



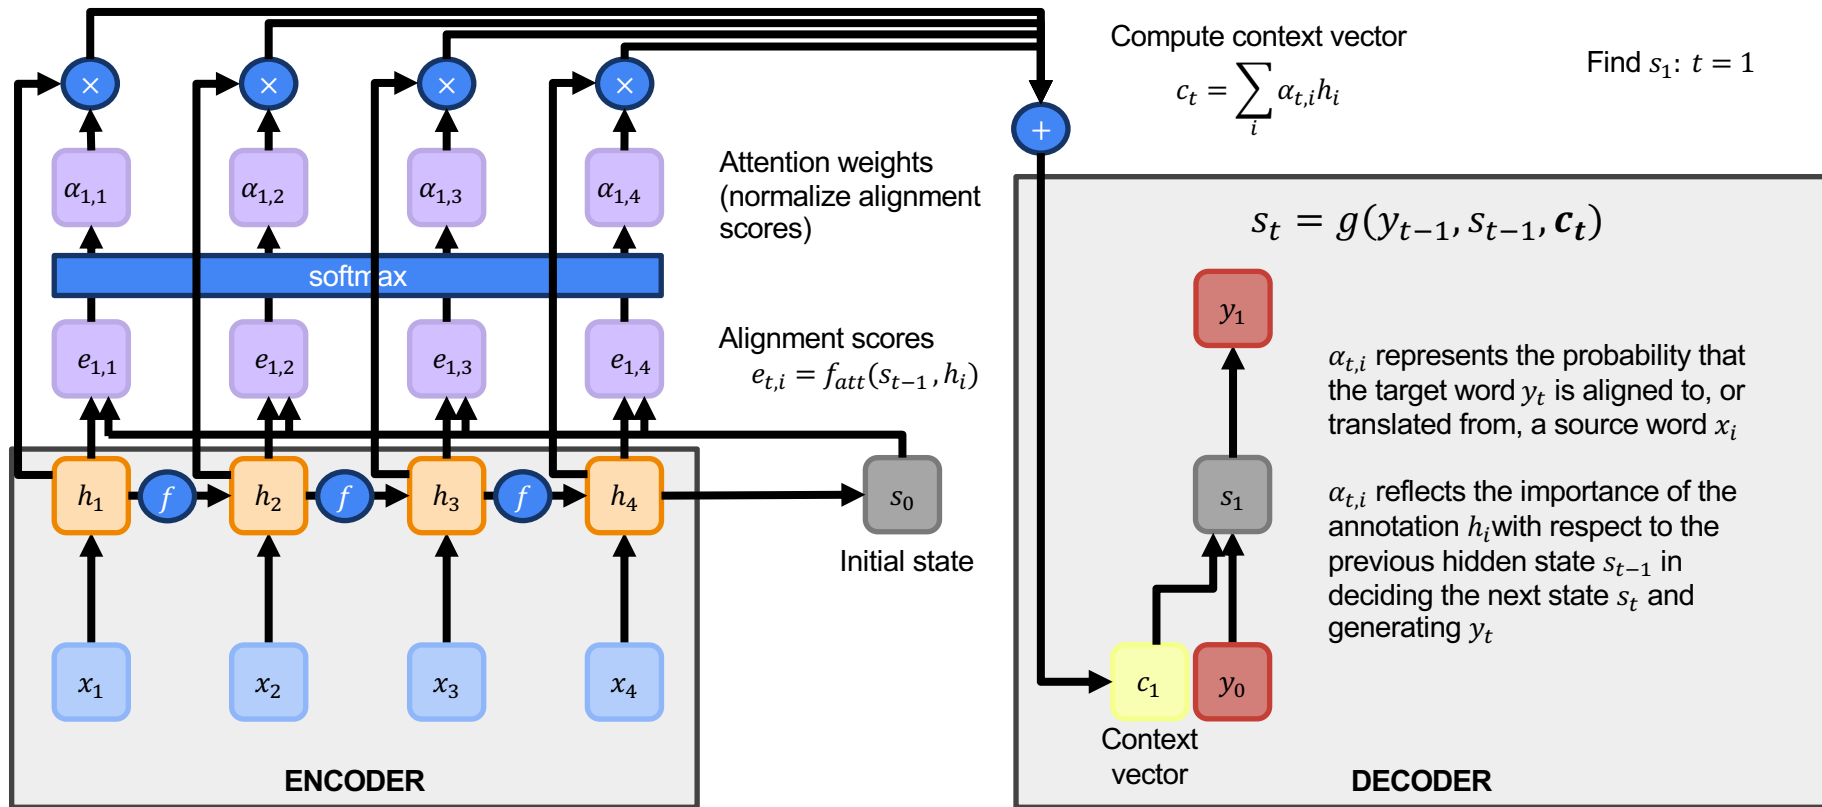
Sequence to Sequence with RNNs + Attention

- ▶ **Idea!** Use a different context vector for each timestep in the decoder

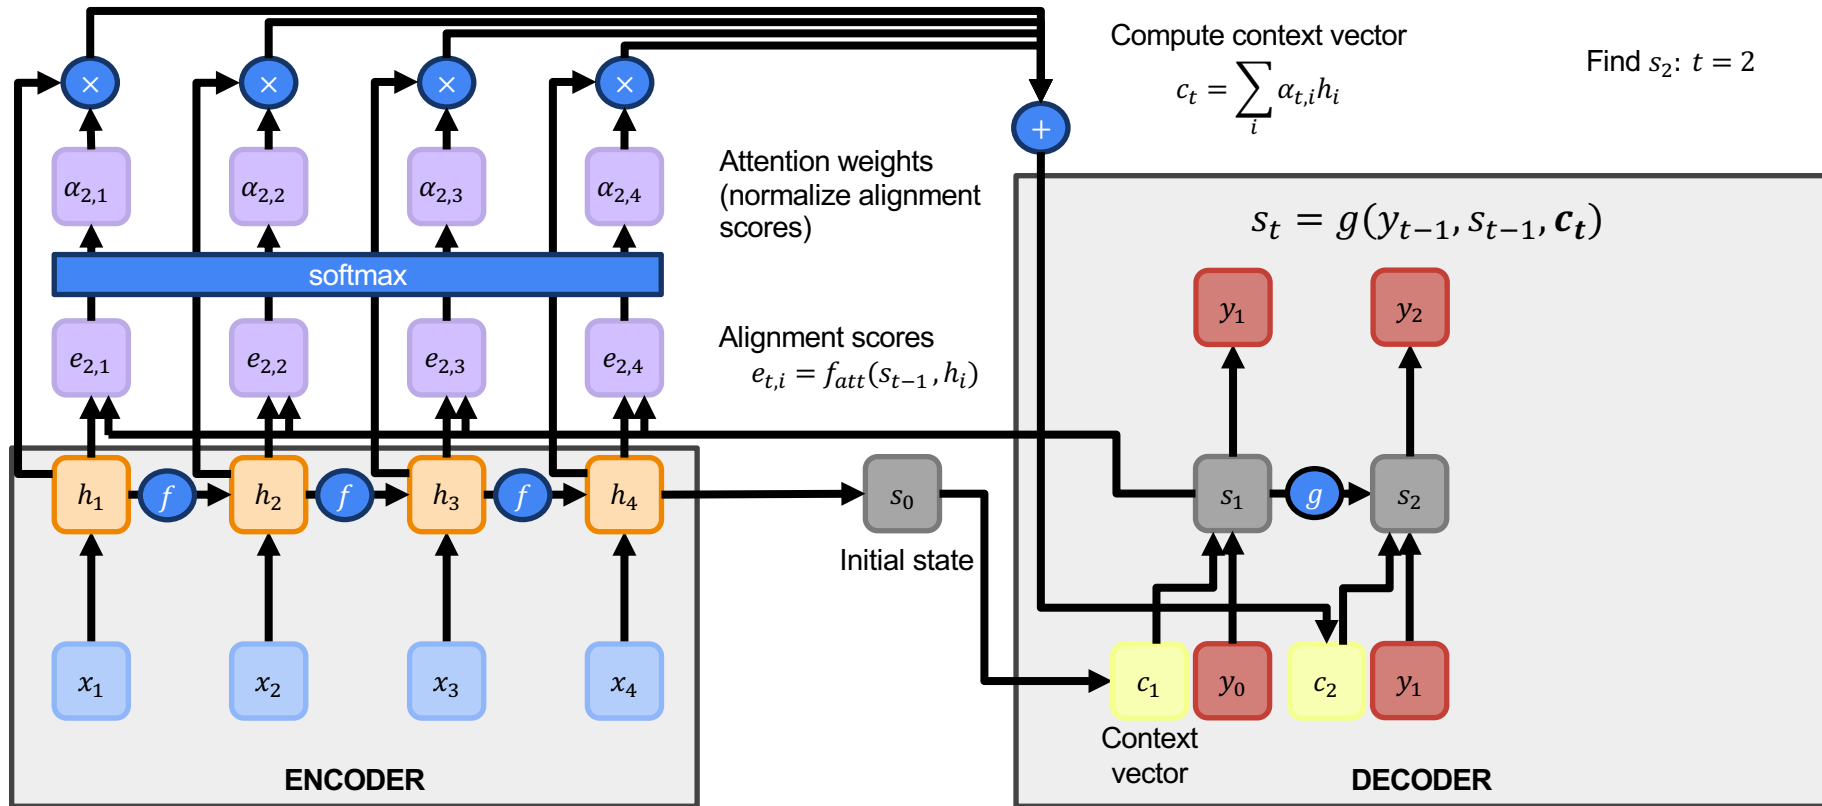
$$s_t = g(y_{t-1}, s_{t-1}, \mathbf{c}_t)$$

- ▶ No more bottleneck through a single vector
- ▶ Craft the context vector so that it “looks at” different parts of the input sequence for each decoder timestep

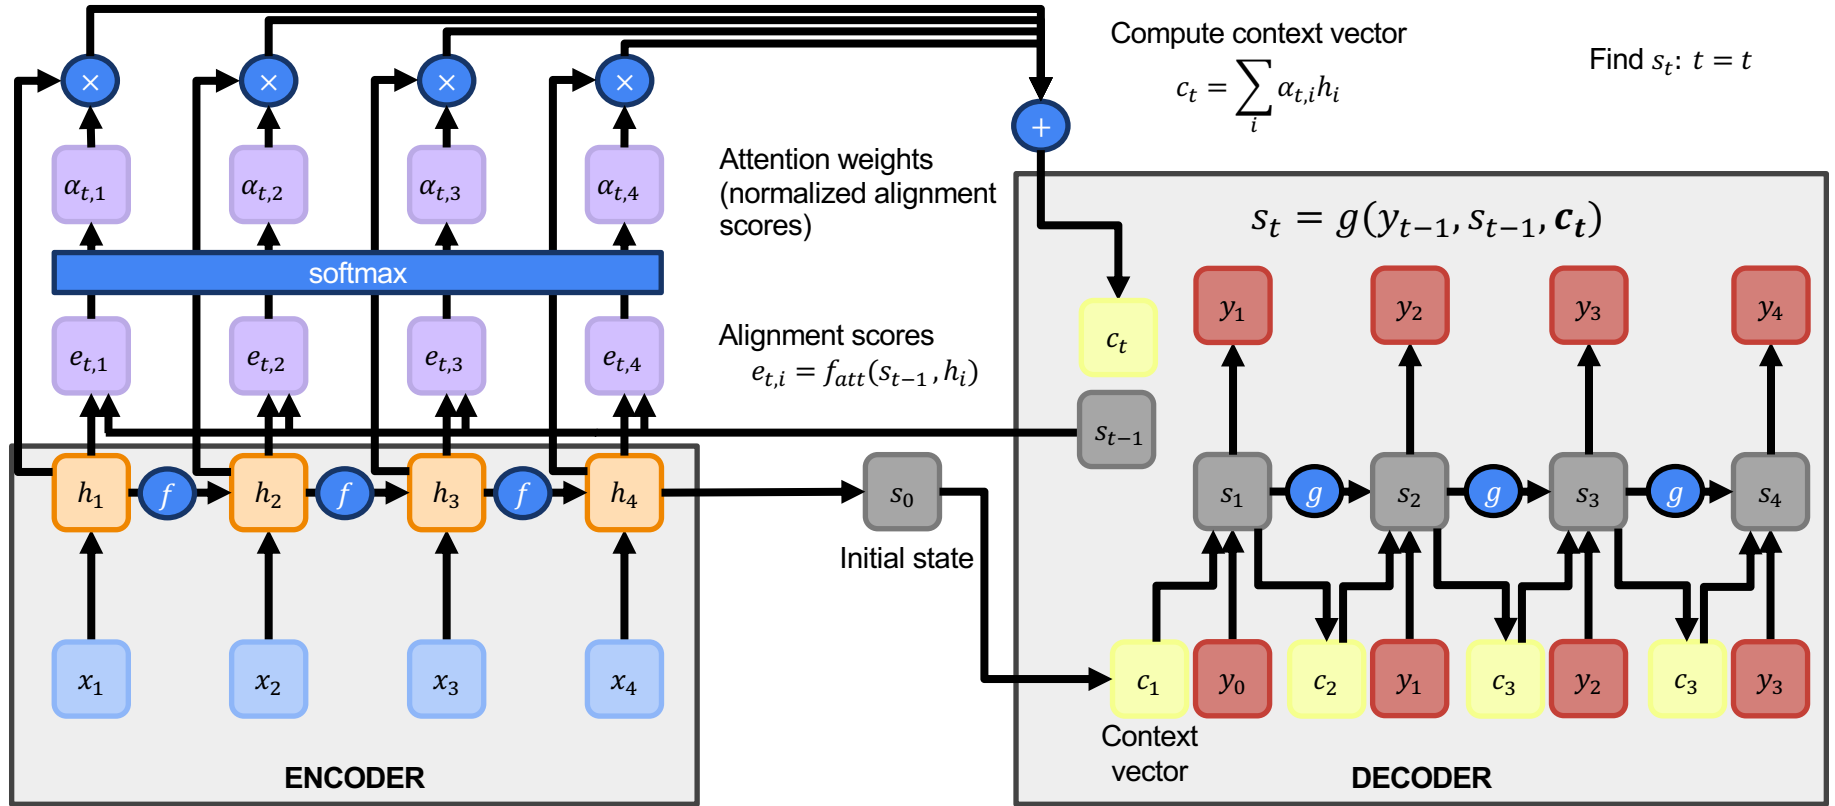
Sequence to Sequence with RNNs + Attention



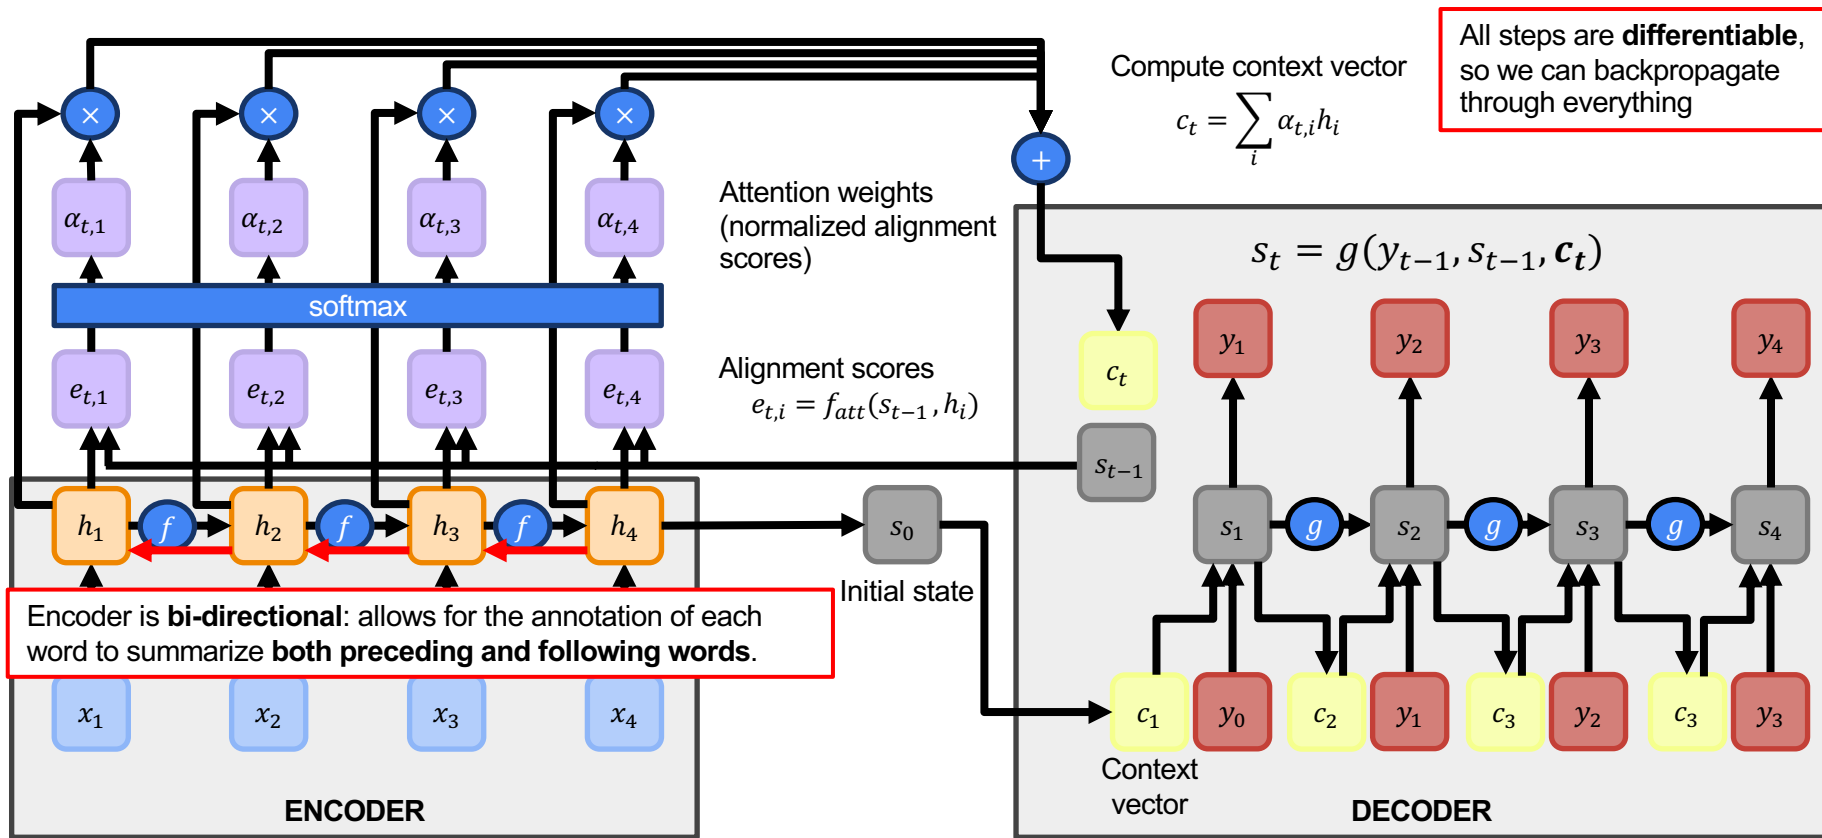
Sequence to Sequence with RNNs + Attention



Sequence to Sequence with RNNs + Attention




Sequence to Sequence with RNNs + Attention



Sequence to Sequence with RNNs + Attention

Application: translation

Each pixel shows the weight $\alpha_{t,i}$ of the annotation of the i -th source word for the t -th target word.




Sequence to Sequence with RNNs + Attention

Application:
text translation

RNN:
RNNenc

RNN + attention:
RNNsearch



Motivating Transformers by Understanding the Limitations of Recurrent Models

Challenge 1: Modeling long-range dependencies

Challenge 2: Optimization due to vanishing and exploding gradients

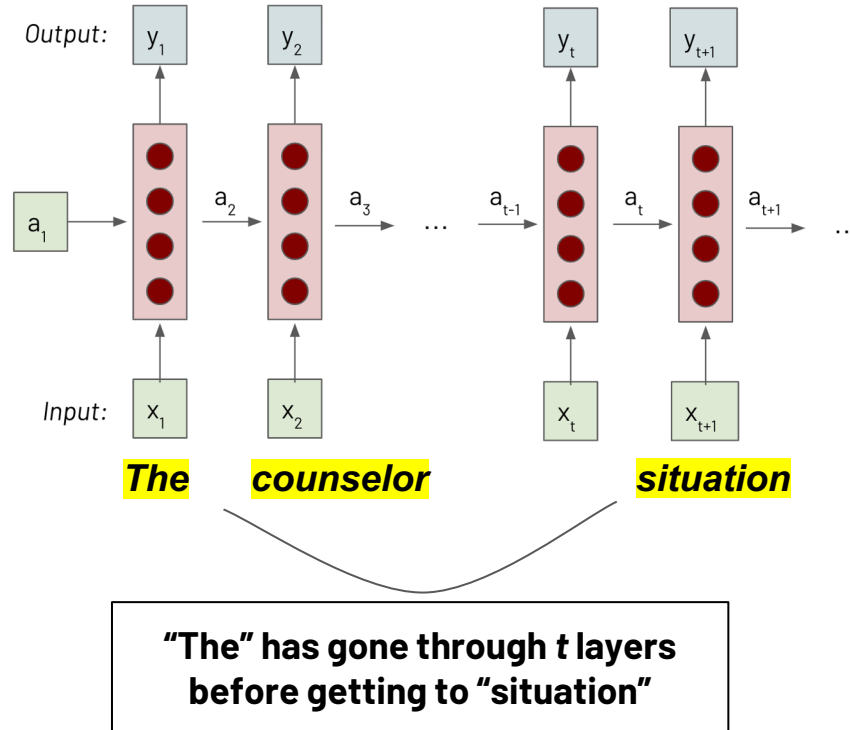
Challenge 3: Slow (sequential, serial) bottleneck

Caveating this discussion: while the challenges we'll discuss originally motivated Transformers, many have continued to make progress on RNNs over the years ([S4](#), [Mamba](#), [Linear Attention](#), [GLA](#), [Based](#), etc.)

Challenge 1: Long Interaction Distances

E.g. “The counselor helped **frame** the situation.”

- Performance degrades as the distance between words increases due to memory constraints: Diluted impact of earlier elements on output as sequence progresses



While RNNs + “Attention” has made some progress, the coupling of the “sequential” structure of RNN with Attention still creates difficulties


Key idea of the

RNN + Attention operation:

All tokens interact with all other tokens' representations

Decoders


Encoders



Challenge 2: RNNs/ LSTMs are difficult to train!

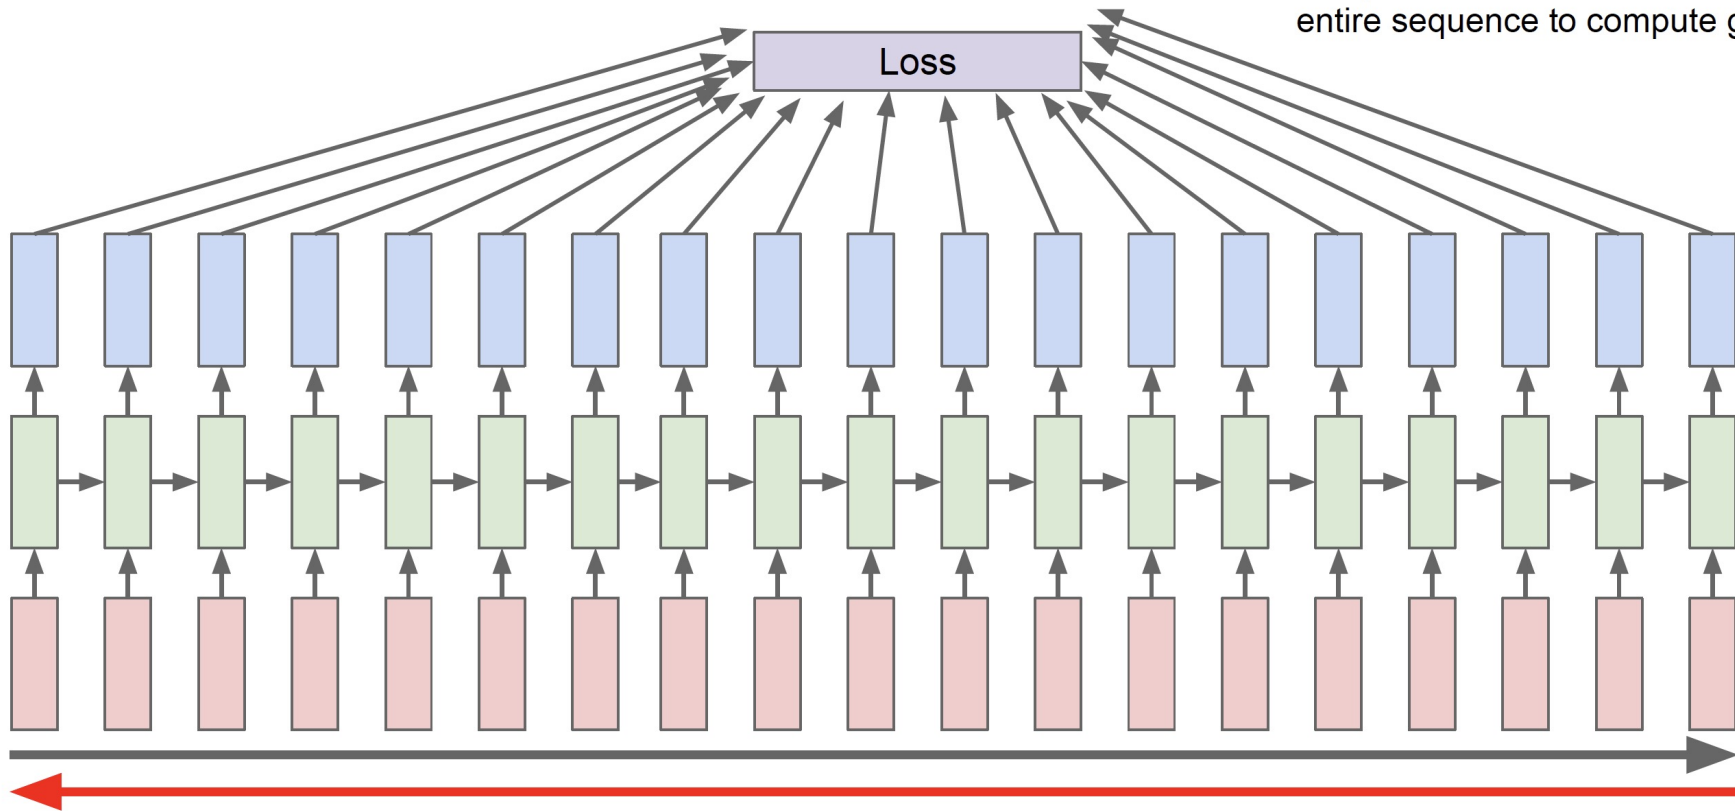
Backpropagation through many timesteps/"layers"...

Recall: backpropagation is about updating the parameters in a way that reduces the loss. We multiply with respect to each set of parameters at each timestep.



Backpropagation through Time

Forward through entire sequence to compute loss, then backward through entire sequence to compute gradient




Challenge 2: RNNs/LSTMs are difficult to train!

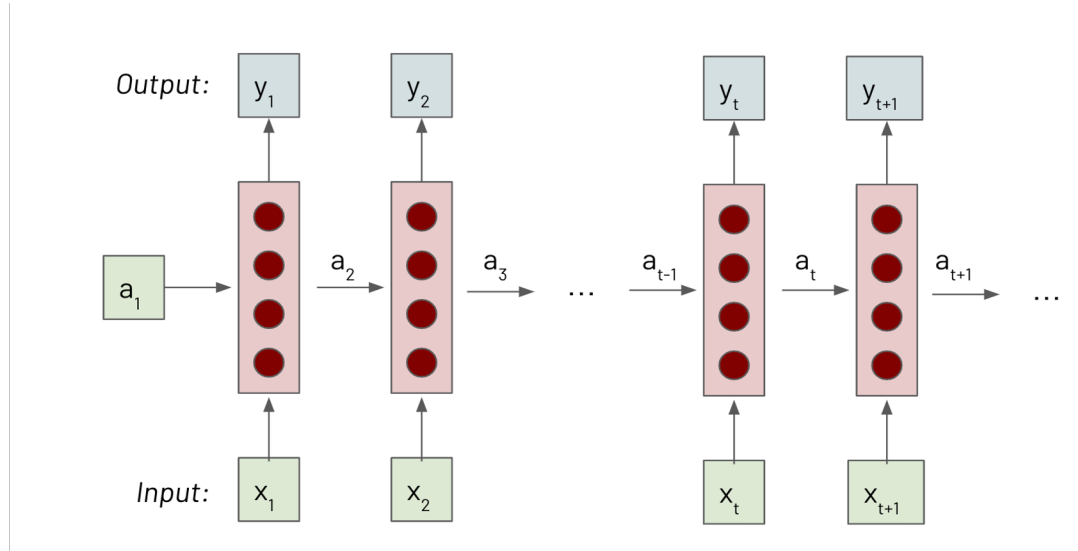
If the value we are multiplying is **large**,
our gradients will grow **exponentially!**

The model becomes **unstable!**

If the value we are multiplying is
small, our gradients will get smaller
each timestep, **going to 0**. The
network **stops learning/learns too
slowly**.



Challenge 3: Parallelizability



Each time step needs to be processed before we can move onto the next step.

Decoupling “Attention” from RNNs


- ▶ **Recall:** attention determines the **importance of elements** to be passed forward in the model.
 - ▶ These weights lets the model pay **attention** to the **most significant parts**
- ▶ **Objective:** a more general attention mechanism not confined to RNNs
 - ▶ We need a modified procedure to:
 1. Determine weights based on context that indicate the elements to attend to
 2. Apply these weights to enhance attended features

Self-Attention and Transformers

- Allows to “focus attention” on particular aspects of the input while generating the output.
- Done by using a set of parameters, called "weights," that determine how much attention should be paid to each input token at each time step.
- These weights are computed using a combination of the input and the current hidden state of the model.

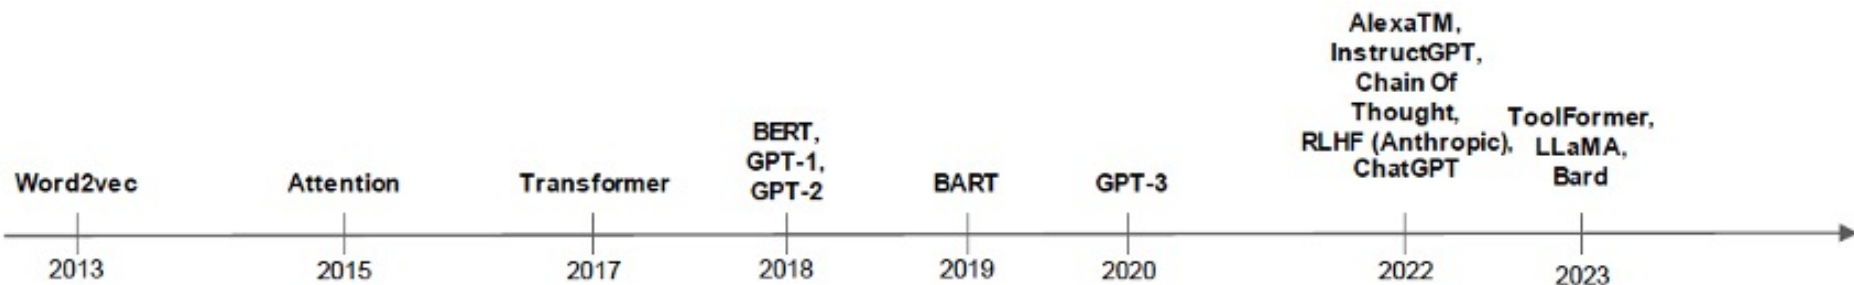
$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

A. Vaswani et al. Attention Is All You Need. NeurIPS 2017.




In encoding the word "it", one attention head is focusing most on "the animal", while another is focusing on "tired". The model's representation of the word "it" thus bakes in some of the representation of both "animal" and "tired".
<https://jalamar.github.io/illustrated-transformer/>

Transformers – the current “standard” for building LLMs and Foundation Models



Slides for video from:
Prof. Jia-Bin Huang
University of Maryland, College Park

<https://www.youtube.com/watch?v=rcWMRA9E5RI>



What?



How?




Why?



Sequence-to-Sequence model


En

How are you?



<start>


你好吃飽嗎 ... ? <end>



ZH


En

How are you?



<start> 你


你好吃飽嗎 ... ? <end>



ZH


En

How are you?



<start> 你


你好吃飽嗎 ... ? <end>



ZH


En

How are you?



transformer


<start> 你好



ZH

En


How are you?



transformer

<start> 你好


你好吃飽嗎...? <end>



ZH

En


How are you?



transformer

<start> 你好嗎


你 好 吃 飽 嗎 ... ? <end>



ZH

En


How are you?



<start> 你 好 嗎


En

How are you?



transformer


<start> 你好嗎？



ZH

En

How are you?




transformer

<start> 你好嗎？


你 好 吃 飽 嗎 ... ? <end>

ZH

AuTo-regressive



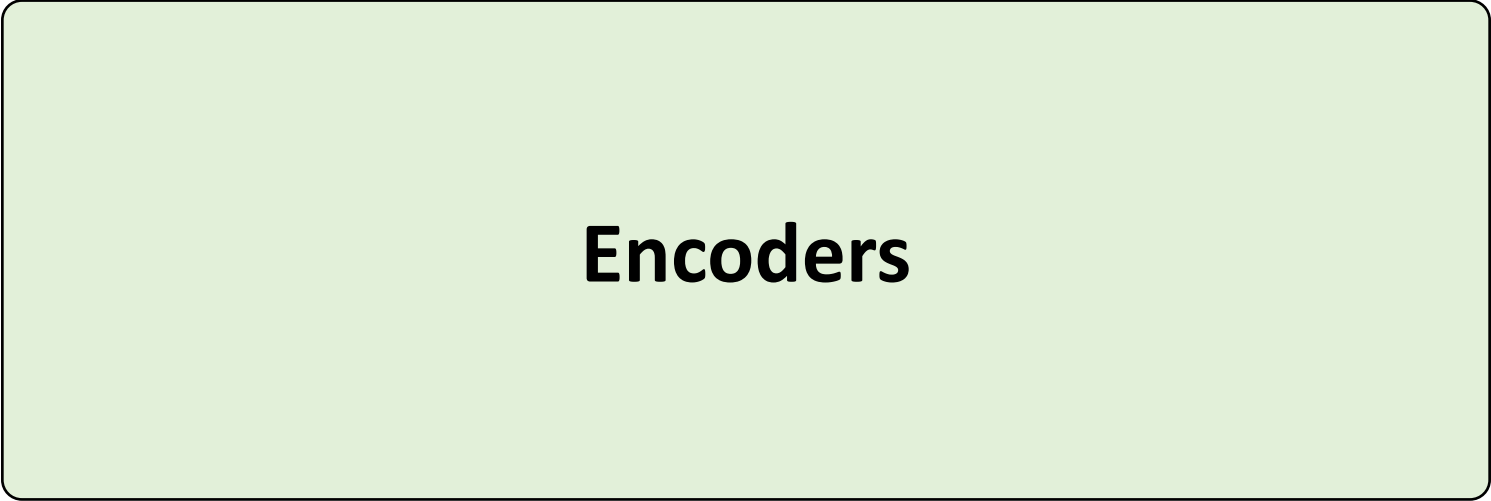
ZH



En

How are you?

