

Any reproduction or distribution of this document, in whole or in part, is prohibited unless permission is granted by the authors

EE-559

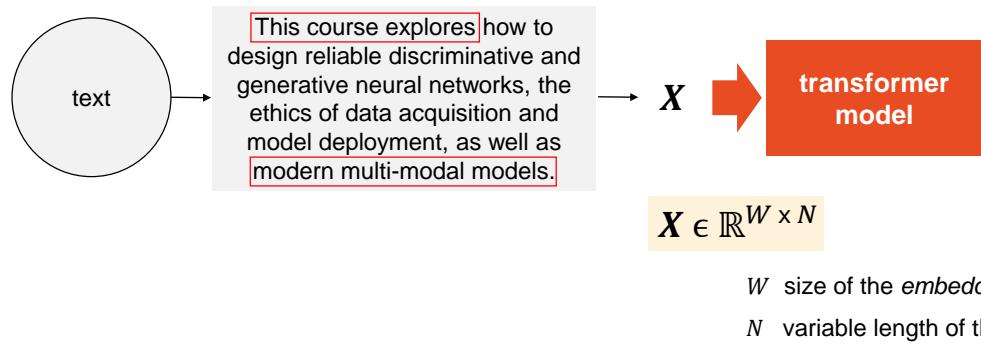
Deep Learning

What's on today?

- **Natural Language Processing**: how to analyse / synthesize language
- **Tokens**: on breaking text into small units
- **Self-attention**: how a model can focus on relevant parts of the input
- **Transformer**: how to use attention in a neural network
- **Encoder model**: how to process text to create a useful representation
- **Decoder model**: how to generate text
- **Encoder-decoder model**: how to map a sequence to a sequence
- **Exercises**: exploring attention and transformers

Natural Language Processing

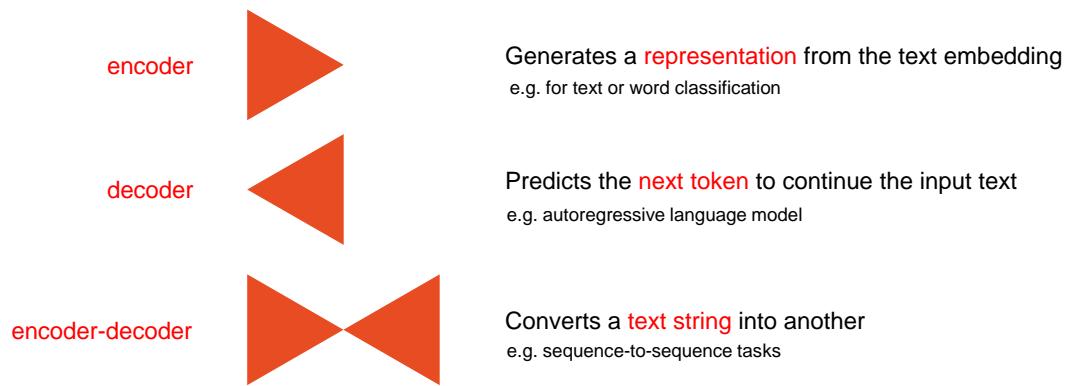
Natural Language Processing (NLP)



Concepts:

Analysis and synthesis of speech and language, syntax, pronouns, connections between words

Transformer models

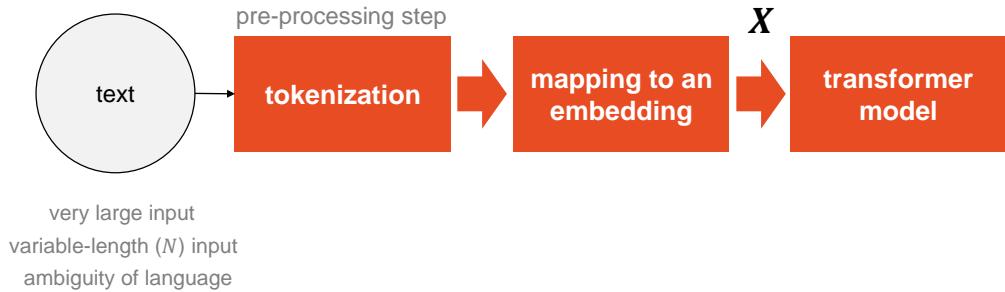


Concepts:

Sentiment analysis, name entity recognition, tokens, generative pre-trained transformers, language translation

