Deep Generative Models: Transformers

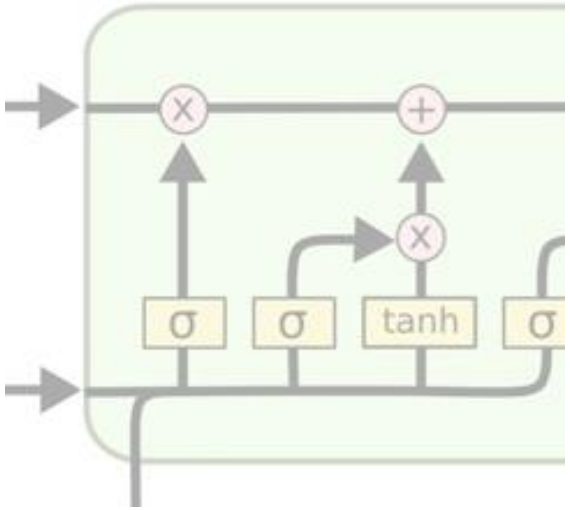
Fall Semester 2026

René Vidal
Director of the Center for Innovation in Data Engineering and Science (IDEAS)
Rachleff University Professor, University of Pennsylvania
Amazon Scholar & Chief Scientist at NORCE



Lc

s!



1990s — **Recurrent Neural Networks**

RNNs, GRUs, LSTMs

2014 — **Simple attention mechanisms**

2017 — **Beginning of transformers**

Attention is all you need

2018 — **Explosion of transformers in NLP**

BERT, GPT-3

2018-2020 — **Explosion into other fields**

Explosion into other fields: ViTs, Alphafold-2.

2021-2022 — **Start of Generative Era**

Codex, Decision Transformers, GPT-X, DALL·E

2025 — **Present Day**

Huge models, more applications: Chat-GPT, GPT-4, Gemini, Llama and open-source LLMs, Whisper, Robotics Transformer, Stable Diffusion, Sora, and so much more...!

Future (?!)



RNN vs Transformers: Why RNNs fall short?

- **RNNS**
  - **Hard to capture long-term dependencies** without modifying the architecture.
  - **Hard to train** due to vanishing and exploding gradients.
  - **Hard to process in parallel** due to sequential nature.

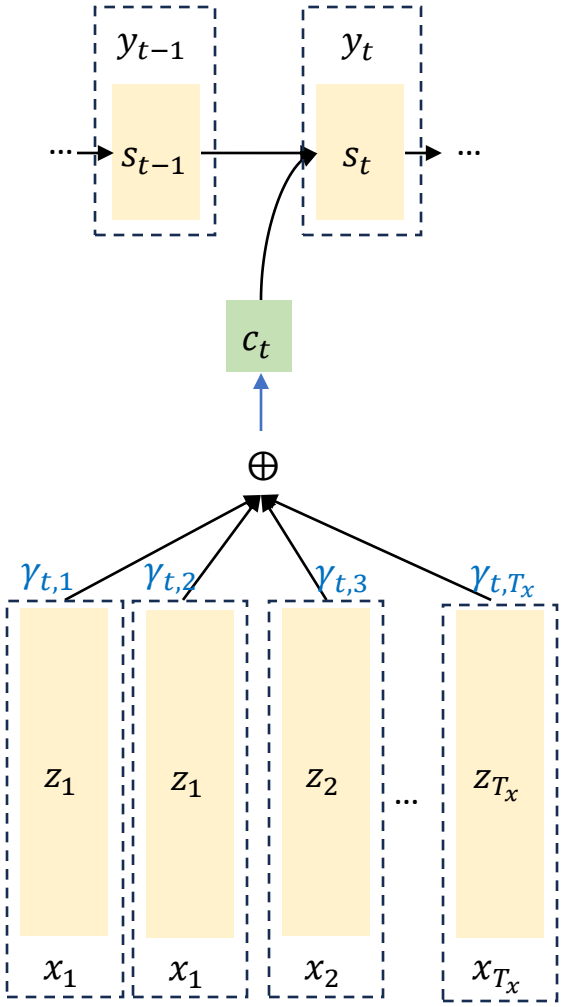
- **Transformers: A non-recurrent solution that solely relies on “attention”:**
  - **No reliance on recurrence:** Transformers capture dependencies across all input *tokens (words)* simultaneously, processing the entire sequence at once. This allows for parallel computation, unlike RNNs that rely on sequential processing.
  - **Captures global dependencies:** The attention mechanism enables modeling of long-range dependencies without the vanishing gradient problem.

Recall the Translate and Align Model in RNNs

- **Problem:** Given source sentence $\mathbf{x} = (x_1, \ldots, x_T)$, build a model that generates target sentence $\mathbf{y} = (y_1, \ldots, y_T)$.
- **Model:** RNN encoder & decoder map source and target sentences to hidden states $\mathbf{z}$ & $\mathbf{s}$, and **context vector** c_t is computed as a weighted sum of the hidden states z_i:

$$e_{t,i} = a(z_i, s_{t-1}) \qquad \gamma_{t,j} = \frac{\exp(e_{t,j})}{\sum_{k=1}^{T_x} \exp(e_{t,k})} \qquad c_t = \sum_{i=1}^{T_x} \gamma_{t,i} z_i$$

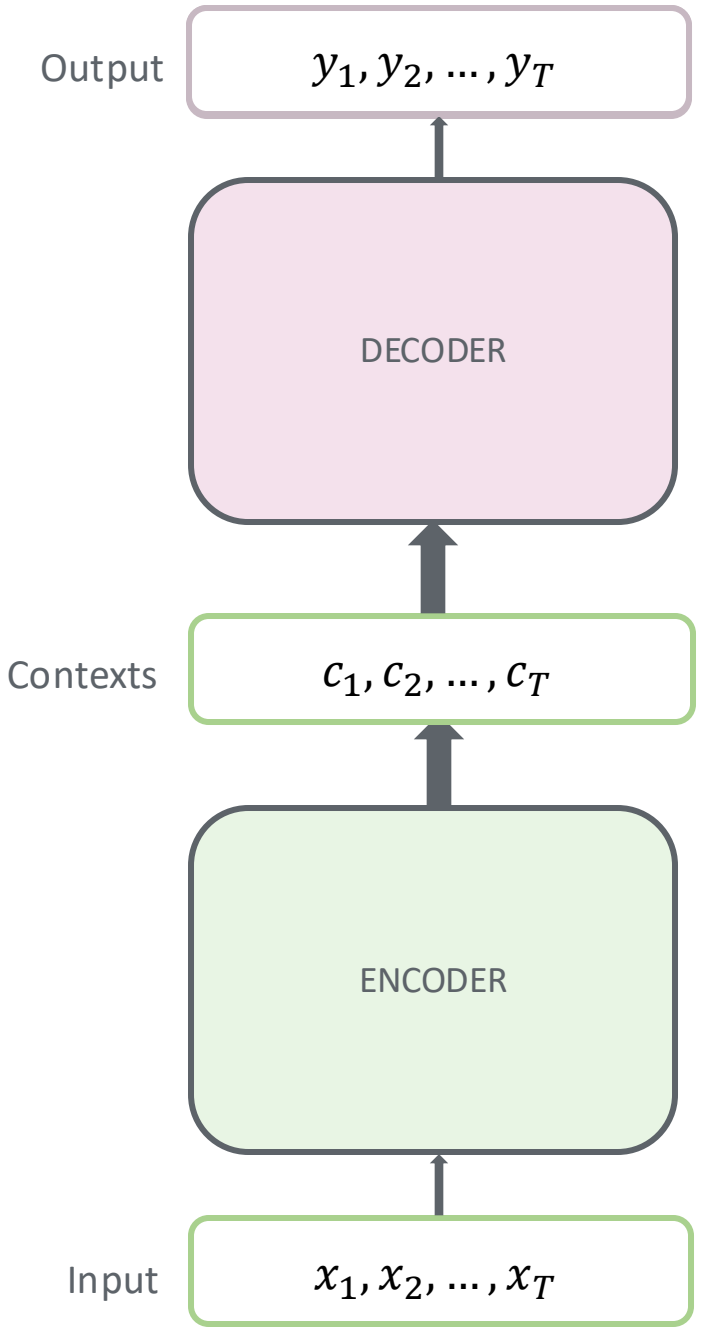
Alignment model | Weights of hidden states | Context vector



- **Attention mechanism**:
  - a is the **alignment model**, typically a feedforward network, and measures how well source word x_i and target word y_t match
  - $\gamma_{t,i}$ is the probability that the target word y_t is aligned to, or translated from, a source word x_i.
  - c_t is the expectation of the hidden state w.r.t. distribution $\gamma_{t,i}$.

From RNNs to Transformers

- Let's keep what is good from Align & Translate:
  - Use encoder to learn latent representation of source sentence
  - Use decoder to learn latent representation of target sentence
  - Align the latent representations of the source/target sentences and form **global contexts**
  - Use decoder to map contexts to target sentences



Transformer

- Let's keep what is good from Align & Translate:
  - Use encoder to learn latent representation of source sentence
  - Use decoder to learn latent representation of target sentence
  - Align the latent representations of the source/target sentences and form **global contexts**
  - Use decoder to map contexts to target sentences

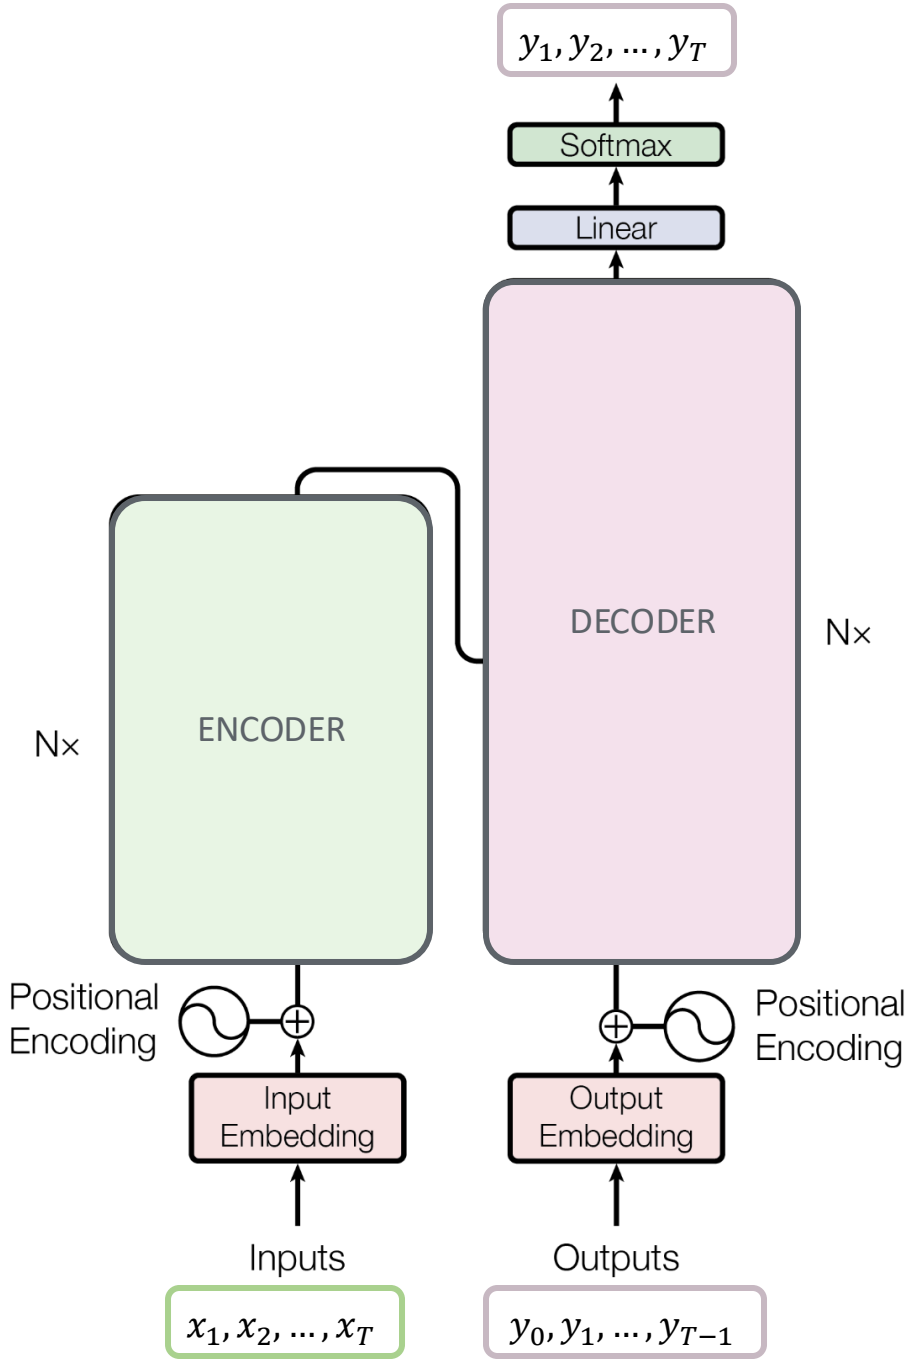


Figure 1: The Transformer - model architecture.