<start> 你好嗎? <end>





Encoders




How are you?

Tokenization



One-hot encoding



TOKEN EMBEDDING

One-hot encoding



cat

$$\begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

dog

$$\begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

bear

$$\begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

cow

$$\begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$


indiv

$$\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

Value 1 at
3687th
entry



TOKEN EMBEDDING



Embedding Space

cat

$$\begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

dog

$$\begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

bear

$$\begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

cow


$$\begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

indiv

$$\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

Value 1 at
3687th
entry

TOKEN EMBEDDING



Embedding Space

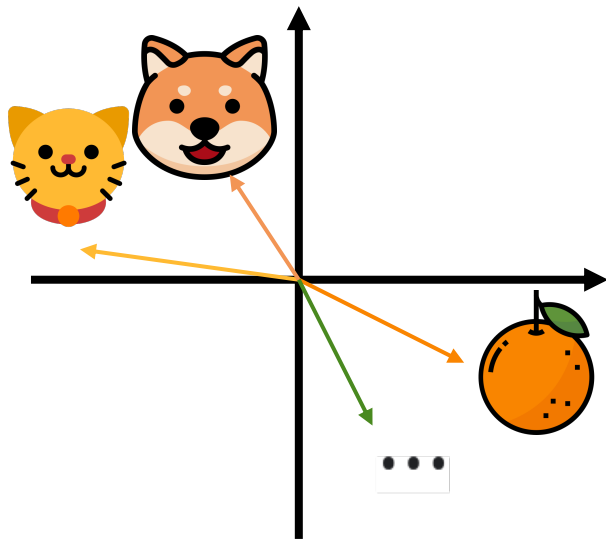
$$d \begin{bmatrix} 0.5 \\ 2.7 \\ 1.2 \\ \vdots \\ 0.2 \end{bmatrix}$$

Embedded token


$$= d \begin{bmatrix} \leftarrow \# \text{ tokens} \rightarrow \\ \mathbf{W} \mathbf{E} \end{bmatrix} \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \text{ dog}$$

Embedding Matrix

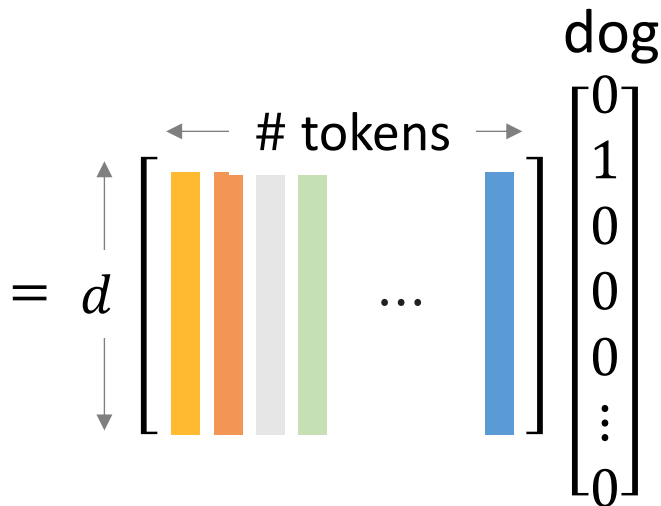
TOKEN EMBEDDING



Embedding Space




Embedded token



Embedding Matrix

TOKEN EMBEDDING



Embedding Space

Apple

I bought an **apple** and an orange.

I bought an **apple** watch.

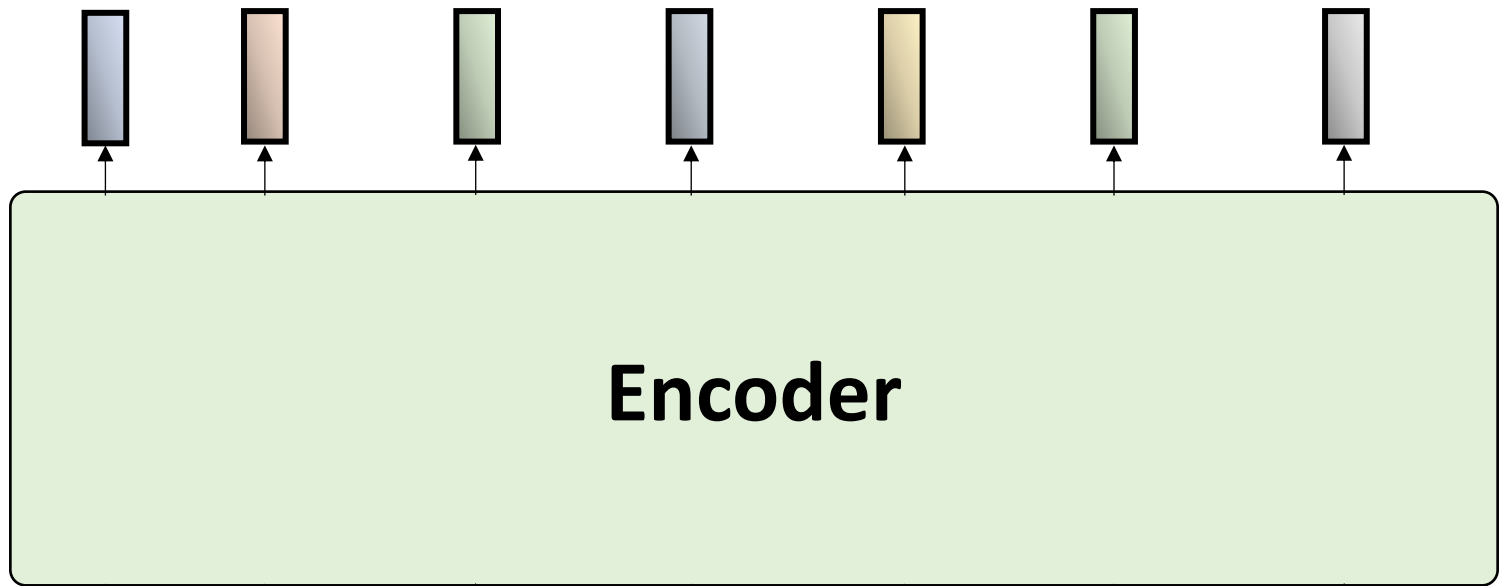
$$d \begin{bmatrix} 0.5 \\ 2.7 \\ 1.2 \\ \vdots \\ 0.2 \end{bmatrix}$$

Embedded token

$$= d \begin{bmatrix} \text{orange} & \text{apple} & \text{watch} & \dots & \text{dog} \end{bmatrix} \begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

tokens

Embedding Matrix



Embedded
Tokens

Token
Embedding

Tokens

I

bought

an

apple

and

an

orange

Encoder

W_E

W_E

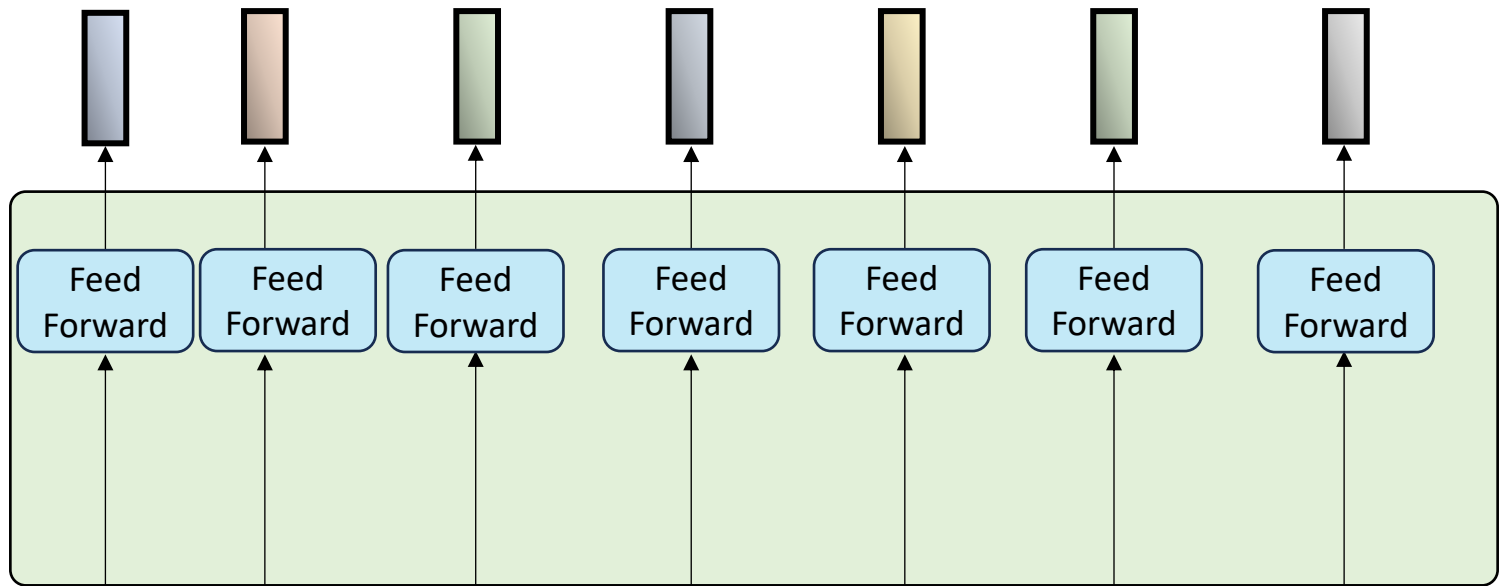
W_E

W_E

W_E

W_E

W_E

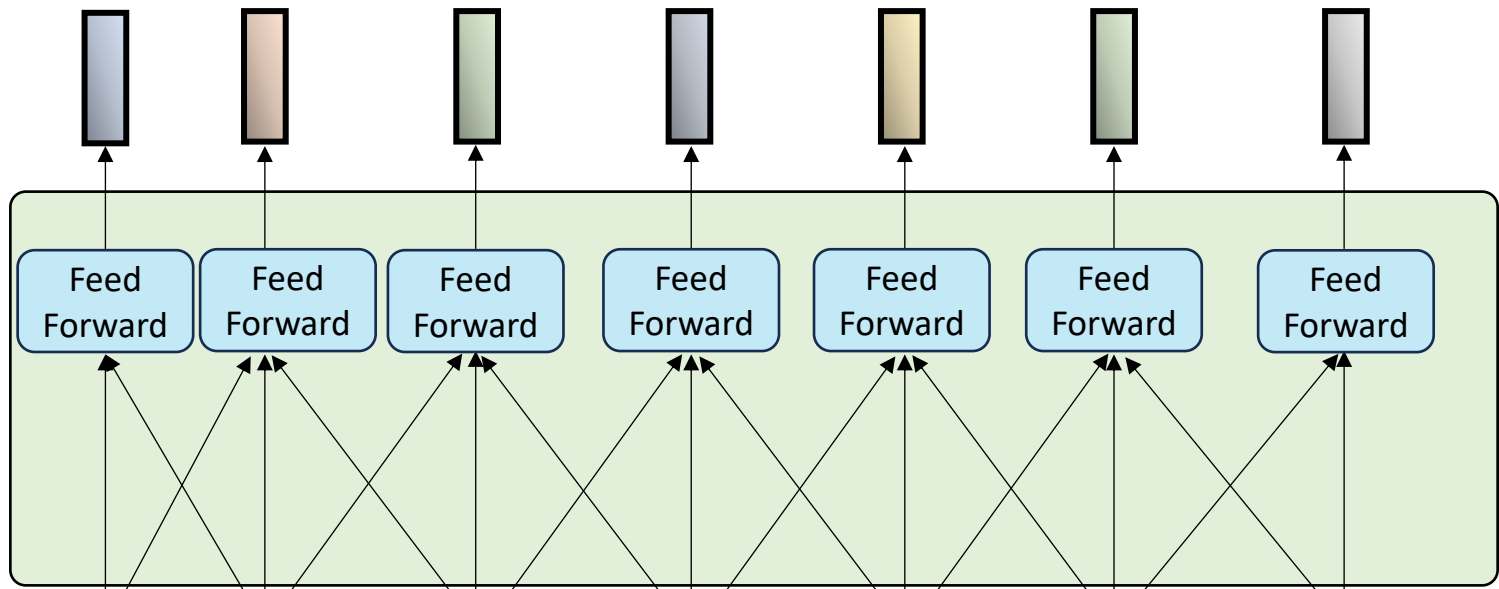


Embedded
Tokens

Token
Embedding

Tokens


I **bought** **an** **apple** **and** **an** **orange**



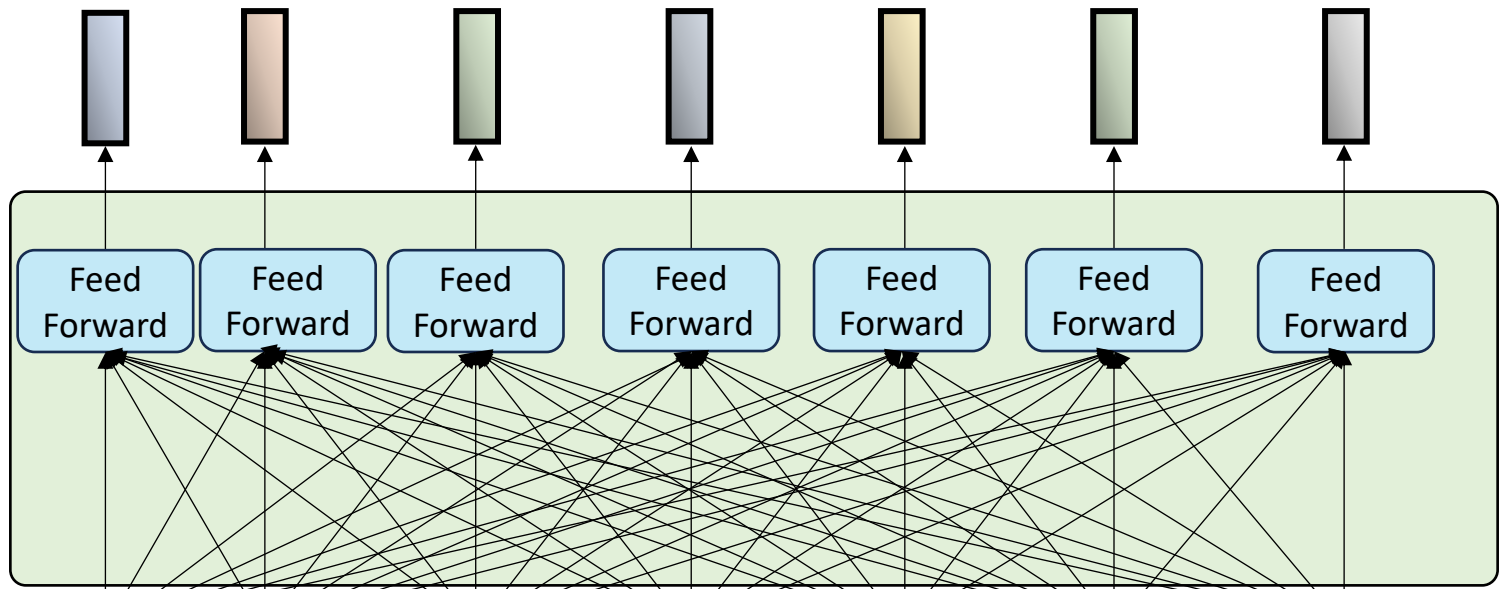
Embedded
Tokens

Token
Embedding

Tokens




I **bought** **an** **apple** **and** **an** **orange**



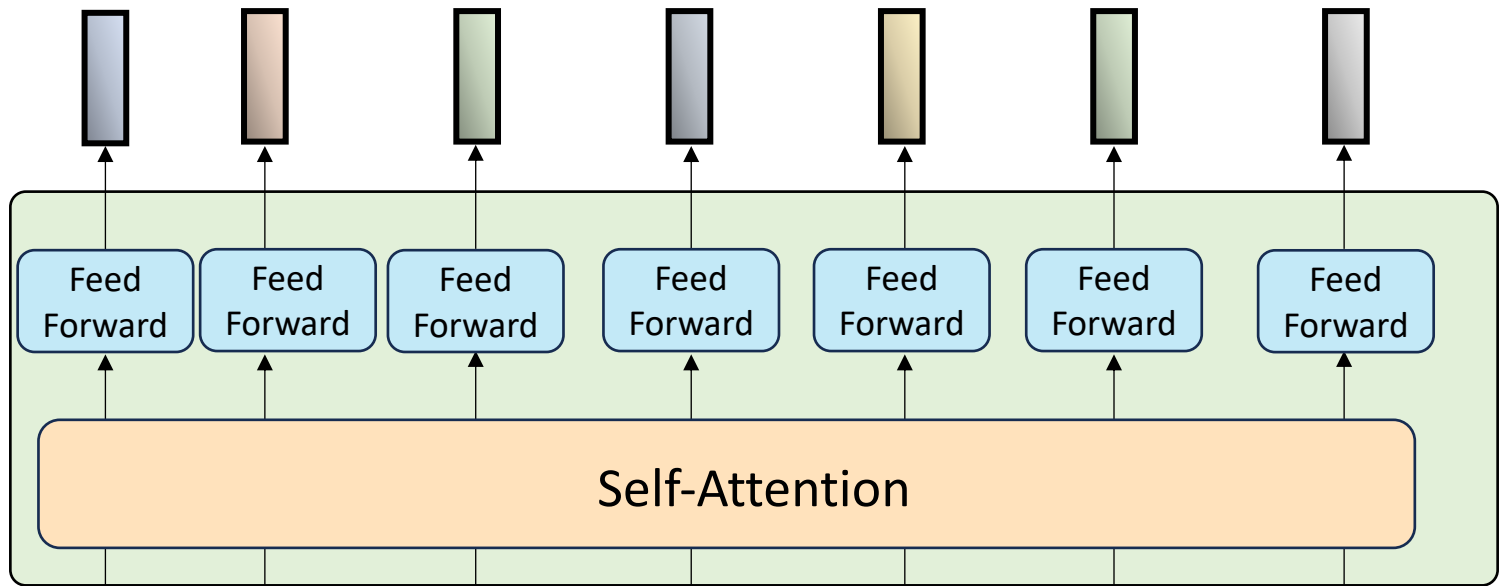
Embedded
Tokens

Token
Embedding

Tokens



I **bought** **an** **apple** **and** **an** **orange**




Embedded
Tokens

Token
Embedding

Tokens



I **bought** **an** **apple** **and** **an** **orange**

Self-Attention



Embedded
Tokens

Tokens



I

bought

an


apple

and

an




orange

Self-Attention



Embedded
Tokens

Tokens



I

bought

an


apple

and

an





orange

Self-Attention



Embedded
Tokens

Tokens



I


bought

an

apple



watch

Self-Attention



Embedded
Tokens

Tokens



I


bought

an


apple

watch

Self-Attention



Embedded
Tokens



Tokens

I


bought

an

apple

watch

Self-Attention



Embedding Space

Embedded
Tokens



Tokens

I

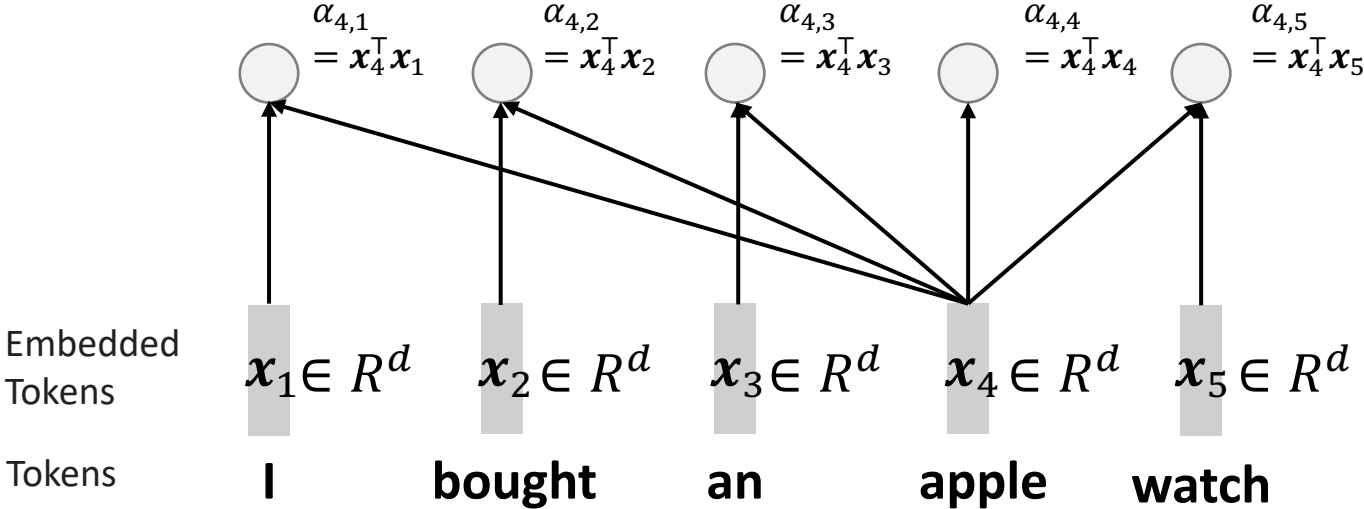
bought

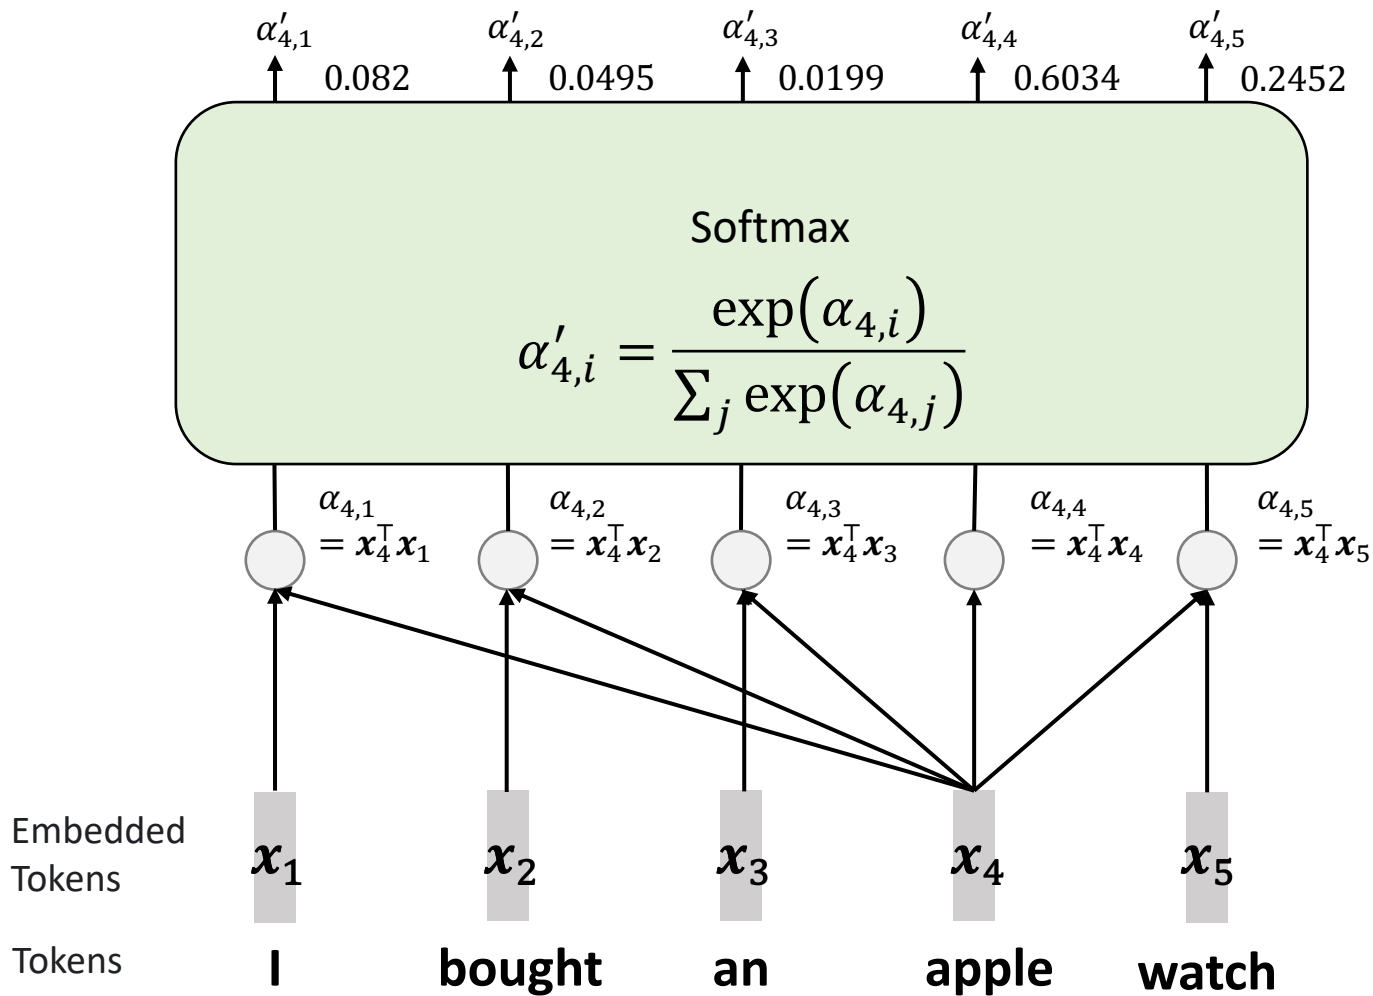
an

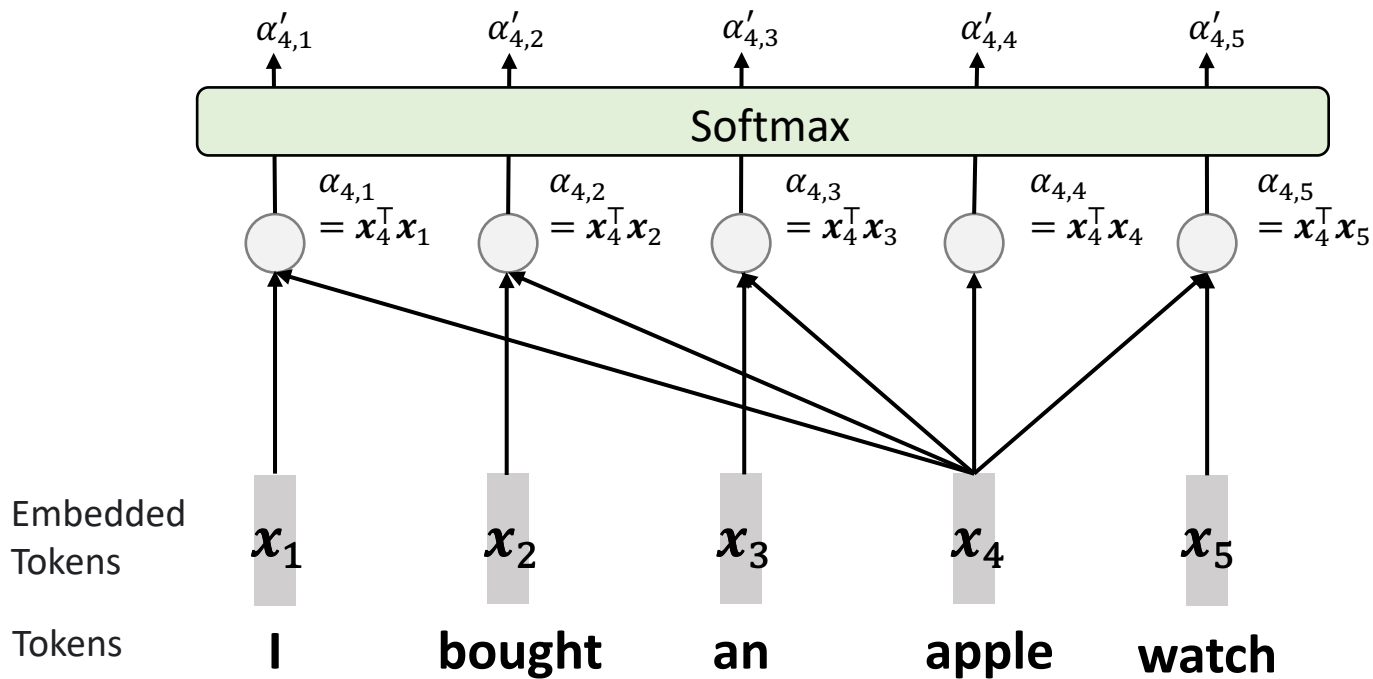
apple

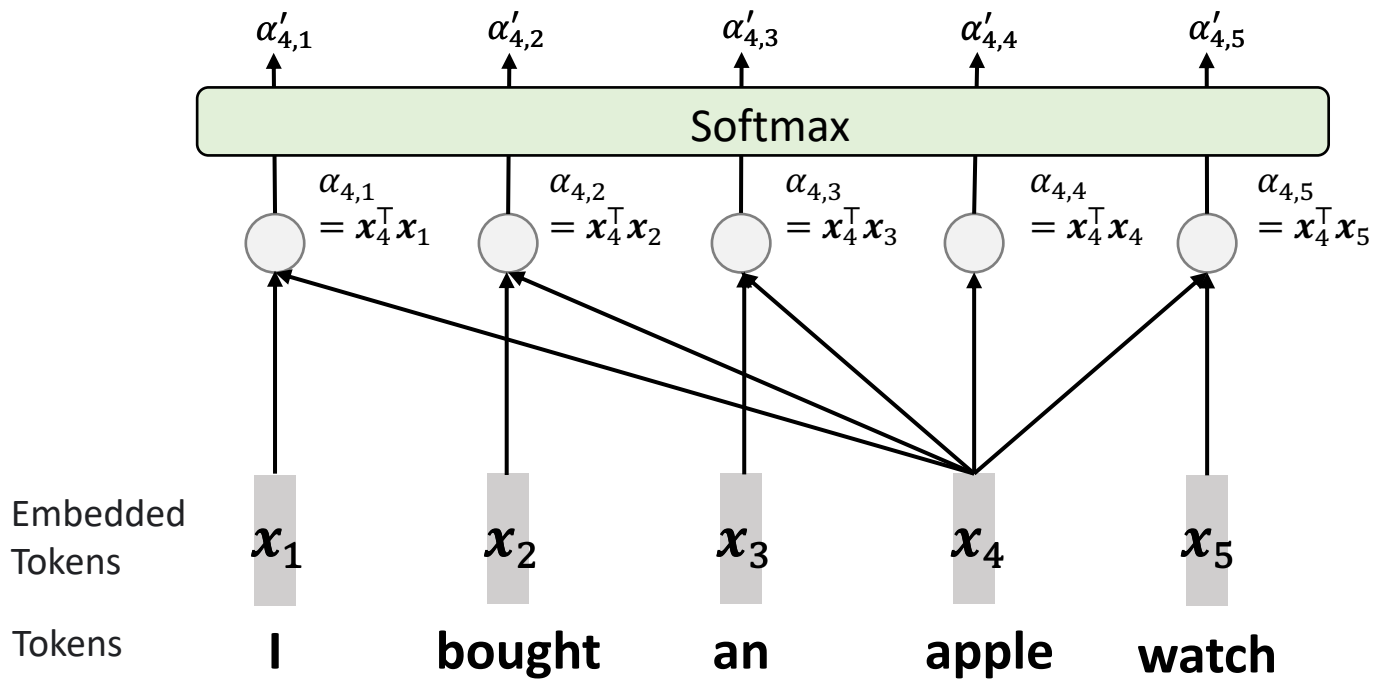
watch

Self-Attention



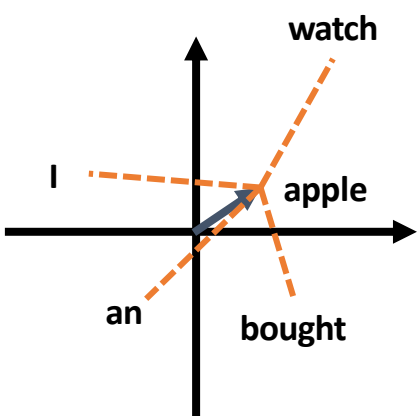






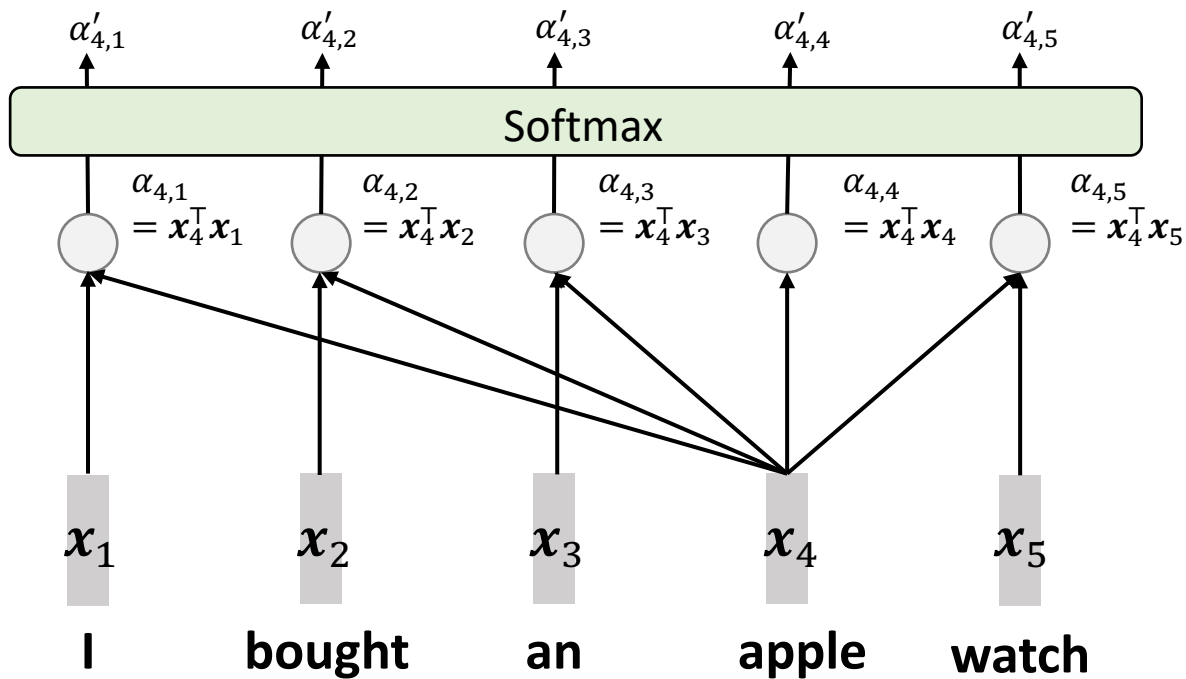
Updated
feature

$$\mathbf{x}'_4 = \alpha'_{4,1}\mathbf{x}_1 + \alpha'_{4,2}\mathbf{x}_2 + \alpha'_{4,3}\mathbf{x}_3 + \alpha'_{4,4}\mathbf{x}_4 + \alpha'_{4,5}\mathbf{x}_5$$