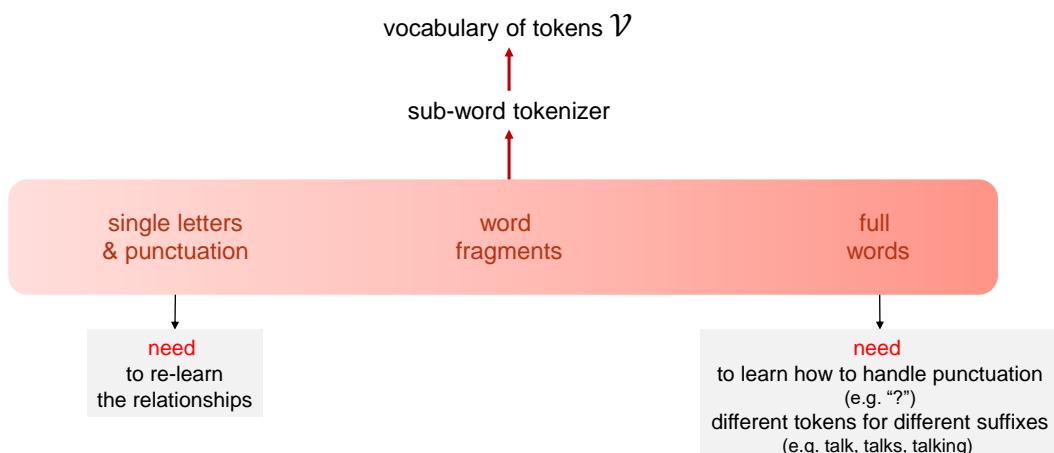
Tokens

Tokenization of text



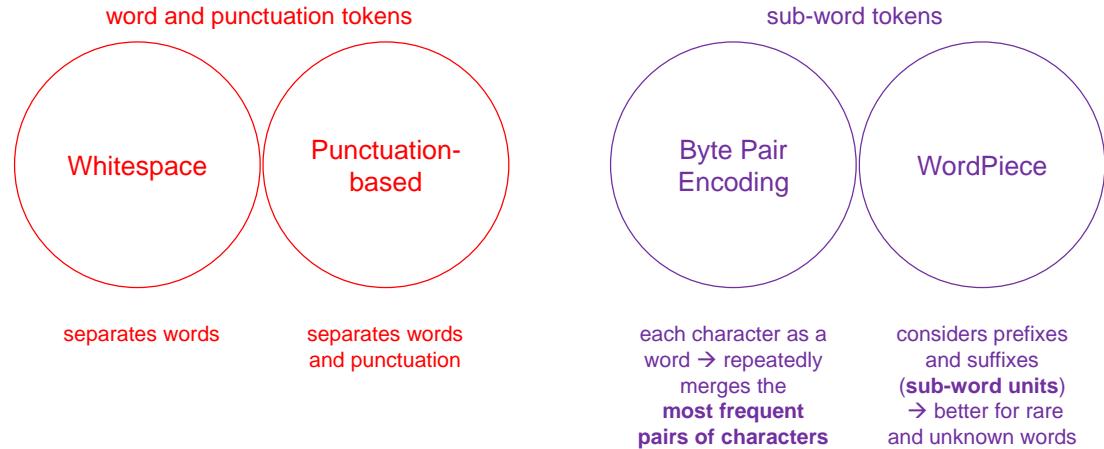
Concepts:
Words and word fragments, sub-word tokenizer, token embedding

Tokenization: choices



Concepts:
Use frequency of commonly occurring sub-strings to merge them (byte pair encoding)

Tokenizers



Concepts: Counting the number of words in a text, identifying frequently occurring character combinations in training text, a tokenizer does not work equally well for all languages

More on units of text

Between words and characters: A Brief History of Open-Vocabulary Modeling and Tokenization in NLP

Sabrina J. Mielke ^{1,2} Zaid Alyafeai ³ Elizabeth Salesky ¹
 Colin Raffel ² Manan Dey ⁴ Matthias Gallé ⁵ Arun Raja ⁶
 Chenglei Si ⁷ Wilson Y. Lee ⁸ Benoit Sagot ^{9*} Samson Tan ^{10*}
BigScience Workshop Tokenization Working Group

¹Johns Hopkins University ²HuggingFace ³King Fahd University of Petroleum and Minerals ⁴SAP

⁵Naver Labs Europe ⁶Institute for Infocomm Research, A*STAR Singapore ⁷University of Maryland

⁸BigScience Workshop ⁹Inria Paris ¹⁰Salesforce Research Asia & National University of Singapore

sjm@sjmielke.com

Abstract

What are the units of text that we want to model? From bytes to multi-word expressions, text can be analyzed and generated at many granularities. Until recently, most natural language processing (NLP) models operated over words, treating those as discrete and atomic tokens, but starting with byte-pair encoding (BPE), subword-based approaches have become dominant in many areas, enabling small vocabularies while still allowing for fast inference. Is the end of the road character-level model or byte-level processing? In this survey, we connect several lines of work from the pre-neural and neural era, by showing how hybrid approaches

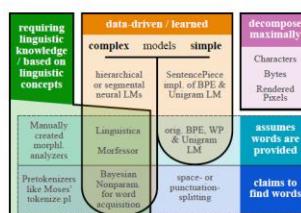
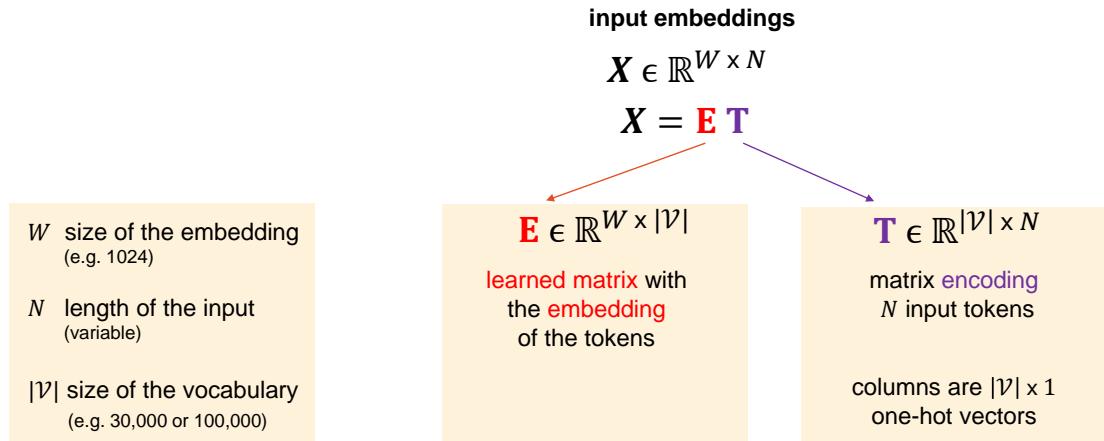