Word to Word Embedding

- First, just like any RNN language tasks, we convert our one-hot vector into embeddings through a word embedding
- Given a sentence, a sequence of one-hot vectors, $\tilde{x} = (\tilde{x}_1, \dots, \tilde{x}_T), \tilde{x}_t \in \{0,1\}^D$
- We obtain the embedding for each word by

$$x_t = \tilde{x}_t^\top \mathbf{E}$$

- Again $\mathbf{E} \in \mathbb{R}^{D \times d}$ is the embedding matrix, and can be pre-trained or learned end-to-end
- In the context of transformers, x_t is also known as a *token*.

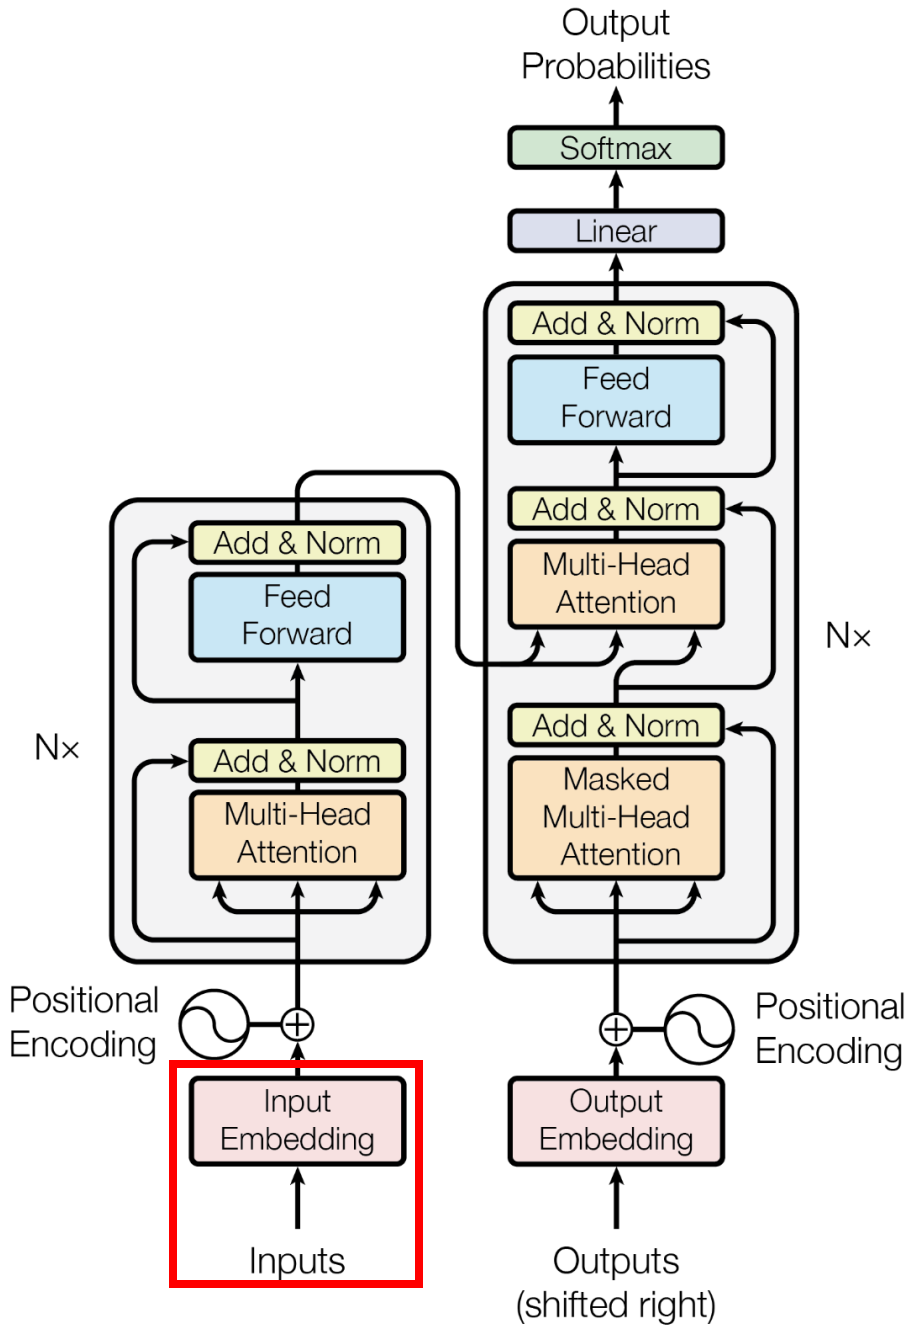


Figure 1: The Transformer - model architecture.

What about the order?

- In RNNs, the recurrence plays a role in telling us the order of the words in a sentence. In transformers, we lose such recurrence.

- Without order, same set of words = same meaning:

  {*I, do, not, like, apples, and, you, like, oranges*}

  {*you, like, apples, and, I, do, not, like, oranges*}

- Need method to encode position of an entity that
  - Outputs a unique encoding for each position
  - Distance between any two positions should be consistent across sentences with different lengths
  - Generalize to longer sentences without any efforts
  - Its values should be bounded

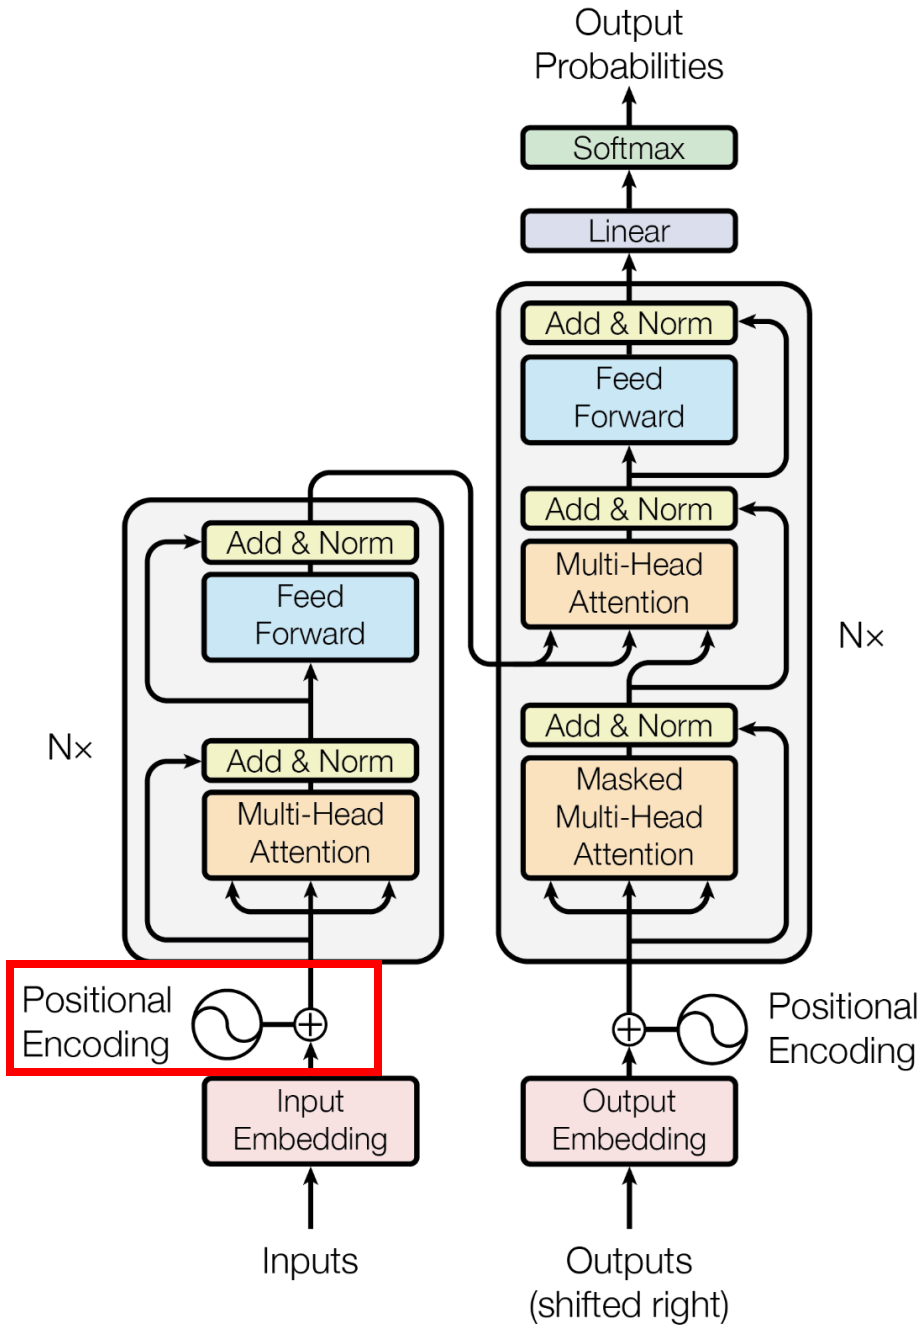
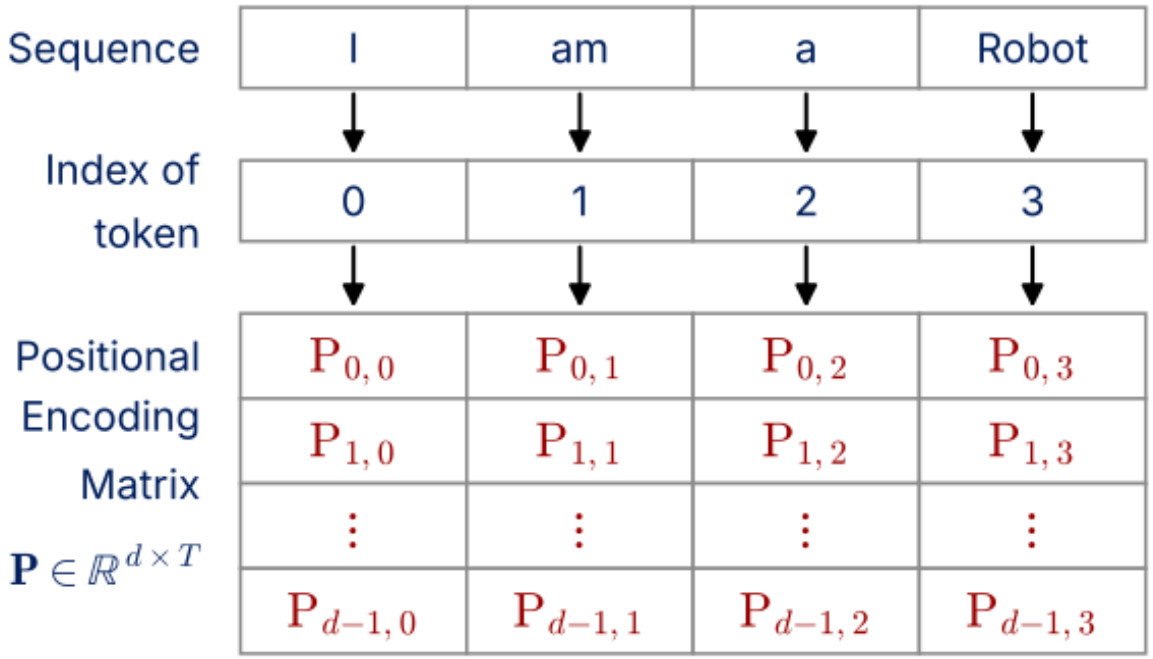


Figure 1: The Transformer - model architecture.