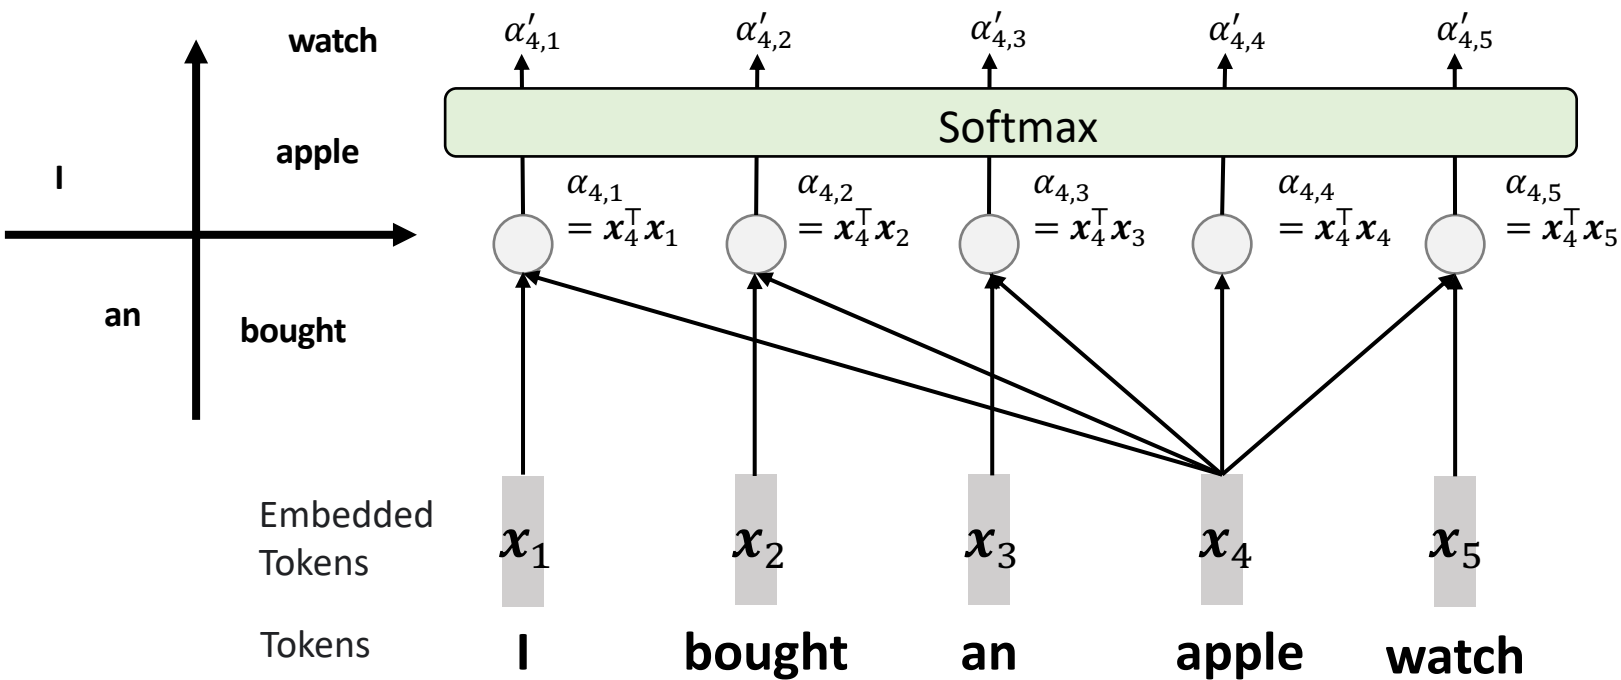
Embedded
Tokens

Tokens



Updated
feature

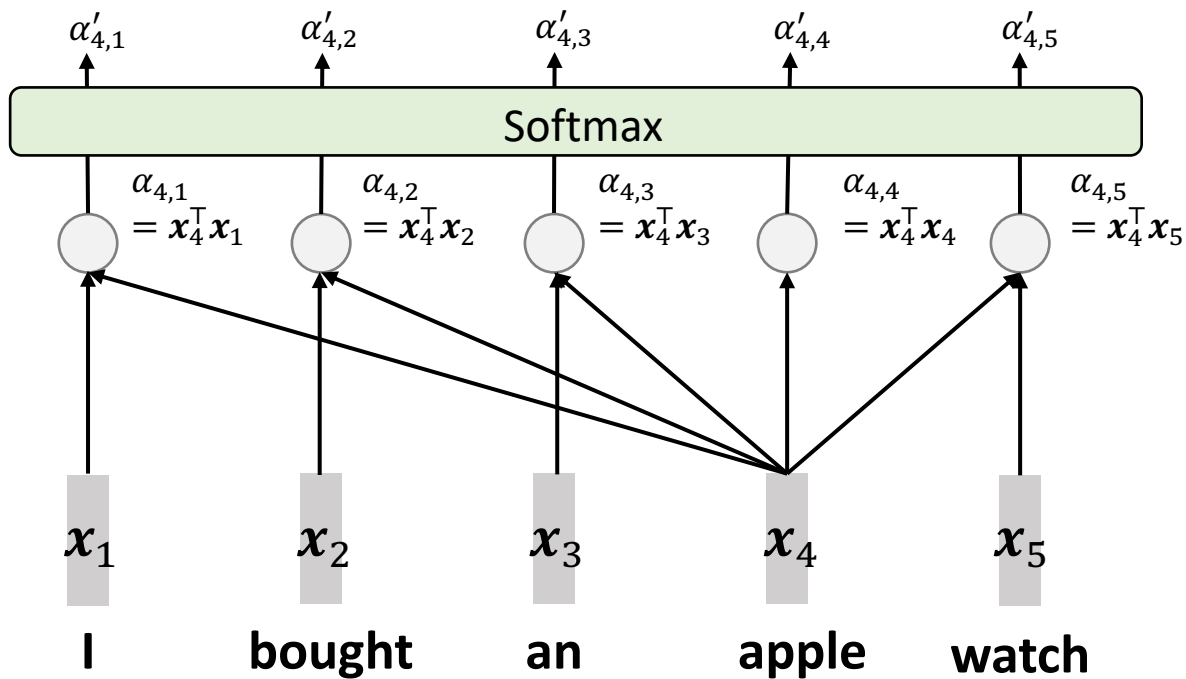
$$\mathbf{x}'_4 = \alpha'_{4,1}\mathbf{x}_1 + \alpha'_{4,2}\mathbf{x}_2 + \alpha'_{4,3}\mathbf{x}_3 + \alpha'_{4,4}\mathbf{x}_4 + \alpha'_{4,5}\mathbf{x}_5$$

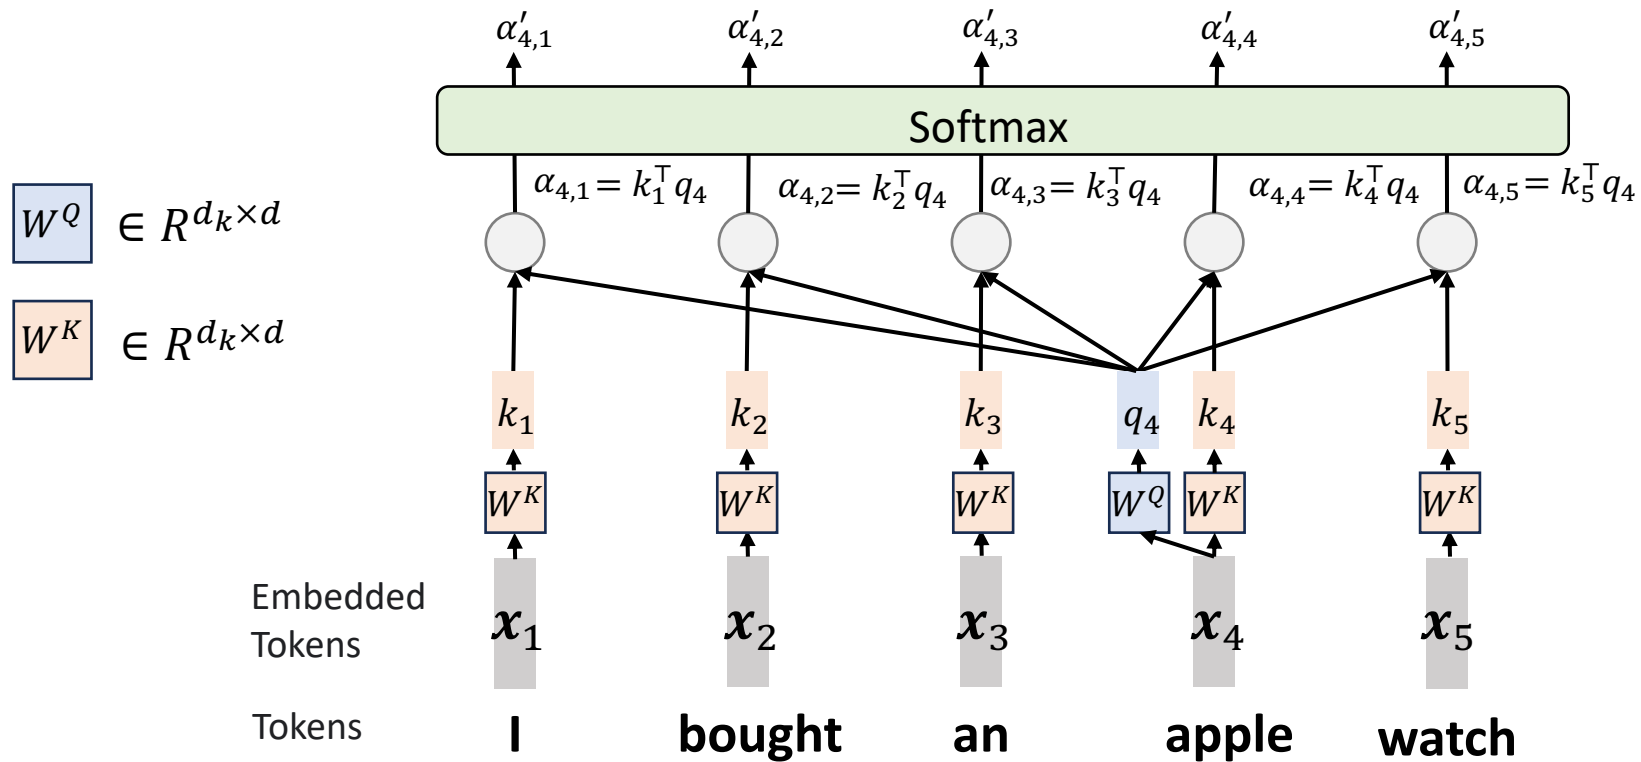


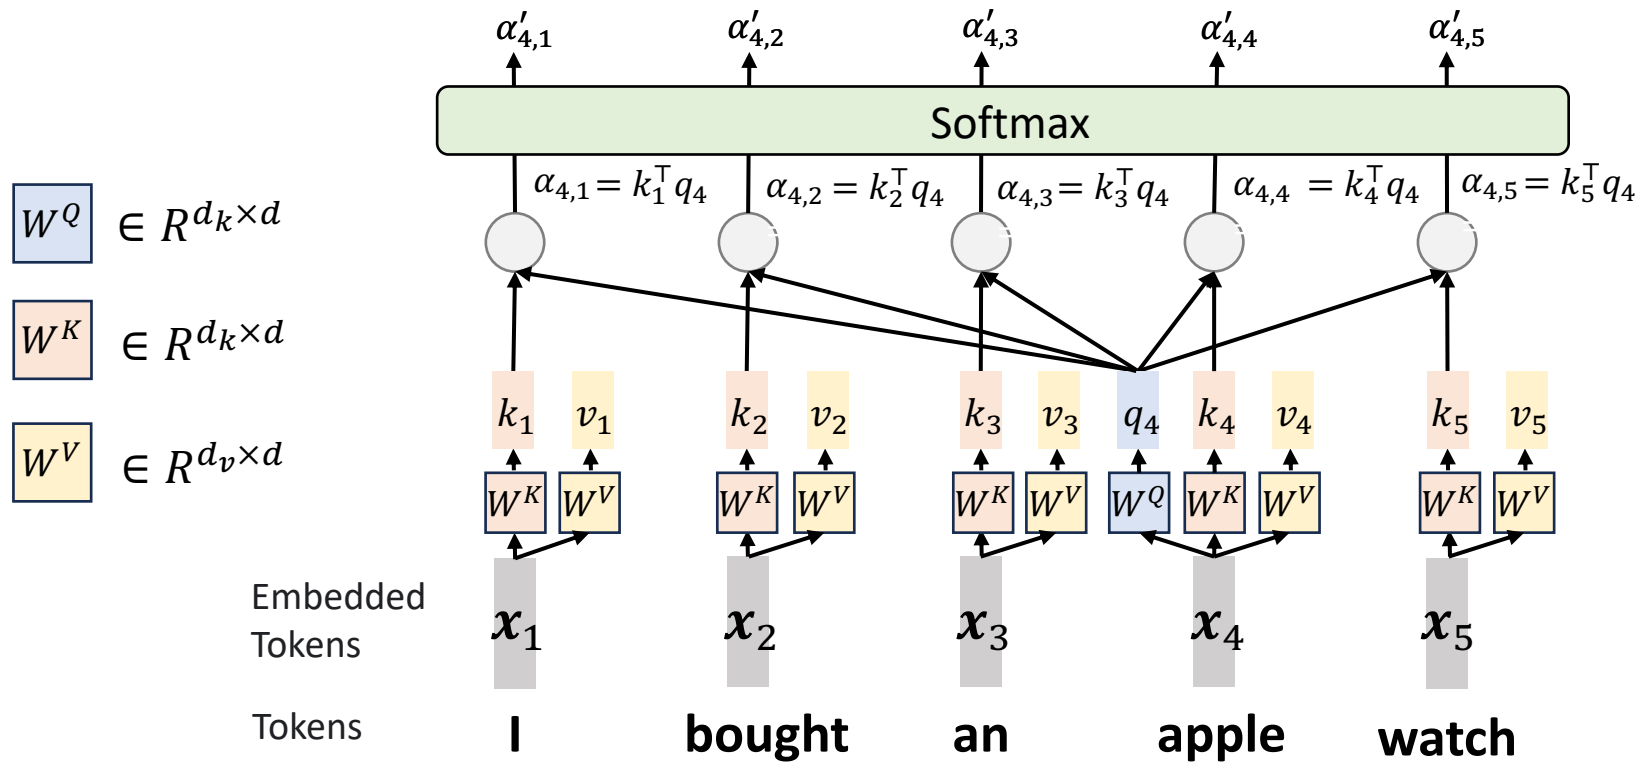
Updated
feature

$$\mathbf{x}'_4 = \alpha'_{4,1}\mathbf{x}_1 + \alpha'_{4,2}\mathbf{x}_2 + \alpha'_{4,3}\mathbf{x}_3 + \alpha'_{4,4}\mathbf{x}_4 + \alpha'_{4,5}\mathbf{x}_5$$

delicious apple

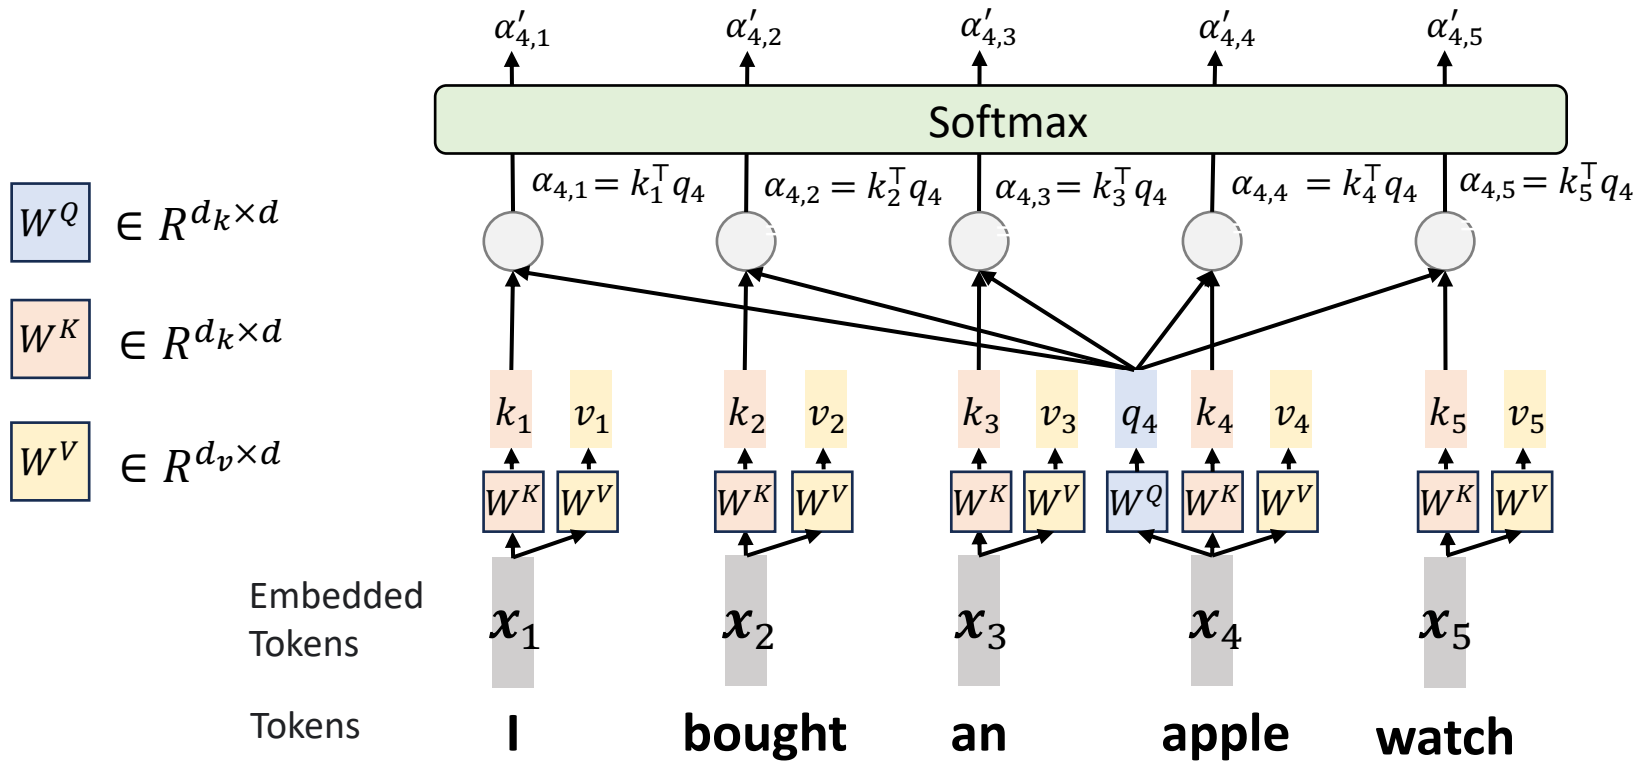






Updated feature

$$\mathbf{x}'_4 = \alpha'_{4,1} v_1 + \alpha'_{4,2} v_2 + \alpha'_{4,3} v_3 + \alpha'_{4,4} v_4 + \alpha'_{4,5} v_5$$



Updated feature

$$\begin{aligned} \mathbf{x}'_4 &= \boxed{W^O} (\alpha'_{4,1} v_1 + \alpha'_{4,2} v_2 + \alpha'_{4,3} v_3 + \alpha'_{4,4} v_4 + \alpha'_{4,5} v_5) \\ &= \sum_i \alpha'_{4,i} \boxed{W^O} \boxed{W^V} \mathbf{x}_i \end{aligned}$$

$$\boxed{W^O} \in R^{d \times d_v}$$

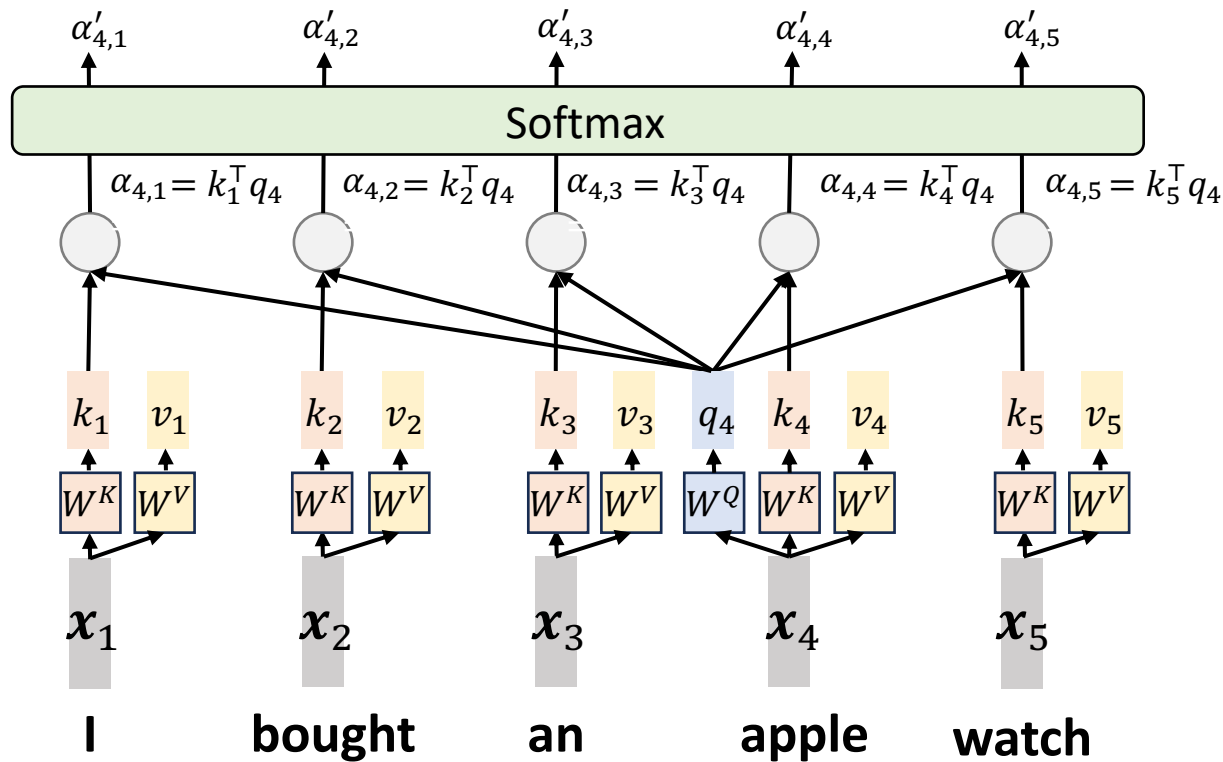
$$\boxed{W^Q} \in R^{d_k \times d}$$

$$\boxed{W^K} \in R^{d_k \times d}$$

$$\boxed{W^V} \in R^{d_v \times d}$$

Embedded Tokens

Tokens



Updated feature

$$\begin{aligned} \mathbf{x}'_4 &= \boxed{W^O} (\alpha'_{4,1} v_1 + \alpha'_{4,2} v_2 + \alpha'_{4,3} v_3 + \alpha'_{4,4} v_4 + \alpha'_{4,5} v_5) \\ &= \sum_i \alpha'_{4,i} \left(\boxed{W^O} \boxed{W^V} \right) \mathbf{x}_i \end{aligned}$$

$$\boxed{W^O} \in R^{d \times d_v}$$

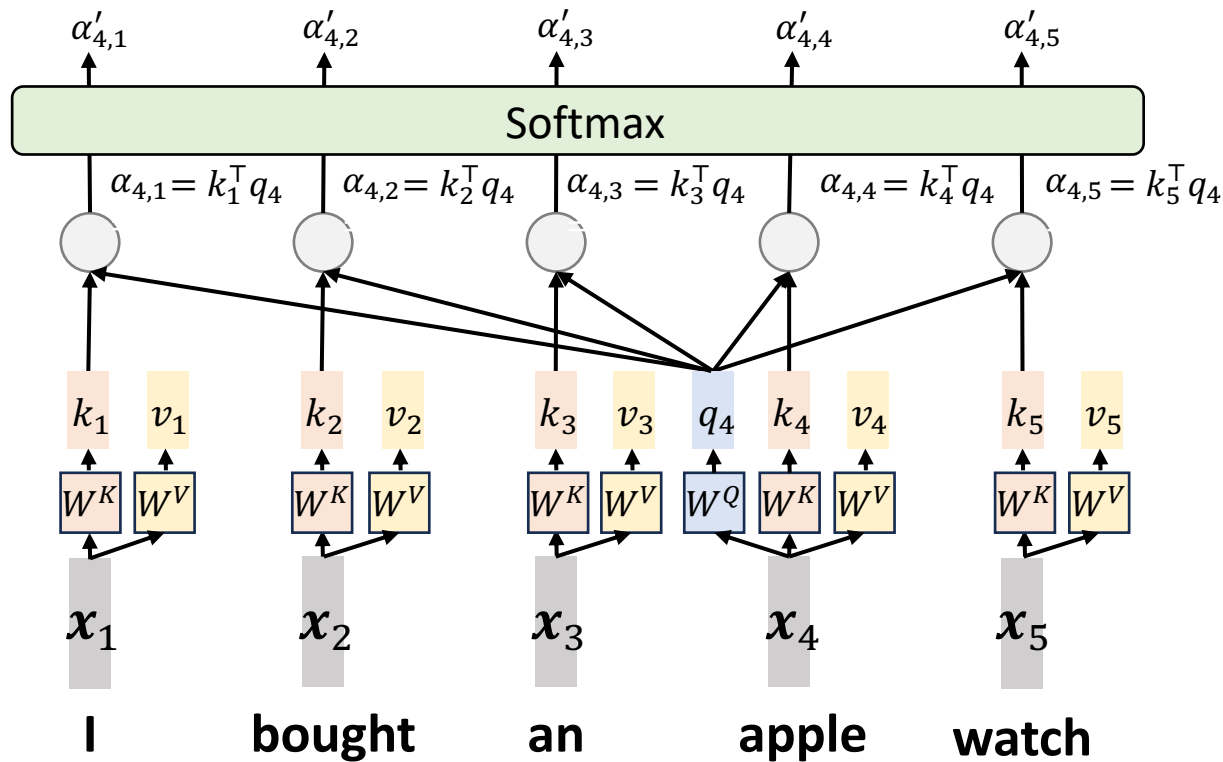
$$\boxed{W^Q} \in R^{d_k \times d}$$

$$\boxed{W^K} \in R^{d_k \times d}$$

$$\boxed{W^V} \in R^{d_v \times d}$$

Embedded Tokens

Tokens



$$W^O \in R^{d \times d_v}$$

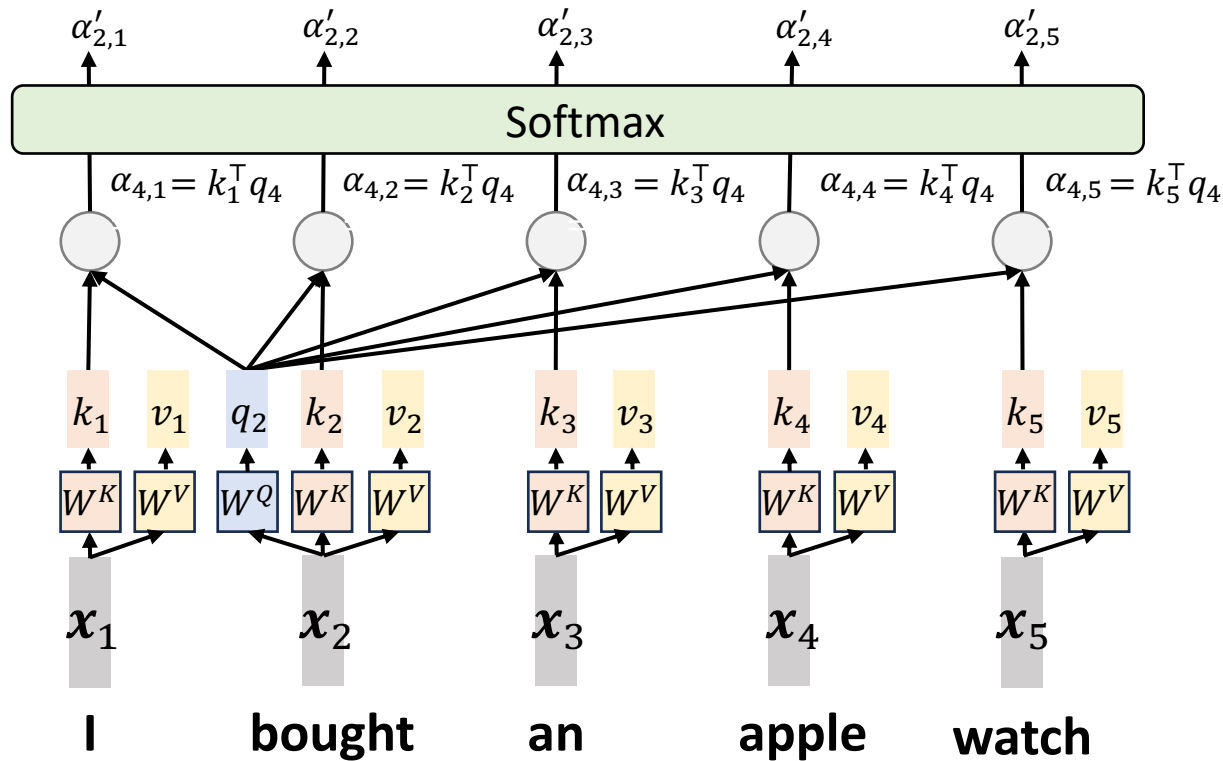
$$W^Q \in R^{d_k \times d}$$

$$W^K \in R^{d_k \times d}$$

$$W^V \in R^{d_v \times d}$$

Embedded
Tokens

Tokens



$$W^O \in R^{d \times d_v}$$

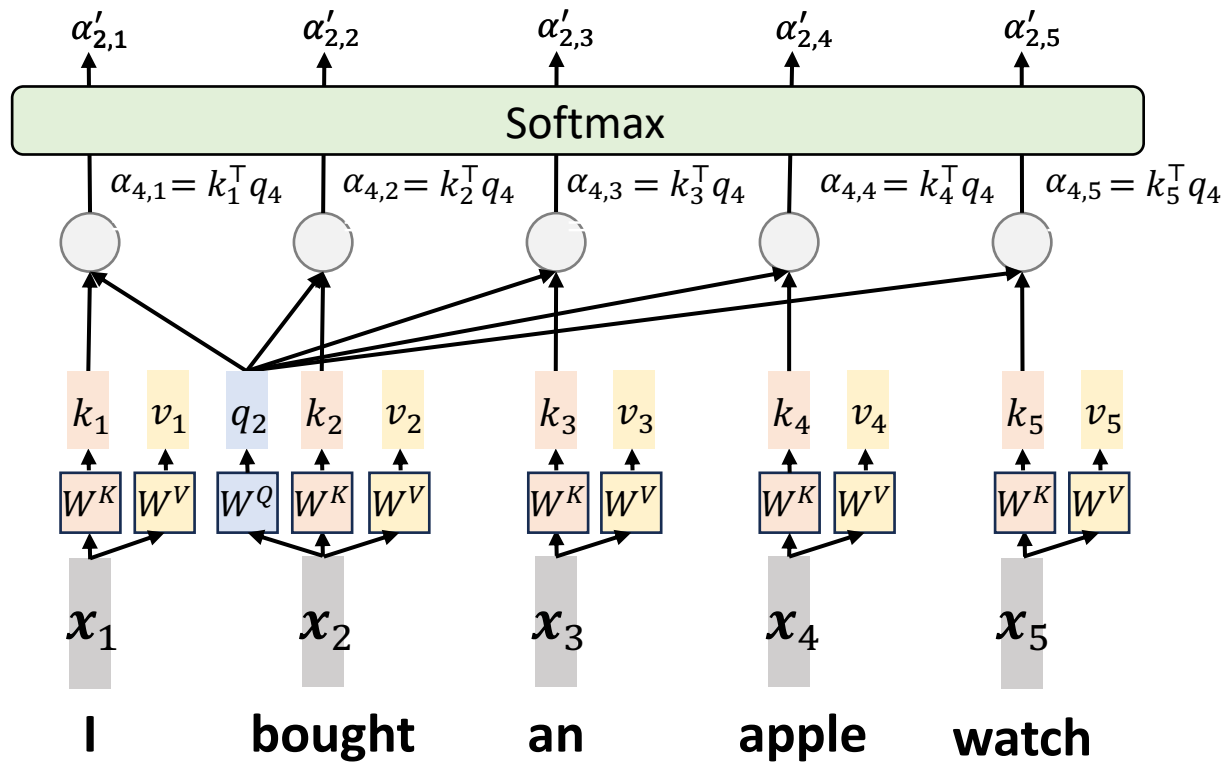
$$W^Q \in R^{d_k \times d}$$

$$W^K \in R^{d_k \times d}$$

$$W^V \in R^{d_v \times d}$$

Embedded
Tokens

Tokens



Updated feature

$$\mathbf{x}'_2 = W^O (\alpha'_{2,1} v_1 + \alpha'_{2,2} v_2 + \alpha'_{2,3} v_3 + \alpha'_{2,4} v_4 + \alpha'_{2,5} v_5)$$

$$W^O \in R^{d \times d_v}$$

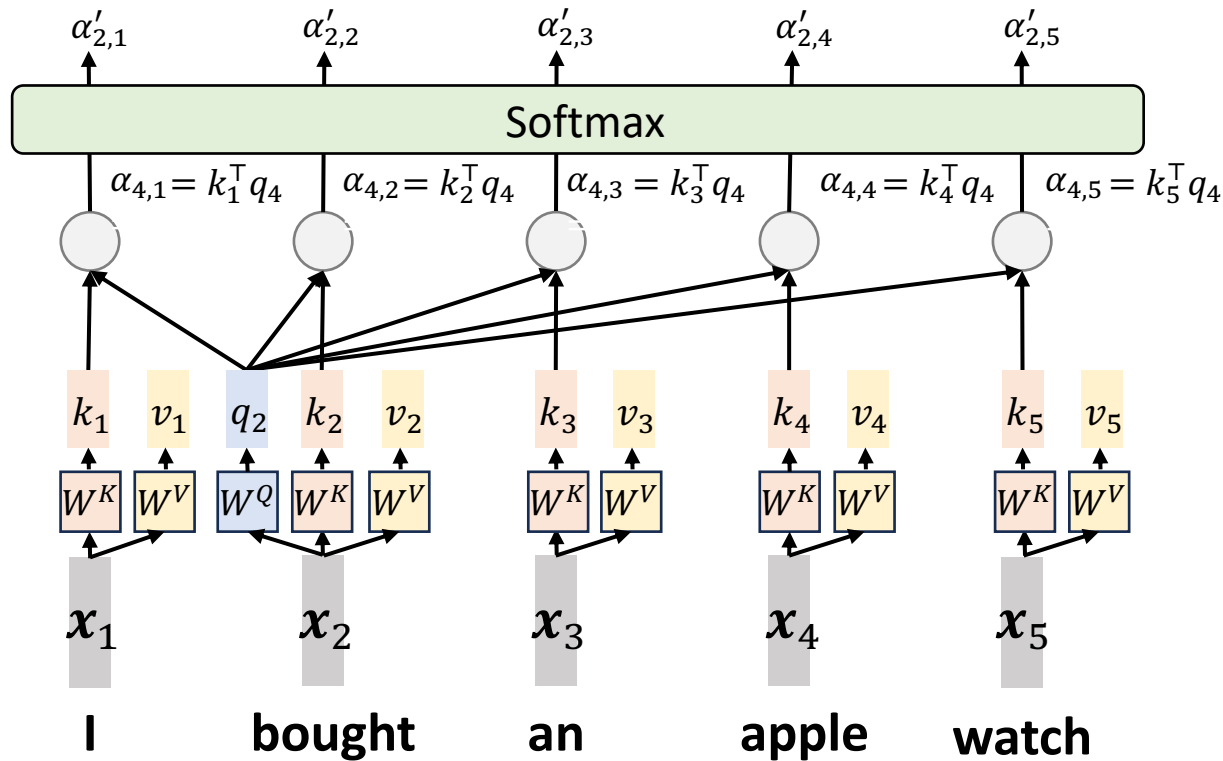
$$W^Q \in R^{d_k \times d}$$

$$W^K \in R^{d_k \times d}$$

$$W^V \in R^{d_v \times d}$$

Embedded Tokens

Tokens



Total weights: 175,181,291,520
Organized into 27,938 matrices



Embedding	$12,288 \times 50,257$ $d_embed * n_vocab$	$= 617,558,016$
Key	$128 \times 12,288 \times 96 \times 96$ $d_query * d_embed * n_heads * n_layers$	$= 14,495,514,624$
Query	$128 \times 12,288 \times 96 \times 96$ $d_query * d_embed * n_heads * n_layers$	$= 14,495,514,624$
Value	$128 \times 12,288 \times 96 \times 96$ $d_value * d_embed * n_heads * n_layers$	$= 14,495,514,624$
Output	$12,288 \times 128 \times 96 \times 96$ $d_embed * d_value * n_heads * n_layers$	$= 14,495,514,624$
Up-projection	$49,152 \times 12,288 \times 96$ $n_neurons * d_embed * n_layers$	$= 57,982,058,496$
Down-projection	$12,288 \times 49,152 \times 96$ $d_embed * n_neurons * n_layers$	$= 57,982,058,496$
Unembedding	$50,257 \times 12,288$ $n_vocab * d_embed$	$= 617,558,016$


$$\alpha_{1,1} = k_1^T q_1$$

$$\alpha_{1,2} = k_2^T q_1$$


$$\alpha_{1,3} = k_3^T q_1$$


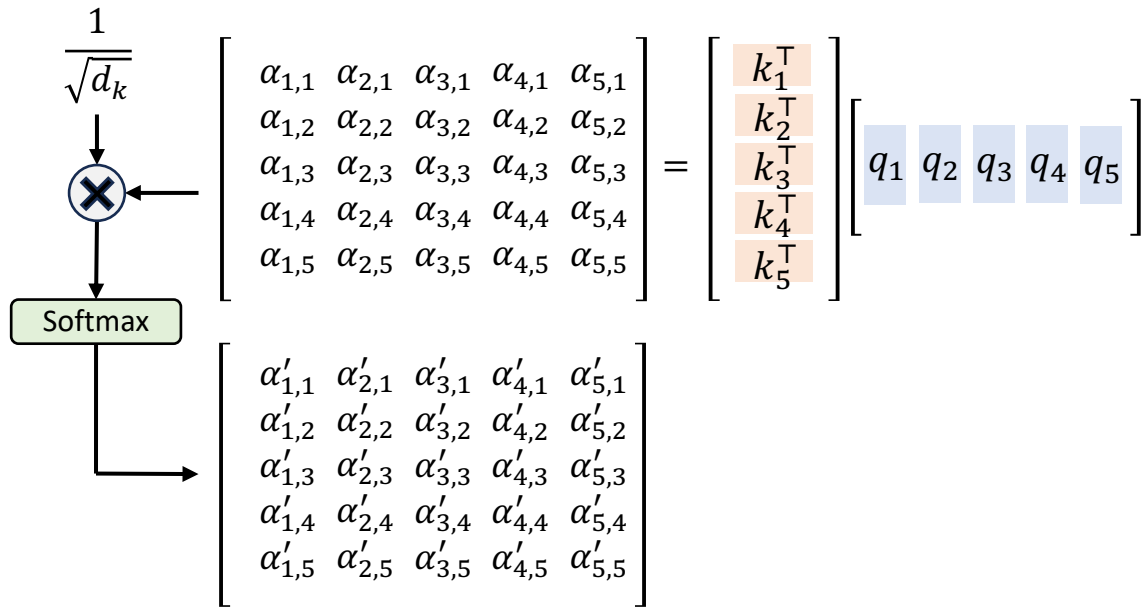
$$\alpha_{1,4} = k_4^T q_1$$


$$\alpha_{1,5} = k_5^T q_1$$



$$\begin{matrix}
 \alpha_{1,1} & \alpha_{2,1} & \alpha_{3,1} & \alpha_{4,1} & \alpha_{5,1} \\
 \alpha_{1,2} & \alpha_{2,2} & \alpha_{3,2} & \alpha_{4,2} & \alpha_{5,2} \\
 \alpha_{1,3} & \alpha_{2,3} & \alpha_{3,3} & \alpha_{4,3} & \alpha_{5,3} \\
 \alpha_{1,4} & \alpha_{2,4} & \alpha_{3,4} & \alpha_{4,4} & \alpha_{5,4} \\
 \alpha_{1,5} & \alpha_{2,5} & \alpha_{3,5} & \alpha_{4,5} & \alpha_{5,5}
 \end{matrix}
 =
 \begin{bmatrix}
 k_1^T \\
 k_2^T \\
 k_3^T \\
 k_4^T \\
 k_5^T
 \end{bmatrix}
 \begin{matrix}
 q_1 & q_2 & q_3 & q_4 & q_5
 \end{matrix}$$