Figure 1: A taxonomy of segmentation and tokenization algorithms and research directions

[arXiv:2112.10508](https://arxiv.org/abs/2112.10508)

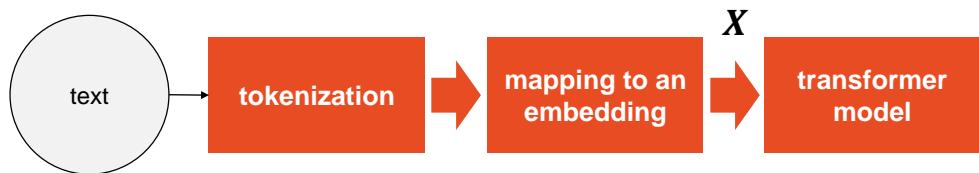
Mapping to an embedding



Concepts:

One-hot vector: the only non-zero element is the entry corresponding to the token, T is sparse

Input embeddings



$$X = [x_1, x_2, \dots, x_m, \dots, x_N]$$

$$x_m \in \mathbb{R}^W$$

Concepts:

Input of length N , embedding of size W , embedding of the token (x_m)

Self-attention

Self-attention

$$\begin{aligned}
 s_n[x_1, x_2, \dots, x_N] &= \sum_{m=1}^N a[x_m, x_n] (\Theta_0^v + \Theta_1^v x_m) \\
 &= \sum_{m=1}^N a[x_m, x_n] v_m
 \end{aligned}$$

self-attention
 (connection between
 word representations)

attention
 (scalar weight)

value
 (linear transformation of each x_m)

$$\begin{aligned}
 a[., x_n] &\geq 0 \\
 \sum_{n=1}^N a[., x_n] &= 1
 \end{aligned}$$

$$\begin{aligned}
 X &\in \mathbb{R}^{W \times N} \\
 x_m &\in \mathbb{R}^W
 \end{aligned}$$

$$\begin{aligned}
 v_m &= \Theta_0^v + \Theta_1^v x_m \\
 &\quad \text{biases} \quad \text{weights}
 \end{aligned}$$

$$\begin{aligned}
 \Theta_0^v &\in \mathbb{R}^W \\
 \Theta_1^v &\in \mathbb{R}^{W \times W}
 \end{aligned}$$

Concepts:

Connections between the inputs, sparse attention coefficients, value (v_m)

Attention

$$\begin{aligned} \mathbf{q}_n &= \Theta_0^q + \Theta_1^q \mathbf{x}_n && \text{query} \\ \mathbf{k}_m &= \Theta_0^k + \Theta_1^k \mathbf{x}_m && \text{key} \end{aligned}$$

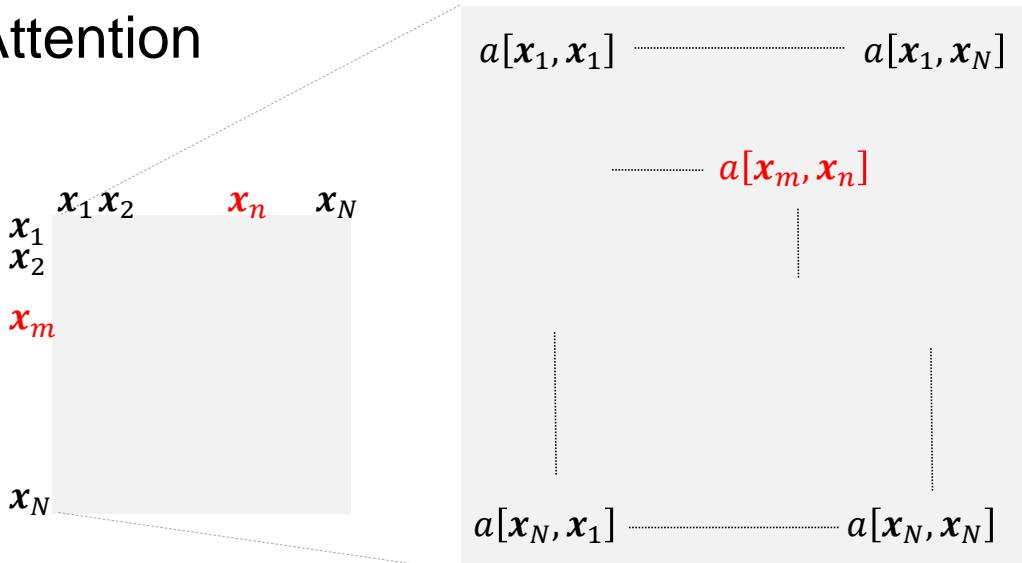
$(\mathbf{k}_m)^T \mathbf{q}_n$ measure of similarity (dot product)

$$\begin{aligned} a[\mathbf{x}_m, \mathbf{x}_n] &= \text{softmax}_m \left((\mathbf{k}_*)^T \mathbf{q}_n \right) \\ &= \frac{e^{(\mathbf{k}_m)^T \mathbf{q}_n}}{\sum_{m'=1}^N e^{(\mathbf{k}_{m'})^T \mathbf{q}_n}} \end{aligned}$$

non-linear due to
the dot product
and the softmax

Concepts: Query and key from information retrieval, dot-product self-attention, measure of similarity between the input \mathbf{x}_n and all the other inputs