Positional Encoding: Why vectors instead of indexes?

- Positional encoding describes the location or position of an entity in a sequence
- Each position is assigned a unique representation



Position encoding matrix for the sequence 'I am a robot'

- Why not just use the index?
  - For long sequences, the indices can grow large in magnitude.
  - If you normalize the index value to lie between 0 and 1, it can create problems for variable length sequences as they would be normalized differently

Positional Encoding: Intuition

- Suppose you want to represent it in binary format:
  - The lowest bit alternates with every number
  - The second-lowest bit alternates every two numbers, and higher bits continue this pattern.
- But using binary values would be a waste of space and doesn't satisfy the consistent distance between any two positions requirement.
- Instead, we can use their continuous counterparts: sinusoidal functions.
- By decreasing their frequencies, we replicate the behavior of binary bits:
  - Higher frequencies alternate more rapidly, similar to the lower bits in binary (e.g., red bits).
  - Lower frequencies alternate more slowly, similar to the higher bits in binary (e.g., orange bits).

0 :	0	0	0	0	8 :	1	0	0	0
1 :	0	0	0	1	9 :	1	0	0	1
2 :	0	0	1	0	10 :	1	0	1	0
3 :	0	0	1	1	11 :	1	0	1	1
4 :	0	1	0	0	12 :	1	1	0	0
5 :	0	1	0	1	13 :	1	1	0	1
6 :	0	1	1	0	14 :	1	1	1	0
7 :	0	1	1	1	15 :	1	1	1	1

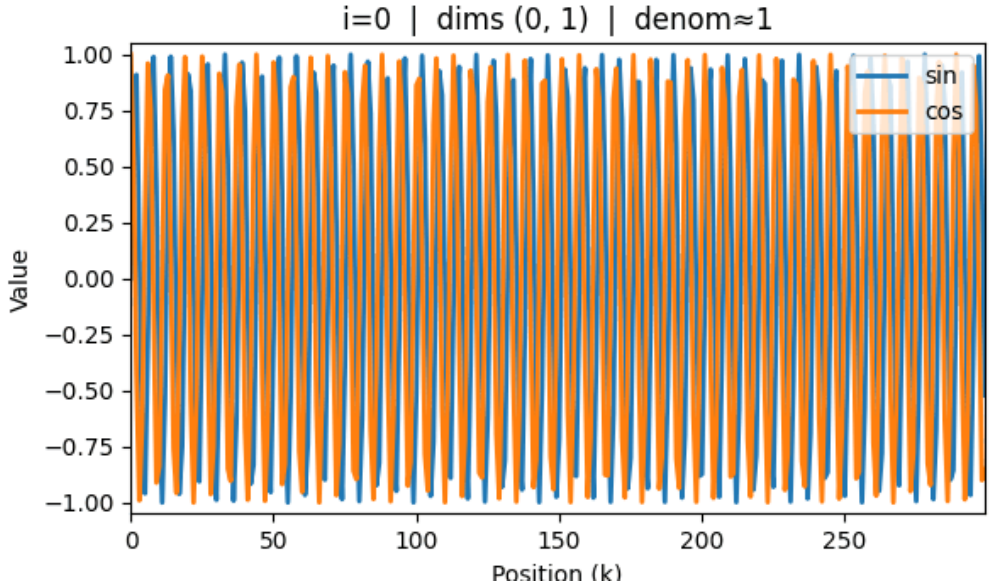
$$\overrightarrow{p_t} = \begin{bmatrix} \sin(\omega_1 t) \\ \cos(\omega_1 t) \\ \\ \sin(\omega_2 t) \\ \cos(\omega_2 t) \\ \\ \vdots \\ \\ \sin\left(\omega_{\frac{d}{2}} t\right) \\ \cos\left(\omega_{\frac{d}{2}} t\right) \end{bmatrix}_{d \times 1}$$

Positional Encoding

- To convey the ordering information, we use **Positional Embeddings** $\mathbf{P} \in \mathbb{R}^{d \times T}$
- In "Attention is All you Need", authors used

$$\mathbf{P}_{2i,k} = \sin\left(\frac{k}{10000^{\frac{2i}{d}}}\right)$$

$$\mathbf{P}_{2i+1,k} = \cos\left(\frac{k}{10000^{\frac{2i}{d}}}\right)$$



$\mathbf{P} \in \mathbb{R}^{d \times T}$		I	am	a	robot
		$t=0$	$t=1$	$t=2$	$t=3$
$i=0$	$j=0$	$\sin(0)$ $=0$	$\sin(1)$ $=0.84$	$\sin(2)$ $=0.91$	$\sin(3)$ $=0.14$
	$j=1$	$\cos(0)$ $=1$	$\cos(1)$ $=0.54$	$\cos(2)$ $=-0.42$	$\cos(3)$ $=-0.99$
$i=1$	$j=2$	$\sin(0)$ $=0$	$\sin(\frac{1}{10})$ $=0.1$	$\sin(\frac{2}{10})$ $=0.2$	$\sin(\frac{3}{10})$ $=0.3$
	$j=3$	$\cos(0)$ $=1$	$\cos(\frac{1}{10})$ $=1$	$\cos(\frac{2}{10})$ $=0.98$	$\cos(\frac{3}{10})$ $=0.96$

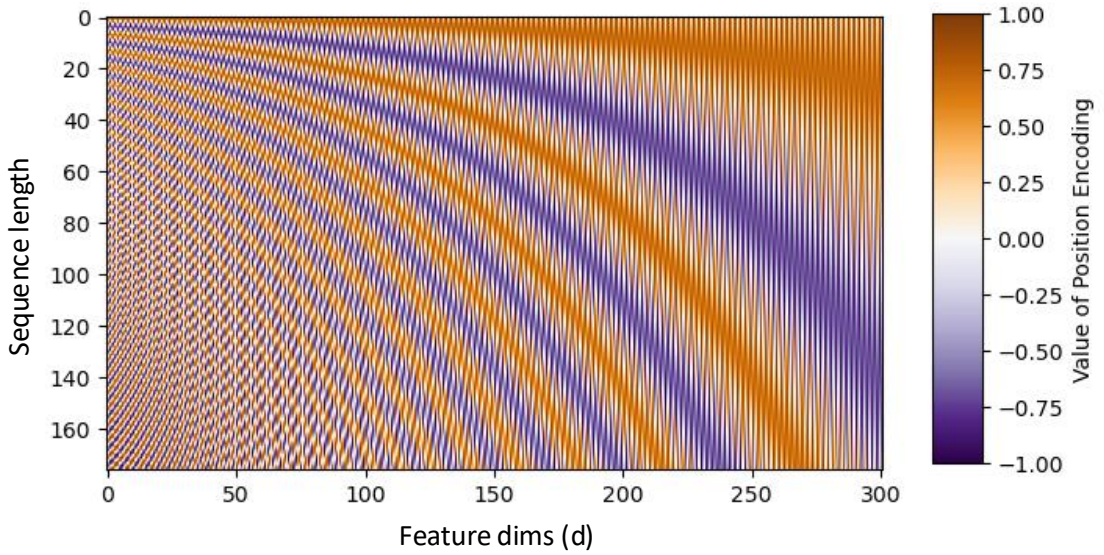
Position encoding matrix for the sequence 'I am a robot'

Positional Encoding

- To convey the ordering information, we use **Positional Embeddings** $\mathbf{P} \in \mathbb{R}^{d \times T}$
- In "Attention is All you Need", authors used

$$\mathbf{P}_{2i,k} = \sin\left(\frac{k}{10000^{\frac{2i}{d}}}\right)$$

$$\mathbf{P}_{2i+1,k} = \cos\left(\frac{k}{10000^{\frac{2i}{d}}}\right)$$



- Let $\mathbf{x} = [x_1, \ldots, x_T] \in \mathbb{R}^{d \times T}$ be the (row) matrix of tokens concatenated together
- The positional embedding gets added to the input directly to the set of tokens:

$$\mathbf{x}^{(0)} = \mathbf{x} + \mathbf{P} \in \mathbb{R}^{d \times T}$$

- We use superscript (0) to denote the input, zero-th layer

Encoder Block

- Just like in the Attend & Align model, we have an encoder that turns input embeddings into hidden embeddings. However, the encoder is not an RNN.
- The main components of an **Encoder Block** are:
  - Multi-Head Attention
  - Layer Normalization
  - Feedforward Neural Networks
  - Skip Connections
- Let's break down the Multi-Head Attention Model!

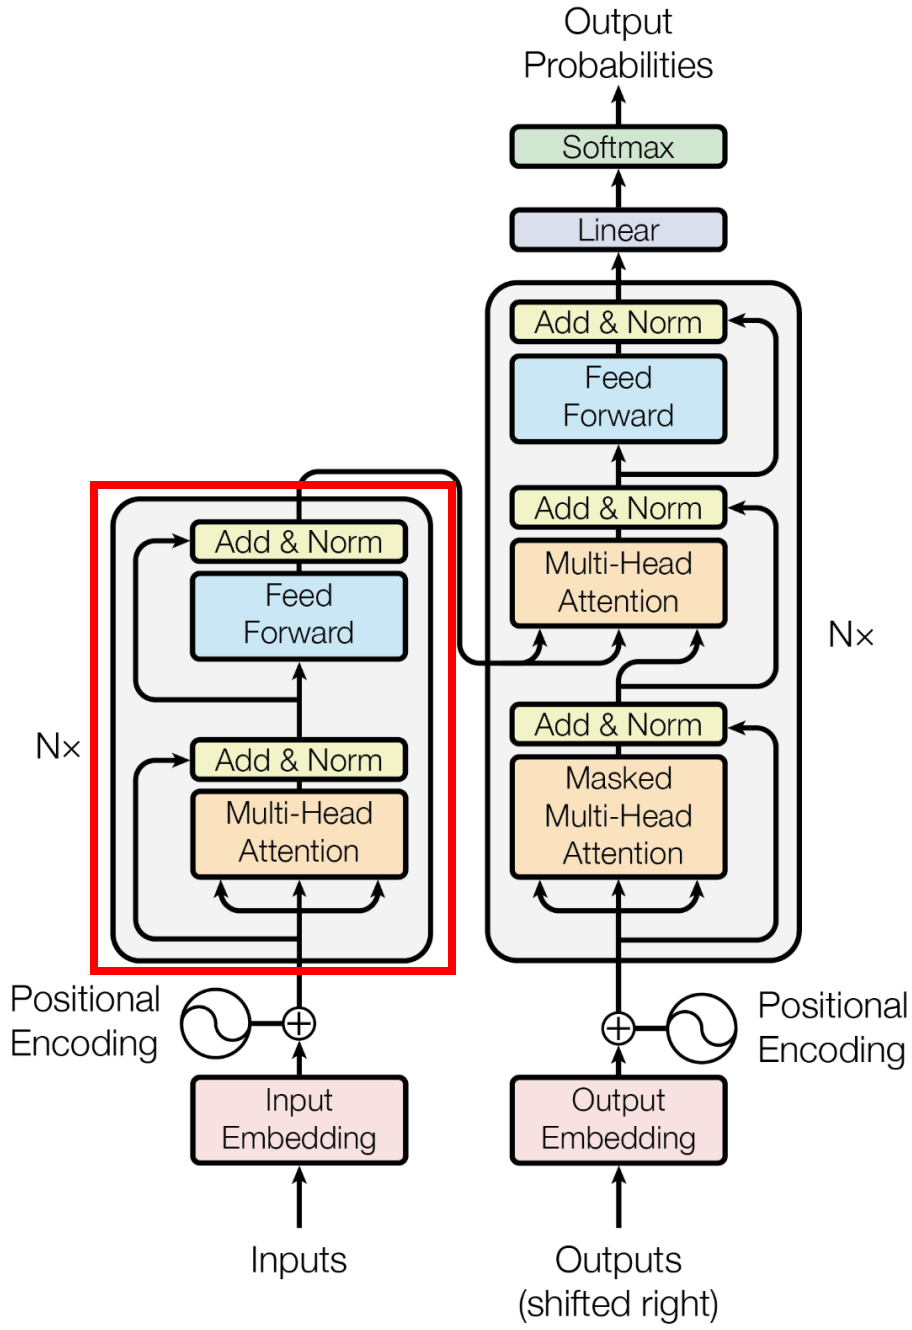
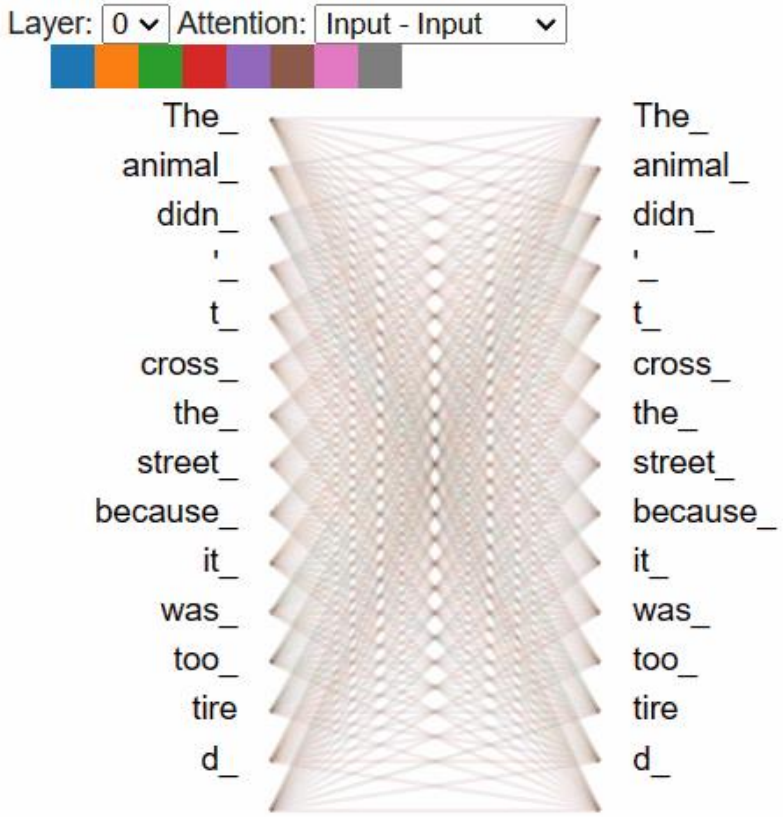