$k = [k^1, k^2, \dots, k^{d_k}]^\top$ $E[k^i] = E[q^i] = 0$
 $q = [q^1, q^2, \dots, q^{d_k}]^\top$ $\text{Var}[k^i] = \text{Var}[q^i] = 1$



$$k^\top q = \sum_{i=1}^{d_k} k^i q^i$$

$\text{Var}[k^\top q] = d_k$












Single-head attention

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

$$Q = W^Q \begin{matrix} x_1 & x_2 & x_3 & x_4 & x_5 \end{matrix}$$
$$K = W^K \begin{matrix} x_1 & x_2 & x_3 & x_4 & x_5 \end{matrix}$$
$$V = W^V \begin{matrix} x_1 & x_2 & x_3 & x_4 & x_5 \end{matrix}$$



Analogy for Q, K, V



- ▶ Library system
- ▶ Imagine you're looking for information on a specific topic (query)
- ▶ Each book in the library has a summary (key) that helps identify if it contains the information you're looking for
- ▶ Once you find a match between your query and a summary, you access the book to get the detailed information (value) you need
- ▶ Here, in Attention, we do a “soft match” across multiple values, e.g. get info from multiple books (“book 1 is most relevant, then book 2, then book 3, etc.”)

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


Single-head attention

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

$$\begin{aligned} Q &= W^Q \begin{matrix} x_1 & x_2 & x_3 & x_4 & x_5 \end{matrix} \\ K &= W^K \begin{matrix} x_1 & x_2 & x_3 & x_4 & x_5 \end{matrix} \\ V &= W^V \begin{matrix} x_1 & x_2 & x_3 & x_4 & x_5 \end{matrix} \end{aligned}$$



...



$$W_i^Q \in R^{d_k \times d}$$

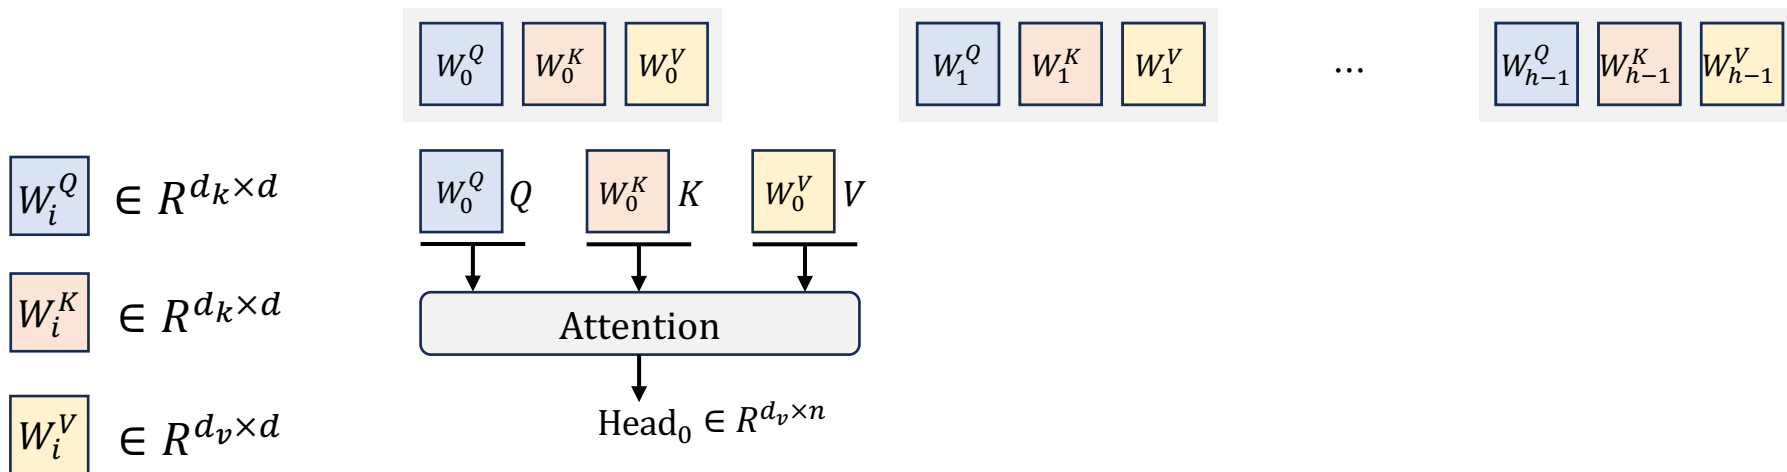
$$W_i^K \in R^{d_k \times d}$$

$$W_i^V \in R^{d_v \times d}$$

Single-head attention

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

$$Q = W^Q \begin{matrix} x_1 & x_2 & x_3 & x_4 & x_5 \end{matrix}$$
$$K = W^K \begin{matrix} x_1 & x_2 & x_3 & x_4 & x_5 \end{matrix}$$
$$V = W^V \begin{matrix} x_1 & x_2 & x_3 & x_4 & x_5 \end{matrix}$$



Single-head attention


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

$$X = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix}$$

$$W_i^Q \in R^{d_k \times d}$$

$$W_i^K \in R^{d_k \times d}$$

$$W_i^V \in R^{d_v \times d}$$



Single-head attention

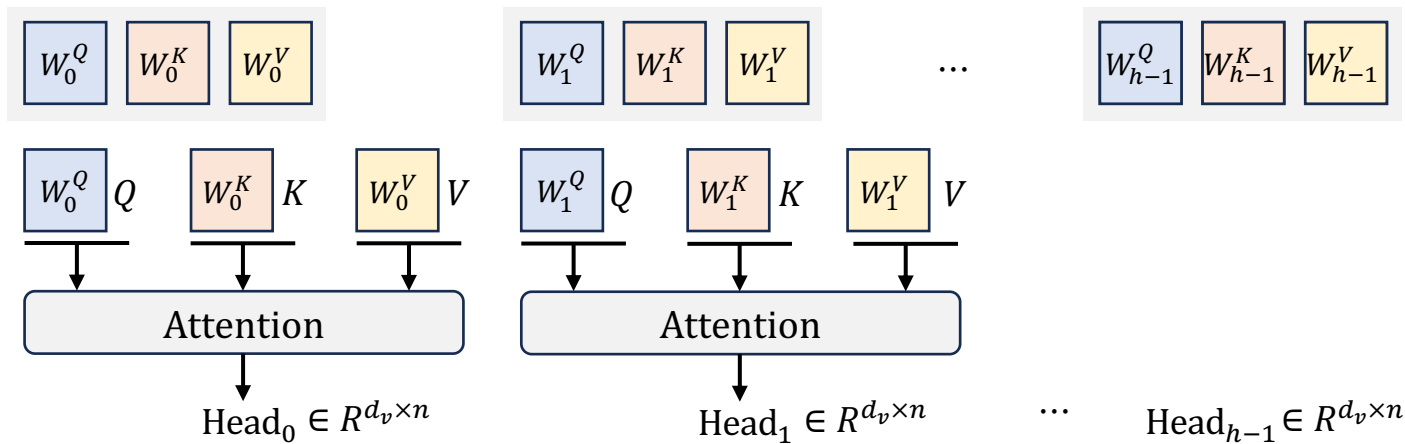
$$\text{Attention}(Q, K, V) = V \text{softmax}\left(\frac{K^T Q}{\sqrt{d_k}}\right)$$

$$Q = W^Q \begin{matrix} x_1 & x_2 & x_3 & x_4 & x_5 \end{matrix}$$
$$K = W^K \begin{matrix} x_1 & x_2 & x_3 & x_4 & x_5 \end{matrix}$$
$$V = W^V \begin{matrix} x_1 & x_2 & x_3 & x_4 & x_5 \end{matrix}$$

$$W_i^Q \in R^{d_k \times d}$$

$$W_i^K \in R^{d_k \times d}$$

$$W_i^V \in R^{d_v \times d}$$



Single-head attention

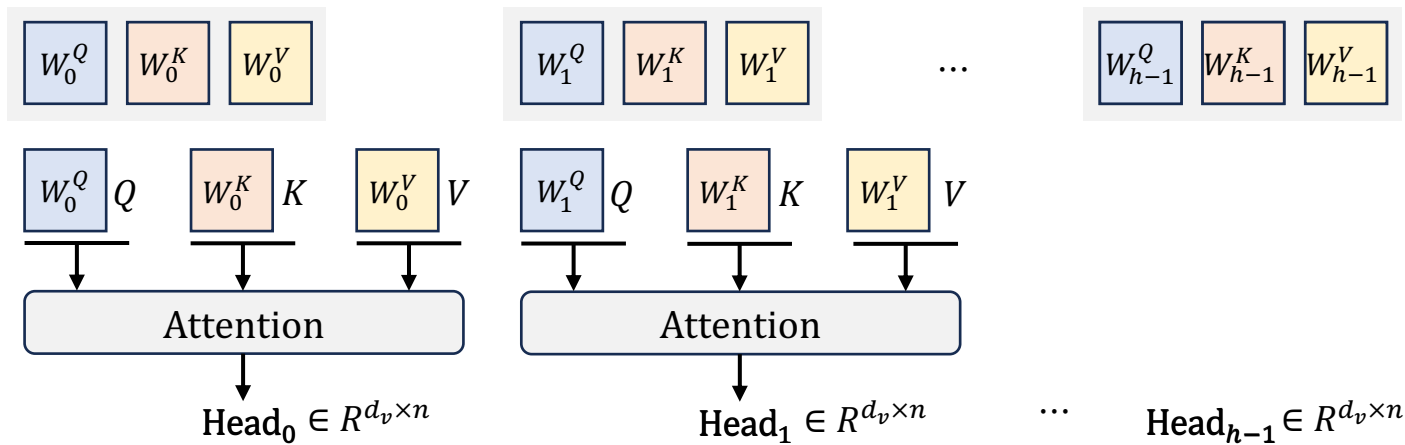
$$\text{Attention}(Q, K, V) = V \text{softmax}\left(\frac{K^T Q}{\sqrt{d_k}}\right)$$

$$Q = W^Q \begin{matrix} x_1 & x_2 & x_3 & x_4 & x_5 \end{matrix}$$
$$K = W^K \begin{matrix} x_1 & x_2 & x_3 & x_4 & x_5 \end{matrix}$$
$$V = W^V \begin{matrix} x_1 & x_2 & x_3 & x_4 & x_5 \end{matrix}$$

$$W_i^Q \in R^{d_k \times d}$$

$$W_i^K \in R^{d_k \times d}$$

$$W_i^V \in R^{d_v \times d}$$



Single-head attention

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



$$Q = W^Q \begin{bmatrix} x_1 & x_2 & x_3 & x_4 & x_5 \end{bmatrix}$$

$$K = W^K \begin{bmatrix} x_1 & x_2 & x_3 & x_4 & x_5 \end{bmatrix}$$


$$V = W^V \begin{bmatrix} x_1 & x_2 & x_3 & x_4 & x_5 \end{bmatrix}$$

Multi-head attention



$$W^O \in R^{d \times hd_v}$$



...



$$W_i^Q \in R^{d_k \times d}$$



$$W_i^K \in R^{d_k \times d}$$



$$W_i^V \in R^{d_v \times d}$$

$$\text{Head}_0 \in R^{d_v \times n}$$

$$\text{Head}_1 \in R^{d_v \times n}$$

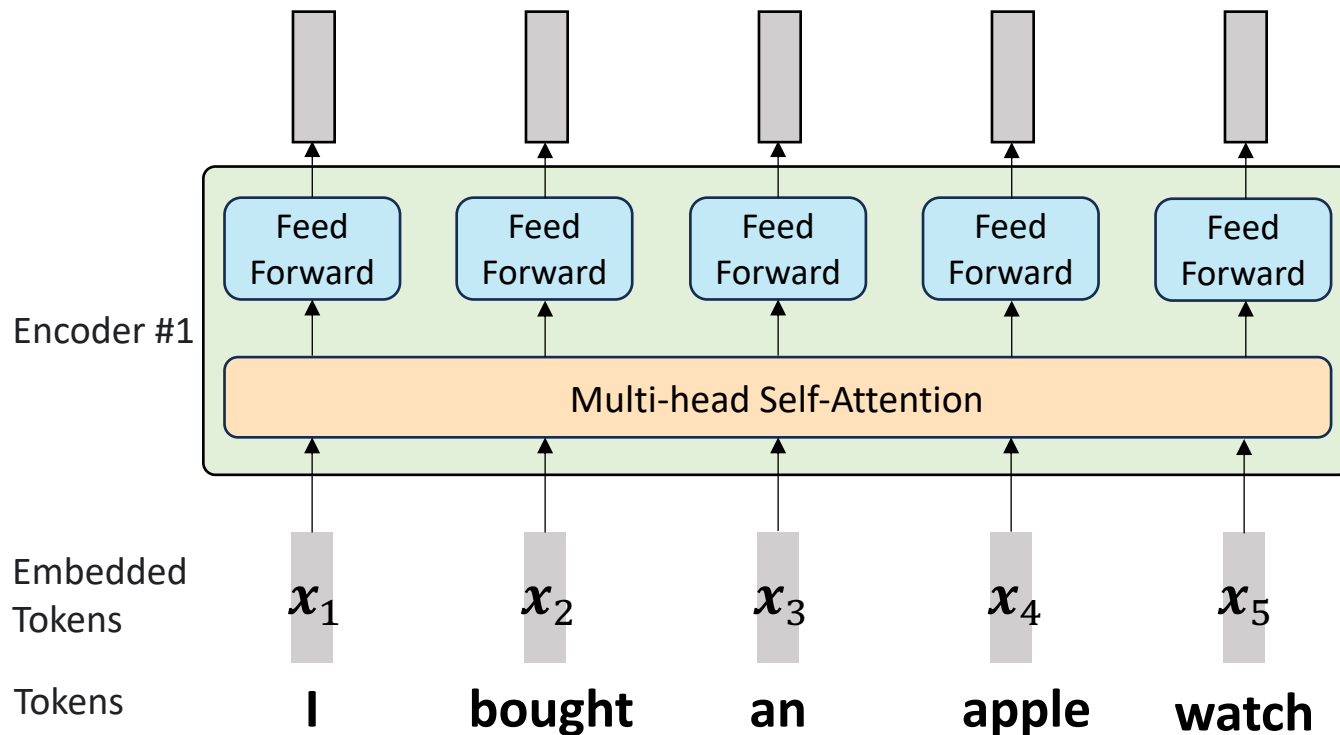
...

$$\text{Head}_{h-1} \in R^{d_v \times n}$$

$$\text{MultiHeadedAttention}(Q, K, V) = W^O \begin{bmatrix} \text{Head}_0 \\ \text{Head}_1 \\ \vdots \\ \text{Head}_{h-1} \end{bmatrix}$$

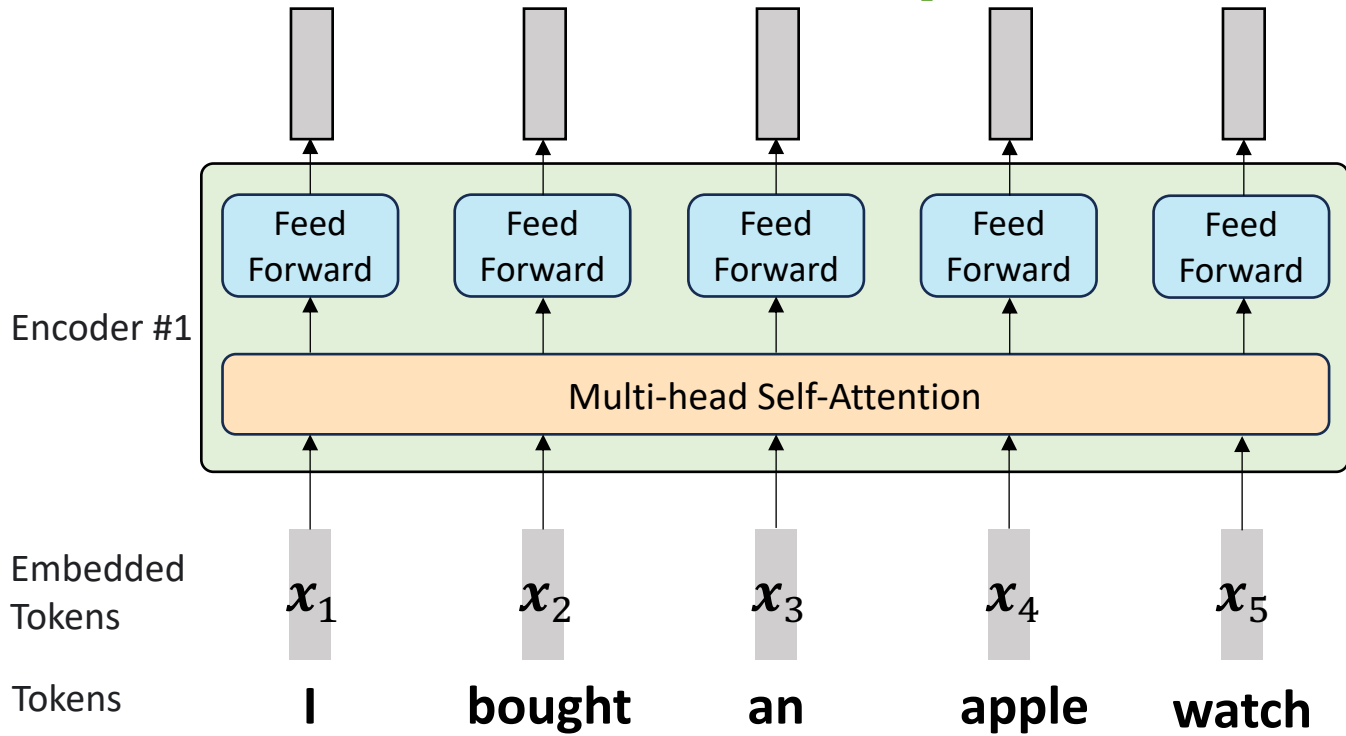

Feed Forward Network (FFN)

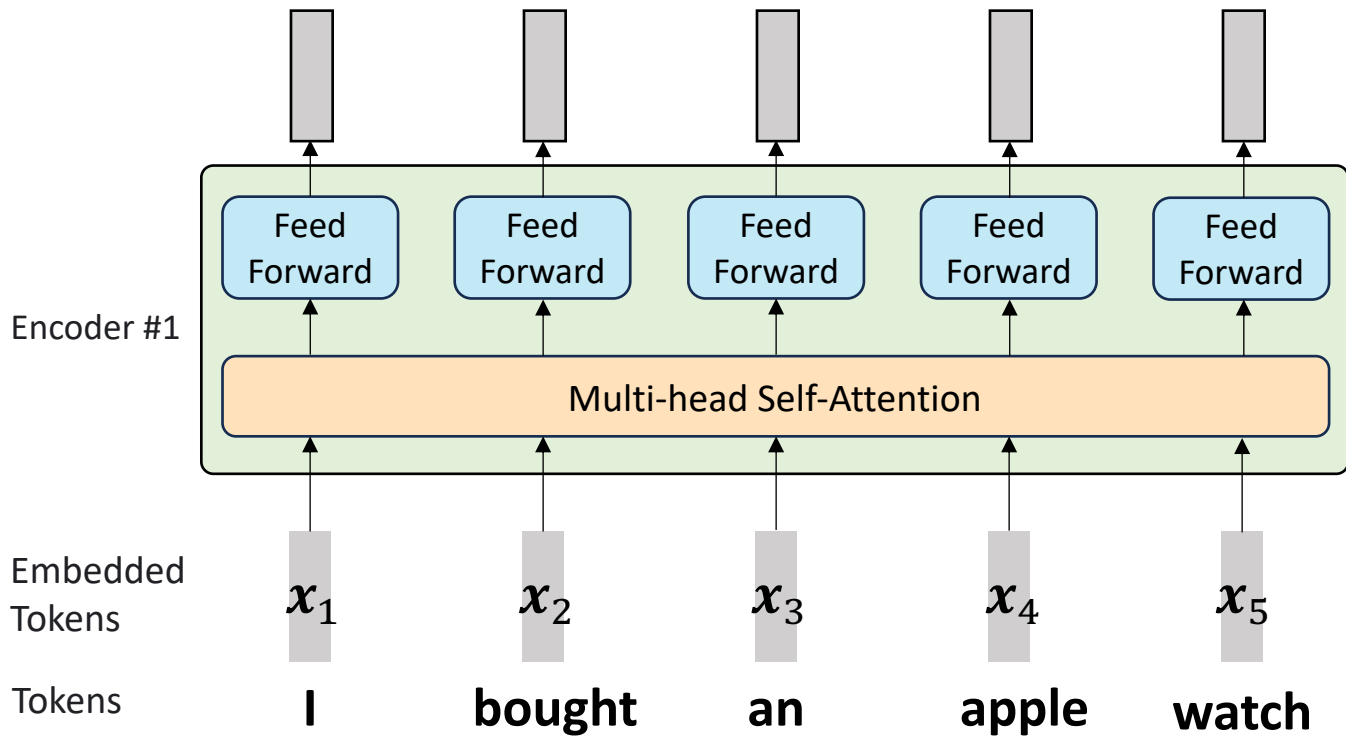
$$FFN(\mathbf{x}) = W_2 \text{ReLU}(W_1 \mathbf{x} + \mathbf{b}_1) + \mathbf{b}_2$$

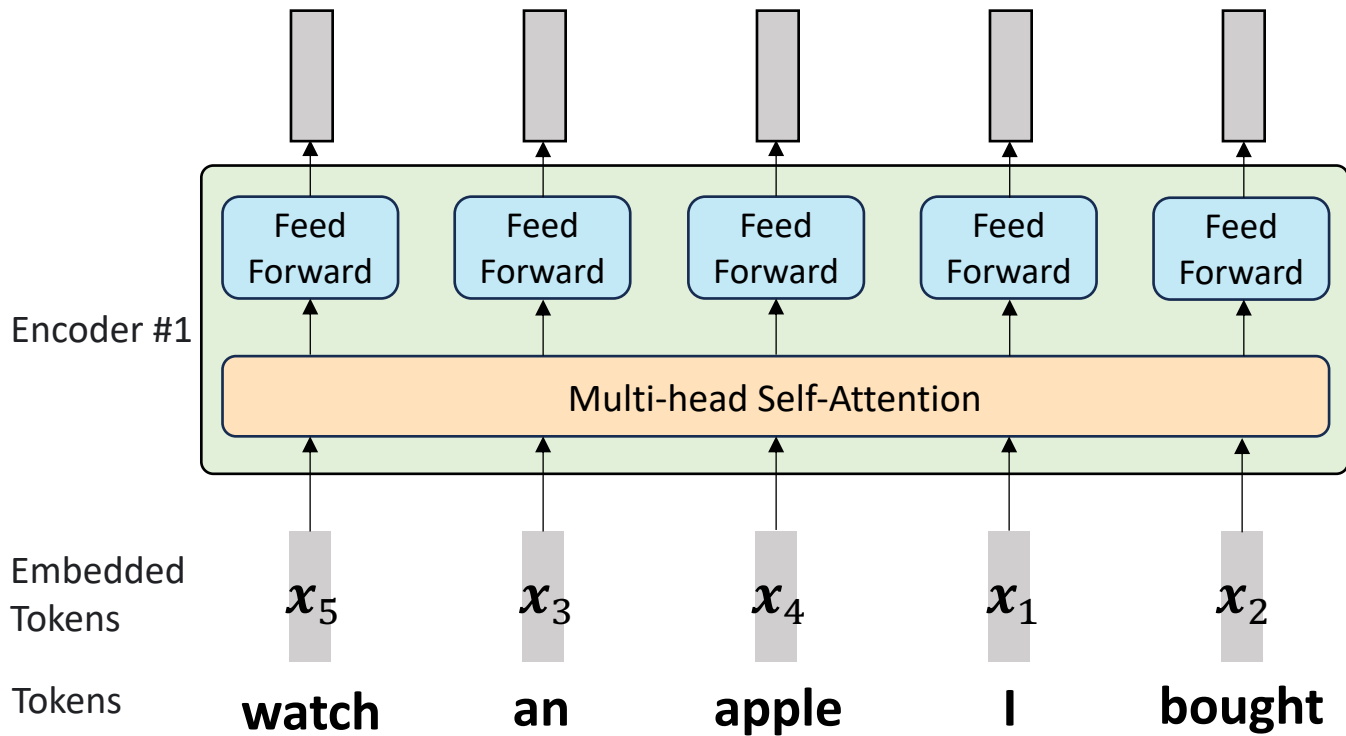


Feed Forward Network (FFN)

$$FFN(x) = W_2 \text{ReLU}(W_1 x + b_1) + b_2$$

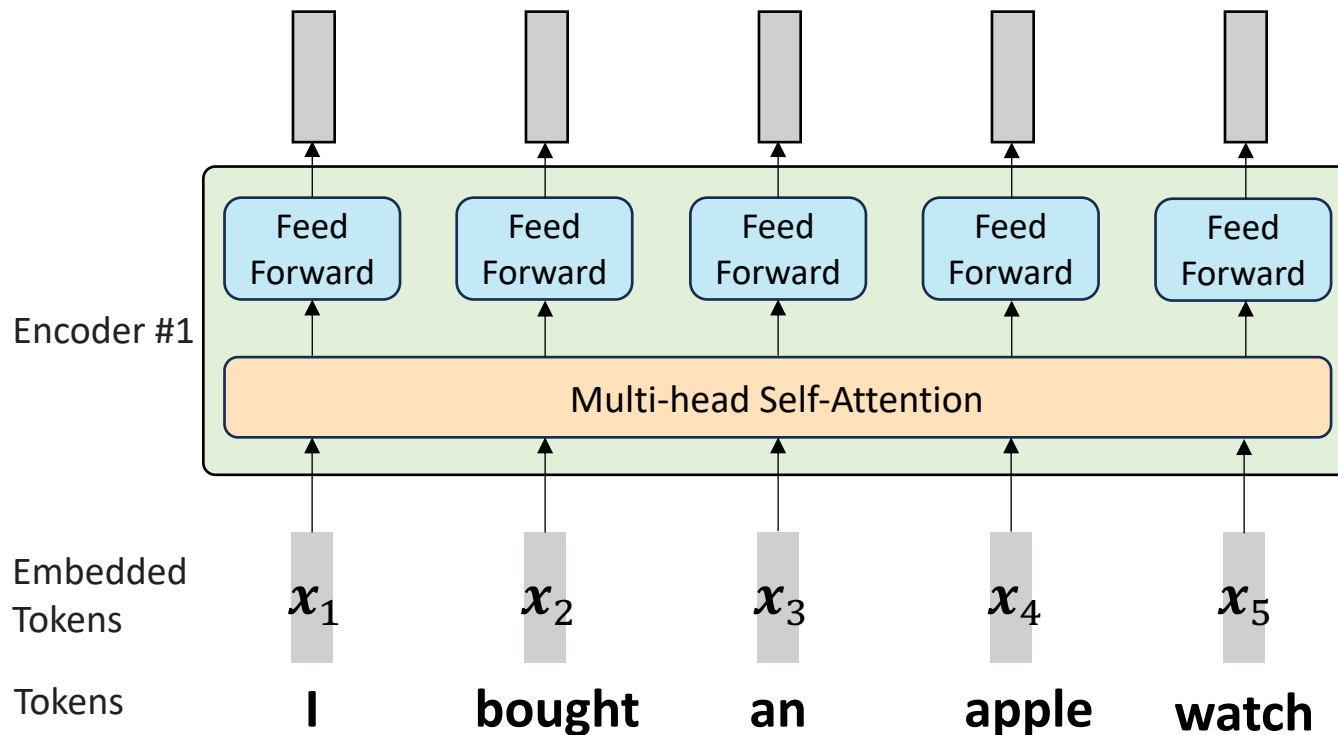








Positional encoding








Positional encoding

Position k	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
Dimension 2^3	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	← Slow oscillating
2^2	0	0	0	0	1	1	1	1	0	0	0	0	1	1	1	1	
2^1	0	0	1	1	0	0	1	1	0	0	1	1	0	0	1	1	
2^0	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	← Fast oscillating


Positional embedding

d



Embedded Tokens

d



Tokens

I

bought

an

apple

watch



Positional encoding









Image: <https://timodenk.com/blog/linear-relationships-in-the-transformers-positional-encoding/>


Positional embedding

d



Embedded Tokens

d



Tokens

I

bought

an

apple

watch

x_1

x_2

x_3

x_4

x_5



Positional encoding

Position k

Angular frequency

$$w_i = N^{-2i/d}$$

$N = 100,000$

$$\begin{bmatrix} \sin(w_0 k) \\ \cos(w_0 k) \\ \sin(w_1 k) \\ \cos(w_1 k) \\ \vdots \\ \vdots \\ \sin\left(\frac{w_{d-1}}{2} k\right) \\ \cos\left(\frac{w_{d-1}}{2} k\right) \end{bmatrix}$$



Normalized Range



Unique identifier, unlimited length




Relative positions as linear transform

$$\begin{bmatrix} \sin(w_i(k + \Delta k)) \\ \cos(w_i(k + \Delta k)) \end{bmatrix} = \begin{bmatrix} \sin(w_i k) \cos(w_i \Delta k) + \cos(w_i k) \sin(w_i \Delta k) \\ \cos(w_i k) \cos(w_i \Delta k) - \sin(w_i k) \sin(w_i \Delta k) \end{bmatrix}$$


Positional embedding

d



Embedded Tokens

d

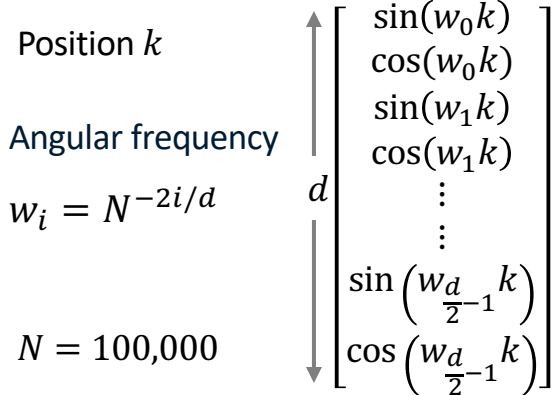


Tokens

I bought an apple watch



Positional encoding



Normalized Range





Unique identifier, unlimited length



Relative positions as linear transform

$$\begin{bmatrix} \sin(w_i(k + \Delta k)) \\ \cos(w_i(k + \Delta k)) \end{bmatrix} = \begin{bmatrix} \sin(w_i k) \cos(w_i \Delta k) + \cos(w_i k) \sin(w_i \Delta k) \\ \cos(w_i k) \cos(w_i \Delta k) - \sin(w_i k) \sin(w_i \Delta k) \end{bmatrix}$$

$$= \begin{bmatrix} \cos(w_i \Delta k) & \sin(w_i \Delta k) \\ -\sin(w_i \Delta k) & \cos(w_i \Delta k) \end{bmatrix} \begin{bmatrix} \sin(w_i k) \\ \cos(w_i k) \end{bmatrix}$$