Attention



Concept:
Modeling data with *long-range* dependencies

Self-attention

$$\begin{array}{ll}
 x_1 & s_{\textcolor{red}{1}}[x_1, x_2, \dots, x_N] = \sum_{m=1}^N a[x_m, x_{\textcolor{red}{1}}] v_m \\
 \vdots & \vdots \\
 x_n & s_{\textcolor{red}{n}}[x_1, x_2, \dots, x_N] = \sum_{m=1}^N a[x_m, x_{\textcolor{red}{n}}] v_m \\
 \vdots & \downarrow \text{scales linearly with } N \\
 x_N & s_{\textcolor{red}{N}}[x_1, x_2, \dots, x_N] = \sum_{m=1}^N a[x_m, x_{\textcolor{red}{N}}] v_m
 \end{array}$$

quadratic dependence on N

Concepts: Value is a linear function of the input, same weights and biases applied to each input, attention weights are *sparse* and are a *non-linear* function of the input, parallelization

$$\begin{aligned}
 X &\in \mathbb{R}^{W \times N} \\
 x_m &\in \mathbb{R}^W
 \end{aligned}$$

Self-attention: matrix form

$$\Theta = \{\Theta_0^v, \Theta_1^v, \Theta_0^q, \Theta_1^q, \Theta_0^k, \Theta_1^k\} \quad \text{shared set of parameters across inputs}$$

$$V(X) = \Theta_0^v \mathbf{1}_N^T + \Theta_1^v X$$

$$Q(X) = \Theta_0^q \mathbf{1}_N^T + \Theta_1^q X \quad \mathbf{1}_N: N\text{-dimensional vector containing only 1s}$$

$$K(X) = \Theta_0^k \mathbf{1}_N^T + \Theta_1^k X$$

$$\begin{aligned}
 S(X) &= V(X) \text{softmax} \left((K(X))^T Q(X) \right) \\
 &= V \text{softmax} \left((K)^T Q \right)
 \end{aligned}$$

scaled dot-product self-attention

$$S(X) = V \text{softmax} \left(\frac{(K)^T Q}{\sqrt{W_q}} \right)$$

dimension of queries (and keys)

Concepts:

Scaled dot-product self-attention: facilitates training, avoids domination of the largest value

Multi-head self-attention

$$V_h(X) = \Theta_{0h}^v \mathbf{1}_N^T + \Theta_{1h}^v X$$

$$Q_h(X) = \Theta_{0h}^q \mathbf{1}_N^T + \Theta_{1h}^q X$$

$$\Theta_h = \{\Theta_{0h}^v, \Theta_{1h}^v, \Theta_{0h}^q, \Theta_{1h}^q, \Theta_{0h}^k, \Theta_{1h}^k\}$$

$$K_h(X) = \Theta_{0h}^k \mathbf{1}_N^T + \Theta_{1h}^k X$$

$$S_h(X) = V_h \text{ softmax} \left(\frac{(K_h)^T Q_h}{\sqrt{W_q}} \right)$$

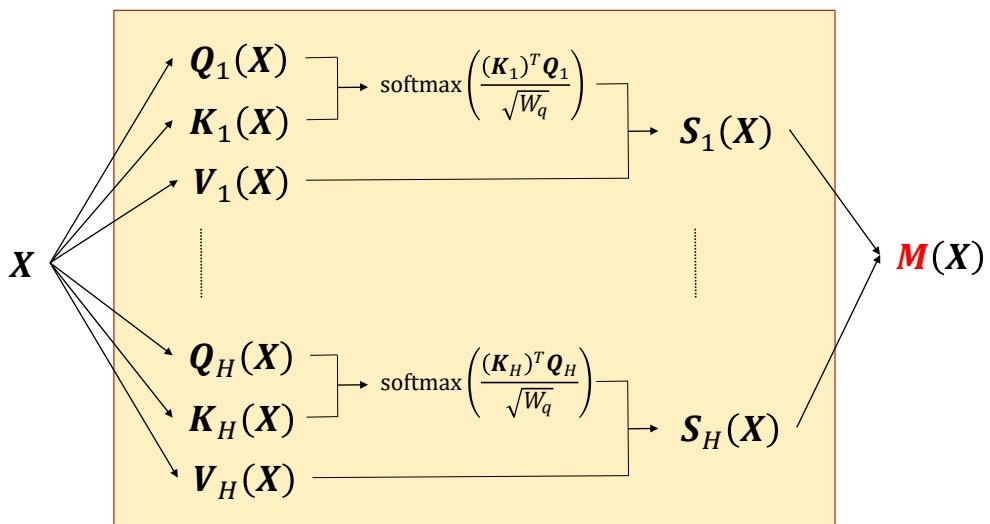
$$M(X) = \Theta_M \left((S_1(X))^T, (S_2(X))^T, \dots, (S_H(X))^T \right)^T$$

Θ_M linear transformation

Concepts:

Self-attention mechanisms (heads) applied in parallel, outputs are vertically concatenated