Figure 1: The Transformer - model architecture.

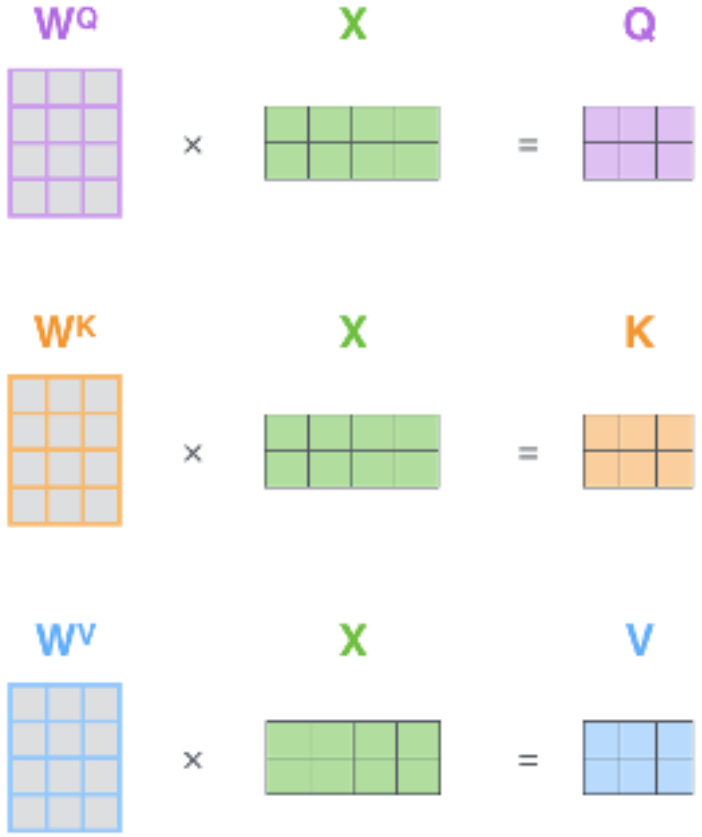
Self-Attention

- Self-attention focuses on important parts of the input by weighing the relevance of each token to others.
  - What does “**it**” in the sentence “The animal didn't cross the street because **it** was too tired.” refer to?
  - Is it referring to the street or to the animal?
- Self-attention allows each token to attend to every other token in the sequence, helping the model capture context and relationships between words.
  - When processing “**it**”, the model uses attention to understand that “**it**” refers to "animal.”
- In RNNs, a hidden state carries context from previous tokens, but attention mechanisms allow direct access to all tokens, without relying on a sequential flow.



Self-Attention

- Given the input embeddings $\mathbf{x} = [x_1, \ldots, x_T]$, we generate three matrices:
  - Query matrix $\mathbf{Q}$
  - Key matrix $\mathbf{K}$
  - Value matrix $\mathbf{V}$
- Input embeddings are transformed into these matrices by multiplying them by three weight matrices $\mathbf{W}^Q, \mathbf{W}^K, \mathbf{W}^V$ that we learn during the training process.
- Analogy for Query, Key, and Value: Library System
  - Imagine you're looking for information on a topic (**query**)
  - Each book has a summary (**key**) to help you identify if it contains relevant information.
  - Once you find a match, you access the book to get the detailed information (**value**) you need.
  - In Attention, we do a "soft match" across multiple books, combine information from each book in proportion to their relevance (e.g., book 1 is most relevant, then book 2, etc.)



Analogy for Query, Key, and Value

Penn Libraries
University of Pennsylvania | Franklin

how to play the violin **Query (Q)**

Bookmarks (0)

Everything | Catalog | Articles+ | Databases | Colenda | Website

Articles+

View and filter 9,343 results

Why I . . . play the violin
Jones, H
Published in BMJ (Online). Volume 376, pp. o322-o322.
Journal Article
2022
South London GP Nicola Weaver tells Helen Jones why she plays in an orchestra

La mariposita ¿La recuerda?: the affective and moral dimensions of professional vision in learning to play the violin
Johnson, S
Published in Mind, culture and activity. Volume 30, Issue 3-4, pp. 209-232.
Journal Article
2023
This is a microanalytic study of a metaphorical story, which is co-constructed by a violin teacher and her emergent bilingual student as part of building the child's professional vision within the art form. The findings center on the multi-functionality of the metaphorical story…

Violin Tutorial I "Violin Master Pro" Teaches People How To Play Violin Like A Master -- Vkoolelite
Published in PRWeb Newswire.
Newspaper Article
2013

How similar is the query to the keys?

Key (K_1) → α_1

Key (K_2) → α_2

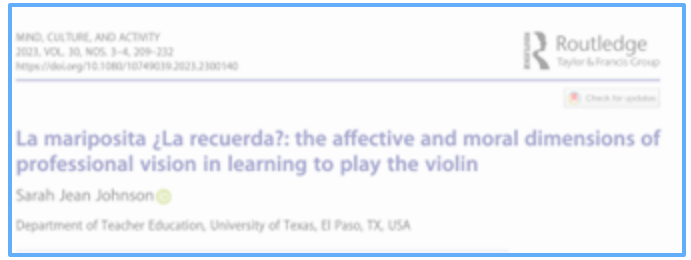
Key (K_3) → α_3

a "soft match" across multiple articles, combining relevant information in proportion to how relevant it is

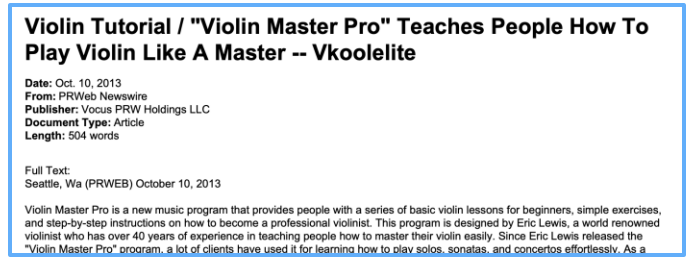
Value (V_1)


CAREERS
Why I . . . play the violin
South London GP Nicola Weaver tells Helen Jones why she plays in an orchestra
Helen Jones

Value (V_2)


Routledge
La mariposita ¿La recuerda?: the affective and moral dimensions of professional vision in learning to play the violin
Sarah Jean Johnson

Value (V_3)


Violin Tutorial / "Violin Master Pro" Teaches People How To Play Violin Like A Master -- Vkoolelite

Date: Oct. 10, 2013
From: PRWeb Newswire
Publisher: Vocus PRW Holdings LLC
Document Type: Article
Length: 504 words

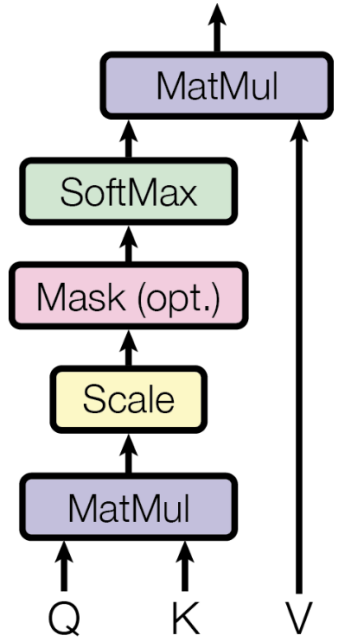
Full Text:
Seattle, Wa (PRWEB) October 10, 2013

Violin Master Pro is a new music program that provides people with a series of basic violin lessons for beginners, simple exercises, and step-by-step instructions on how to become a professional violinist. This program is designed by Eric Lewis, a world renowned violinist who has over 40 years of experience in teaching people how to master their violin easily. Since Eric Lewis released the "Violin Master Pro" program, a lot of clients have used it for learning how to play solos, sonatas, and concertos effortlessly. As a

Self-Attention

- Calculate the attention score by taking the dot product of $\mathbf{Q}$ and $\mathbf{K}^\top$.
- Divide the scores by $\sqrt{d_k}$, where d_k is the dimension of the hidden embedding, to ensure the variance of the dot product does not grow with d_k, leading to unstable attention mechanism.
- Apply the softmax function to the scaled scores, turning them into probabilities.
- Multiply softmax scores by $\mathbf{V}$ to obtain the final attention output.
- The self-attention, thus, is defined as:

$$\mathrm{SA}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \mathrm{softmax}\left(\frac{\mathbf{Q}^\top \times \mathbf{K}}{\sqrt{d_k}}\right)\mathbf{V}$$



- The term “self” comes from the fact that $\mathbf{Q}, \mathbf{K}, \mathbf{V}$ are all derived from the same input sequence $\mathbf{x} = [x_1, \ldots, x_T]$

Multi-Head Self-Attention (MSA)

- **Multi-head Self Attention (MSA)** extends Self-Attention by introducing multiple independent attention heads, each focusing on different types of relationships.