Tokens **I** **bought** **an** **apple** **watch**

Embedded
Tokens

\uparrow
 d
 \downarrow

\mathbf{x}_1

\mathbf{x}_2

\mathbf{x}_3

\mathbf{x}_4

\mathbf{x}_5

Tokens

I

bought

an

apple

watch

Position

$k = 1$

$k = 2$

$k = 3$

$k = 4$

$k = 5$

\uparrow
 d
 \downarrow

\mathbf{P}_1

\mathbf{P}_2

\mathbf{P}_3

\mathbf{P}_4

\mathbf{P}_5

d

$$\begin{bmatrix} \sin(w_0 k) \\ \cos(w_0 k) \\ \sin(w_1 k) \\ \cos(w_1 k) \\ \vdots \\ \vdots \\ \sin\left(w_{\frac{d}{2}-1} k\right) \\ \cos\left(w_{\frac{d}{2}-1} k\right) \end{bmatrix}$$

\mathbf{x}_i

concat

\mathbf{P}_i

\mathbf{x}_i

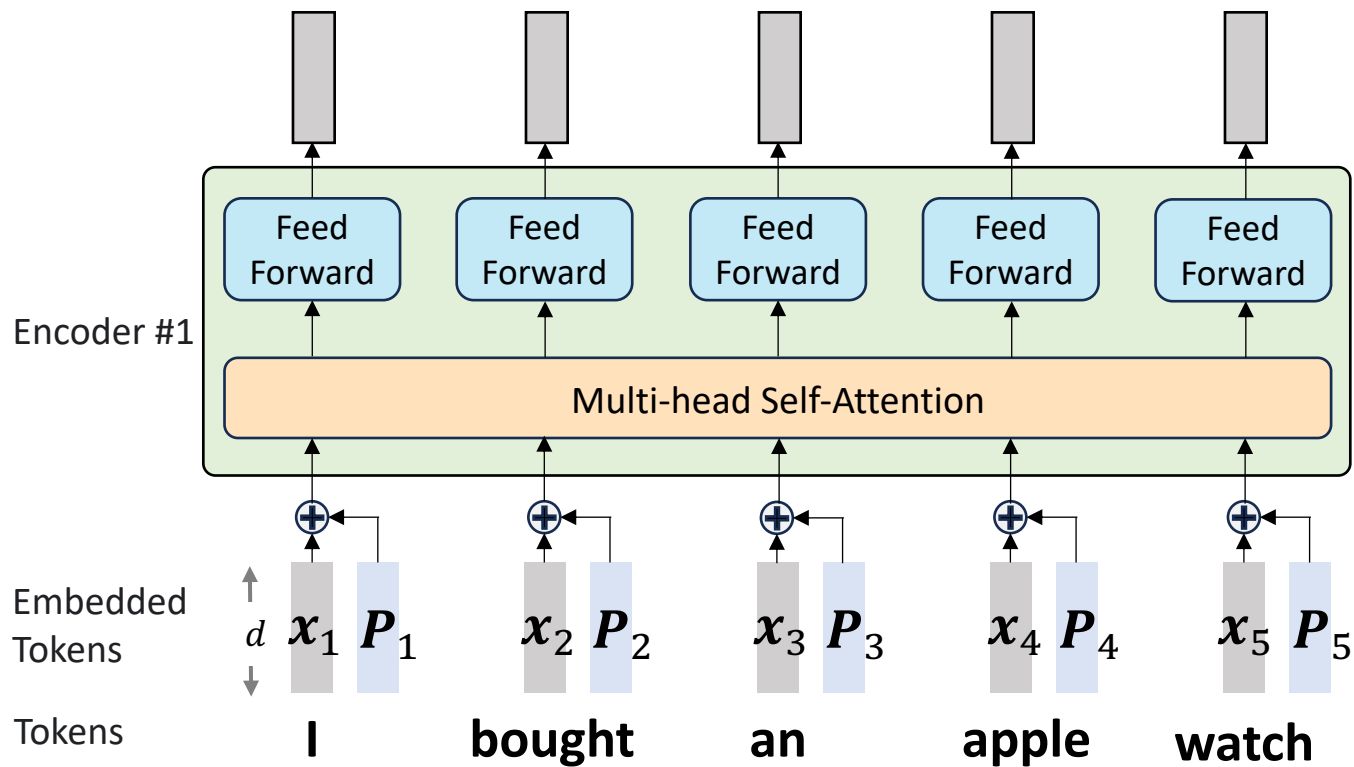
MLP

\mathbf{P}_i

\mathbf{x}_i

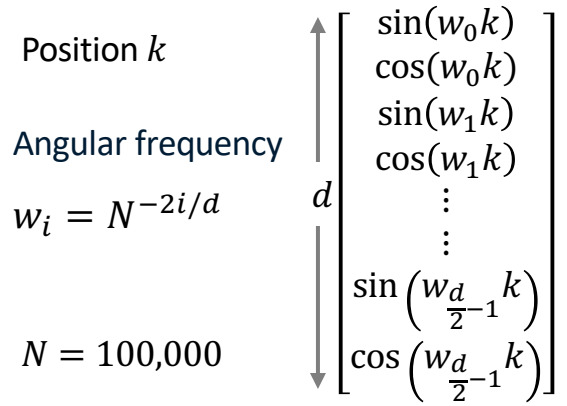
+

\mathbf{P}_i





Positional encoding



Sinusoidal positional encoding

Relative positional encoding

KERPLE

RoPE






CoPE

NoPE


YaRN

FIRE

Positional embedding



Embedded Tokens



Tokens

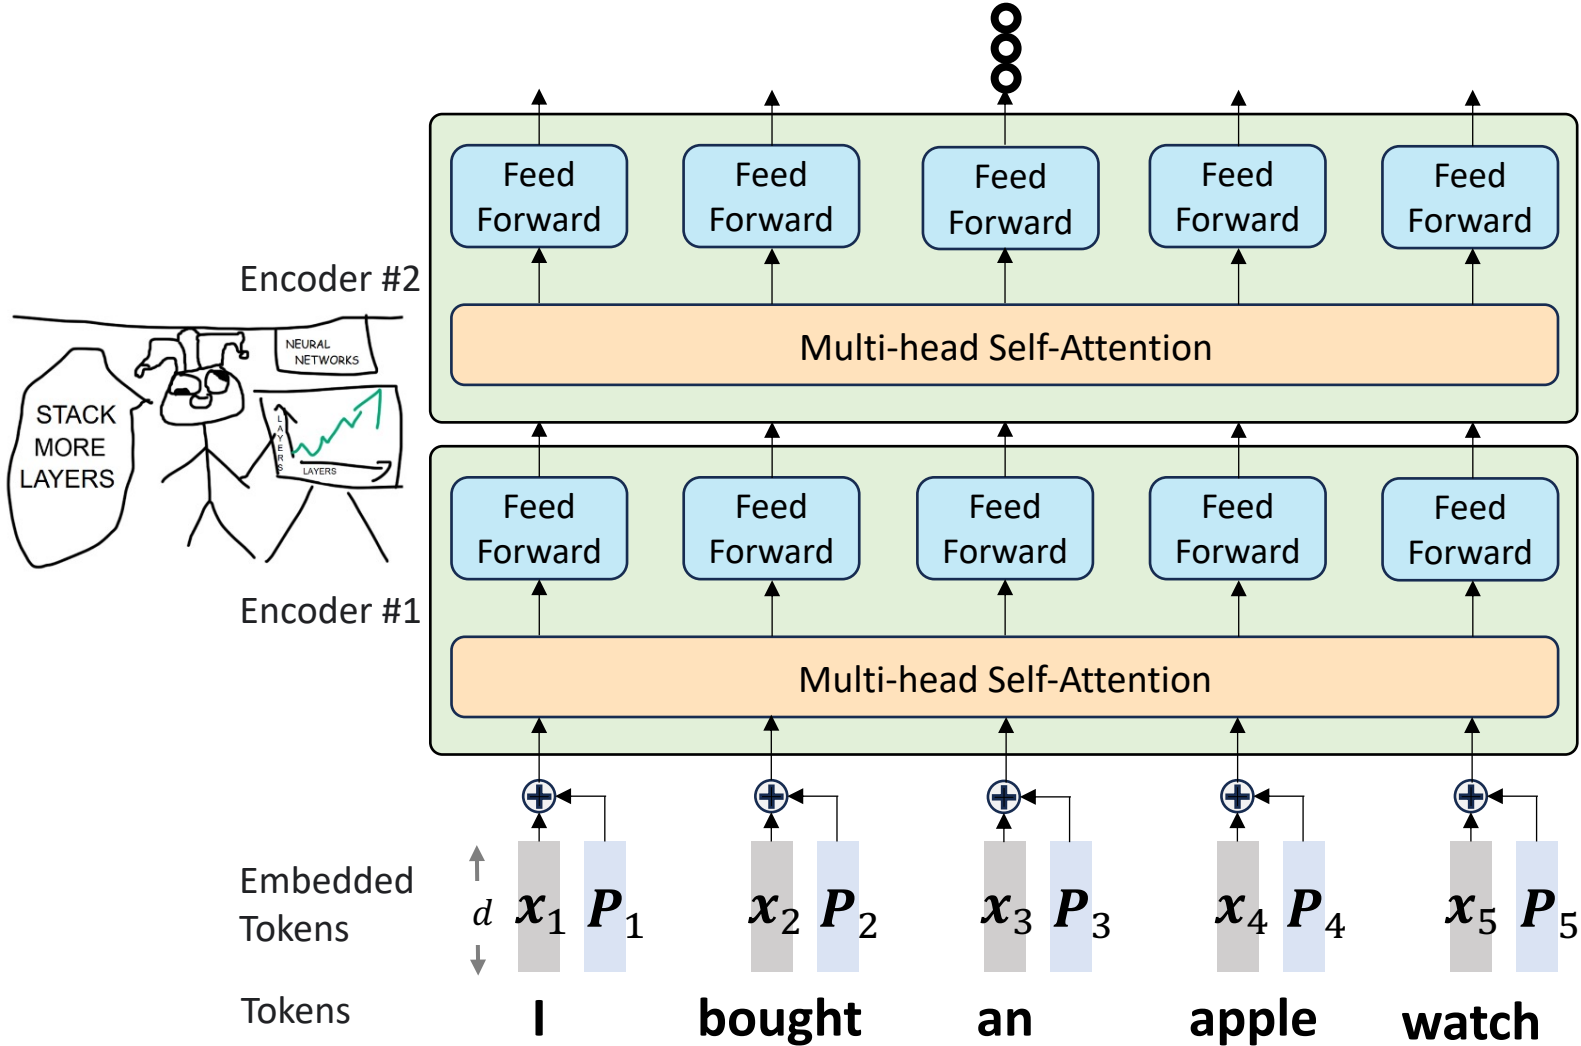
I

bought

an

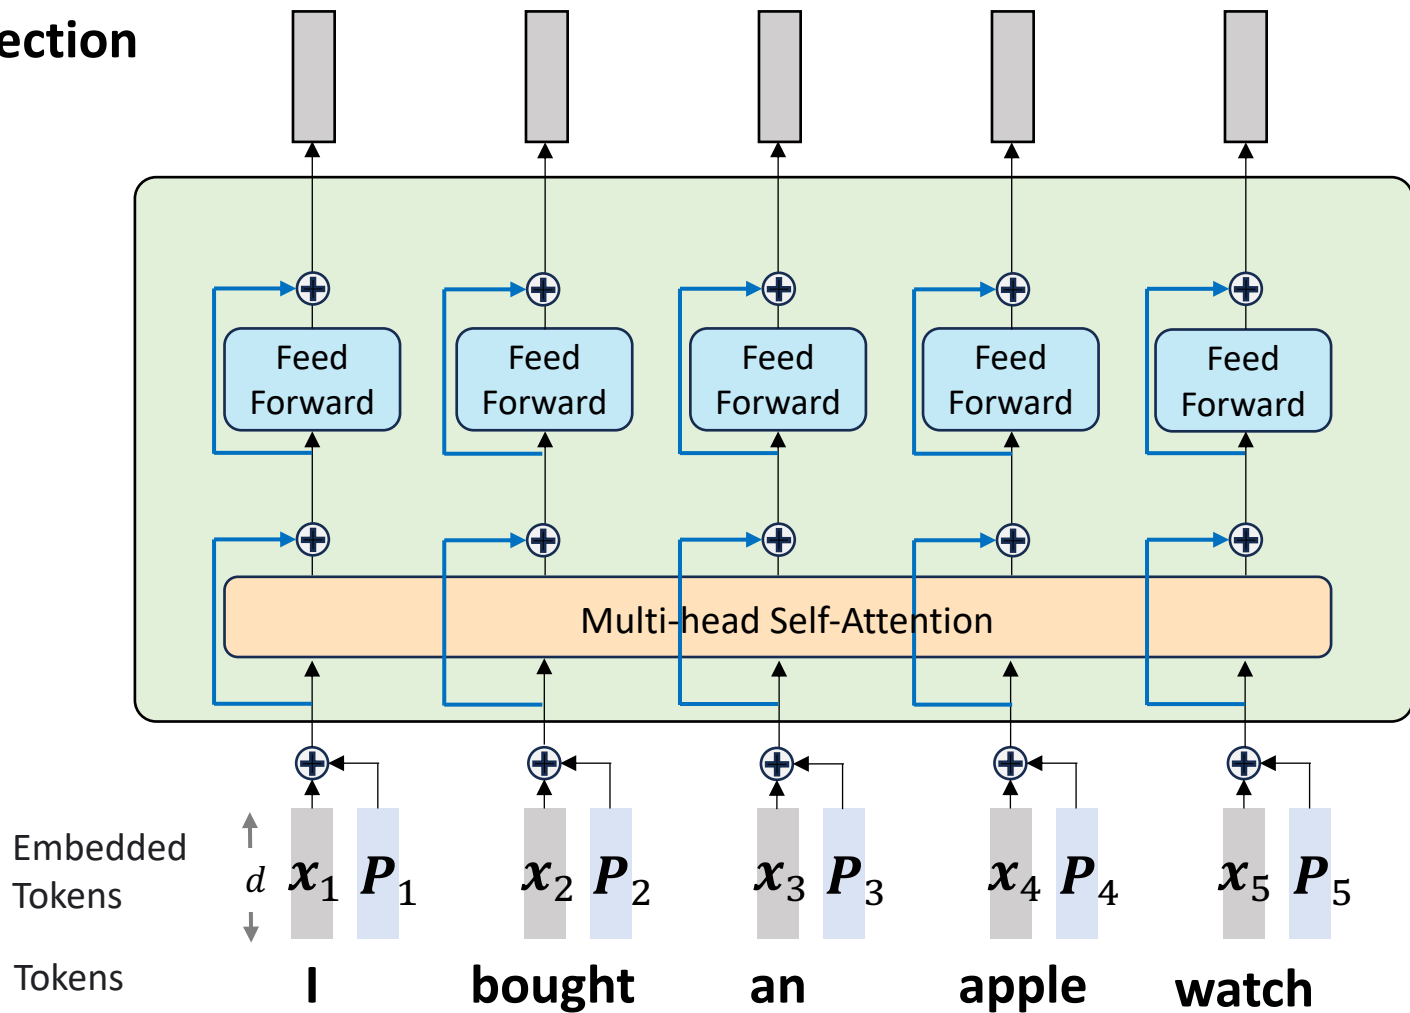
apple

watch



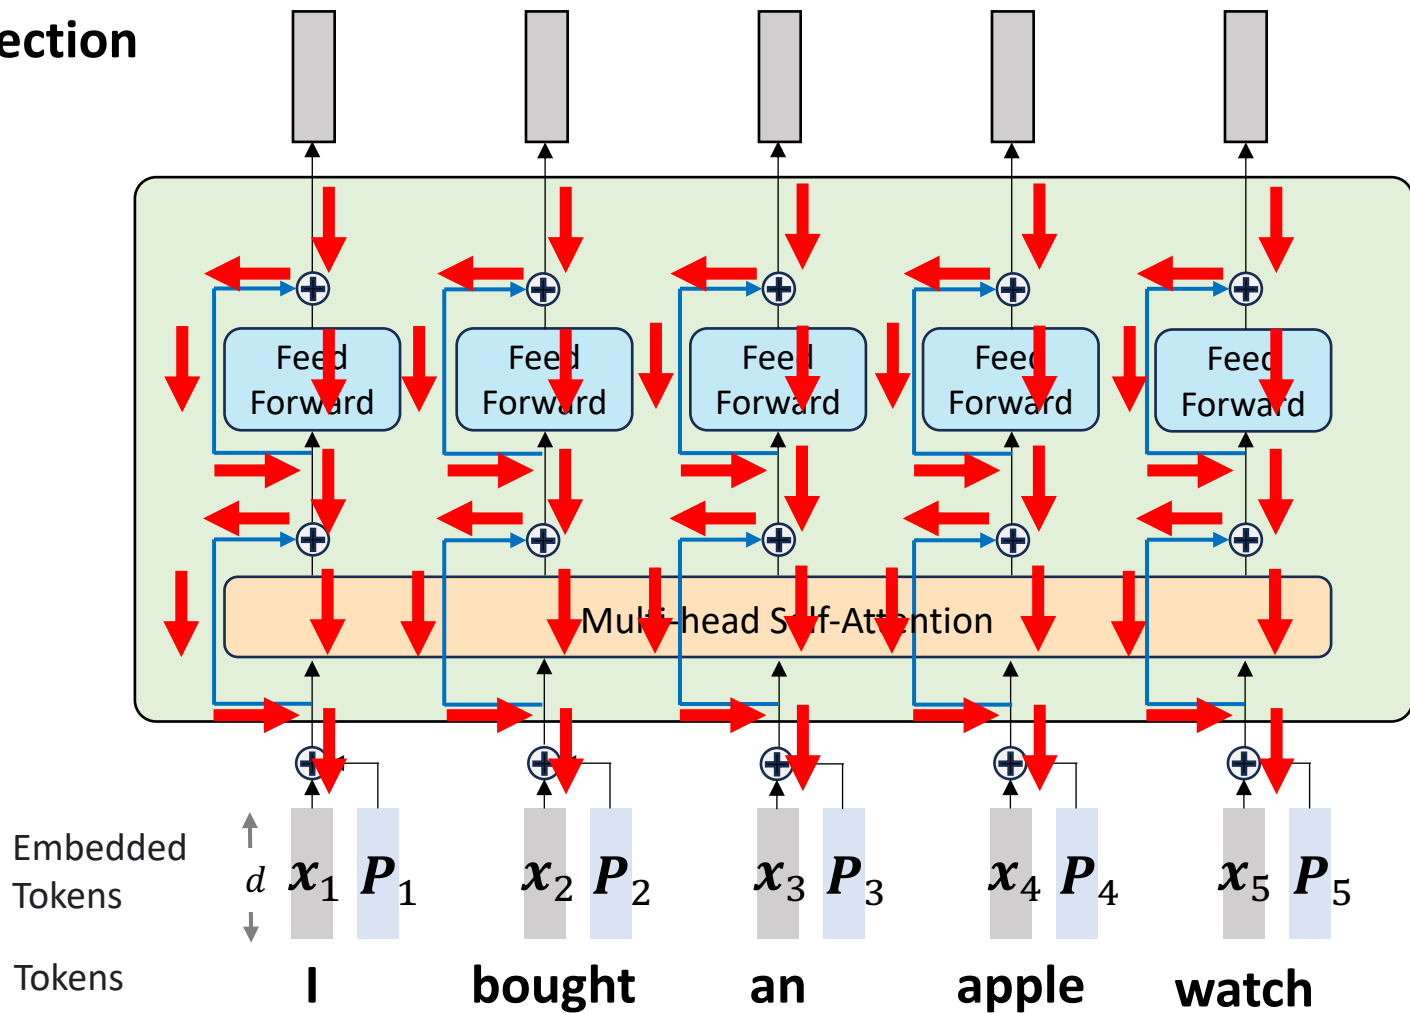


Residual connection





Residual connection

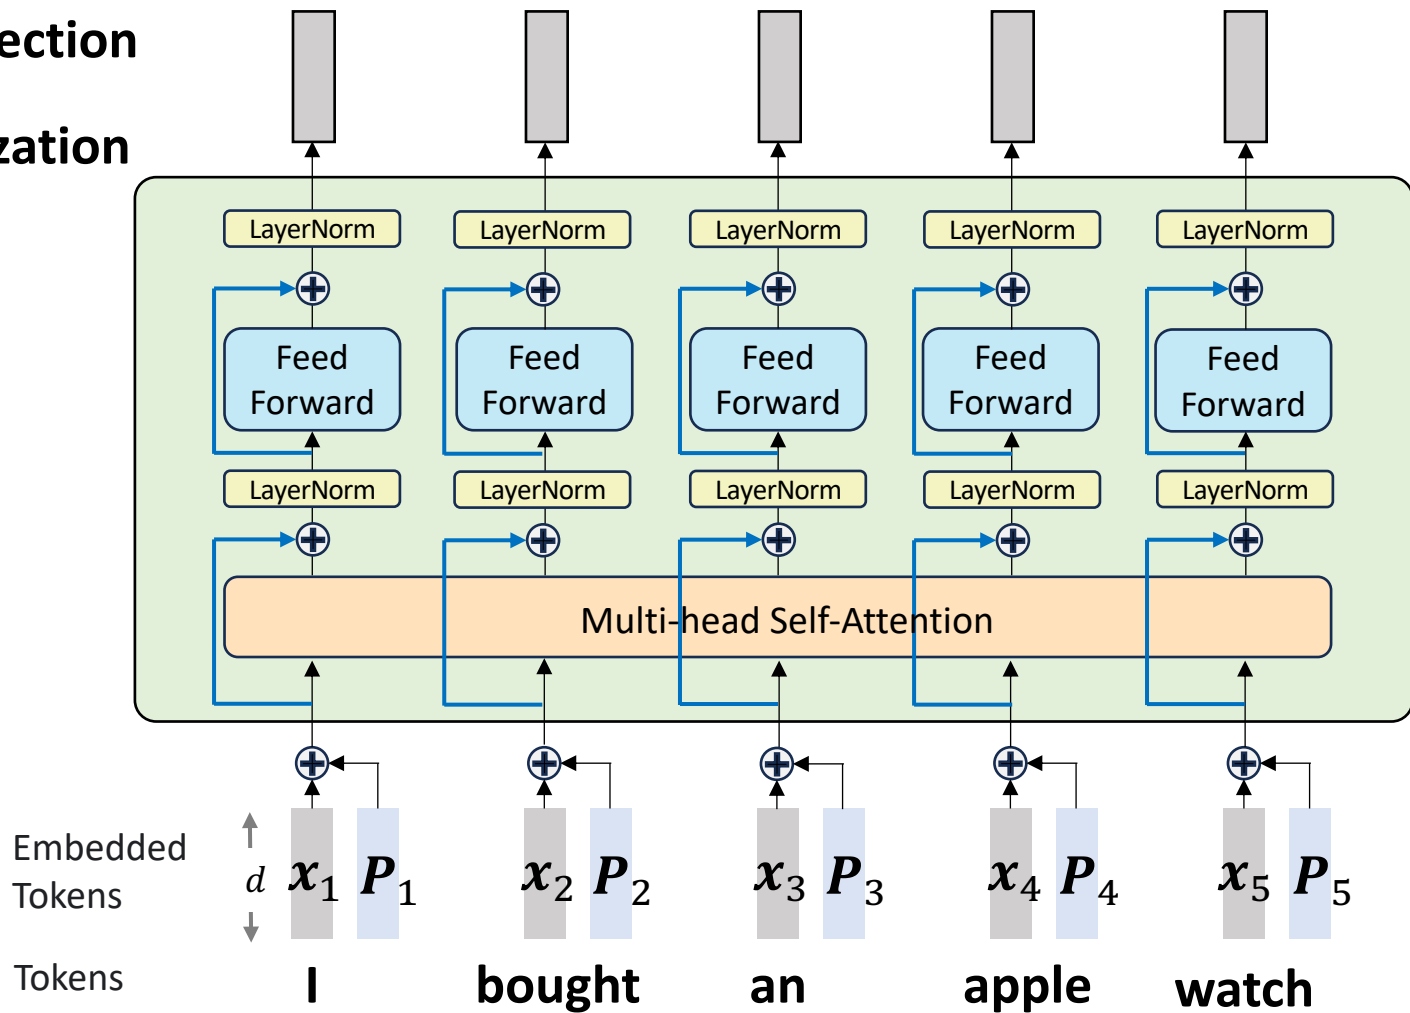




Residual connection



Layer normalization





Residual connection



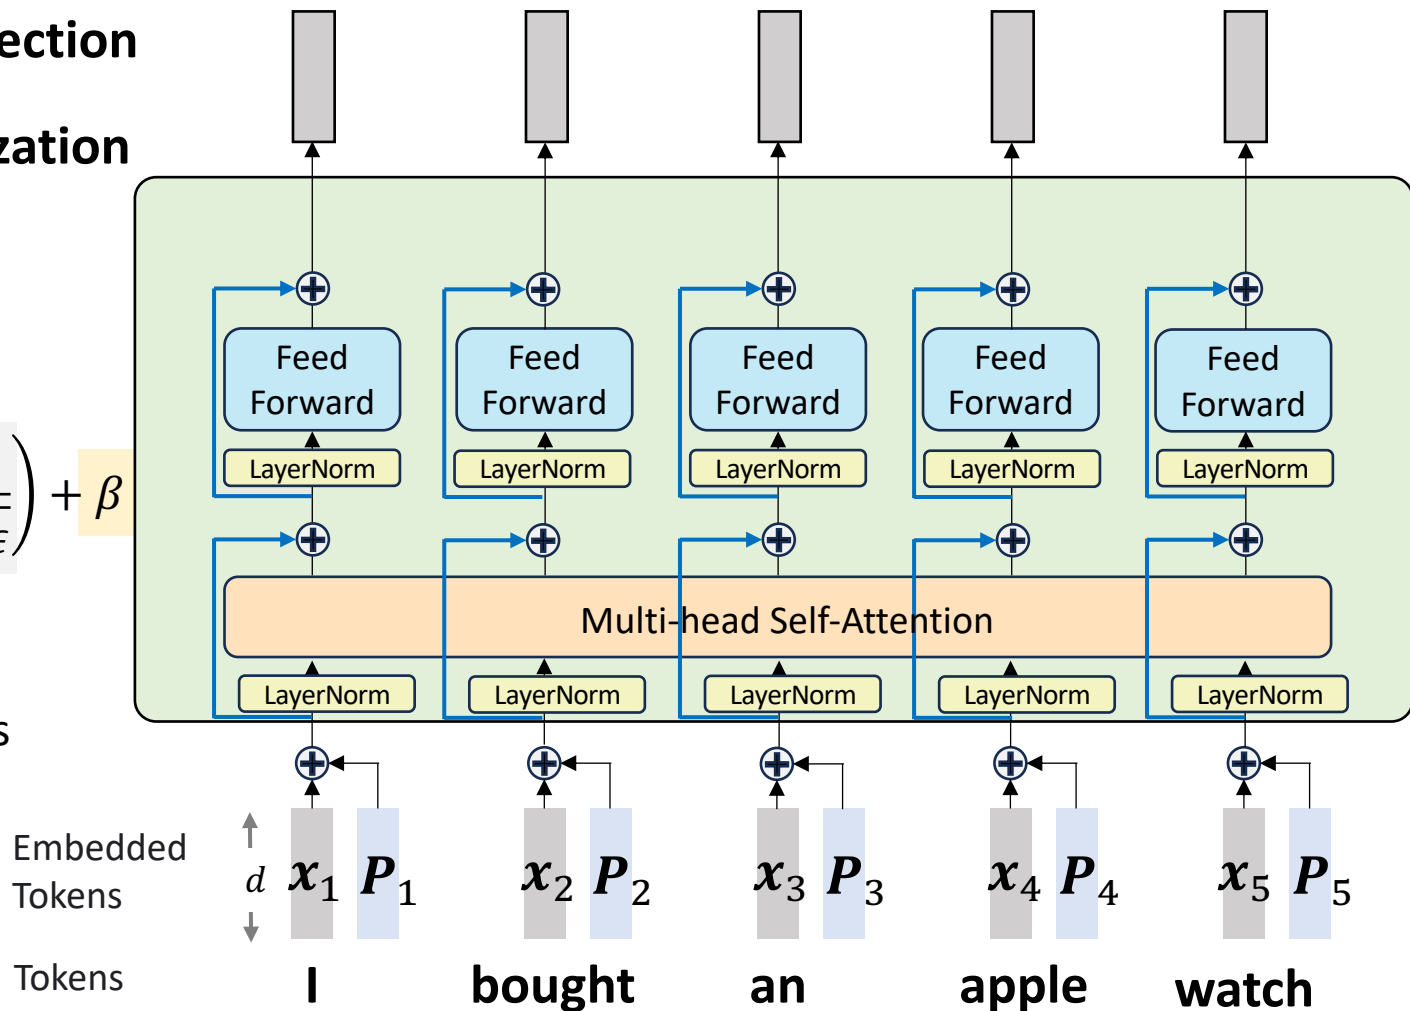
Layer normalization

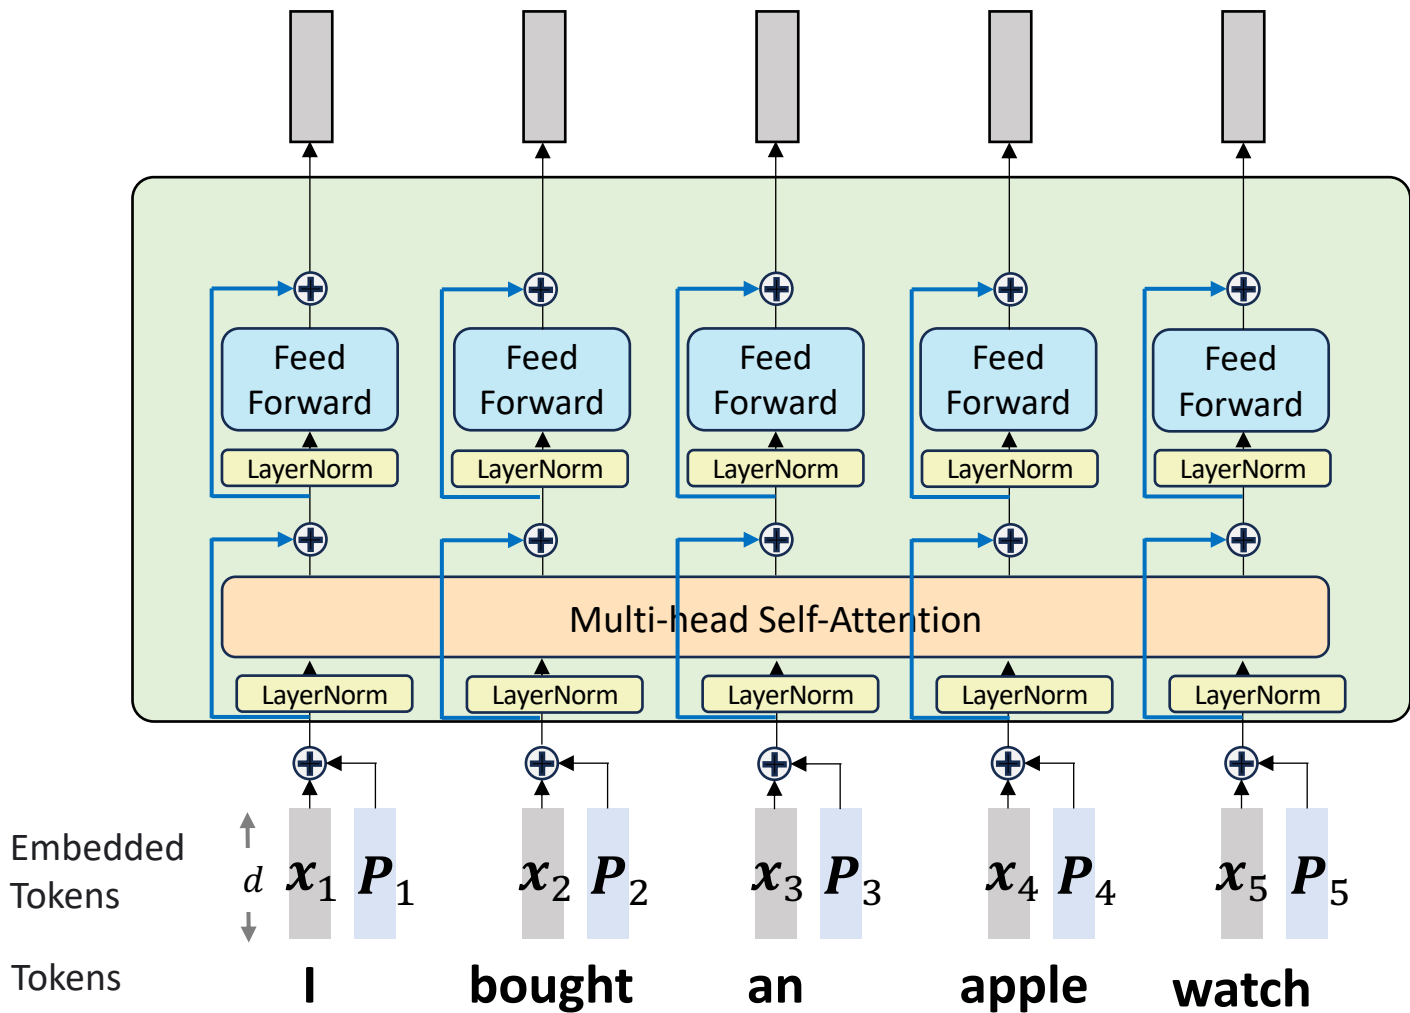
LayerNorm(\mathbf{x}) =

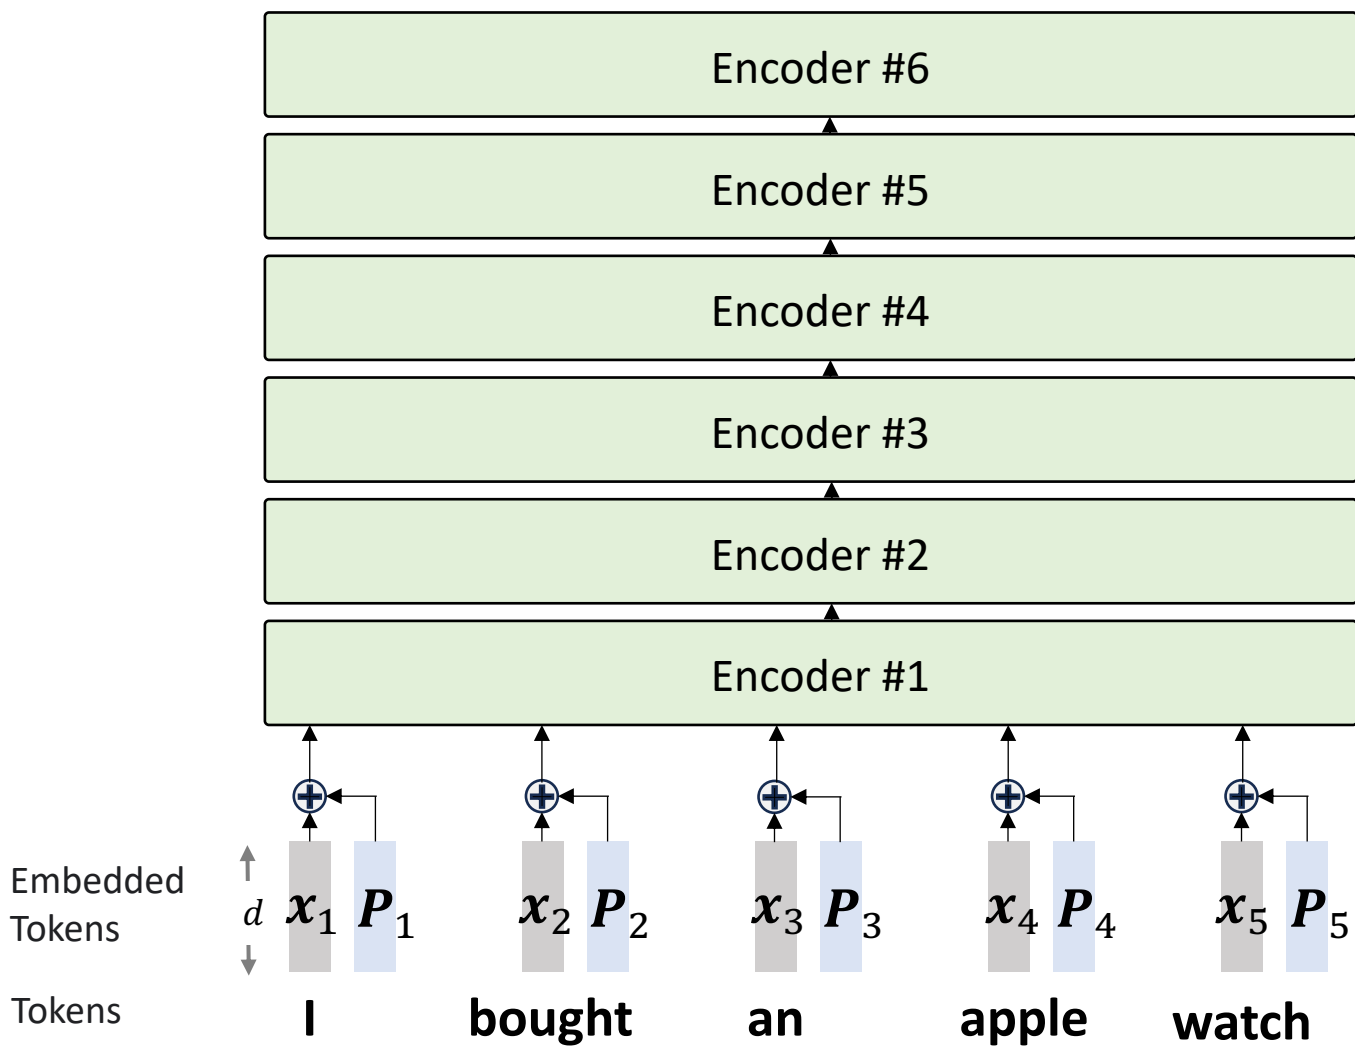
$$\gamma \left(\frac{\mathbf{x} - \text{mean}(\mathbf{x})}{\sqrt{\text{Variance}(\mathbf{x}) + \epsilon}} \right) + \beta$$