Multi-head self-attention



Transformer

Attention Is All You Need

Ashish Vaswani*
 Google Brain
 avaswani@google.com

Noam Shazeer*
 Google Brain
 noam@google.com

Niki Parmar*
 Google Research
 nikip@google.com

Jakob Uszkoreit*
 Google Research
 usz@google.com

Llion Jones*
 Google Research
 llion@google.com

Aidan N. Gomez* †
 University of Toronto
 aidan@cs.toronto.edu

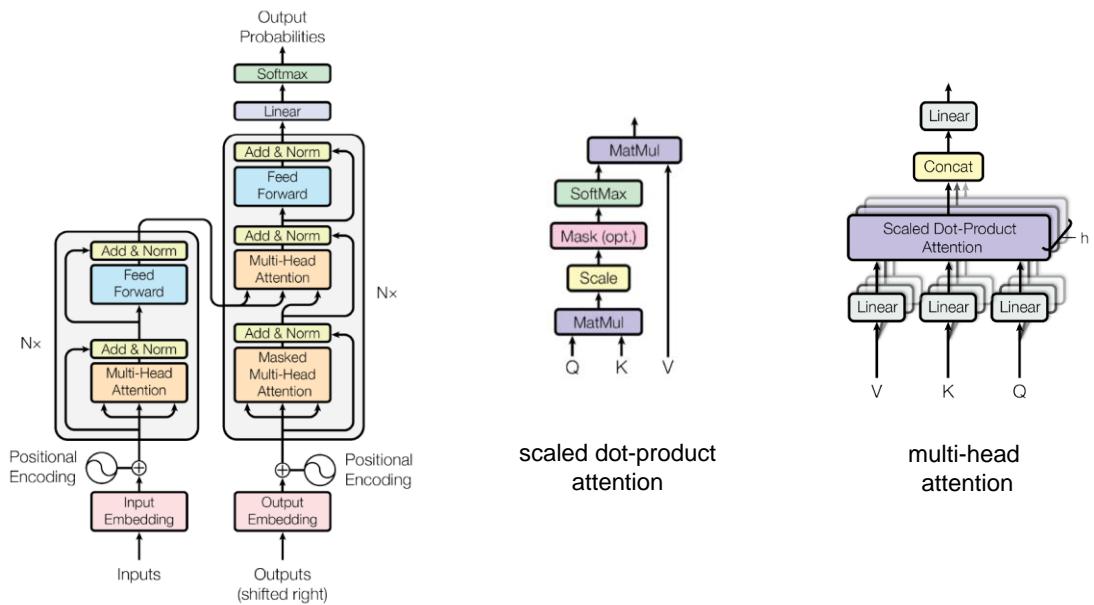
Lukasz Kaiser*
 Google Brain
 lukaszkaiser@google.com

Illia Polosukhin* ‡
 illia.polosukhin@gmail.com

Abstract

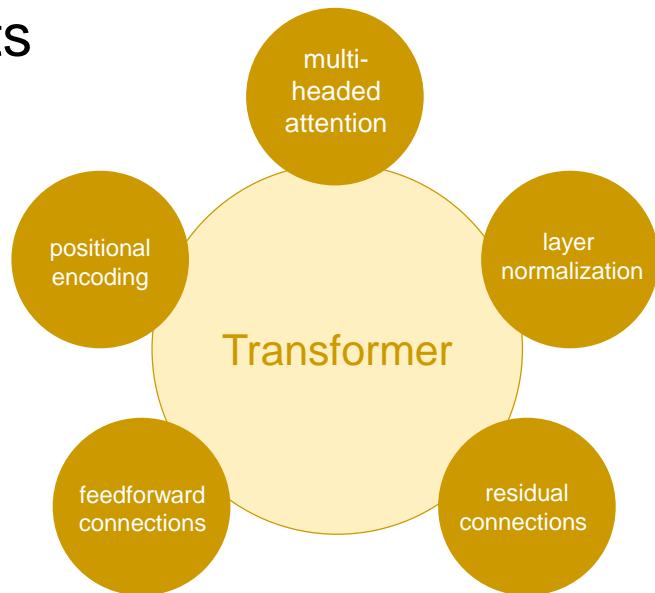
The dominant sequence transduction models are based on complex recurrent or convolutional neural networks that include an encoder and a decoder. The best performing models also connect the encoder and decoder through an attention mechanism. We propose a new simple network architecture, the Transformer, based solely on attention mechanisms, dispensing with recurrence and convolutions entirely. Experiments on two machine translation tasks show these models to be superior in quality while being more parallelizable and requiring significantly less time to train. Our model achieves 28.4 BLEU on the WMT 2014 English-to-German translation task, improving over the existing best results, including ensembles, by over 2 BLEU. On the WMT 2014 English-to-French translation task, our model establishes a new single-model state-of-the-art BLEU score of 41.8 after training for 3.5 days on eight GPUs, a small fraction of the training costs of the best models from the literature. We show that the Transformer generalizes well to other tasks by applying it successfully to English constituency parsing both with large and limited training data.

[arXiv:1706.03762](https://arxiv.org/abs/1706.03762)



[arXiv:1706.03762](https://arxiv.org/abs/1706.03762)

