

Large Language Models on Supercomputers


A brief overview

Speaker: Simeon Harrison
Trainer at EuroCC Austria

EuroCC

Fully funded EU project

- EuroCC is EU-funded international initiative aimed to support the uptake of AI and High-Performance Computing (HPC) in Europe
- Set up of 32 National Competence Centres (NCCs) across Europe
- EuroCC Austria is one of them
- Service Provider for AI, HPC and HPDA



Need More Compute-Power?

LUMI

- Fastest supercomputer in Europe and the fifth fastest globally.
- Sustained computing power (HPL) is 380 petaflops
- Over 262 000 AMD EPYC CPU cores
- Equipped with AMD Radeon Instinct MI250X GPUs

<https://www.lumi-supercomputer.eu/>

Leonardo

- Second fastest supercomputer in Europe and the sixth fastest globally.
- Sustained computing power (HPL) is 239 petaflops
- Intel new gen Sapphire Rapids 56 cores
- Equipped with custom NVIDIA A100 SXM6 64GB GPUs

<https://leonardo-supercomputer.cineca.eu/>

European HPC Landscape

EuroHPC JU systems

Different access modes:
[Calls for Proposals](#)

EuroHPC development access:
[Opportunity to test the system](#)

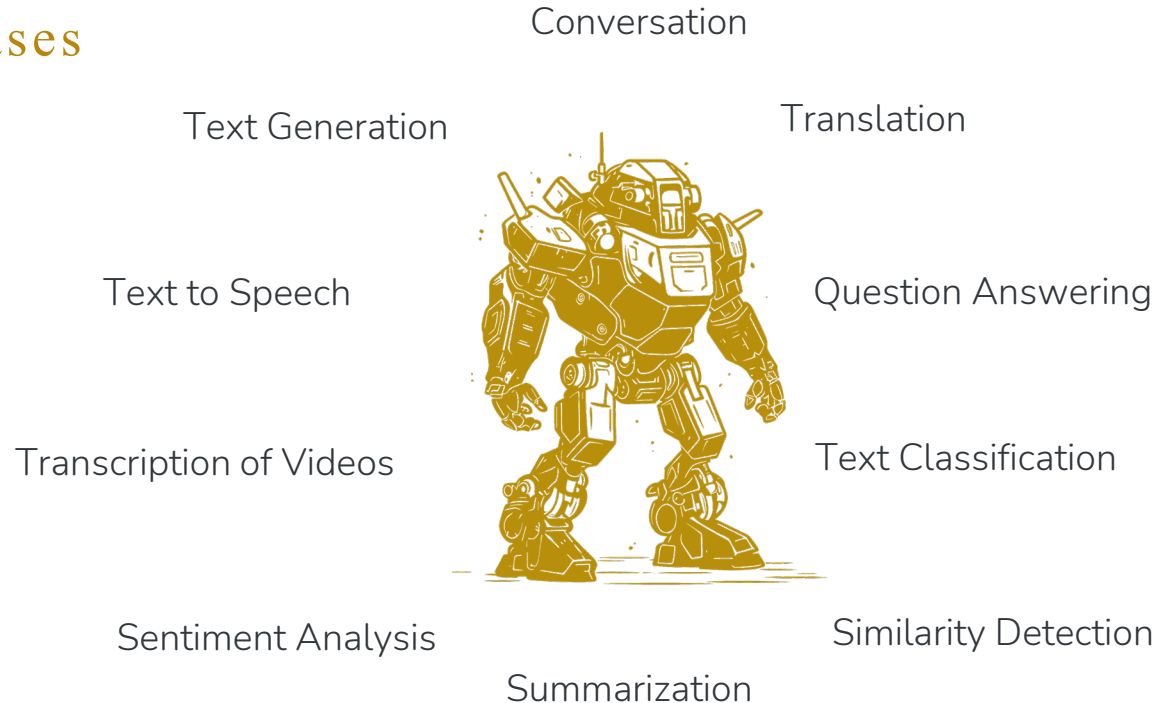
Applicants can request a small number of node hours to get acquainted with the supercomputers to further develop their software.




What can LLMs be used for?

Many different use-cases

- Made possible by the transformer architecture
- Choose your model according to the use-case
- Is there a pre-trained model for your use-case?