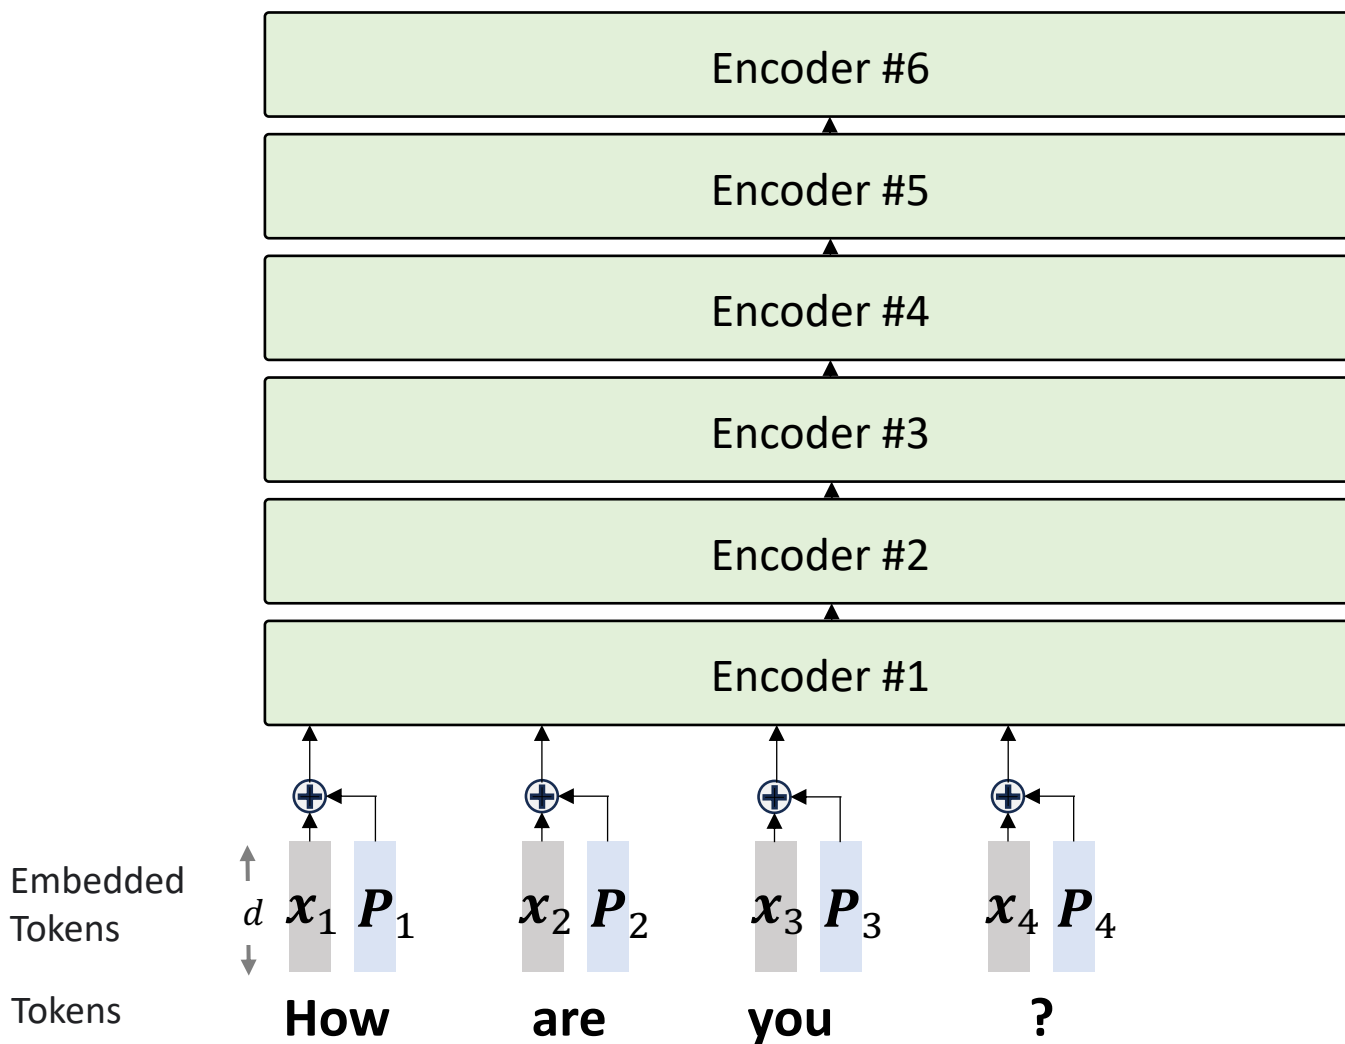
$\gamma, \beta \in \mathbb{R}$


Learnable parameters
















⋮



Embedded
Tokens

d



Tokens

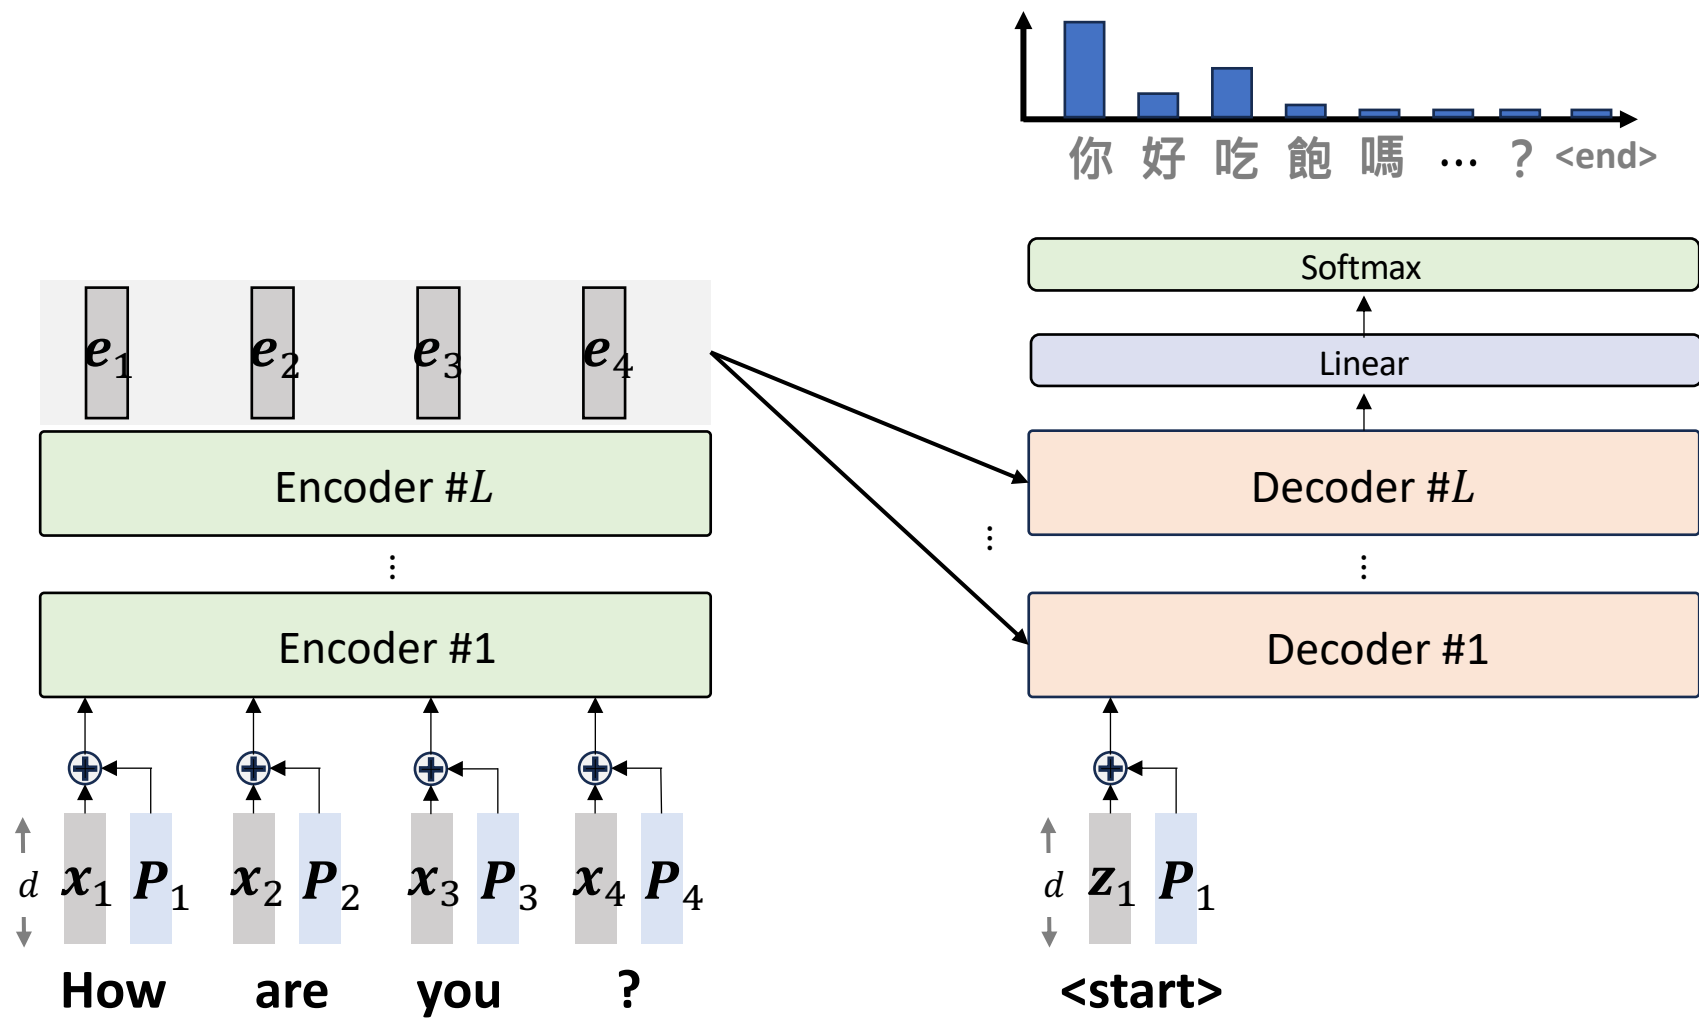
How

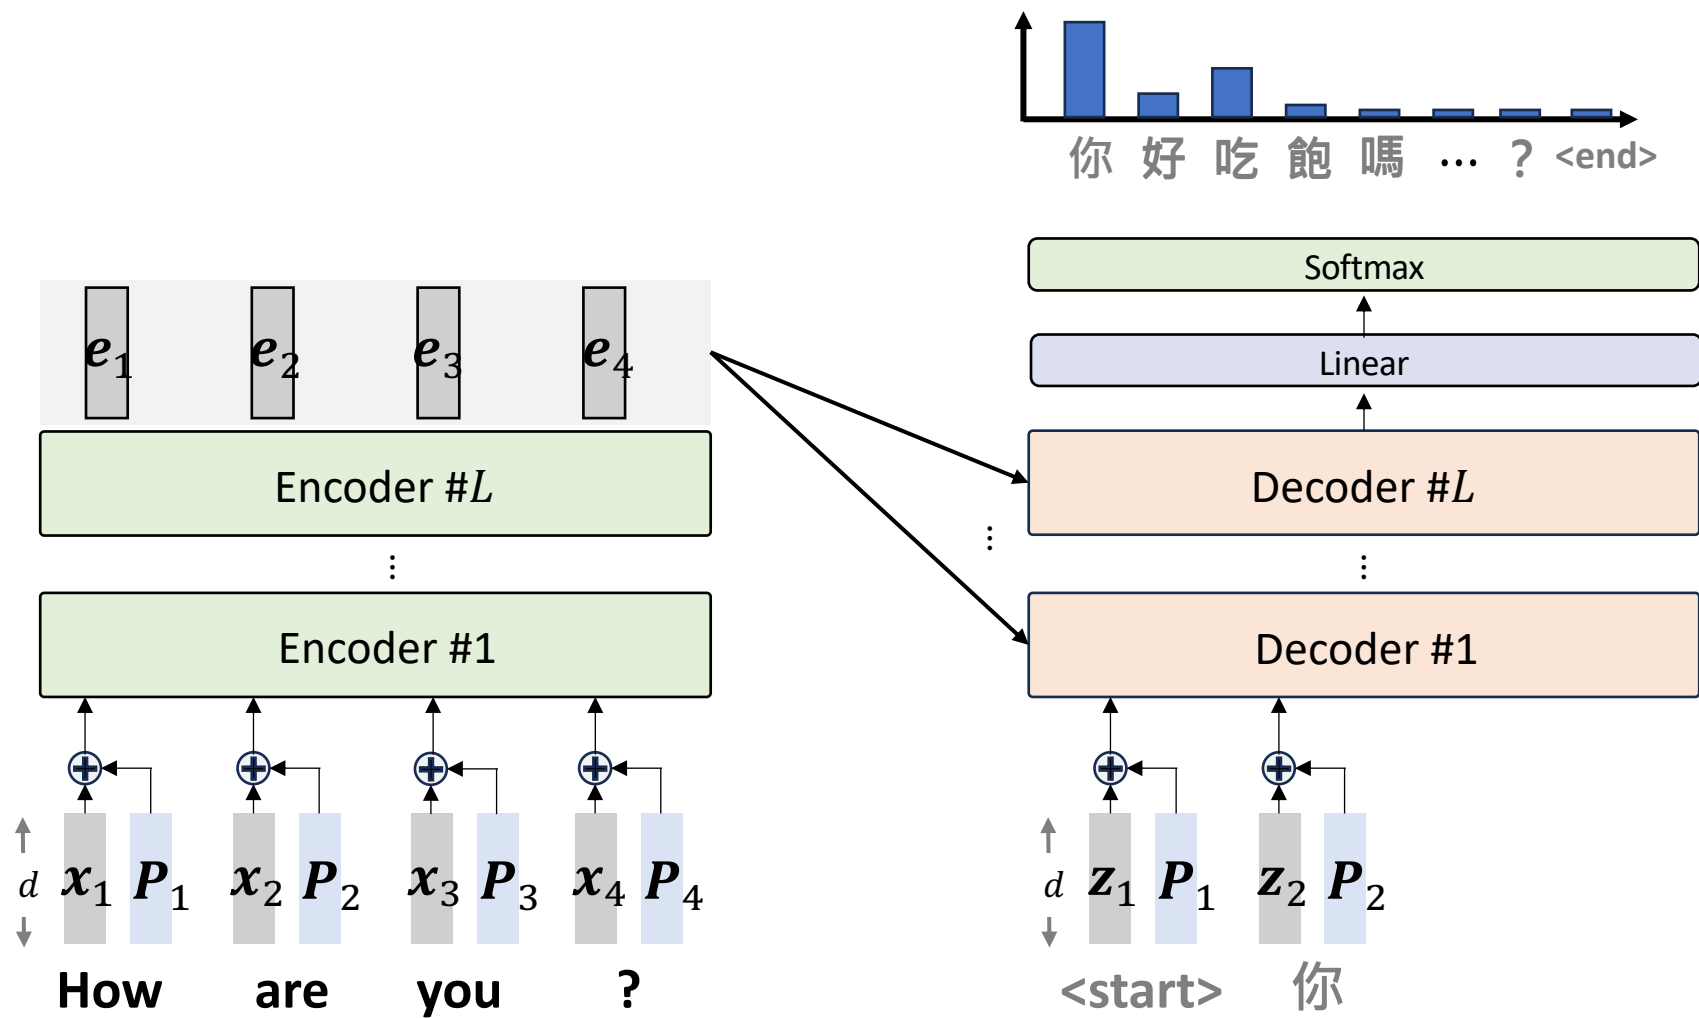
are

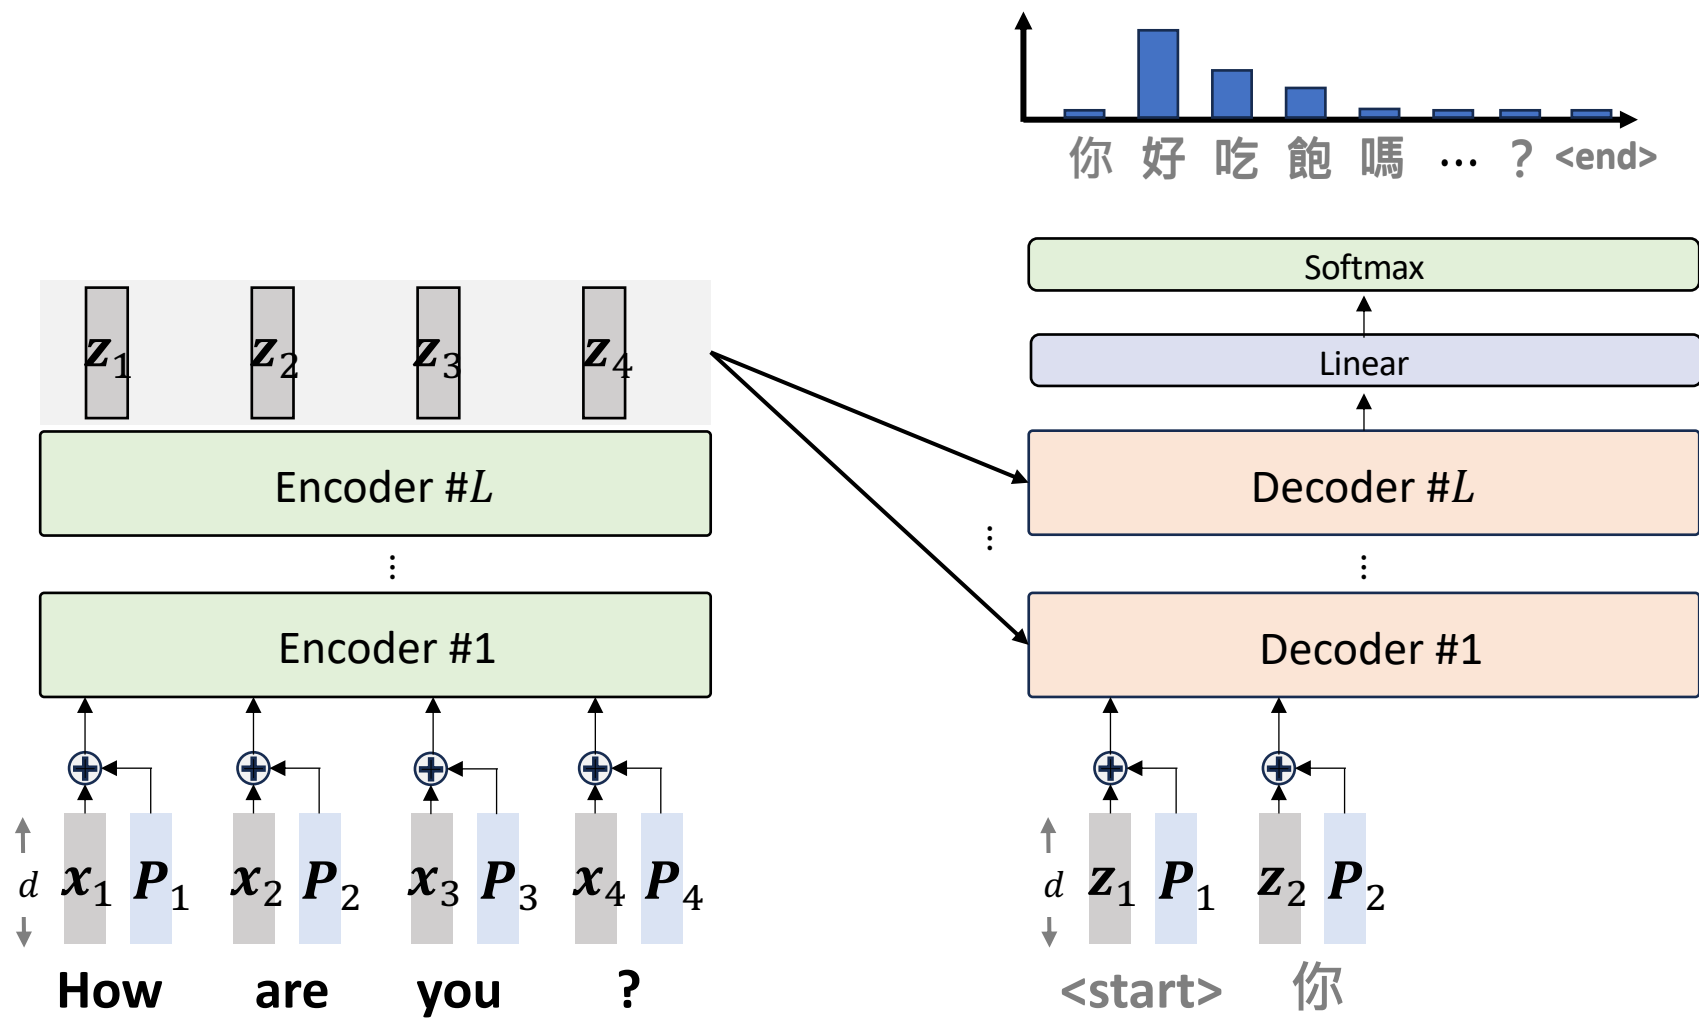
you

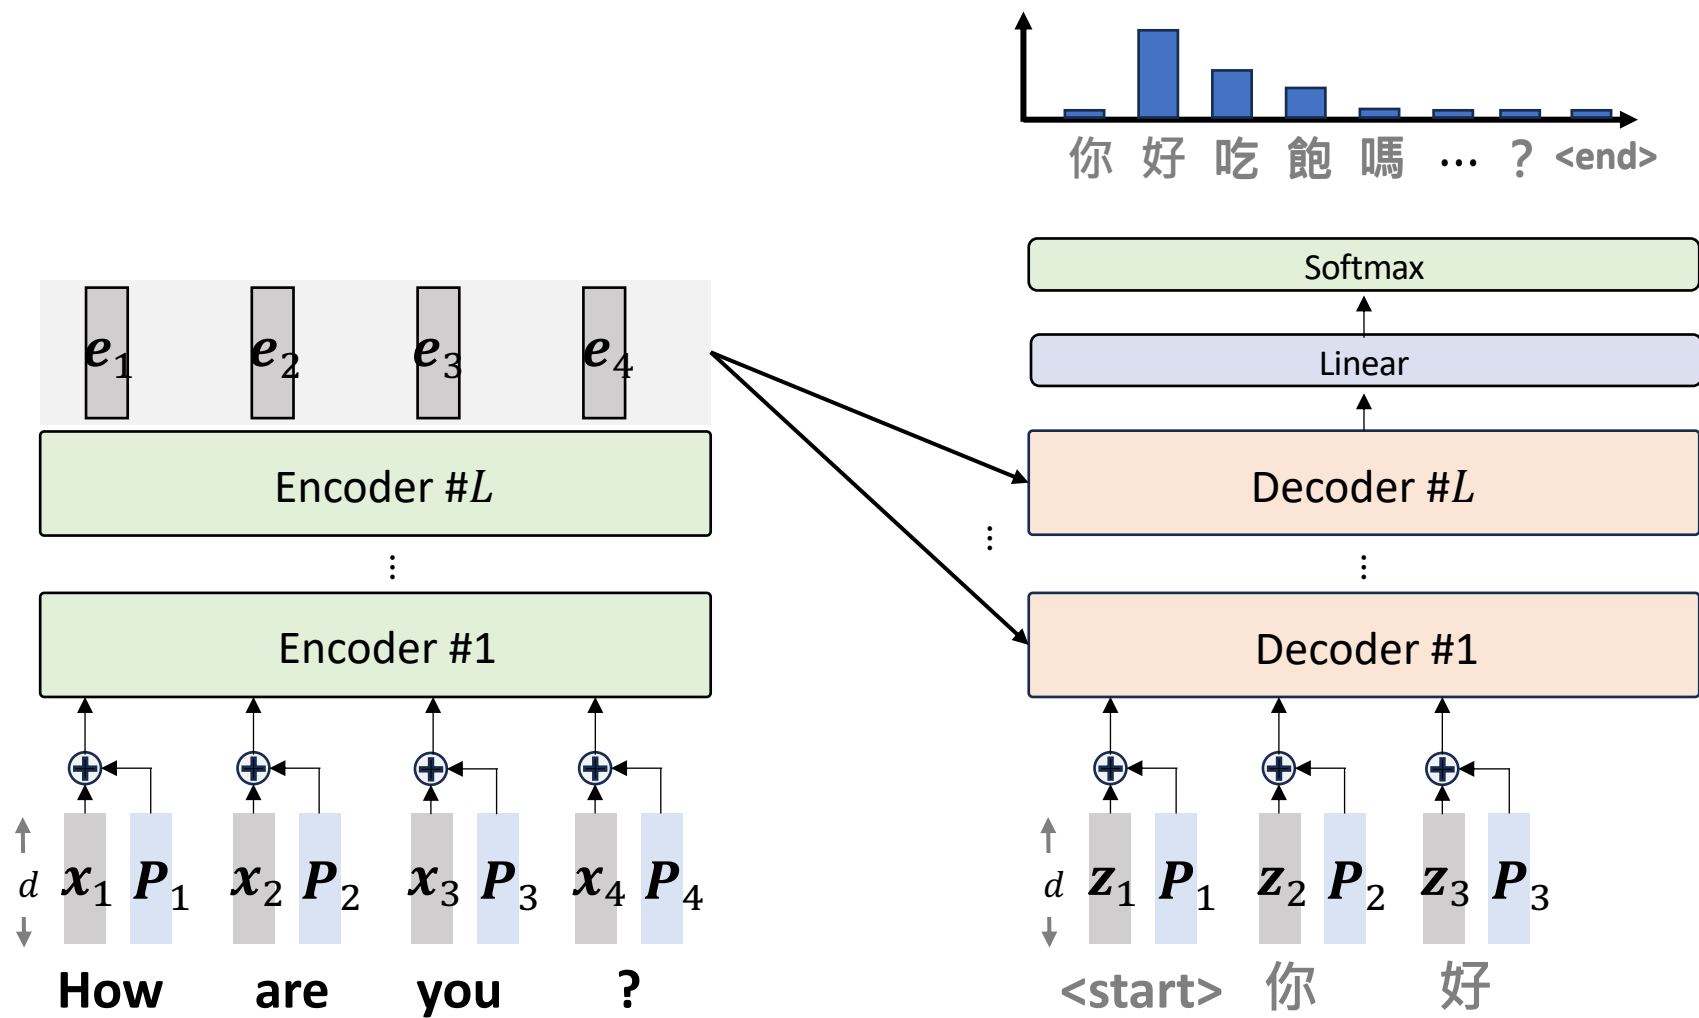
?