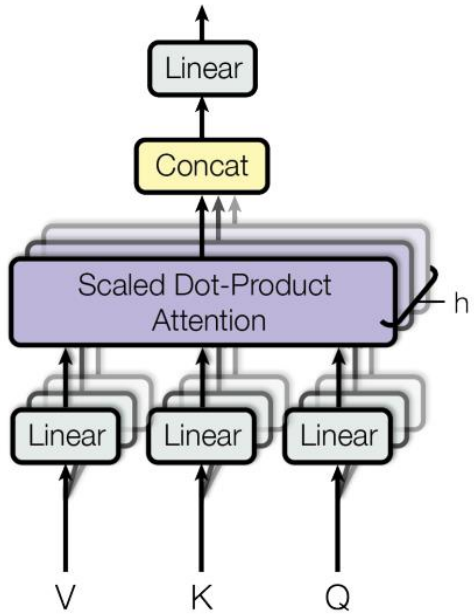


- Each head has its own set of weight matrices:

$$\mathrm{MSA}(\mathbf{Q},\mathbf{K},\mathbf{V}) = \mathbf{W}_O \begin{bmatrix} \mathrm{SA}(\mathbf{Q}_1,\mathbf{K}_1,\mathbf{V}_1) \\ \vdots \\ \mathrm{SA}(\mathbf{Q}_h,\mathbf{K}_h,\mathbf{V}_h) \end{bmatrix}$$

$$\mathbf{Q}_i = \mathbf{W}_i^Q \mathbf{x} \qquad \mathbf{K}_i = \mathbf{W}_i^K \mathbf{x} \qquad \mathbf{V}_i = \mathbf{W}_i^V \mathbf{x}$$

where $\mathbf{W}_i^Q, \mathbf{W}_i^K, \mathbf{W}_i^V \in \mathbb{R}^{d_h \times d}$ are weight matrices for the query, key, and value of each head $i = 1, \ldots, h$, and $\mathbf{W}_O \in \mathbb{R}^{d \times (h \cdot d_h)}$ is the weighting matrix for fusing all attention heads.



Residual Connection & Layer Normalization

- **Residual Connection:** combines input with output of a sub-layer (either self-attention or feed forward)

$$Output = LN\big(\mathbf{x} + SubLayer(\mathbf{x})\big)$$

  - It allows the gradients to flow through the network directly, bypassing non-linear transformations.

- **LayerNorm:** normalizes x_i across feature dimension

$$LN(x_i) = \gamma \cdot \frac{x_i - \mu_i}{\sigma_i} + \beta$$

  - Parameters $\gamma, \beta \in \mathbb{R}^d$ are learned for each layer, and

$$\mu_i = \frac{1}{d}\sum_{k=1}^{d} x_{i,k}\,, \qquad \sigma_i^2 = \frac{1}{d}\sum_{k=1}^{d}\big(x_{i,k} - \mu\big)^2$$

  - This ensures consistent scaling across layers, leading to more stable training.

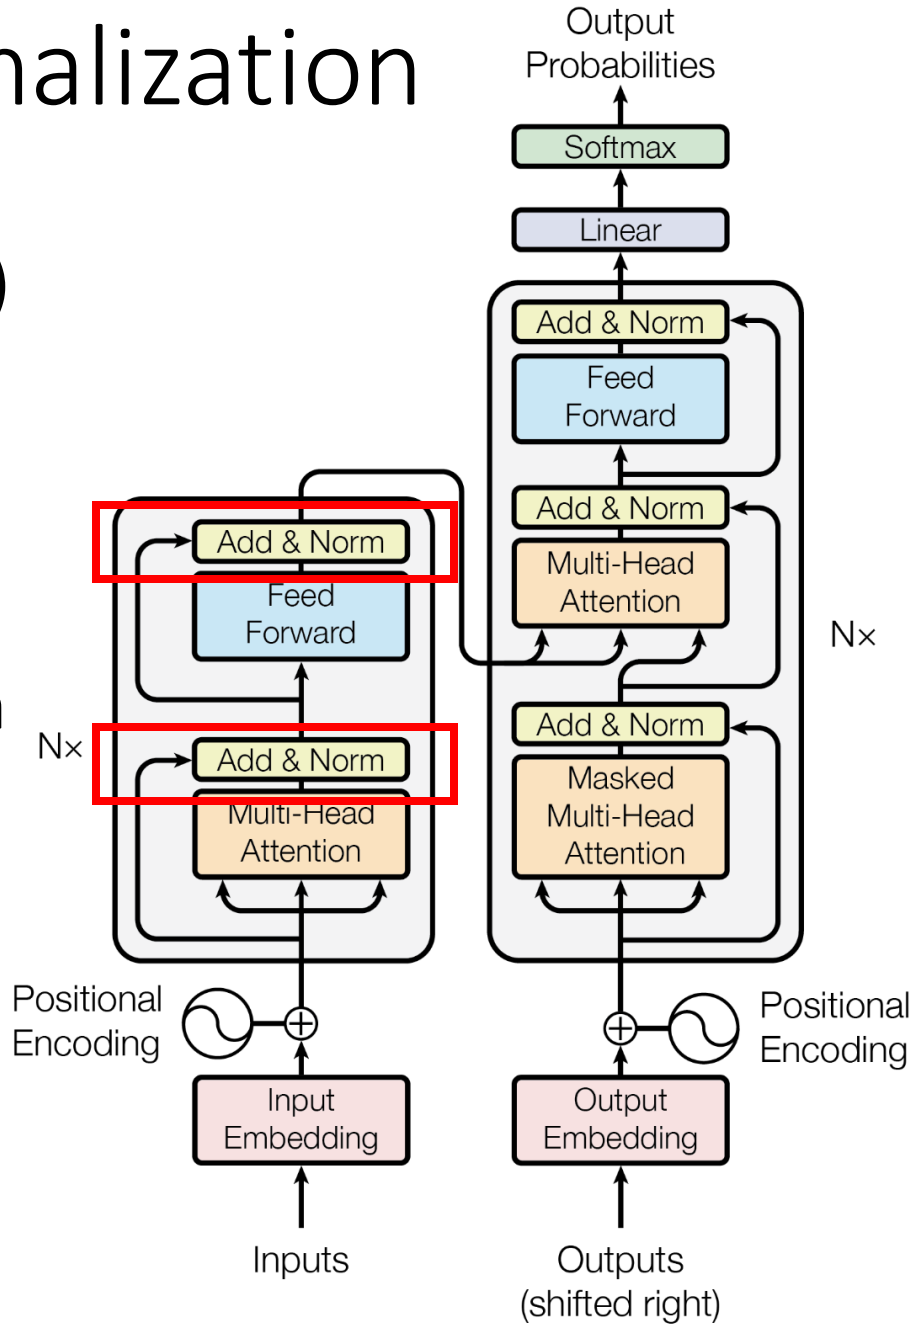


Figure 1: The Transformer - model architecture.

Encoder Block Summarized

- Putting everything together mathematically, the encoder block can be described by

$$\hat{\mathbf{x}}^{(l)} = \mathrm{LN}\left(\mathrm{MSA}\left(\mathbf{x}^{(l-1)}, \mathbf{x}^{(l-1)}, \mathbf{x}^{(l-1)}\right) + \mathbf{x}^{(l-1)}\right)$$
$$\mathbf{x}^{(l)} = \mathrm{LN}\left(\mathrm{FFN}\left(\hat{\mathbf{x}}^{(l)}\right) + \hat{\mathbf{x}}^{(l)}\right)$$

  where FFN is a feed forward neural network and LN denotes Layer Norm

- Note that the input and output dimension of the encoder block is the same: $\mathbb{R}^{d \times T}$
- We can stack encoder blocks together to make it *deeper*
- The output is like the input: a collection of tokens, but **in context with other tokens**

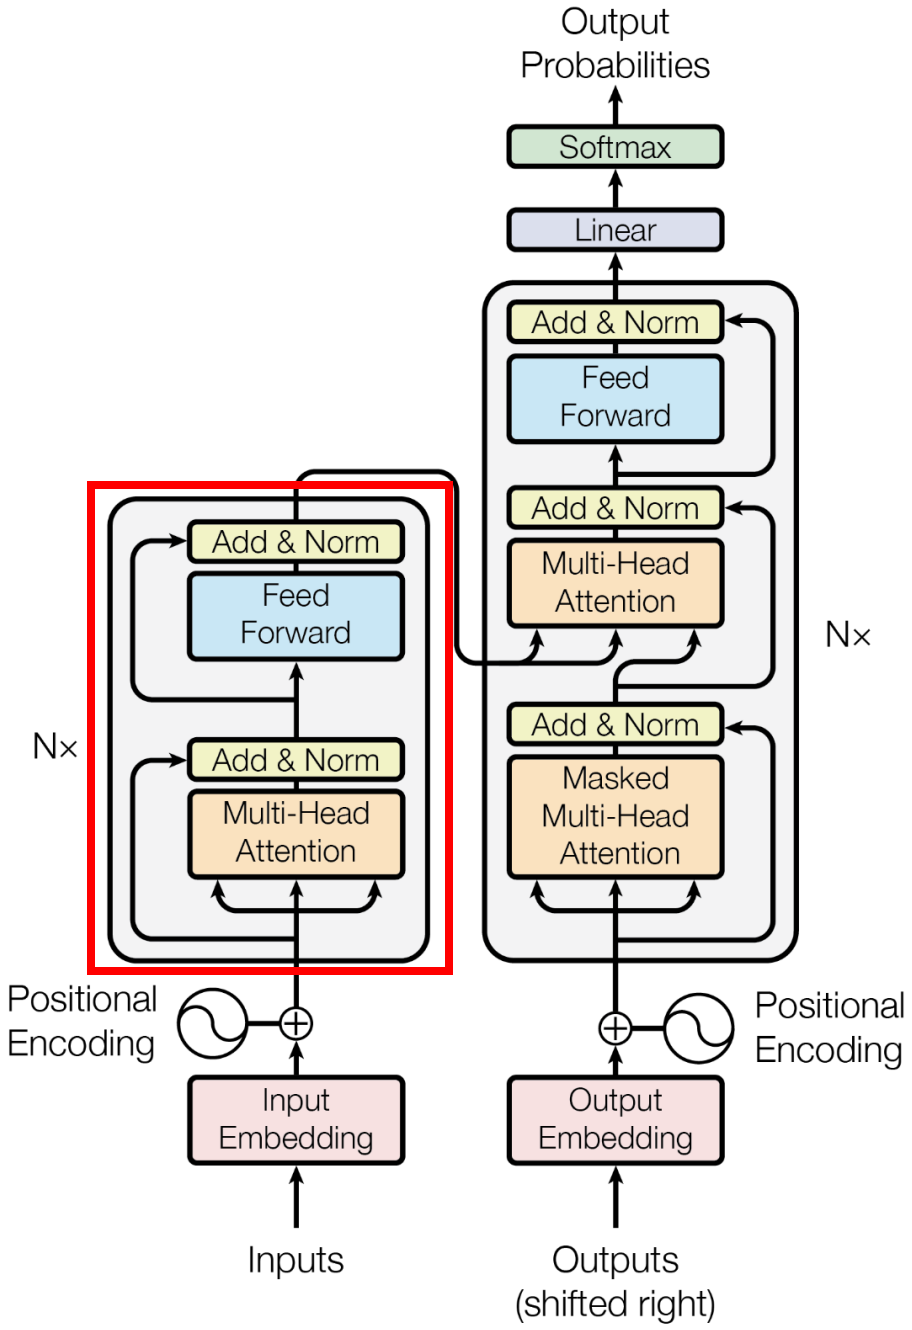
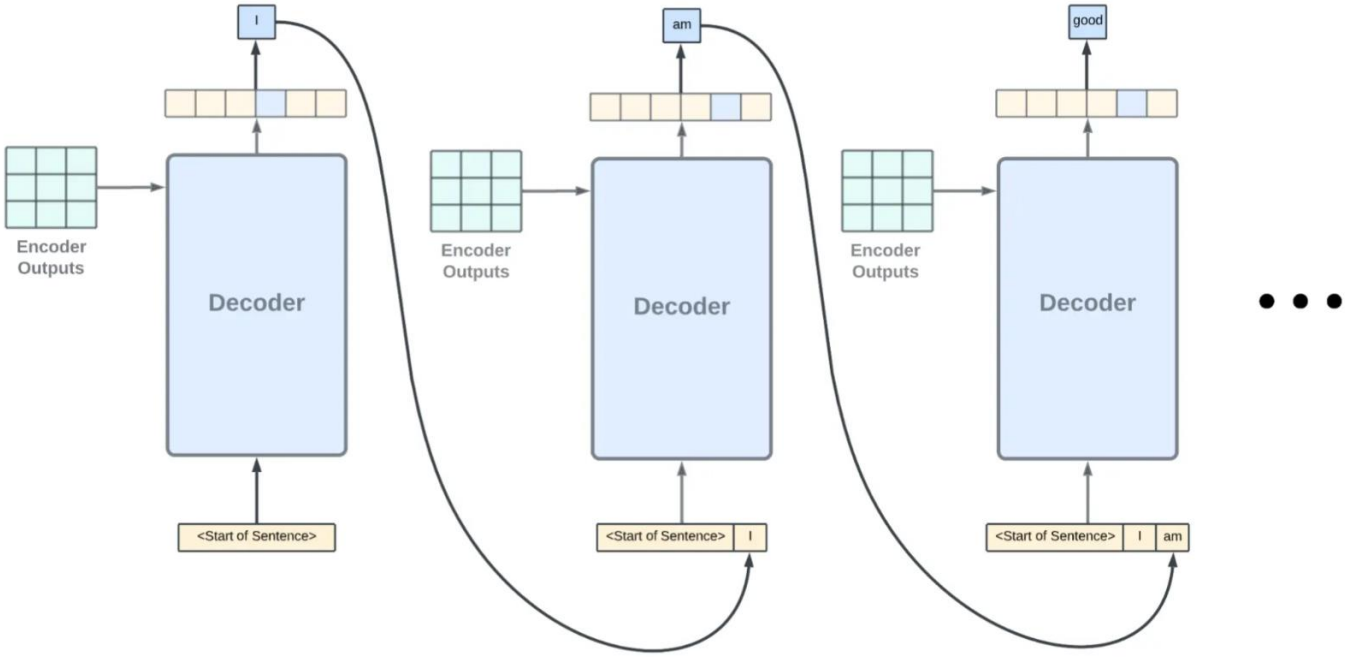