Key components



Positional encoding

$$\mathbf{X} \quad \xrightarrow{\quad} \quad \mathbf{X} + \Pi$$

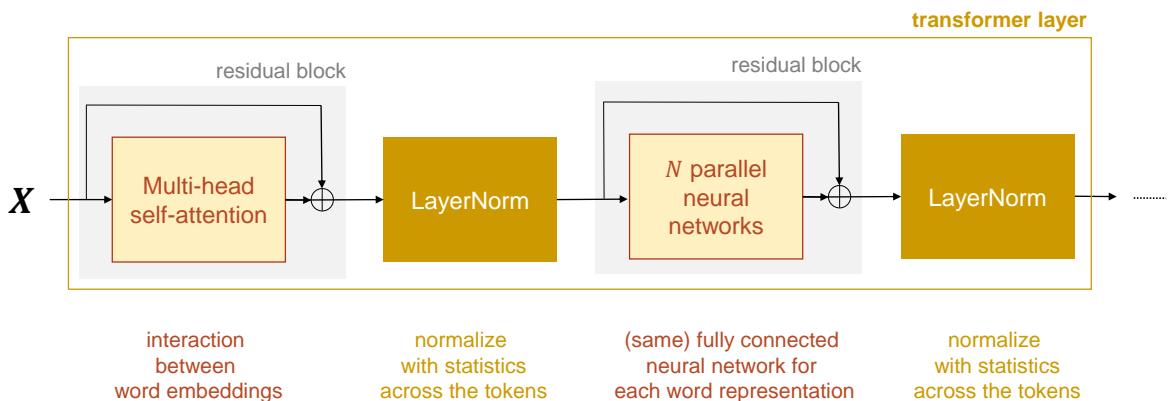
Π

- each **column** of Π is **unique** (information about the *position* in the input sequence)
- added to the input only or at every layer
- learned or “**hand-crafted**”

Concept:

Self-attention disregards the order of the tokens in a sentence

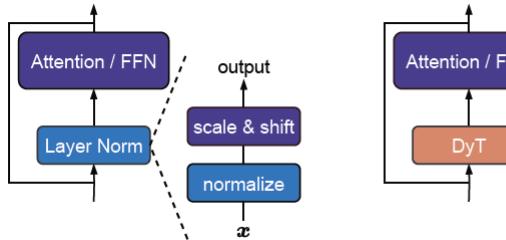
Transformer layer



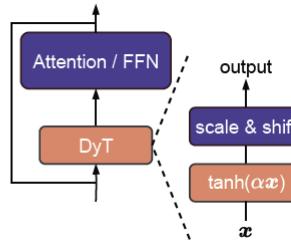
Concepts:

Output has the same size as the input ($W \times N$), residual block: output added to the original input

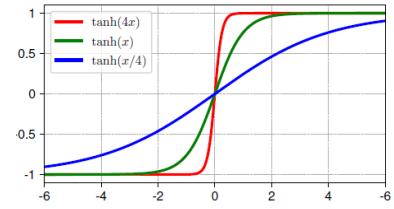
Layer normalization or Dynamic Tanh?



original transformer block
with **layer normalization**



block with **Dynamic Tanh**
(DyT) layer

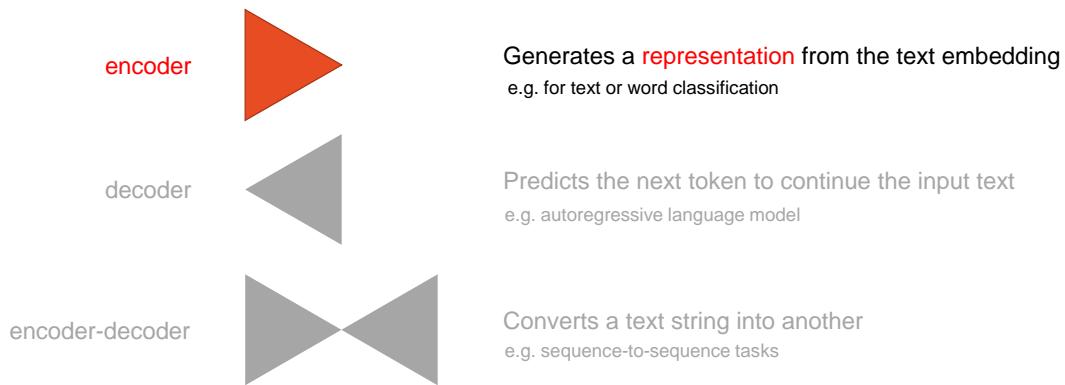


[arXiv:2503.10622](https://arxiv.org/abs/2503.10622)

Concept:
Transformers without normalization

Encoder model

Transformer models



Encoder model example: BERT

$P = 340,000,000$ parameters

$W = 1,024$ size of the embedding

$|\mathcal{V}| = 30,000$ size of the vocabulary (number of tokens)

24 transformers (transformer layers)

$H = 16$ heads for the self-attention of each transformer

$W_v = W_q = W_k = 64$ dimension of the value, query, key

Concept: Bidirectional Encoder Representations from Transformers (BERT):
jointly conditioning on both left and right context (in all layers)

Training BERT

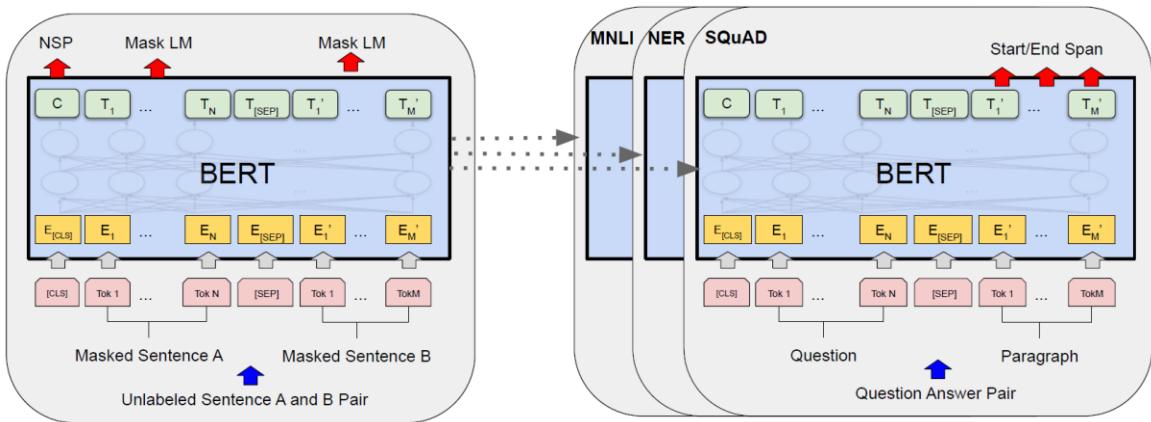


to learn general information
about the statistics of language
(on a **large corpus** of text)

adapted to solve a **task**
using a **smaller body**
of supervised training data

Concepts: Missing word prediction, understanding syntax,
(statistical) understanding about the world, fine-tune with extra layer after the transformer network