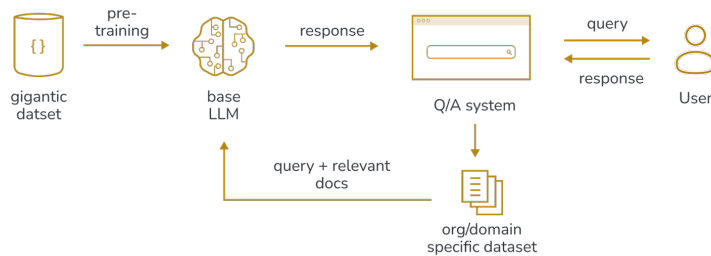
How can you influence LLMs?



How can you use LLMs with your data?


RAG:

Retrieval Augmented Generation



- Ideal for tapping into company's knowledge DBs
- Minimises hallucinations by grounding response on retrieved evidence
- Can quickly adapt to changing data
- Makes it easier to interpret result

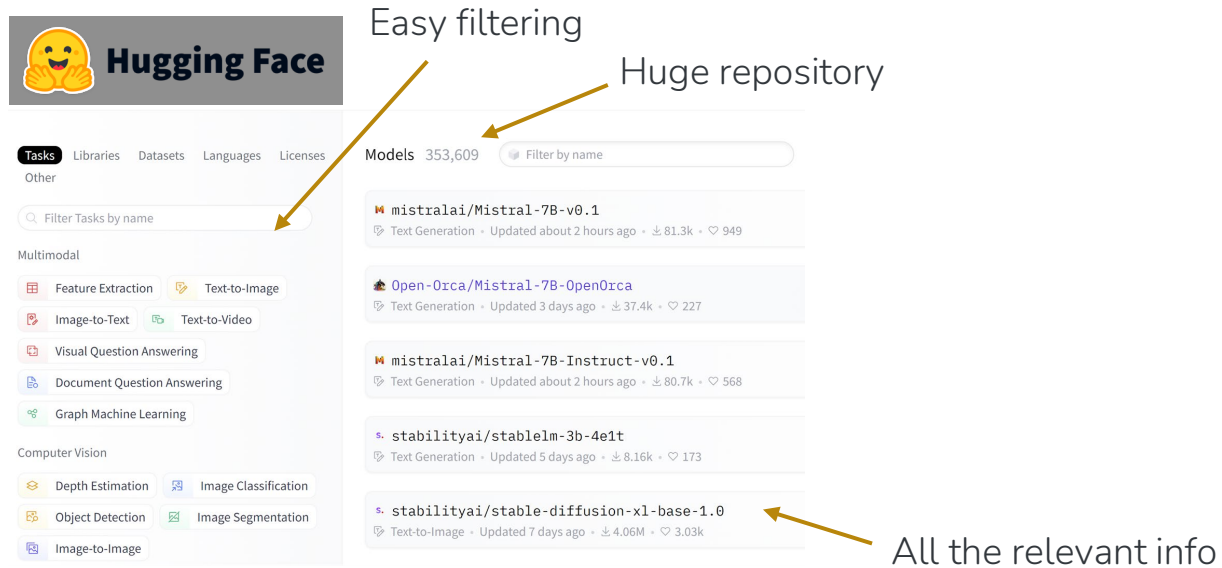
Finetuning



- Ideal if plenty of labelled data is available
- Teaches model domain specific vocabulary
- Company's writing/answer style is „baked“ into model through fine-tuned parameters

Transformer Models

Spoilt for Choice at <https://huggingface.co/>



The screenshot shows the Hugging Face Models page. Annotations include:

- Easy filtering**: Points to the 'Filter Tasks by name' search bar.
- Huge repository**: Points to the 'Models 353,609' header.
- All the relevant info**: Points to the details of the 'stabilityai/stable-diffusion-xl-base-1.0' model, which includes the model name, task type (Text-to-Image), update time (7 days ago), size (4.06M), and likes (3,033).



Pick the Right Model

mistralai/Mistral-7B-Instruct-v0.2   like 1.04k



Text Generation



Transformers



PyTorch



Safetensors

mistral

finetuned

conversational

 arxiv:2310.06825



License: apache-2.0



Model card



Files and versions



Community

61



Edit model card

Model Card for Mistral-7B-Instruct-v0.2

The Mistral-7B-Instruct-v0.2 Large Language Model (LLM) is an improved instruct fine-tuned version of [Mistral-7B-Instruct-v0.1](#).

For full details of this model please read our [paper](#) and [release blog.post](#).