Figure 1: The Transformer - model architecture.

Decoder Block

- Now, we are going to switch gears into the decoder blocks
- At a high level,
  - During inference, the decoder will take in a <BOS> (beginning of sentence) token as input, and recursively predict the next word until the <EOS> (end of sentence) token is predicted
  - Just like our previous methods for machine translation, the decoder should take in *context* from the encoder to predict what the next token should be



Decoder Block: Attention Layers

- What is the input?
  - Token embeddings (shifted right): the decoder sees previous target only.
  - Positional encoding: same as the encoder.
- In the Encoder, each block consists of only *one* Multi-Head Self-Attention layer.
- In the Decoder, each block consists *two layers*:
  - The first one is a Masked Multi-Head Self-Attention with tokens from input.
    - *Allows each token to attend to previous ones in the sequence.*

$$\hat{\mathbf{y}}^{(l)} = \mathrm{LN}\left(\mathrm{MaskedMSA}\left(\mathbf{y}^{(l-1)}, \mathbf{y}^{(l-1)}, \mathbf{y}^{(l-1)}\right) + \mathbf{y}^{(l-1)}\right)$$

  - But, what does the "Masked" in Masked Multi-Head Attention mean?

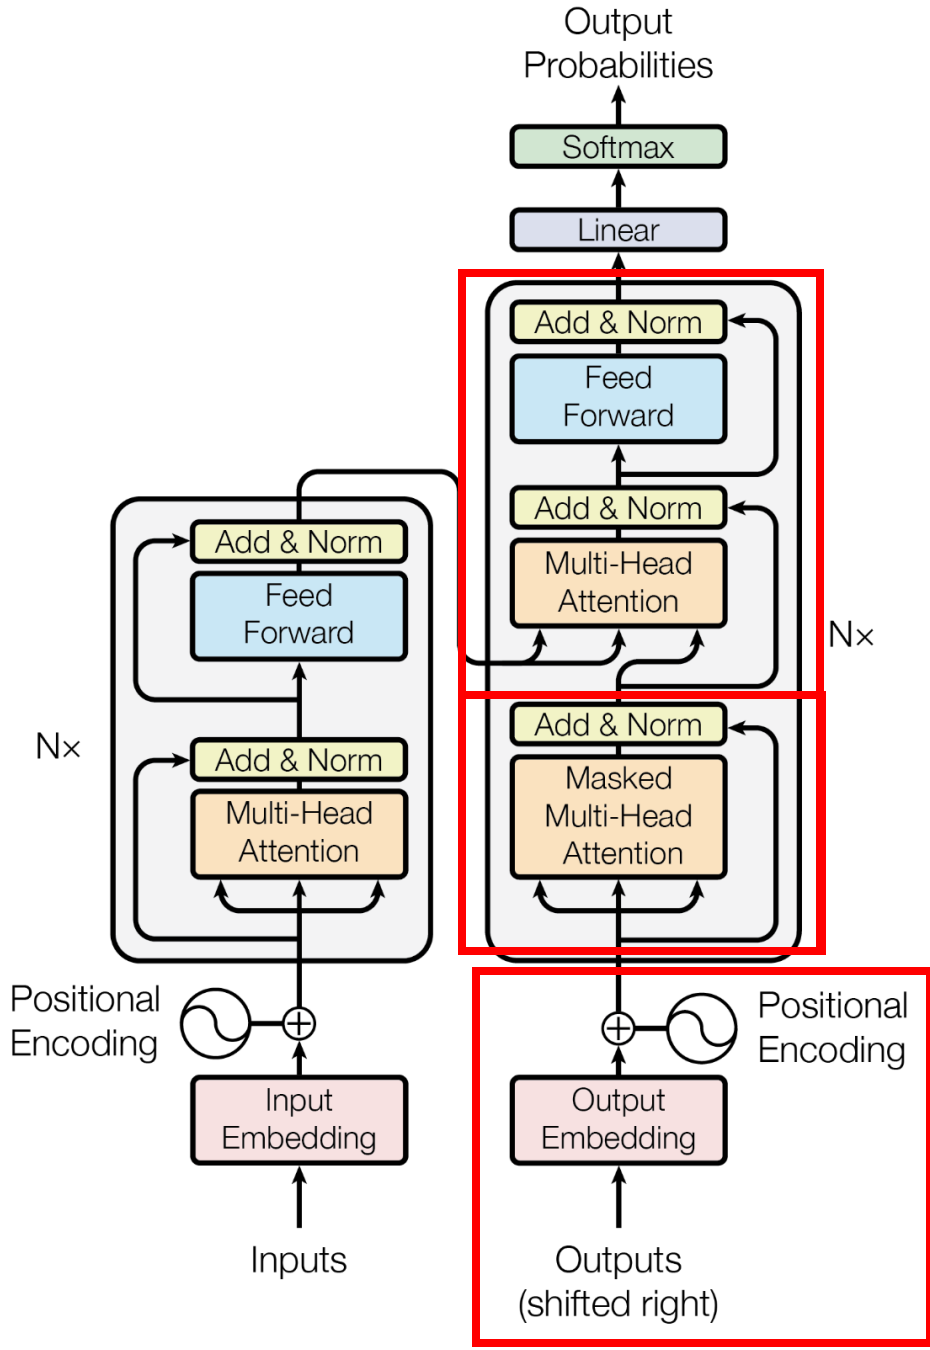


Figure 1: The Transformer - model architecture.

Decoder Block: Masked MSA

- Just liked MSA, **MaskedMSA** calculates attention scores using a scaled dot-product of query and key vectors and normalizes these scores with a *softmax* function to obtain attention weights.

- During training, **MaskedMSA** applies masks on the attention matrices. This is important to preserve the autoregressive property, where each token is predicted based on the preceding tokens only.

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

M is a mask that applies $-\infty$ to all **future positions** so they don't contribute to the *softmax* output.

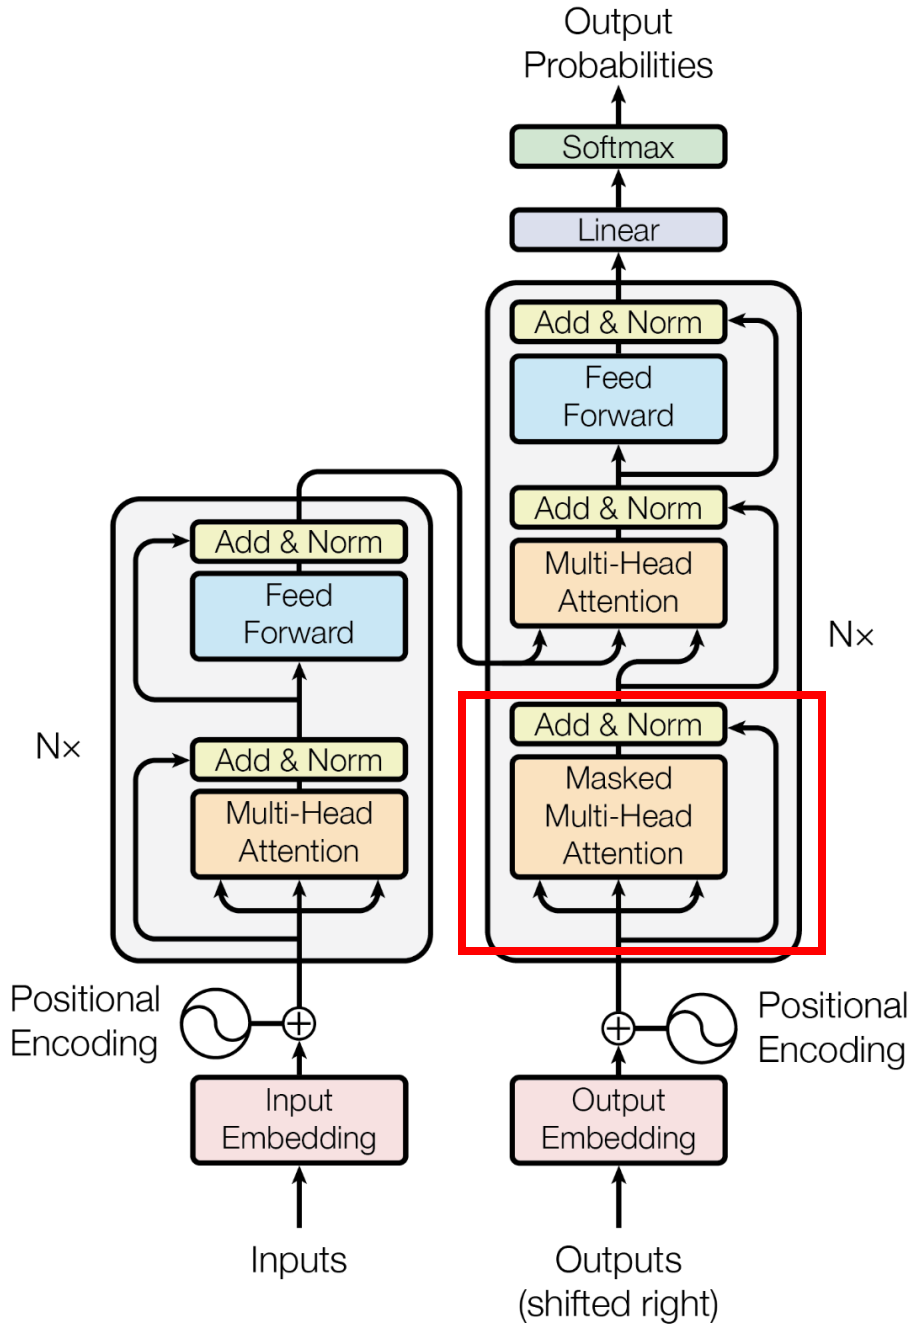


Figure 1: The Transformer - model architecture.