BERT: pre-training and fine-tuning

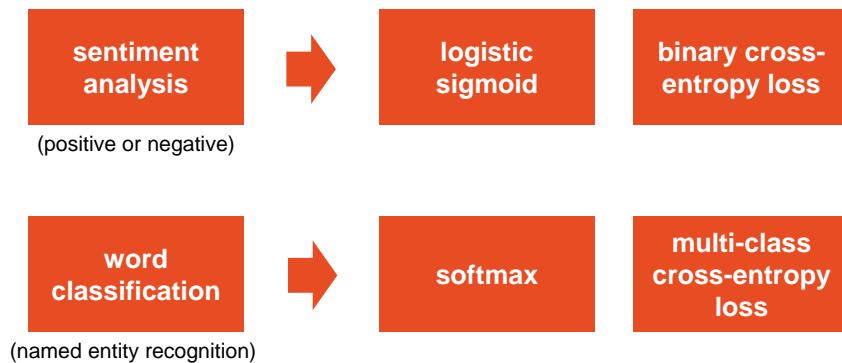


pre-training

fine-tuning
(all parameters)

[arXiv:1810.04805](https://arxiv.org/abs/1810.04805)

BERT: tasks



Decoder model

Transformer models



Decoder model example: GPT3

$P = 175,000,000,000$ parameters

$W = 12,288$ size of the embedding

$|\mathcal{V}| = 50,257$ size of the vocabulary (number of tokens)

96 transformers

$H = 96$ heads for the self-attention of each transformer

$W_v = W_q = W_k = 128$ dimension of the value, query, key

Context:

GPT4 has got $P = 1,800,000,000,000$ parameters

GPT3: next token prediction

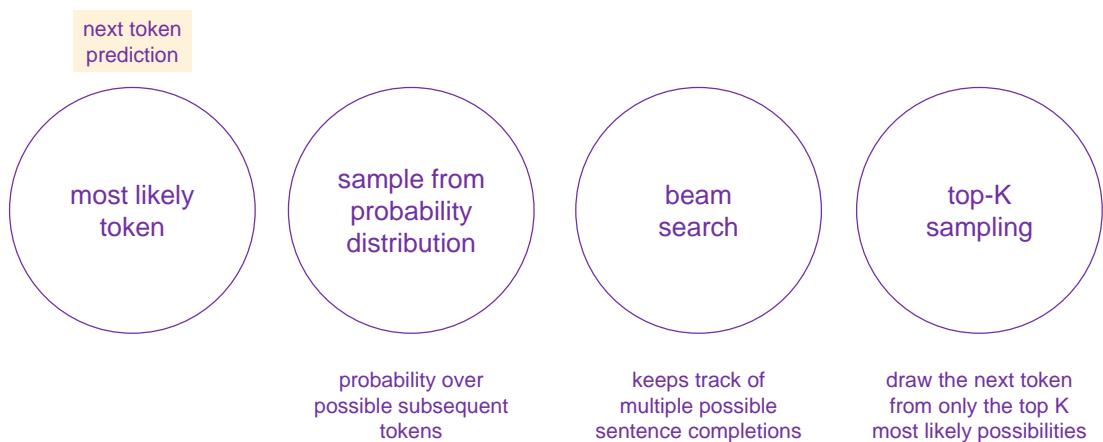
Each token t_n only interacts with previous tokens (**masked self-attention**)

$$P(\underbrace{t_1, t_2, \dots, t_{N-1}, t_N}_{\text{sentence}}) = P(t_1) \prod_{n=2}^N P(t_n|t_1, t_2, \dots, t_{n-1})$$

New extended sequence: fed back to the encoder for the **next token prediction**

Concepts: Autoregressive language model, predict the next token w/o access to the future, connection between maximizing the log probability of the tokens (loss function) and the next token prediction task

GPT3: generative model



The Slido logo, which consists of the word "slido" in a lowercase, bold, sans-serif font.

To understand transformers, was today's lecture more useful than prompting a large language model?

① Start presenting to display the poll results on this slide.