Prepare your Data

Garbage in – garbage out

- Most underrated aspect of AI
- Most time consuming aspect of AI. Time spent in data preparation reflects in the quality of the product
- For fine-tuning you need labelled data
- Remember, that you are going to change the models parameters with your data

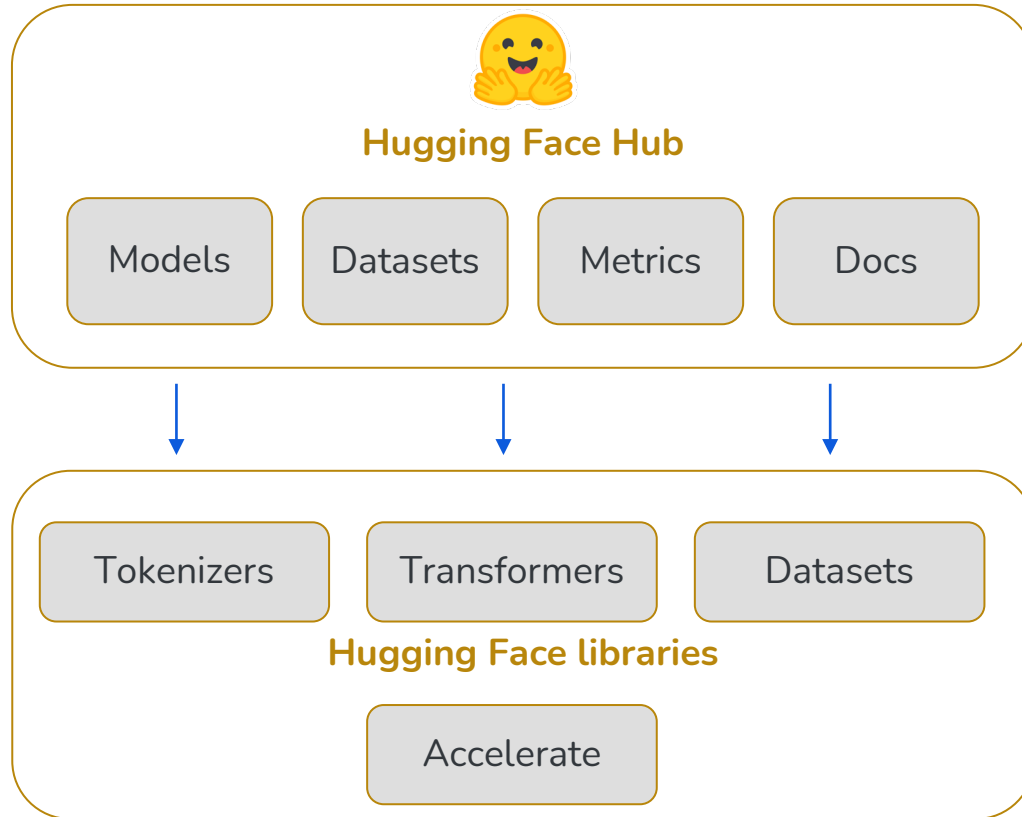


The Hugging Face Ecosystem

From 🤗 Transformers to the 🤗 Hub

Speaker: Simeon Harrison
Trainer at EuroCC Austria

Hub and libraries



Libraries

Tokenizers

Very fast at tokenizing text, thanks to the Rust backend.

Takes care of pre- and postprocessing, such as normalization and transformation to required format.

You can load tokenizers just like you load models.

Transformers

Provides a standardized interface to a wide range of transformer models.

Supports PyTorch, TensorFlow and JAX and makes switching between them easy.

Provides task-specific heads to fine-tune LLMs on downstream tasks

Libraries

Datasets

Simplifies the process of downloading and transforming data into required format.

Interface to thousands of datasets that can be found on the hub.

Interoperable with most common data wrangling tools and frameworks such as Pandas, NumPy etc.

Accelerate

Adds a layer of abstraction to training loops with the goal of making code portable (local laptop to cluster)

Makes training on multi-GPU, multi-node easier to accomplish, as you can switch between parallelization strategies.

Transformer Anatomy

Attention is really all you need?