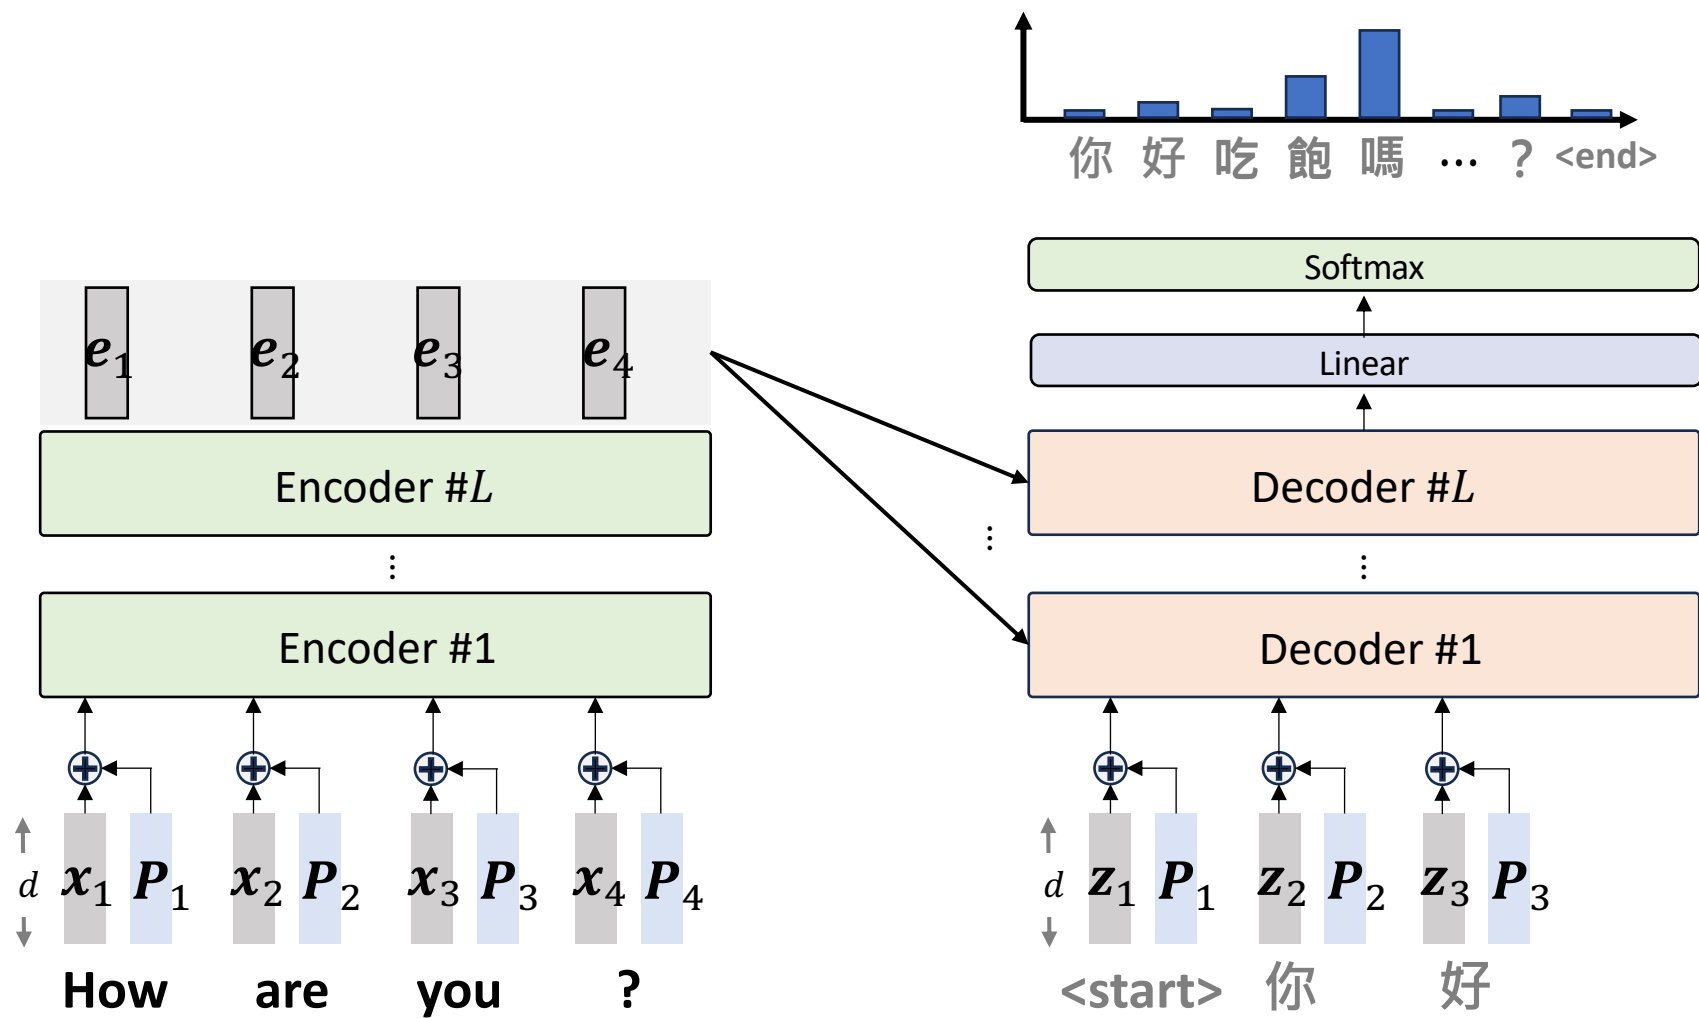


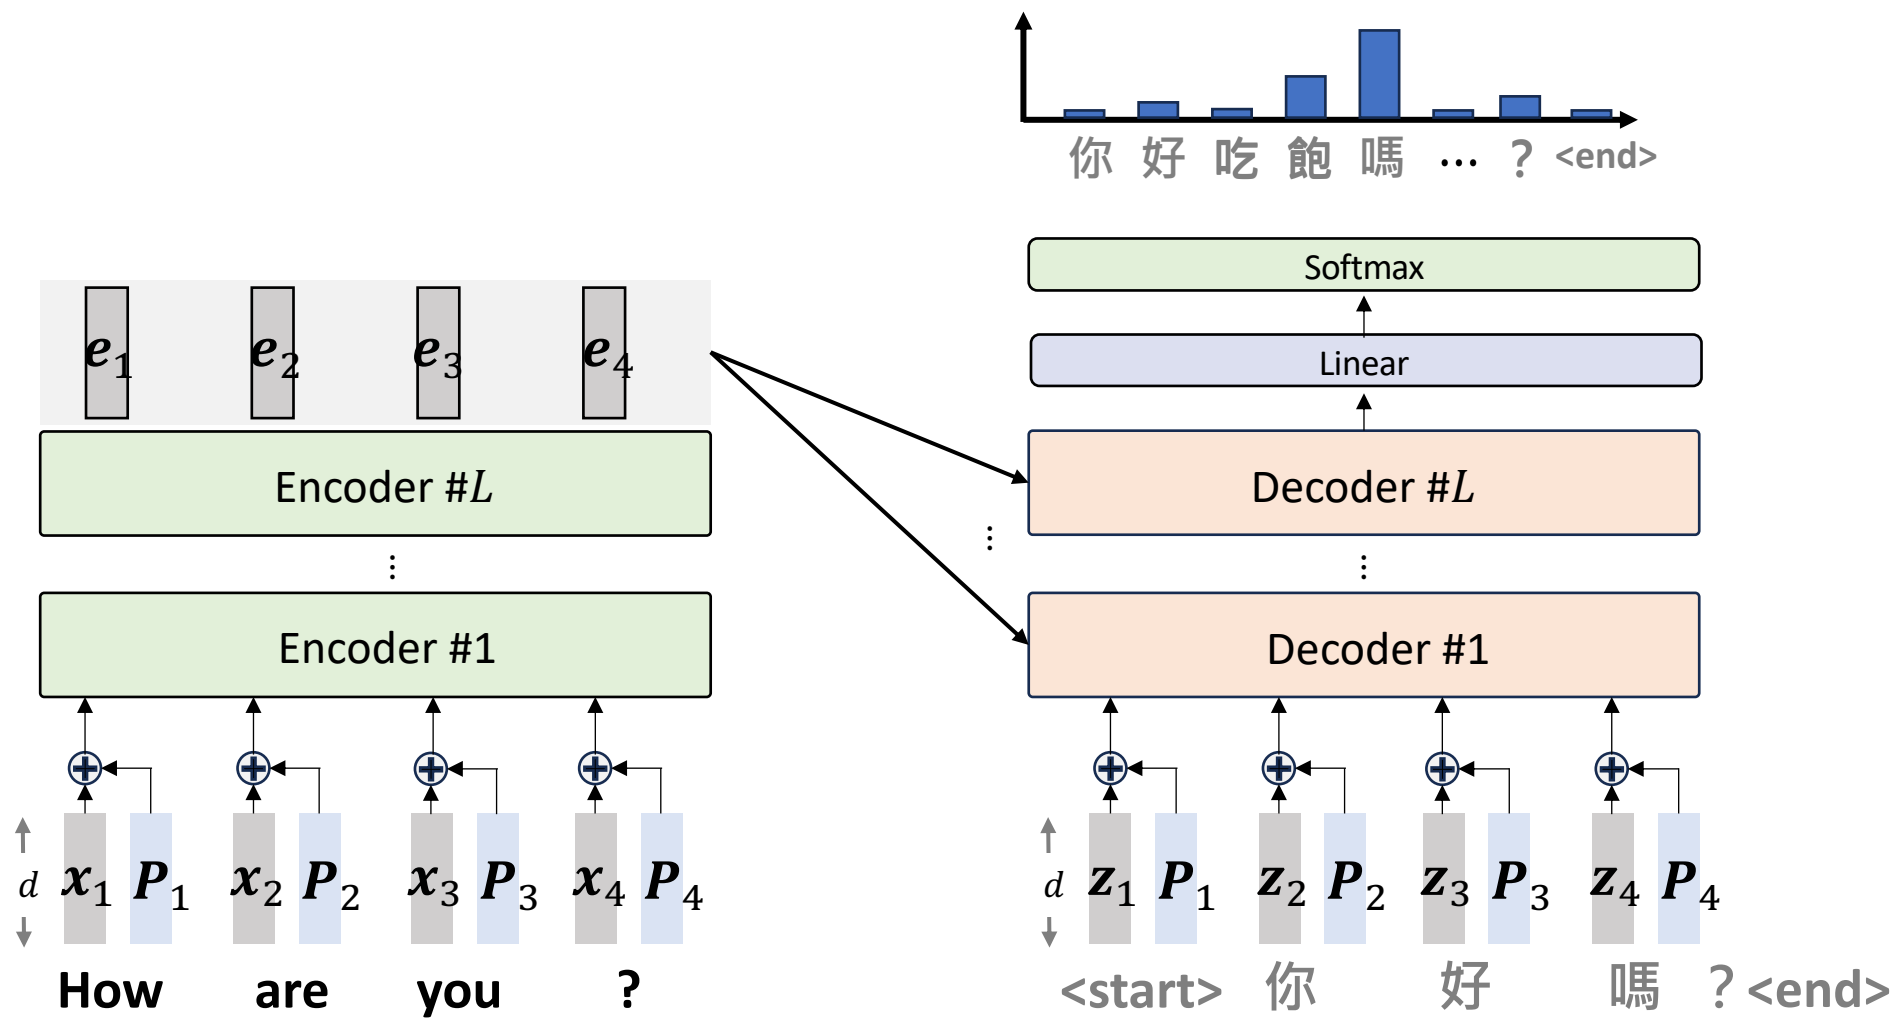


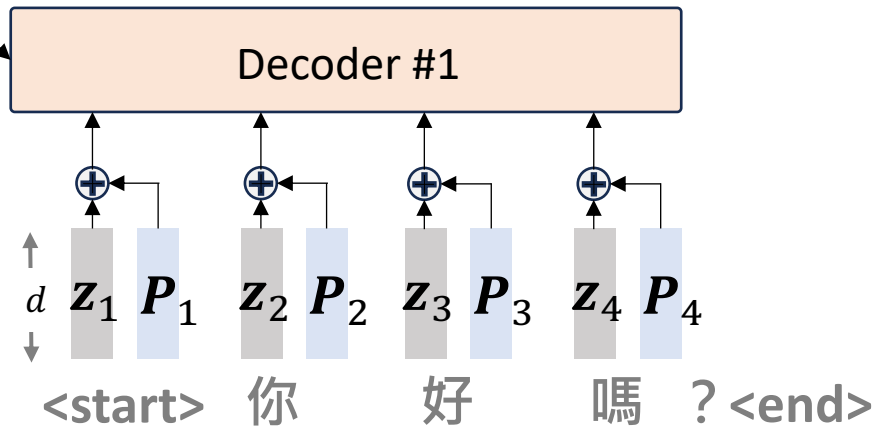
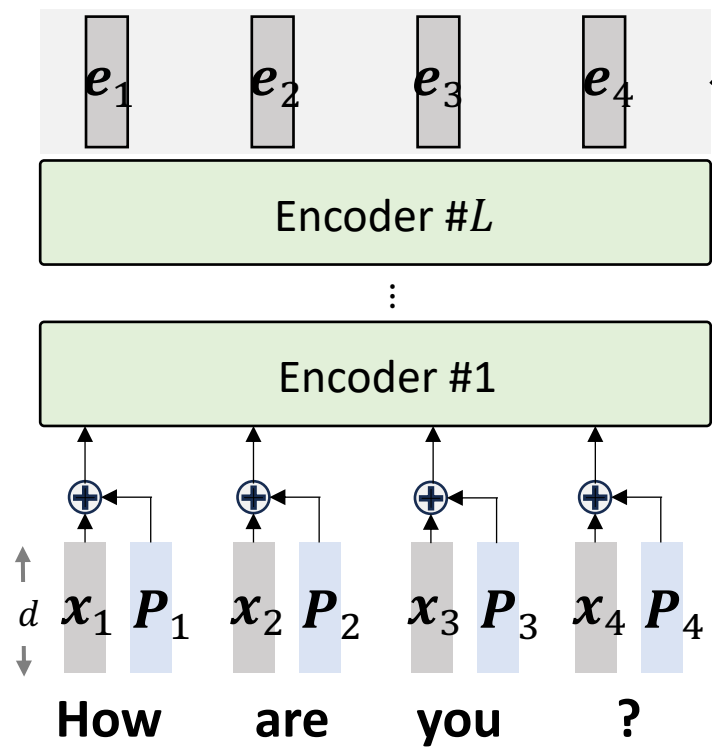


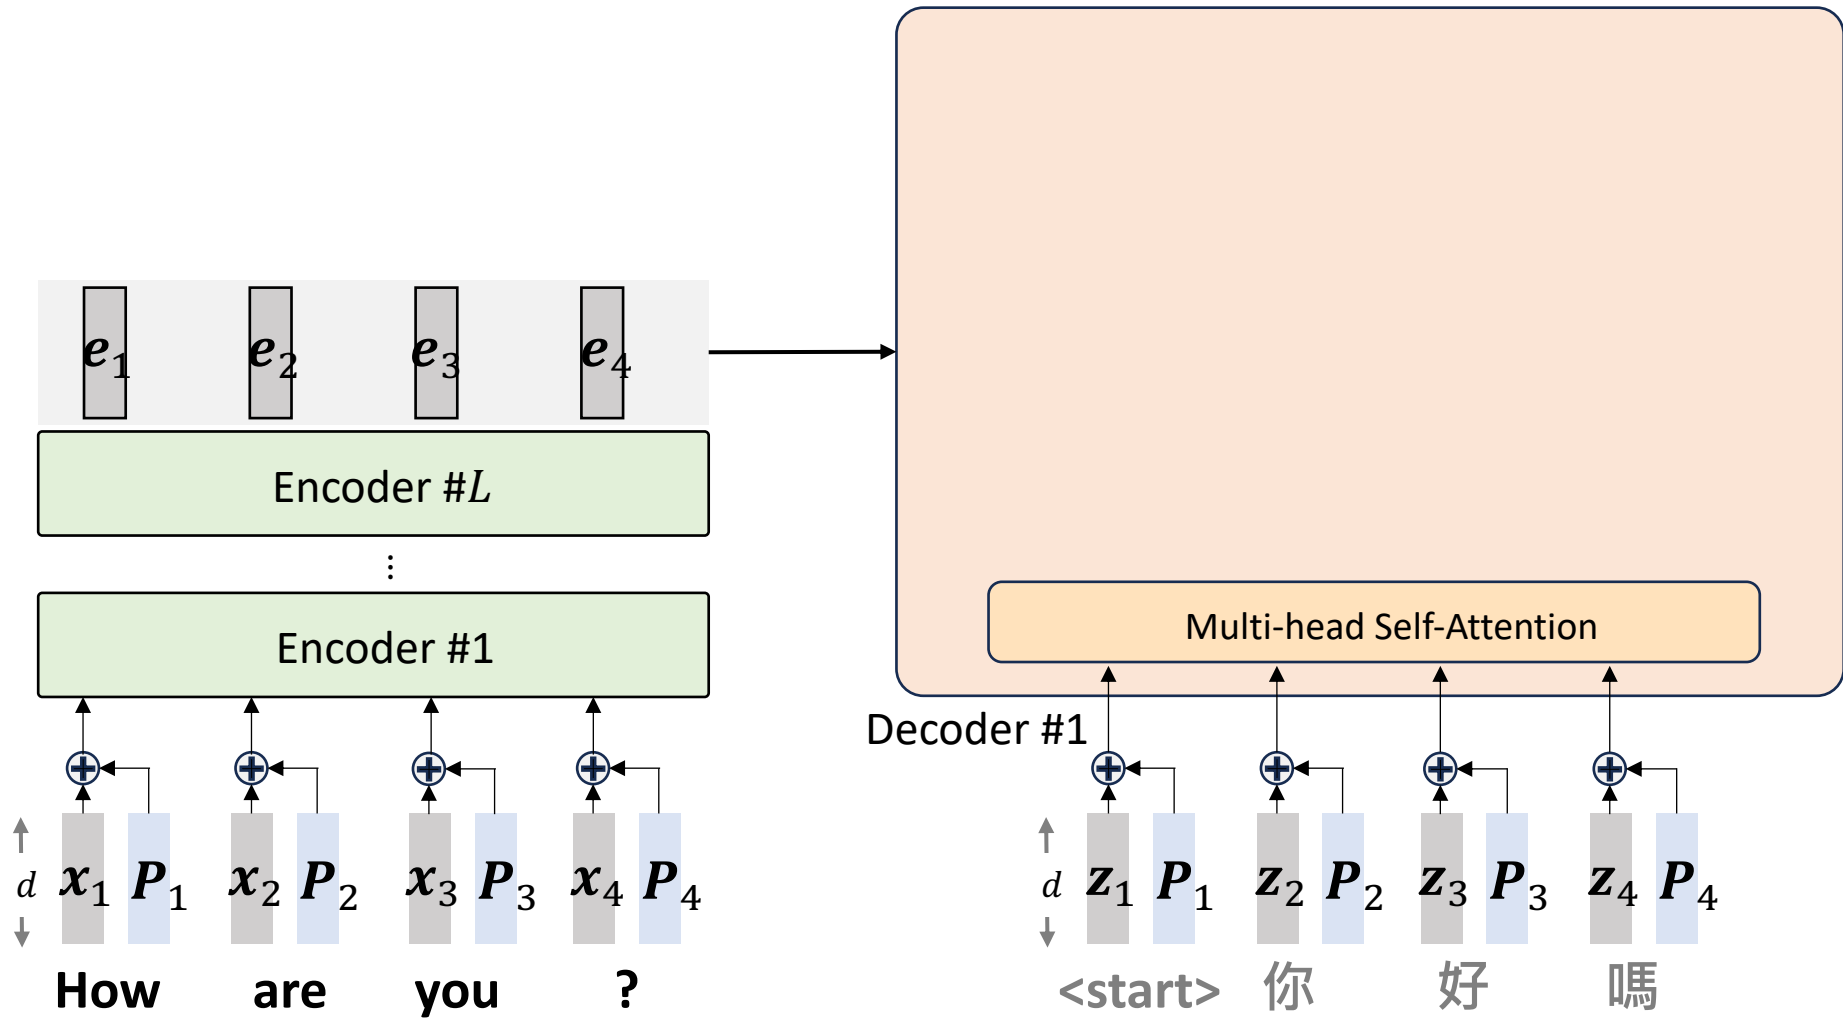


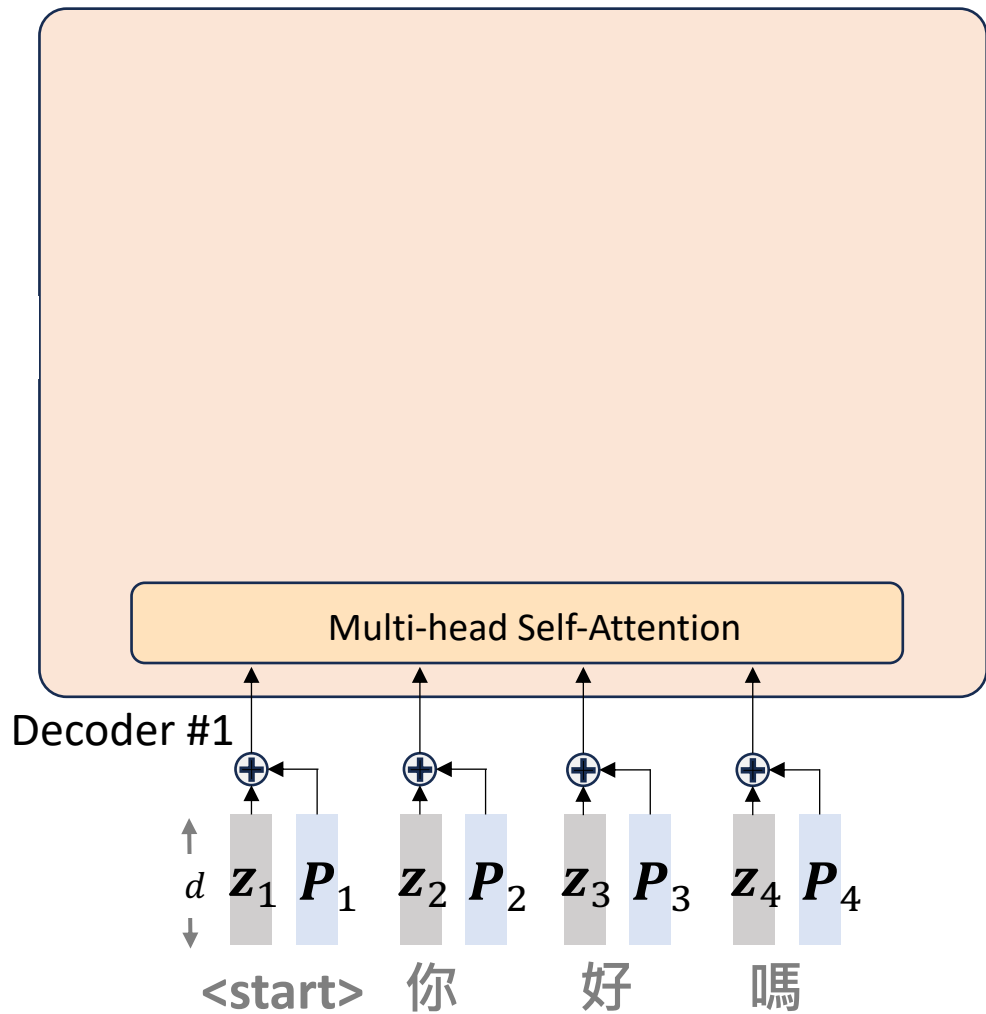


















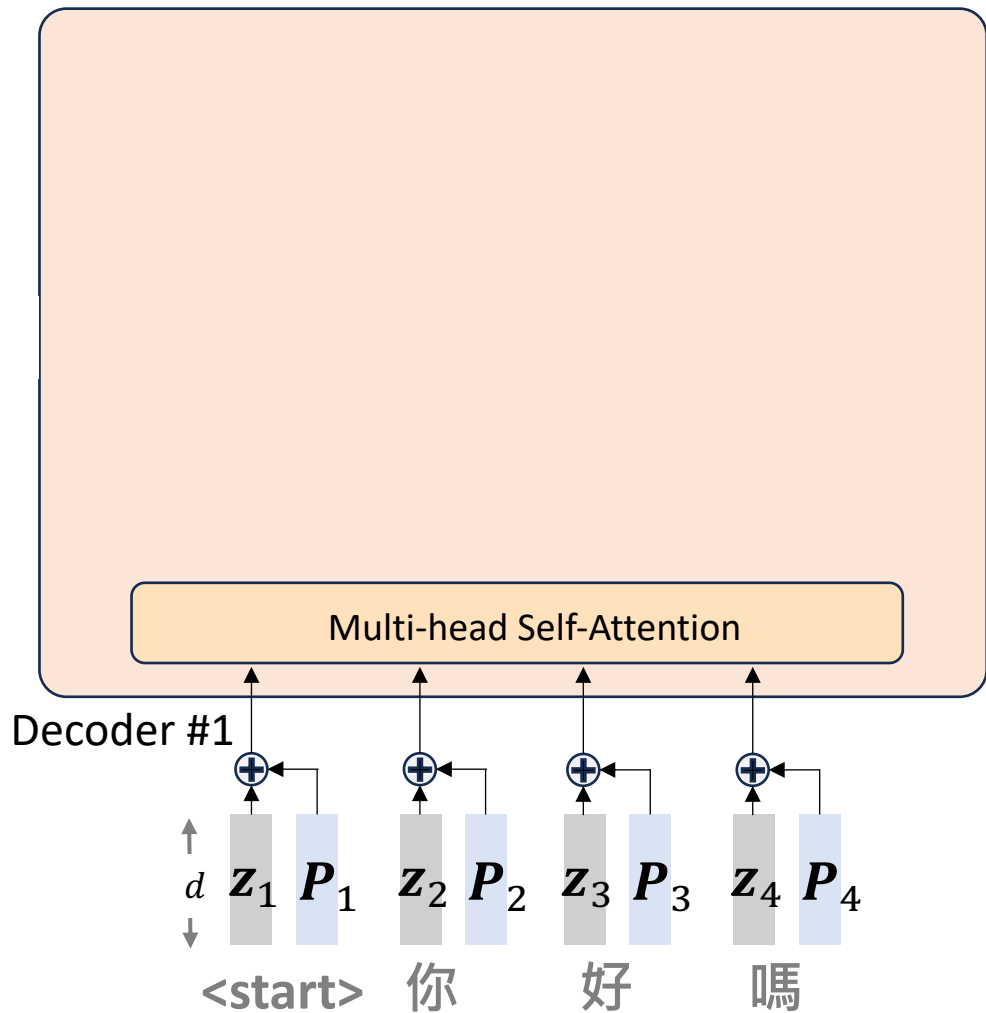



$$\text{Attention}(Q, K, V) = V \text{softmax}\left(\frac{K^T Q}{\sqrt{d_k}}\right)$$

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

$M =$

0	0	0	0	0
$-\infty$	0	0	0	0
$-\infty$	$-\infty$	0	0	0
$-\infty$	$-\infty$	$-\infty$	0	0
$-\infty$	$-\infty$	$-\infty$	$-\infty$	0



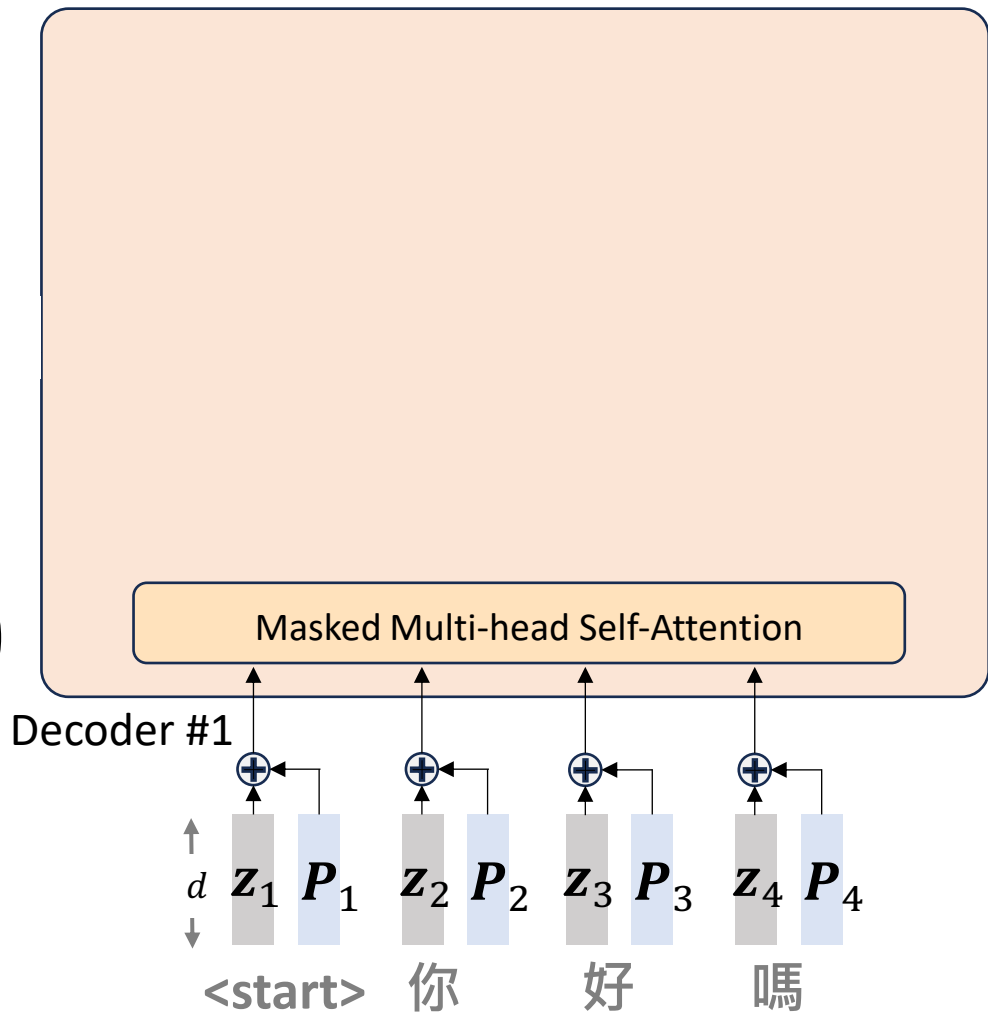


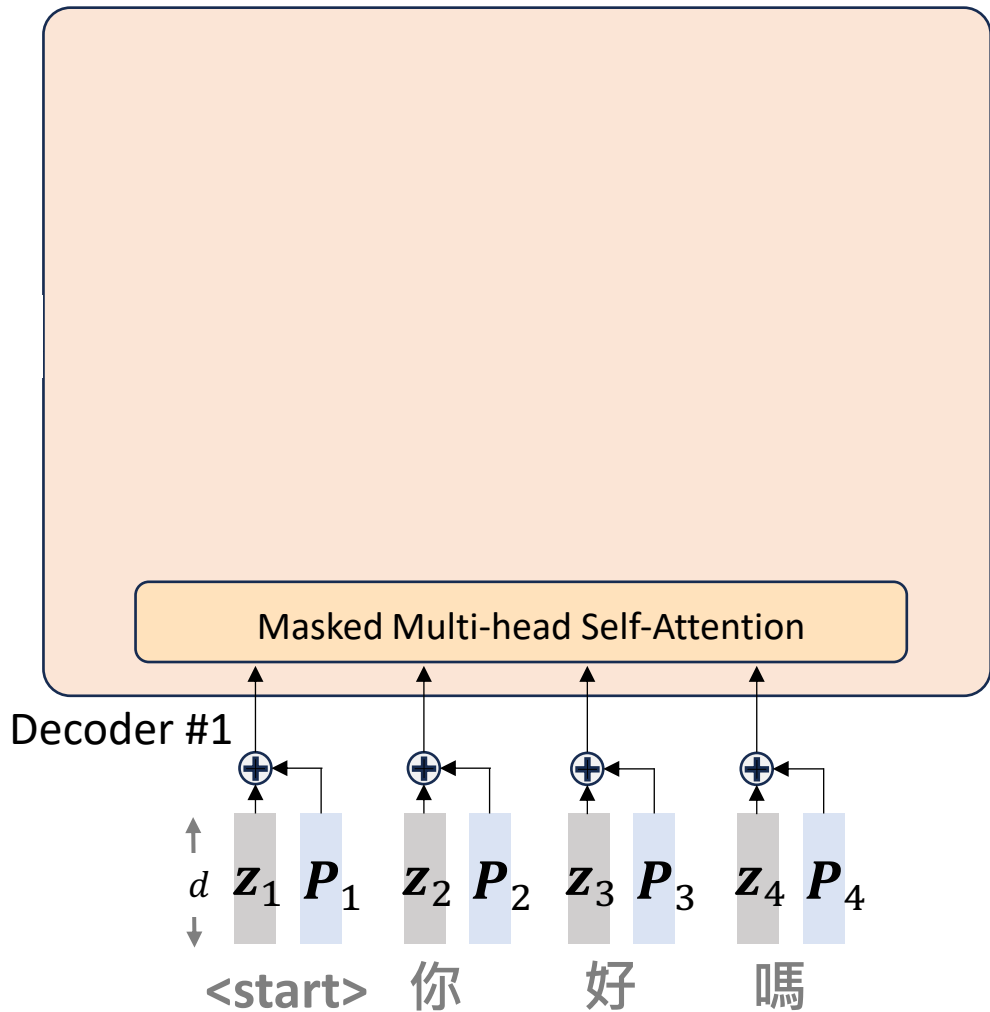


$$\text{Attention}(Q, K, V) = V \text{softmax}\left(\frac{K^T Q}{\sqrt{d_k}}\right)$$

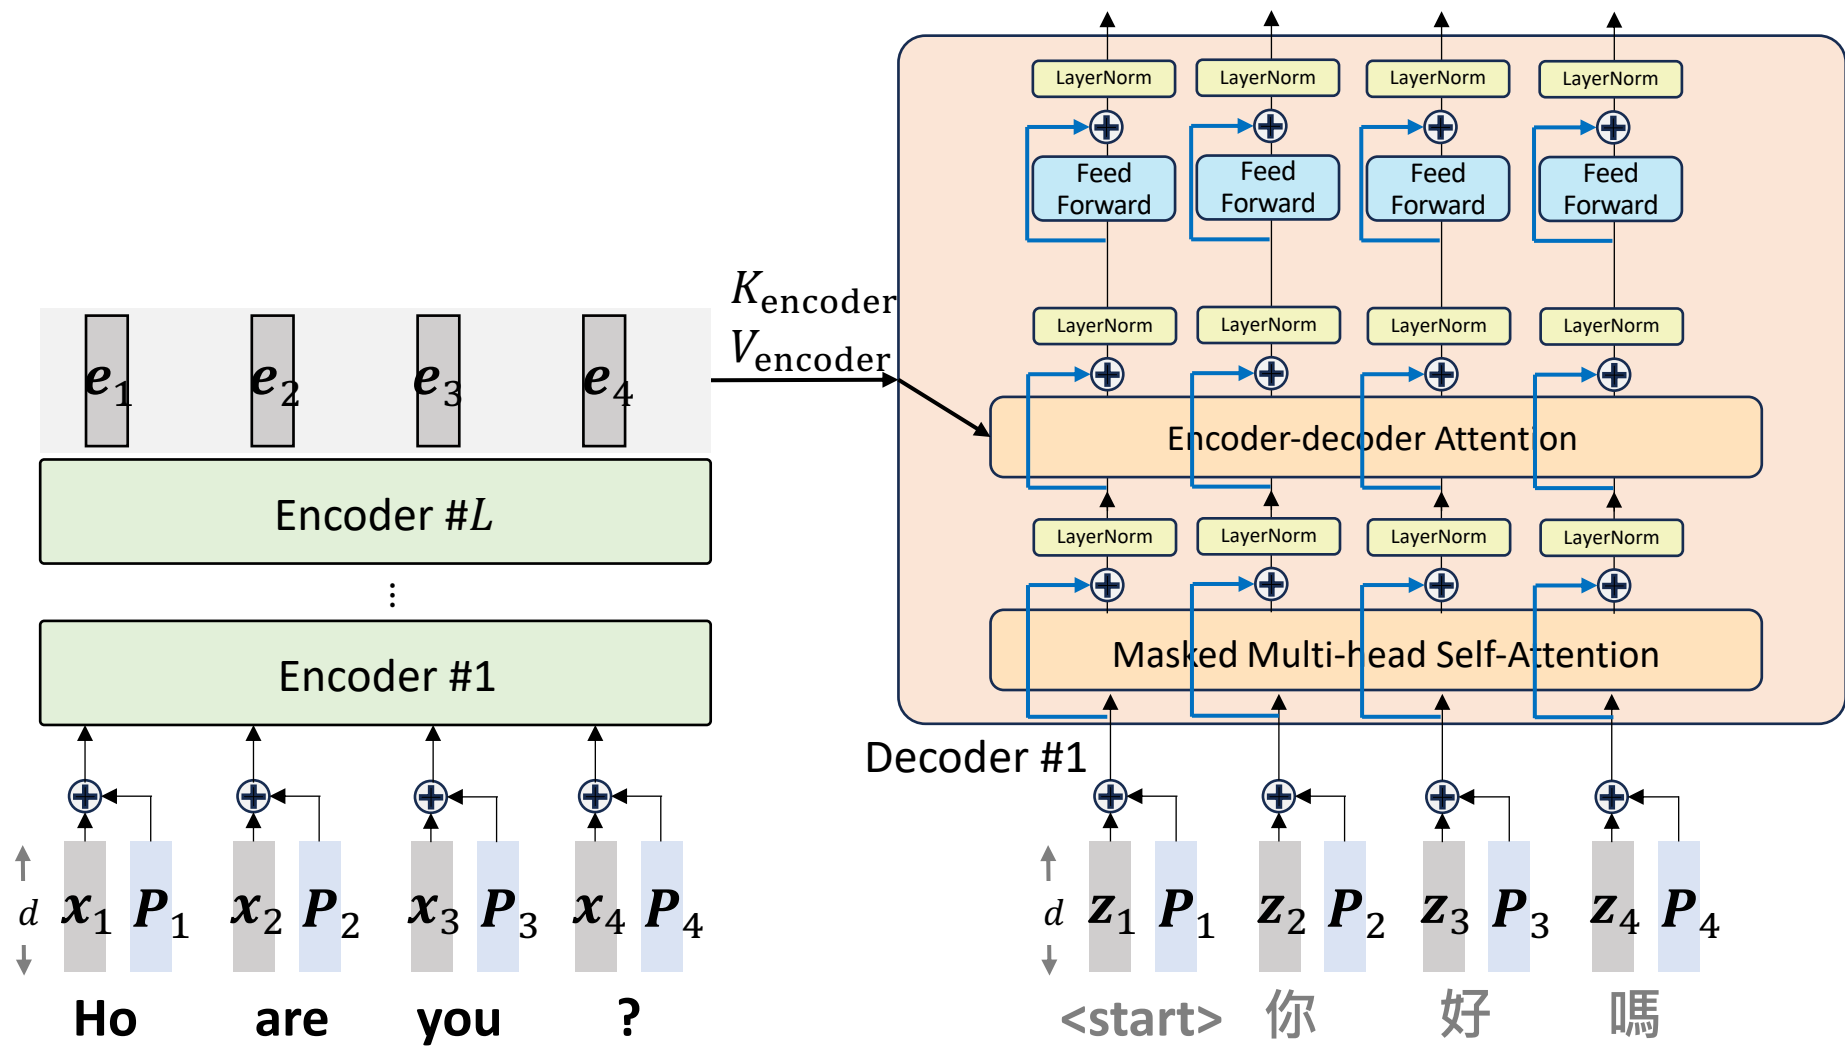
$$\text{MaskedAttention}(Q, K, V) = V \text{softmax}\left(\frac{K^T Q}{\sqrt{d_k}} + M\right)$$