Decoder Block: Attention Layers

- In the Encoder, each block consists of only *one* Multi-Head Self-Attention (MSA) layer.
- In the Decoder, each block consists *two layers*:
  - The Second one is a Multi-Head *Cross* Attention (MCA).
  - MCA applies the same mechanism as MSA, but the *queries*, *keys*, and *values* come from different sources.
  - *Key* and *value* matrices come from the output $\mathbf{x}^{(N)}$ of the encoder, and *query* matrix from the previous MSA.

$$\tilde{\mathbf{y}}^{(l)} = \mathrm{LN}\left(\mathrm{MCA}\left(\hat{\mathbf{y}}^{(l-1)}, \mathbf{x}^{(N)}, \mathbf{x}^{(N)}\right) + \hat{\mathbf{y}}^{(l-1)}\right)$$

  - This allows the decoder to focus on relevant part of encoded input

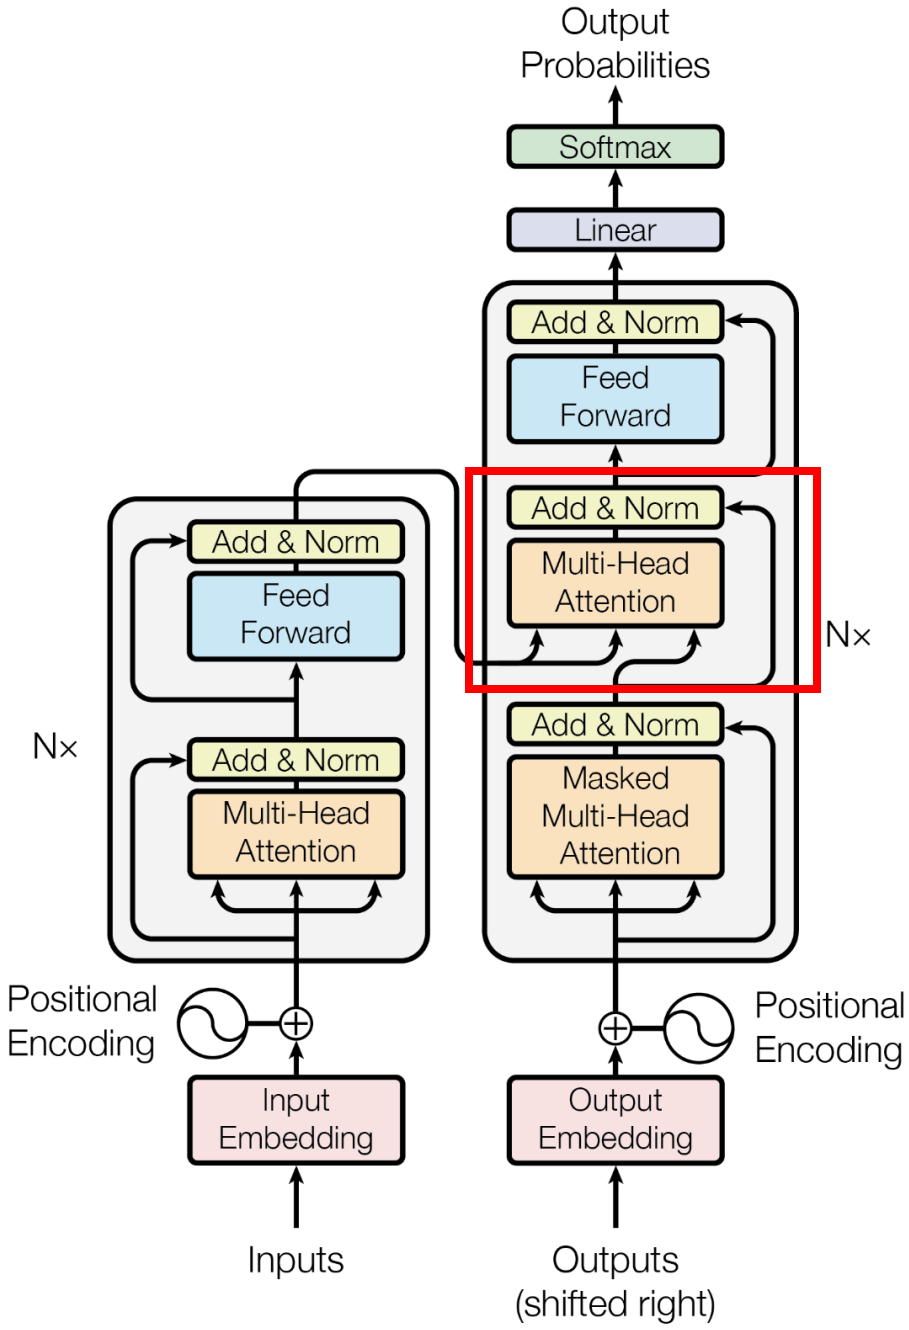


Figure 1: The Transformer - model architecture.

Decoder Block: Summarized

- Summarizing a forward pass of the Decoder Block, along with Layer Norms and Feedforward Networks like the Encoder:

$$\hat{\mathbf{y}}^{(l)} = \mathrm{LN}\left(\mathrm{MaskedMSA}\left(\mathbf{y}^{(l-1)}, \mathbf{y}^{(l-1)}, \mathbf{y}^{(l-1)}\right) + \mathbf{y}^{(l-1)}\right)$$
$$\tilde{\mathbf{y}}^{(l)} = \mathrm{LN}\left(\mathrm{MCA}\left(\hat{\mathbf{y}}^{(l-1)}, \mathbf{x}^{(N)}, \mathbf{x}^{(N)}\right) + \hat{\mathbf{y}}^{(l-1)}\right)$$
$$\mathbf{y}^{(l)} = \mathrm{LN}\left(\mathrm{FFN}\left(\tilde{\mathbf{y}}^{(l)}\right) + \tilde{\mathbf{y}}^{(l)}\right)$$

- After the decoder, the output is projected via a linear layer and softmax to obtain probabilities.
  - The final decoder output $\mathbf{y}^{(N)} \in \mathbb{R}^{d \times T}$ is mapped to vocabulary logits:

$$\mathbf{z} = \mathbf{W}_L \mathbf{y}^{(N)} + \mathbf{b}_L, \qquad \mathbf{W}_L \in \mathbb{R}^{V \times d}$$

  - Finally, the model predicts the probability distribution:

$$\mathbf{p} = \mathrm{softmax}(\mathbf{z})$$

    - Each row $\mathbf{p}_t \in \mathbb{R}^V$ represents model's predicted probability over the vocabulary for the next word: $\mathbf{p}_t = p_\theta(y_t | y_{<t}, \mathbf{x})$.

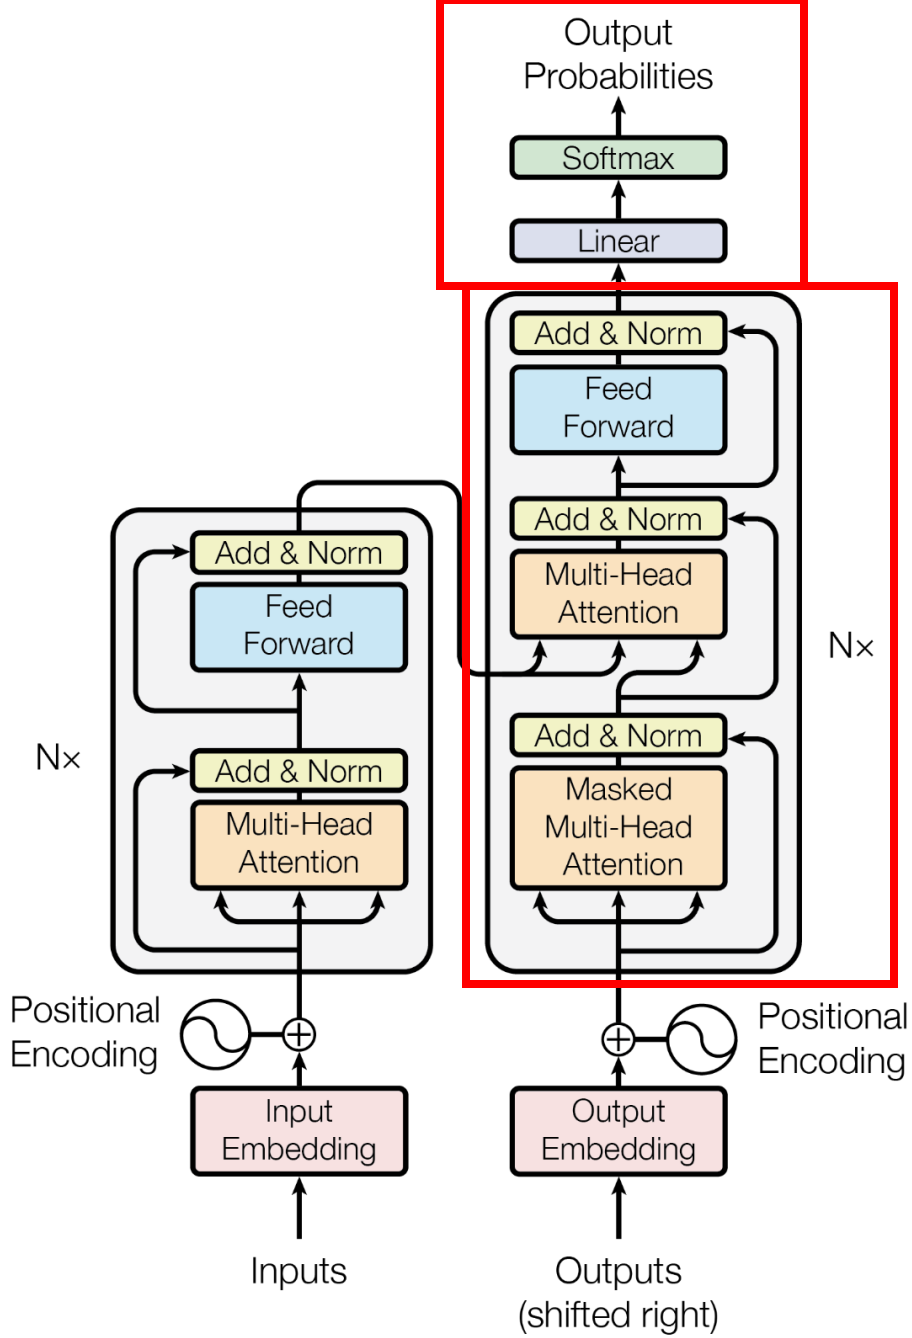


Figure 1: The Transformer - model architecture.

Training the Transformer

- Training the Transformer follows the same intuition with other Seq2Seq models.
- Autoregressive Sequence Modeling:
  - The decoder uses masked self-attention so that each token predicts the next word without looking at future tokens.
  - The model defines the conditional probability of the target sequence $\mathbf{y}$ given the source $\mathbf{x}$ as follows:

$$p_\theta(\mathbf{y}|\mathbf{x}) = \prod_{t=1}^{T} p_\theta(y_t|y_{<t}, \mathbf{x})$$

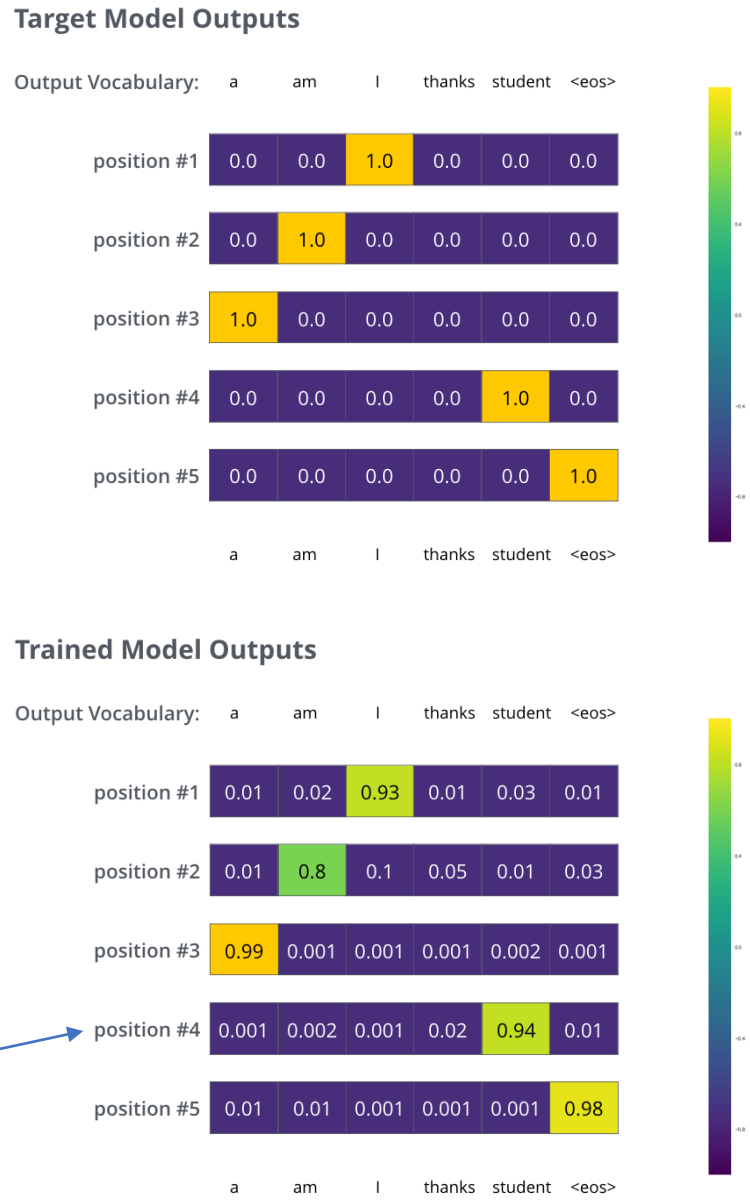
$p_\theta(y_t|y_{<t}, \mathbf{x})$ corresponds to the decoder's *softmax* output for the next token.