Speaker: Simeon Harrison
Trainer at EuroCC Austria

Transformer Anatomy

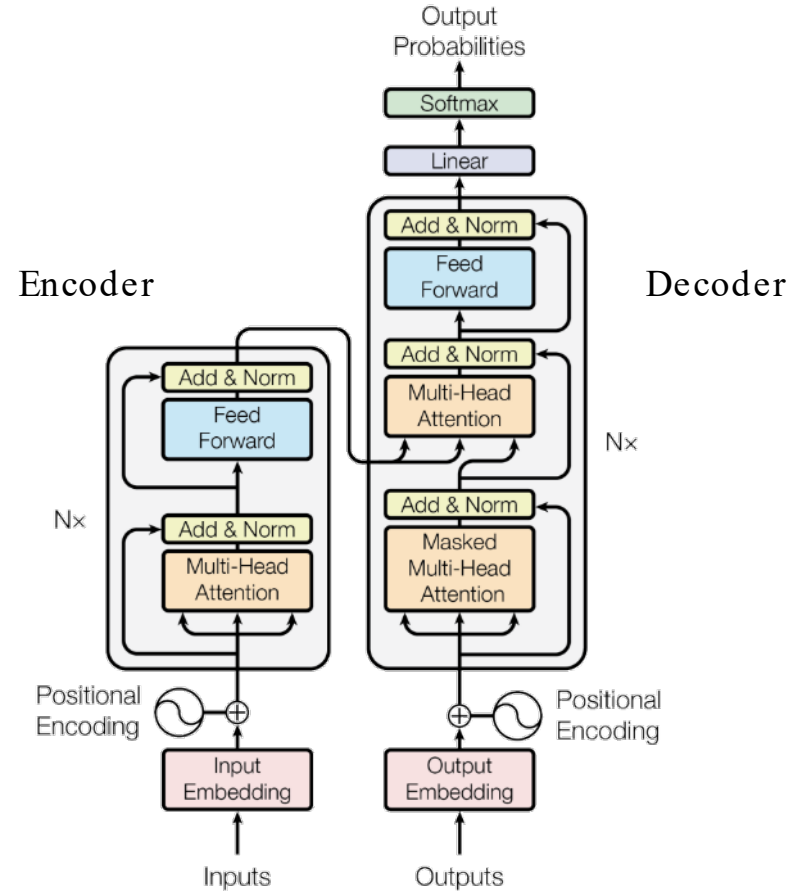
The original architecture

A transformer consists of an encoder and/or decoder block.

Words (tokens) are input as numerical representations (embeddings).

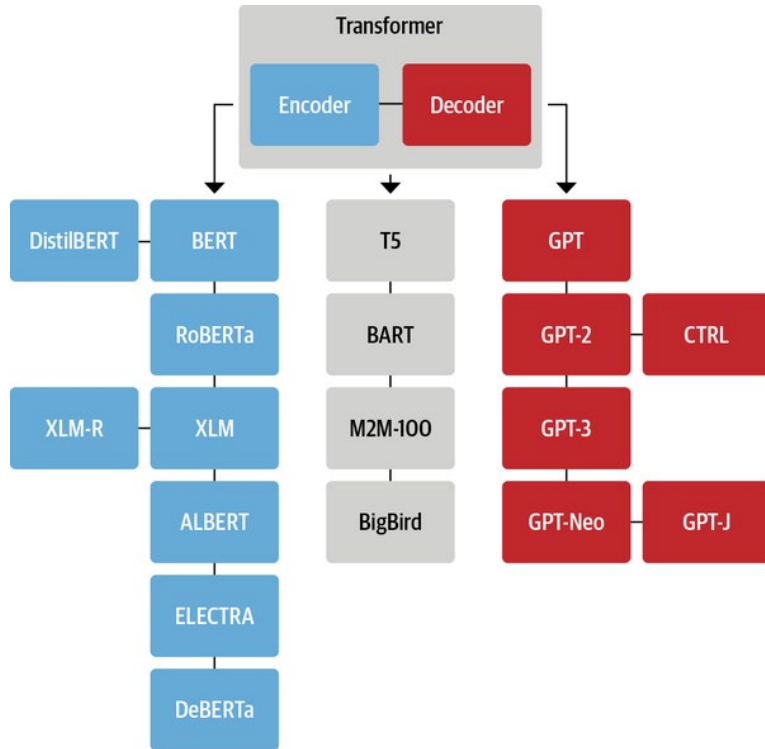
About 1/3 of all parameters are in the multi-head attention blocks

About 2/3 of all parameters are in the feed forward networks (also known as multi layer perceptron)



Source: "Attention Is All You Need", Vaswani et al.

Transformer Family



Encoder only:

These models excel at text classification, named entity recognition, and question answering

Decoder only:

Very good at predicting the next word in a sequence, therefore mostly used for text generation

Encoder-Decoder:

These models are often used for machine translation or summarization tasks.

Context Is All You Need

Embeddings

Here, we will refer to “word” instead of “token”, as it makes the content easier to explain.

A word embedding comes as a multi dimensional vector (e.g. 12.000 dim).