$M =$

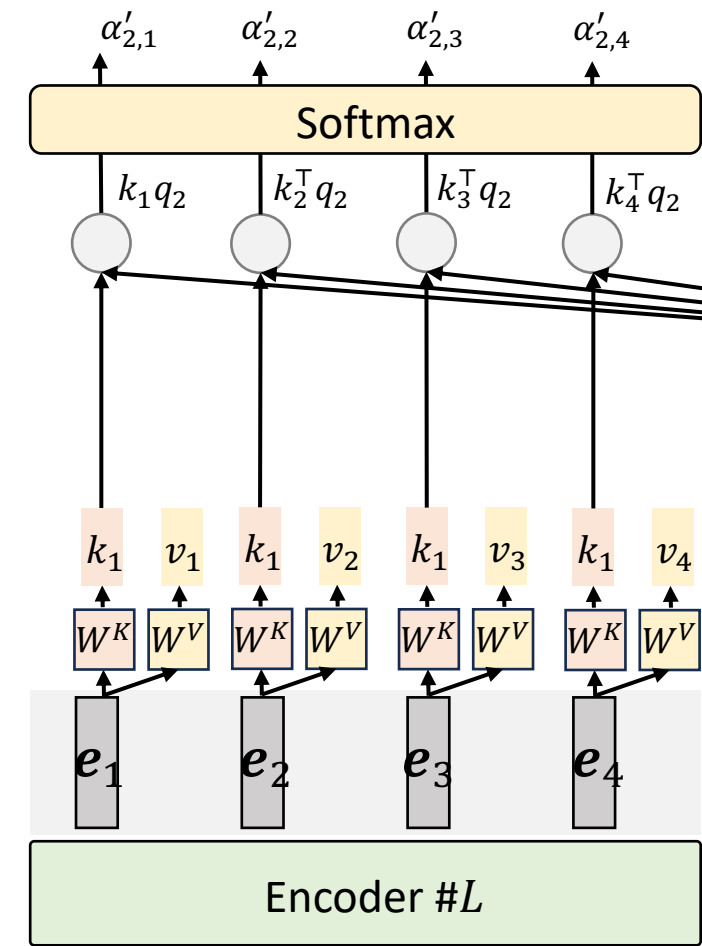
0	0	0	0	0
$-\infty$	0	0	0	0
$-\infty$	$-\infty$	0	0	0
$-\infty$	$-\infty$	$-\infty$	0	0
$-\infty$	$-\infty$	$-\infty$	$-\infty$	0











$$z'_2 = W^o (\alpha'_{2,1} v_1 + \alpha'_{2,2} v_2 + \alpha'_{2,3} v_3 + \alpha'_{2,4} v_4)$$

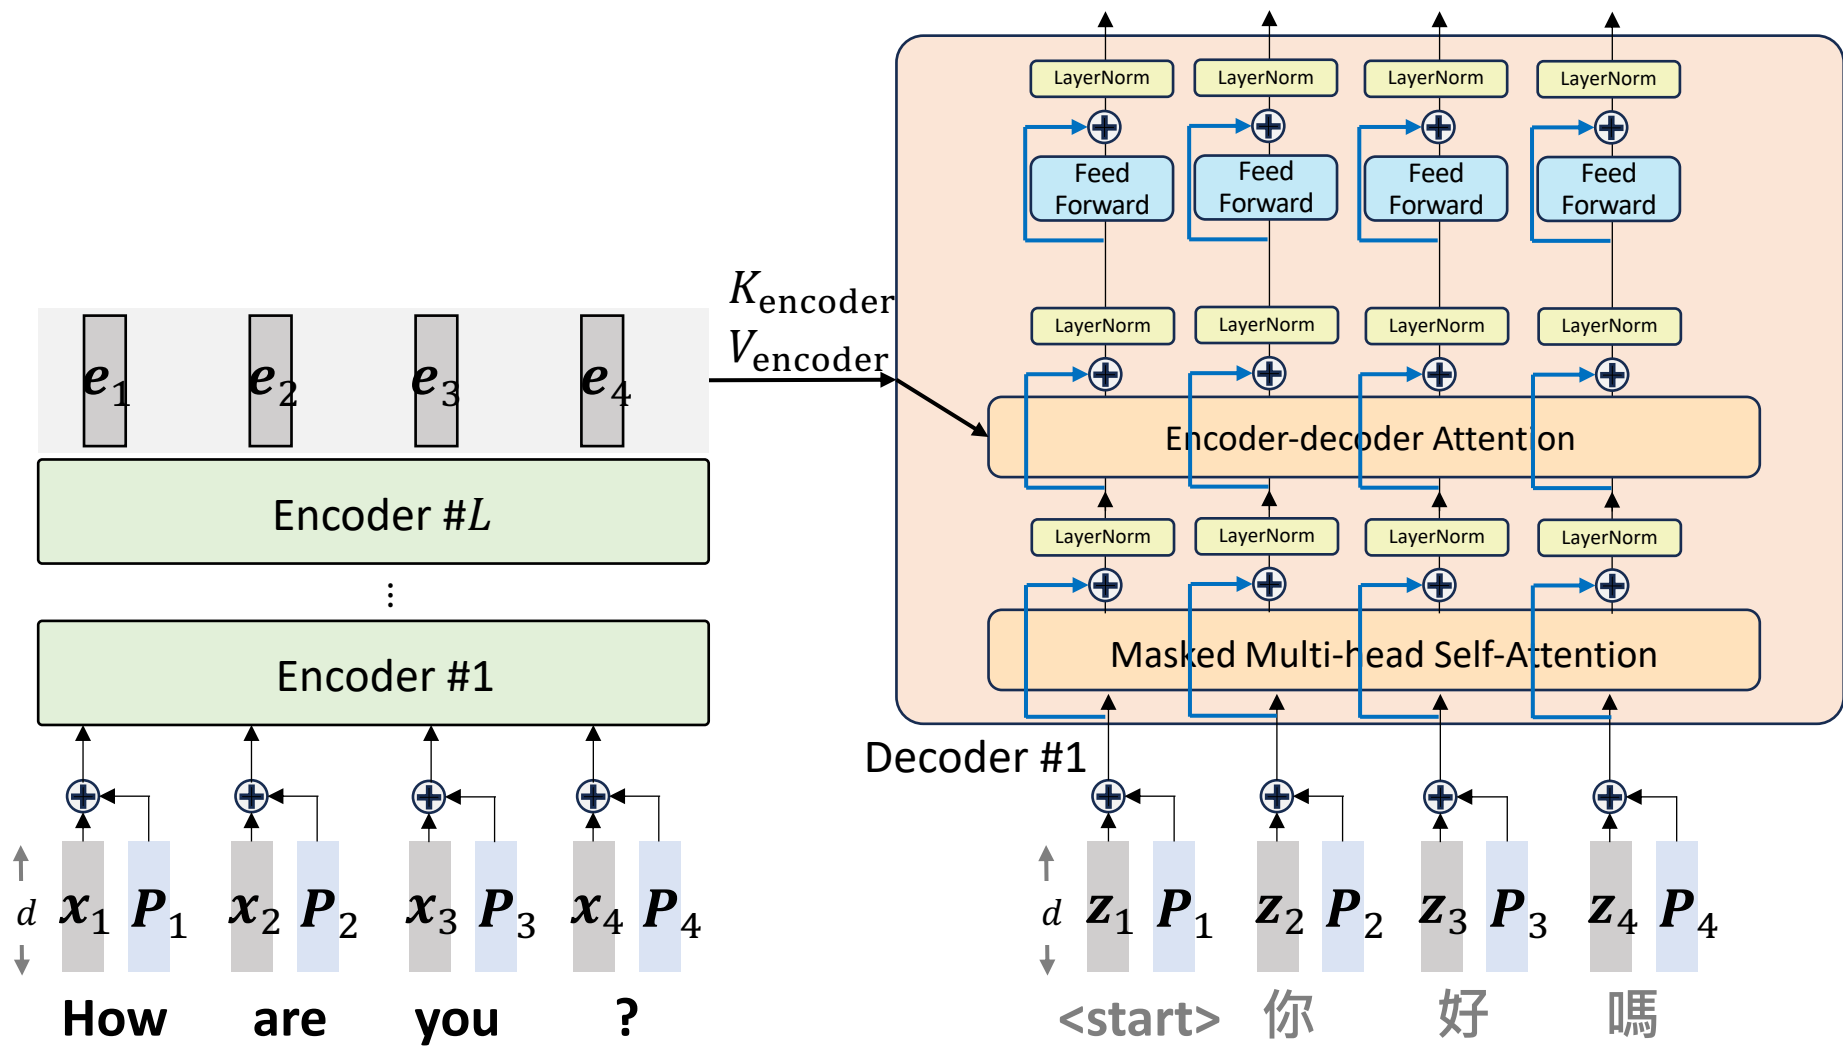
Cross-attention

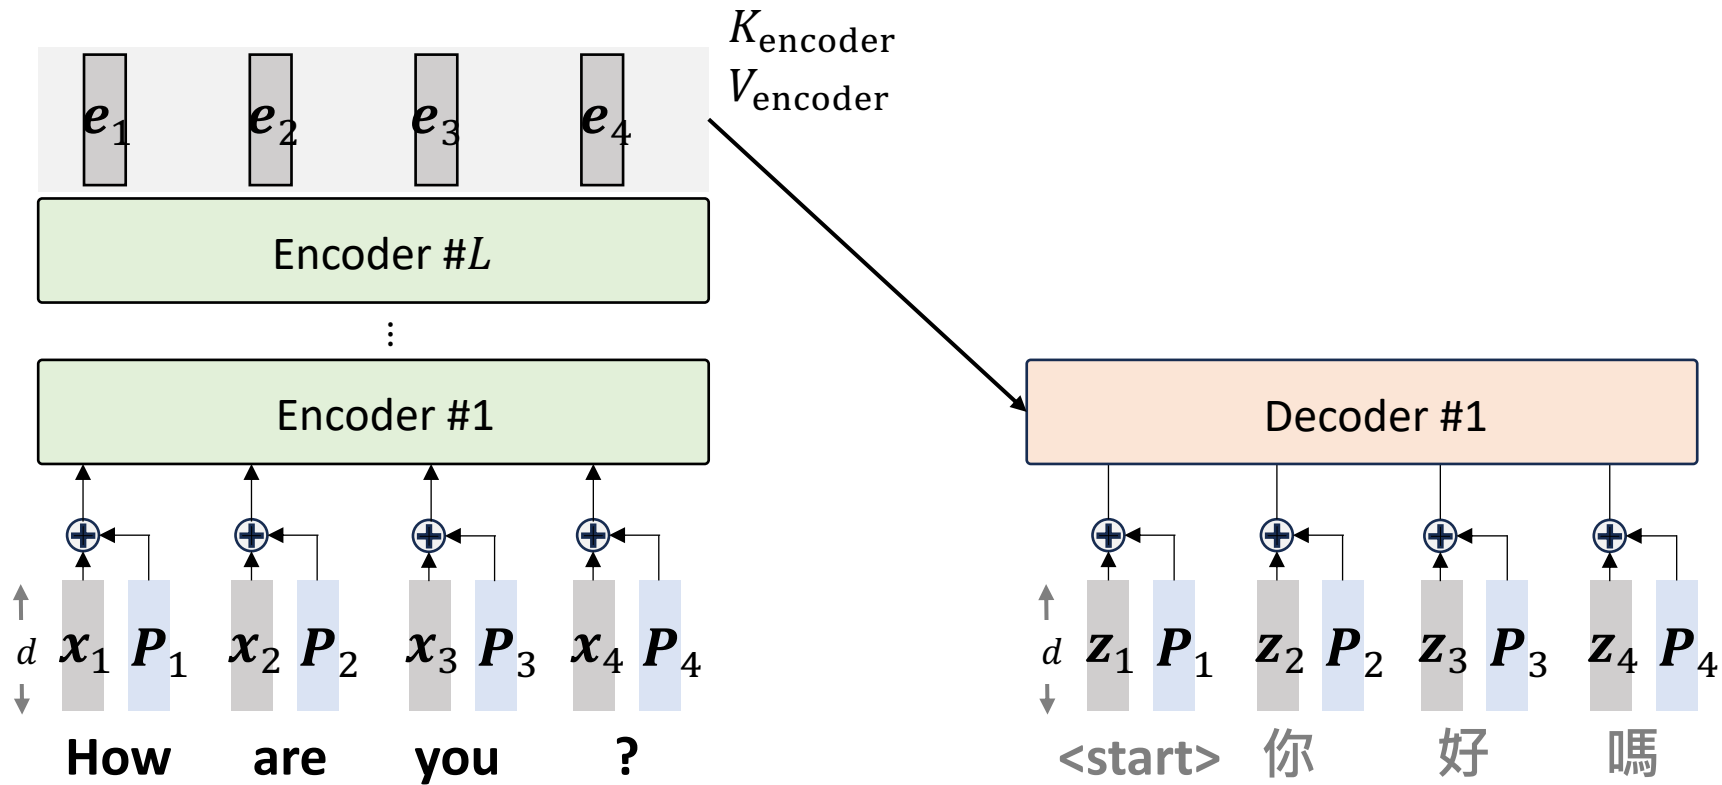
Encoder-decoder attention

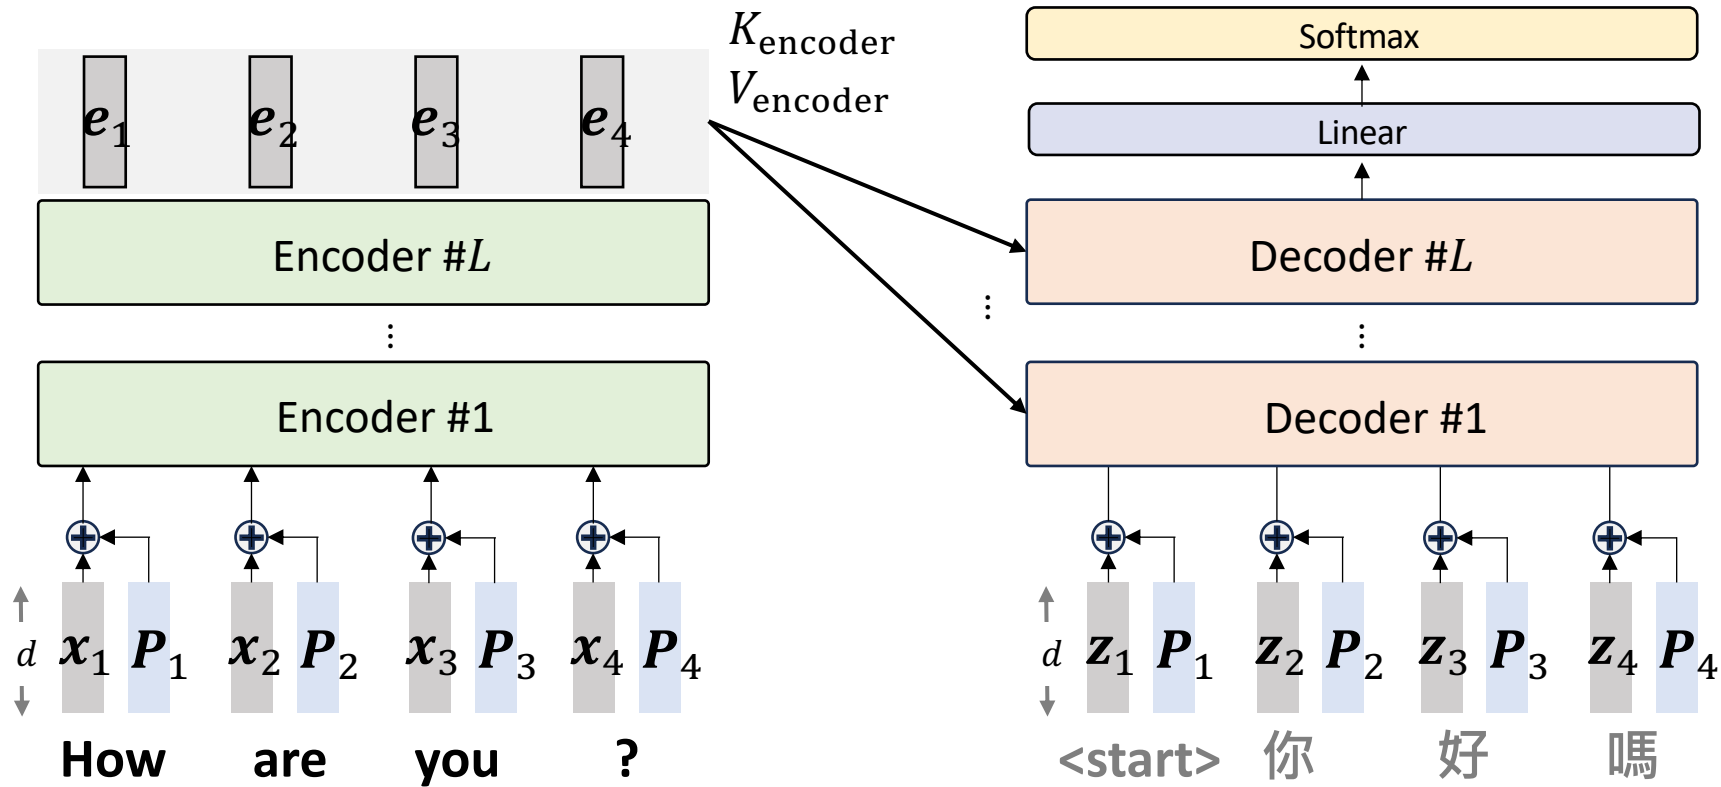
Decoder #1

<start> 你

(ignore the scaling $1/\sqrt{d_k}$ here for simplicity)







Transformer & Multi-Head Attention

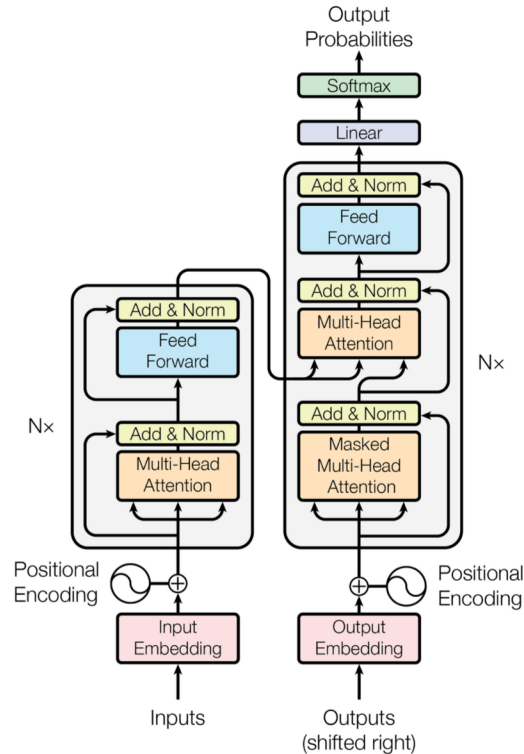


Figure 1: The Transformer - model architecture.

Summary: Attention and Transformers

- ▶ Attention weights used to compute the context vector, which is a weighted sum of the input at different positions
- ▶ Context vector is used to update the hidden state of the model, which is used to generate the final output
- ▶ "Pay attention" to different parts of the input, depending on the task at hand → more accurate and natural-sounding output, esp. when working with longer inputs (e.g. paragraphs)

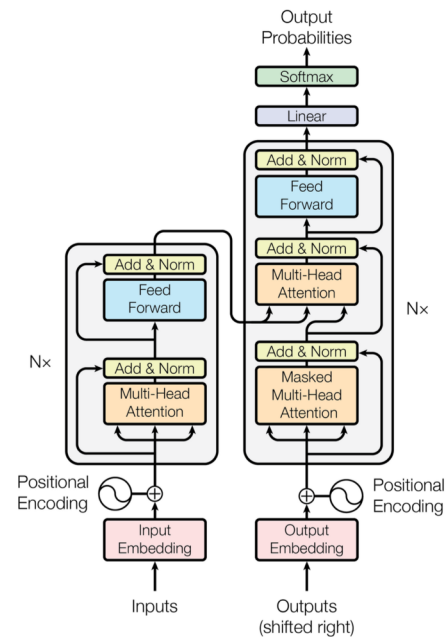


Figure 1: The Transformer - model architecture.

Ways Attention was used in the *original* Transformer Architecture

- **Encoder-decoder cross-attention**

- Allow decoder layers to attend all parts of the latent representation produced by the encoder
- Pull context from the encoder sequence over to the decoder

- **Self-attention in the encoder**

- Allow the model to attend to all positions in the previous encoder layer
- Embeds context about how elements in the sequence relate to one another

- **Masked self-attention in the decoder**

- Allow the model to attend to all positions in the previous decoder layer up to and including the current position (during auto-regressive process)
- Prevent forward looking bias by stopping leftward information flow during training
- Also embed context about how elements in the sequence relate to one another

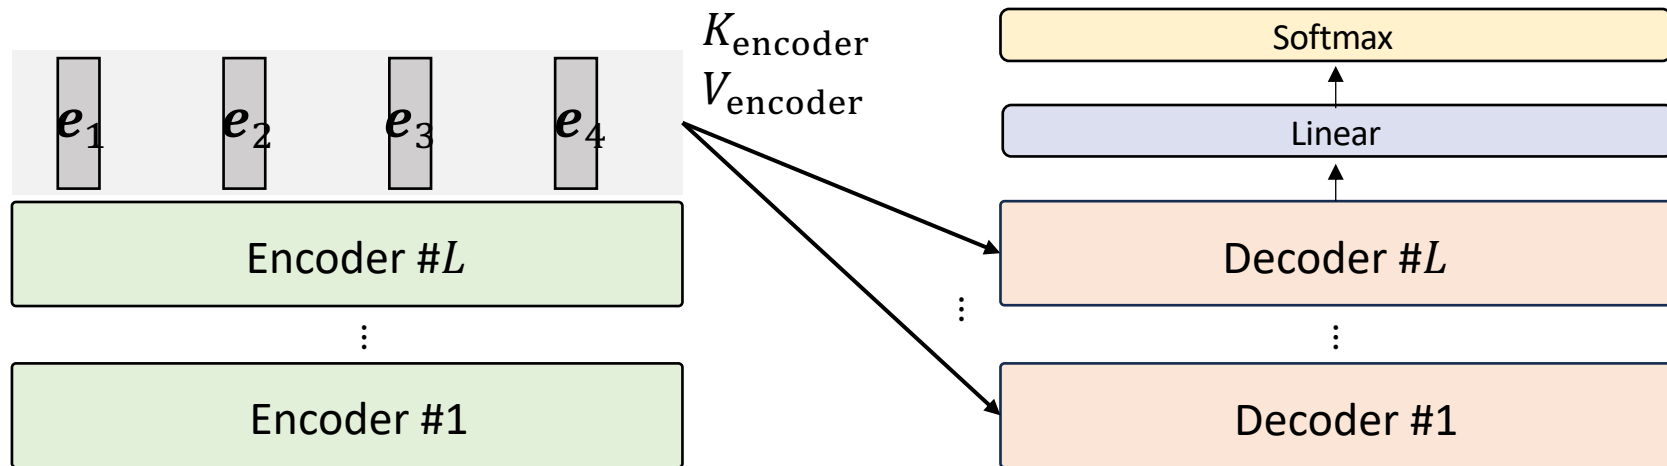
Encoder-Decoder Transformer

Examples:


Attention is all you need, T5, BART.

Good for:

Machine translation, summarization. QA
(when input/target are sufficiently different)



Encoder-Decoder Transformer



Examples:

Attention is all you need, T5, BART.

Good for:

Machine translation, summarization. QA
(when input/target are sufficiently different)

Encoder-only Transformer




Examples:

BERT, RoBERTa, DeBERTa, X-BERT

Good for:

Classification, sequence tagging, sentiment analysis

(Understand text, but not generate them)



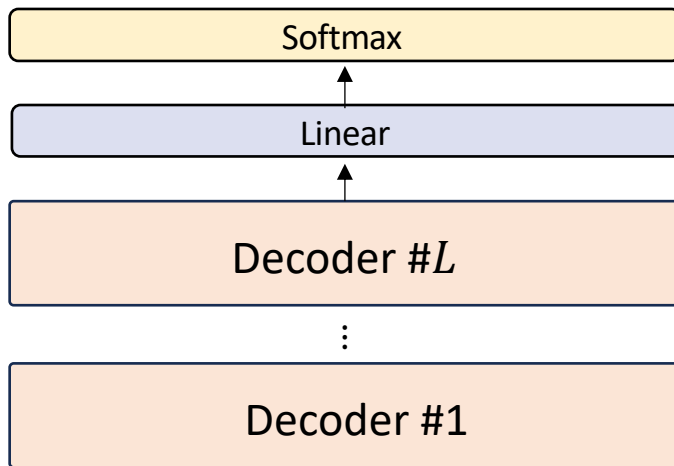
Decoder-only Transformer

Examples:

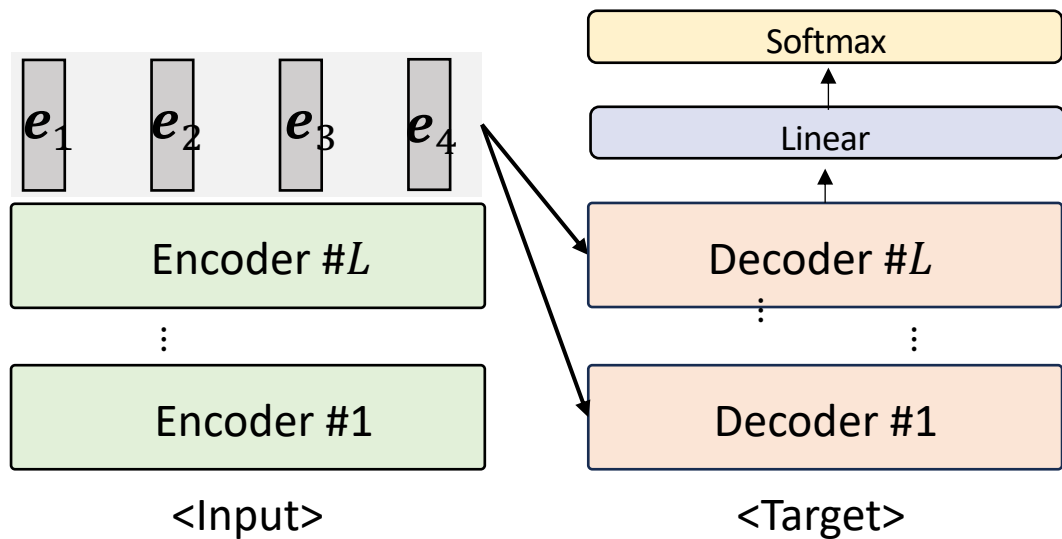
GPT-X (OpenAI), PaLM (Google), LLaMA (Meta)
BLOOM (BigScience)

Good for:

Text generation, multi-round conversation

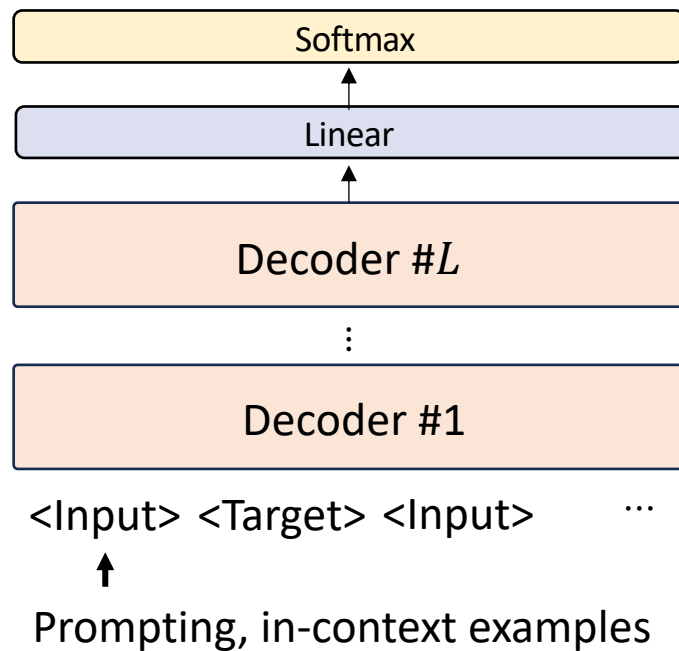


Encoder-Decoder Transformer



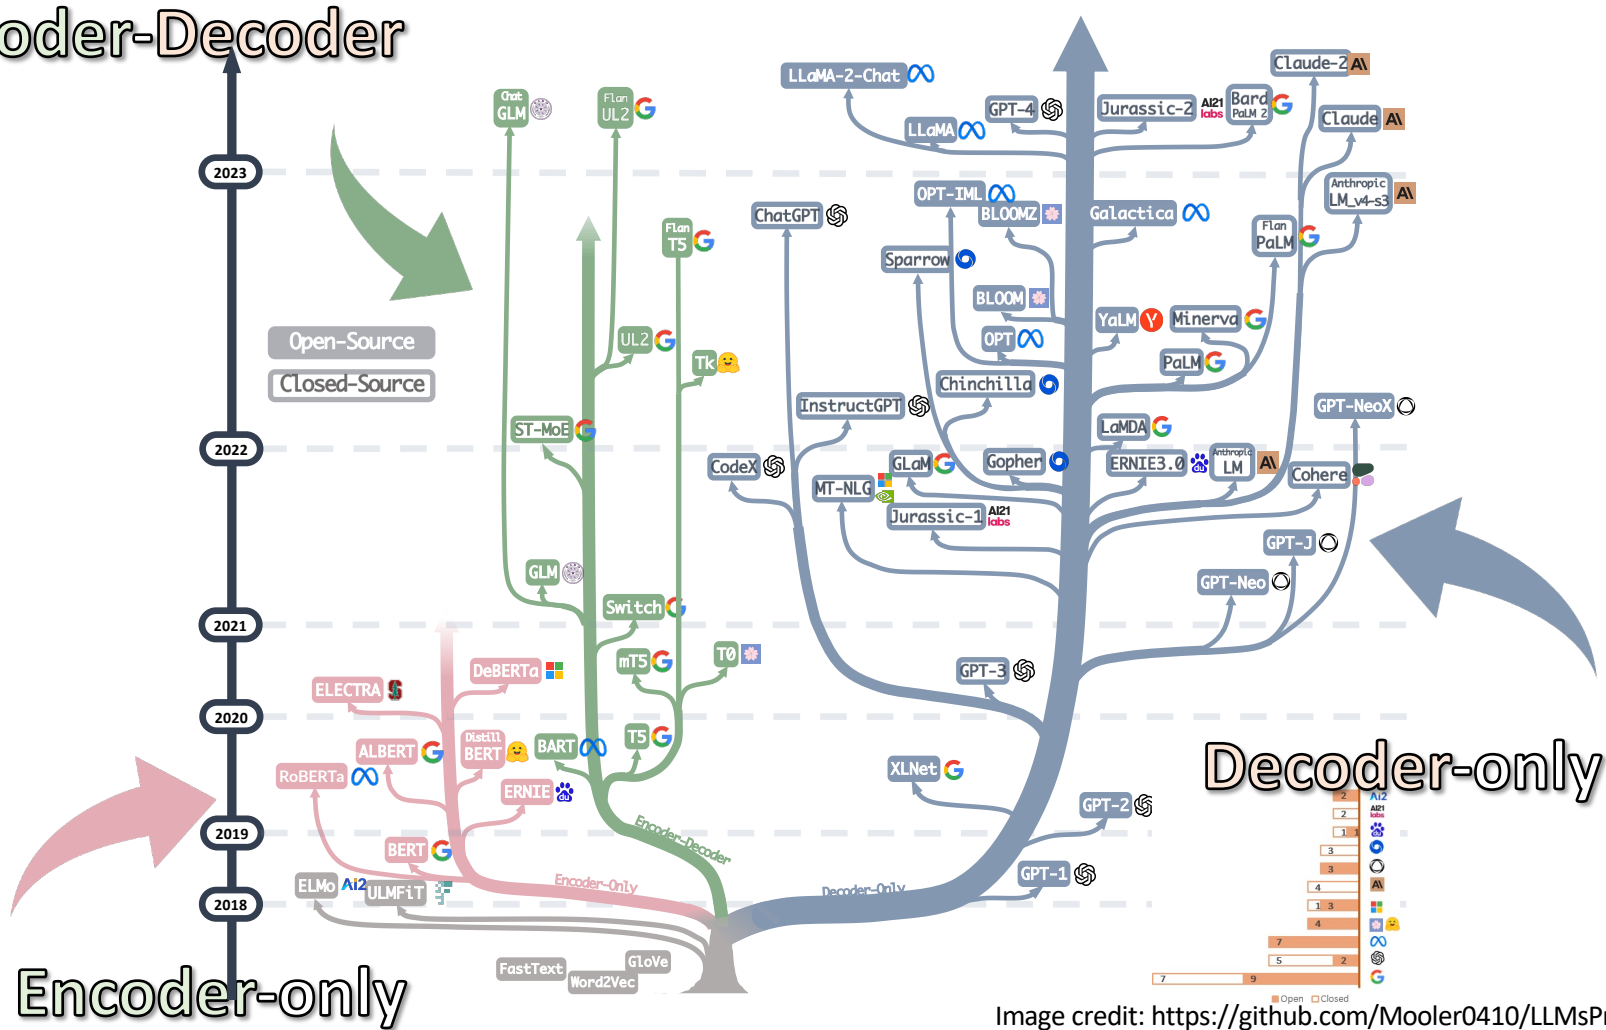
Different parameters for **encoder**/**decoder**

Decoder-only Transformer



Shared parameters

Encoder-Decoder



Encoder-only

Decoder-only

Transformers vs. RNNs

Challenges with RNNs	Transformers
<ul style="list-style-type: none">● Long range dependencies● Gradient vanishing and explosion● Large # of training steps● Sequential/recurrence → can't parallelize● Complexity per layer: $O(n*d^2)$	<ul style="list-style-type: none">● Can model long-range dependencies● No gradient vanishing and explosion● Fewer training steps● Can parallelize computation!● Complexity per layer: $O(n^2*d)$

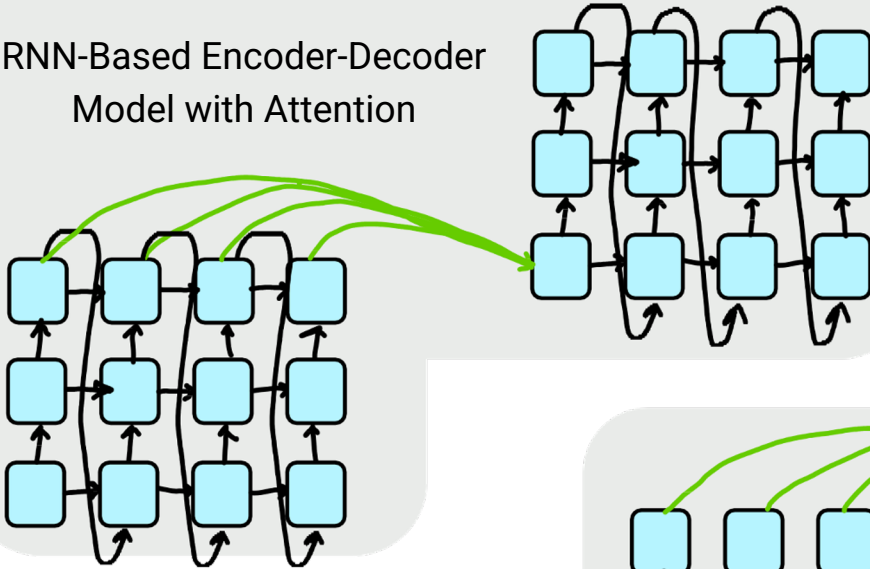
- ▶ When sequence length (n) \ll representation dimension (d), the complexity per layer is lower for a Transformer model compared to RNN models ; no true for real-world LLMs

Differences in Attention Mechanism of RNN vs. Transformer

Feature	RNN with Attention (Bahdanau et al. 2015)	Transformer
Attention Type	Additive (Bahdanau) Attention	Scaled Dot Product Attention
Alignment	Based on decoder hidden state and encoder hidden states	Based on dot-product of query and keys (global attention)
Efficiency	Processes sequences step-by-step	Parallel processing of all positions
Context	Weighted sum of encoder hidden states at each step	Attends to all encoder positions for every output
Self-Attention	Not used	Self-attention in both encoder and decoder

Computational Dependencies for Recurrence vs. Attention

RNN-Based Encoder-Decoder Model with Attention



Transformer Advantages:

- # unparallelizable operations does not increase with sequence length.
- Each word interacts with each other, so maximum interaction distance is $O(1)$.

Transformer-Based Encoder-Decoder Model