  - The model is trained to minimize the cross-entropy between the predicted distribution and the ground truth.

$$\mathcal{L} = -\sum_{(\mathbf{x},\mathbf{y})\in\mathcal{D}} \log p_\theta(\mathbf{y}|\mathbf{x})$$

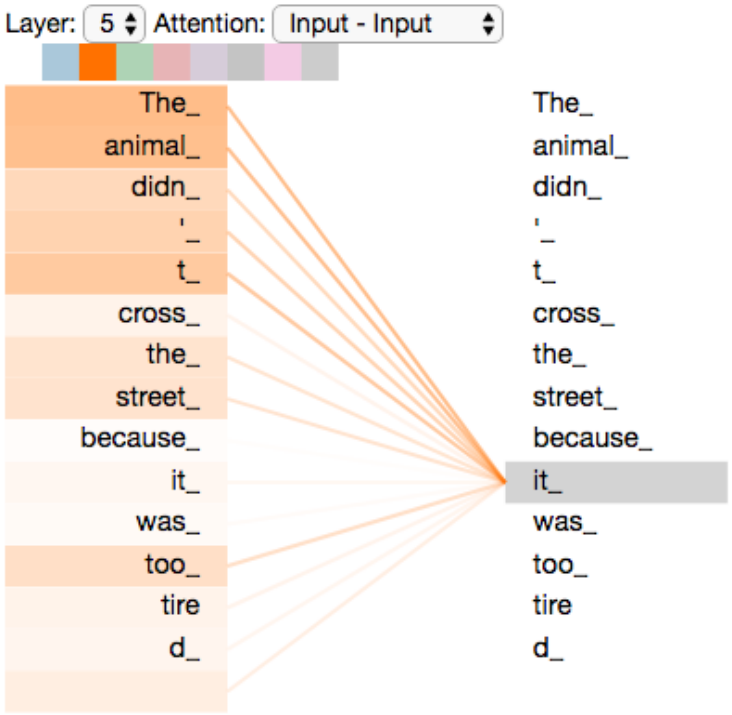
This Is $p_\theta(y_4|y_{<4}, \mathbf{x})$

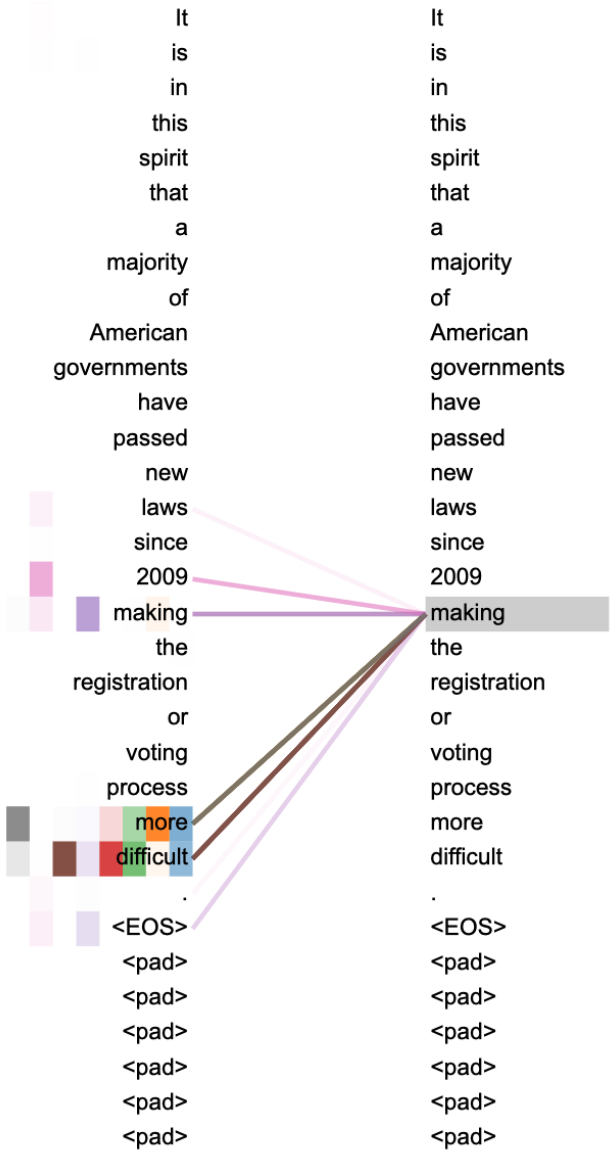
For example, when translating
"Soy un estudiante" to "I am a student"



Attention Visualization: Long distance dependency

- Earlier we saw the sentence: “The animal didn't cross the street because **it** was too tired.”
- What does “**it**” in this sentence refer to? The visualization of self-attention shows the association of “it” with beginning parts like “The animal”.

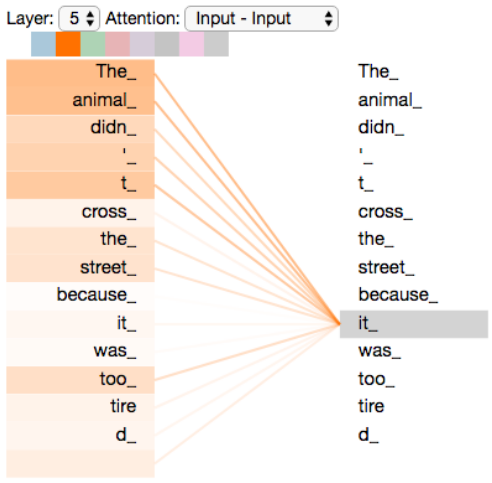


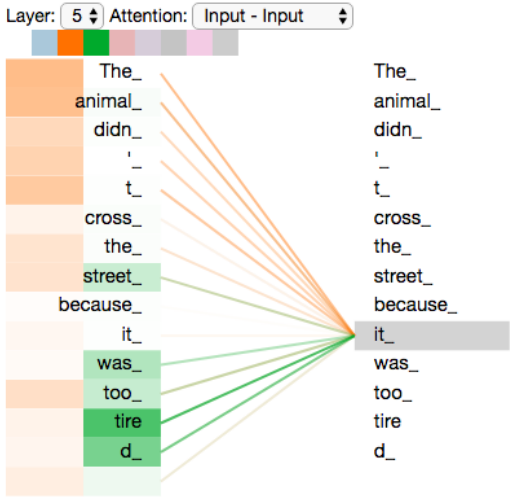


- On the right we see another visualization showing how different words in a longer sentence relate to each other.
- Check out this interactive visualization.

Attention: Attention from Different Heads

- Attention heads can specialize to capture various dependencies, such as syntactic and semantic relationships.
- This allows the model to attend to different types of causalities between words in a sentence.





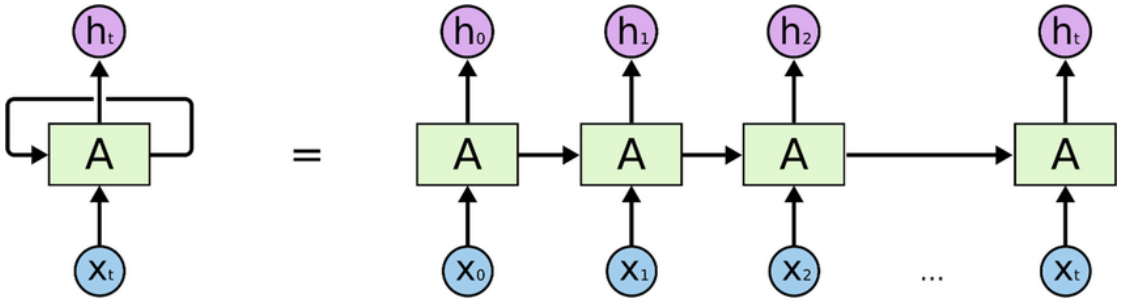




RNNs vs. Transformers

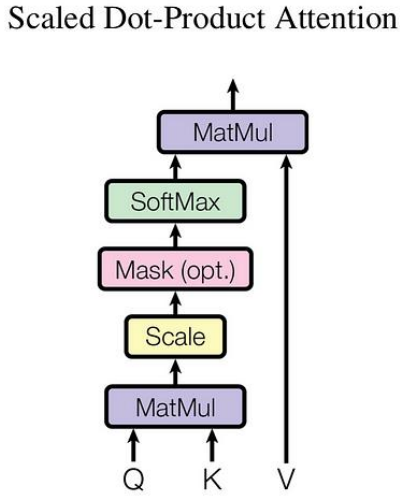
Recurrent Neural Network

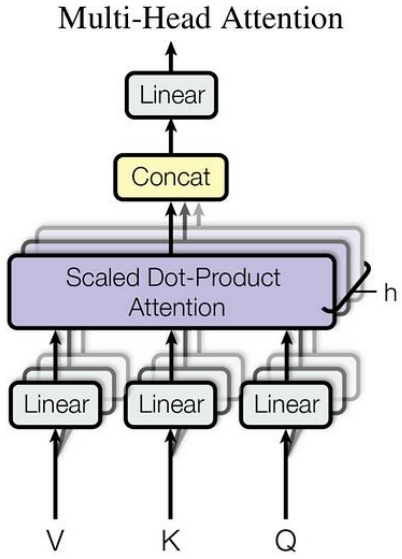
- Handle sequential data
- Learn sequential dependencies
- Each time step depends on the previous one



Transformers

- Handle sequential data
- Learn sequential dependencies
- Use self-attention to capture global context





RNNs vs. Transformers

Recurrent Neural Network	Transformers

• (-) Learning long-range dependences is challenging due to recurrent structure
  • Can be aided by specialized architectures like LSTM and GRU
  • Suffer from training issues such as vanishing gradient
	• (+) Attention mechanism better captures long-range dependences
	  • Able to handle both global context and local context
	  • No vanishing gradient issues

• (-) Hard to scale up because each time step depends on the previous one	• (+) Processes tokens in parallel, makes it efficient for training on GPUs
• (+) Usually smaller number of parameters, does not require lots of data to train	• (-) Usually large number of parameters, requires lots of data to train

Upcoming Lectures: Iterations of Transformers

Natural Language Processing

- **BERT (Bidirectional Encoder Representations from Transformers)**
- GPT (Generative Pre-trained Transformer)
- RoBERTa (Robustly Optimized Bert Pre-training)
- T5 (Text-to-Text Transfer Transformer)

Computer Vision

- **ImageGPT**
- **Vision Transformer**
- Swin Transformer, Pyramid Vision Transformer