The initial word embedding in all of the examples of the word „mole“ is the same.



The European mole is a mammal



Take a biopsy of the mole

$$6.02 \times 10^{23}$$

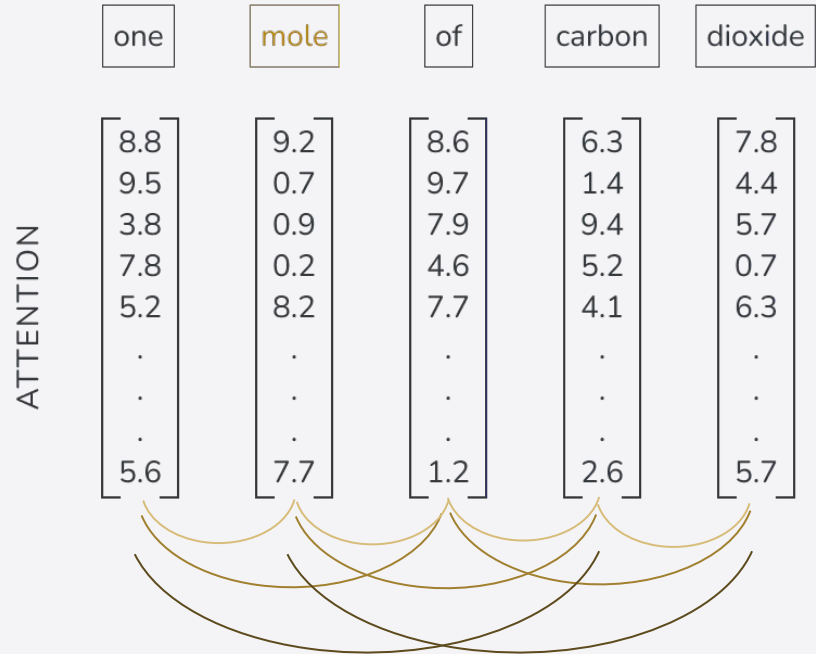
One mole of carbon dioxide

Context Is All You Need

Attention

The word „mole“ should be represented by a **unique vector** in the embedding space, depending on its **context**.

An **attention** block should **compute the vectors that you need to add** to the original, generic vector to get it to the correct, meaningful, rich representation, depending on the context in which the word is used.




Context Is All You Need

Lion

We associate the word „lion“ with a big cat, living wild on the African continent.

We probably imagine a majestic predator with a big mane.

The embedding of the word „lion“ is a vector with a certain length and direction within the embedding space.

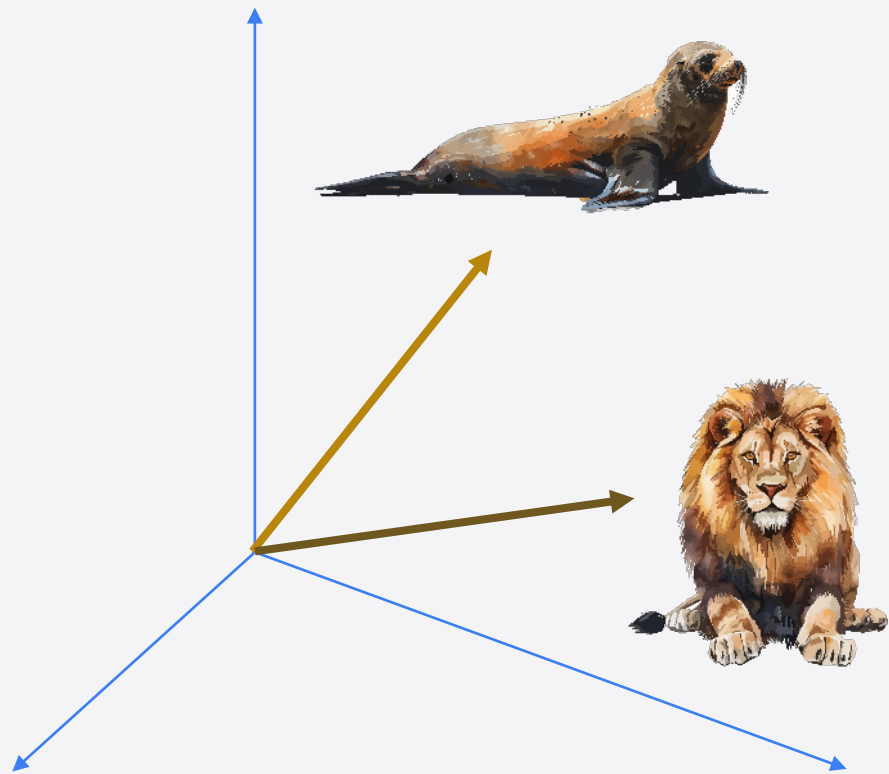


Context Is All You Need

Sea Lion

However, as soon we add the word „sea“ in front of „lion“ we imagine a totally different animal.