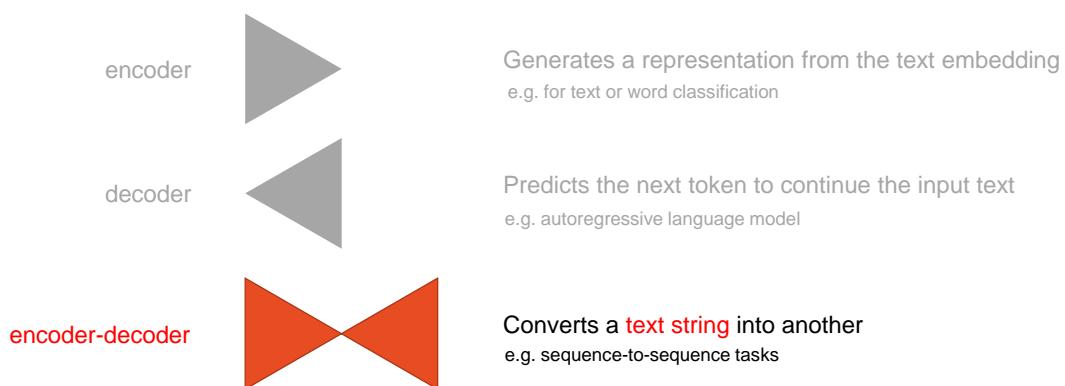
The Slido logo, which consists of the word "slido" in a lowercase, bold, sans-serif font.

In terms of learning about transformers, how did today's lecture compare to using a large language model?

① Start presenting to display the poll results on this slide.

Encoder-decoder model

Transformer models



Example: machine translation

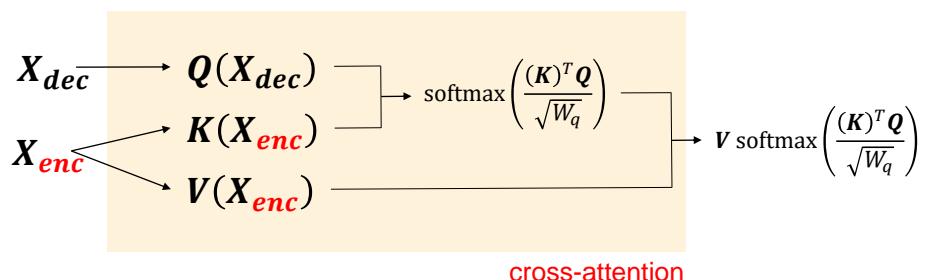
decoder layers also attend to the output of the encoder



each output token is *conditioned*
on the previous output tokens
and the source sentence

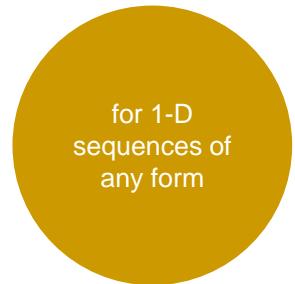
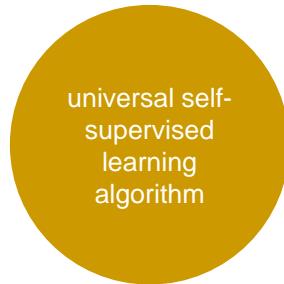
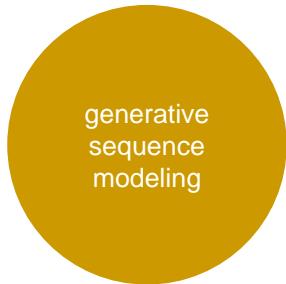
Concepts: Sequence-to-sequence task, information about the source language (encoder) and the target language statistics (decoder)

Encoder-decoder attention: cross attention



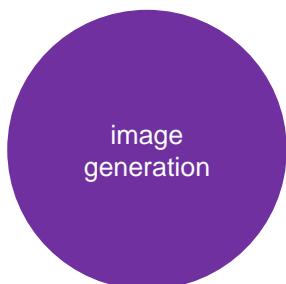
Concepts: Cross-attention (encoder-decoder attention), queries computed from the decoder embeddings, keys and values computed from the encoder embeddings

Transformers

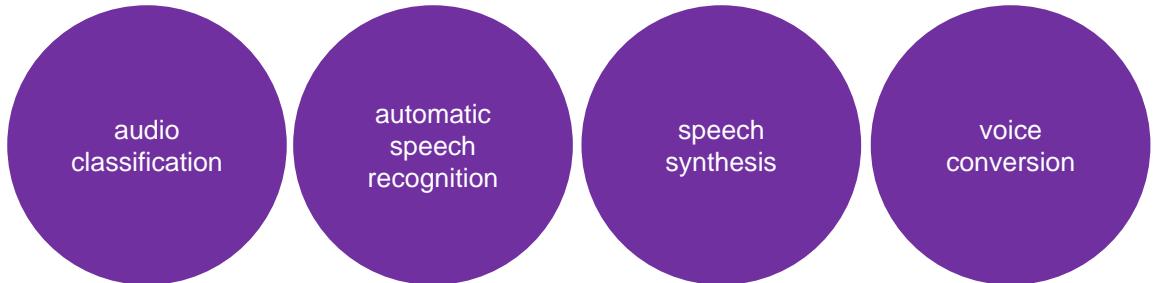


Concepts: Transformer as *self-supervised* learning algorithm, sequences of *bytes*, trained to maximize the likelihood (mode covering), can be applied to any data type

Transformers for visual tasks



Transformers for audio tasks



Transformers for audio-visual tasks



Exercises

Today's exercises

Practice

You will become familiar with fine-tuning a **transformer** for binary text classification

Marked (*last one for Submission 1!*)

Includes **positional encoding**, scaled dot-product **attention** and multi-head **attention**

The Slido logo, which consists of the word "slido" in a lowercase, sans-serif font with a green outline.

**What questions do you have
regarding the mini-project?**

ⓘ Start presenting to display the poll results on this slide.

What did we learn today?

- Natural Language Processing
- Tokens
- Self-attention
- Transformer
- Encoder model
- Decoder model
- Encoder-decoder model
- Exercises

EE-559

Deep Learning

andrea.cavallaro@epfl.ch