The same goes for the embedding. The attention mechanism needs to update the direction and length of the vector so that it represents the animal in question correctly.




Context Is All You Need

Sea Lion Cuddly Toy

The context depends on more than just the immediate words to the left and right.


The embedding of „sea lion cuddly toy“ will certainly be very different of just „lion“.

In order to achieve that the vector for „lion“ needs to attend to all the other words in the input (context size).




Self Attention


the European mole is a mammal




Self-attention




take a biopsy of the mole



Self-attention



Self Attention



Attention! Queries, Keys and Values

The Attention Score

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

Q.....query matrix

K.....key matrix

V.....value matrix

d.....dimension of (smaller) query-key space

Attention! Queries, Keys and Values

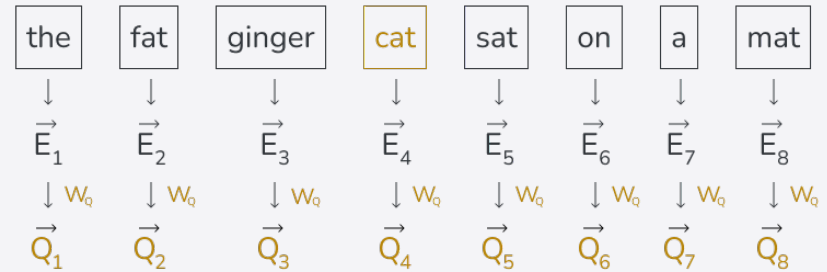
Query

The query vector is obtained by multiplying a matrix full of trainable parameters with the embedding vector of that word (token).

Each word has its own query vector:

$$\vec{Q}_i = W_Q \vec{E}_i$$

It maps the embedding vector to a much smaller dimensional space (e.g. 128 dimensions)



Imagine it like this:

One particular attention head is focussing on nouns. The query vector is like looking for labels with „adjective“ on them to better understand the noun.

In reality, this is much more abstract.

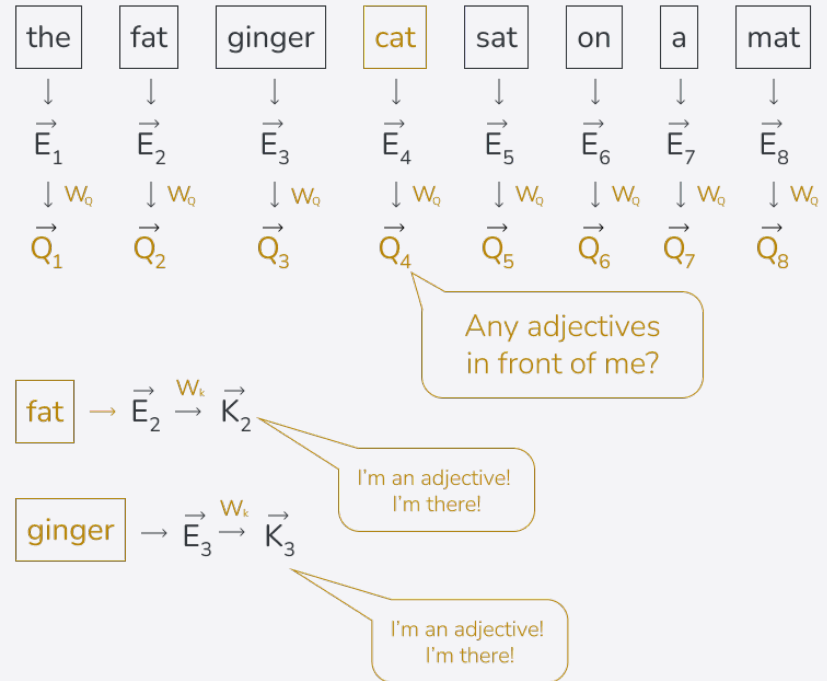
Attention! Queries, Keys and Values

Key

The key vector is also obtained by multiplying a matrix full of trainable parameters with the embedding vector of that word (token). Each word also has its own key vector: $\vec{K}_i = W_K \vec{E}_i$

Imagine it like this:

We are still dealing with the particular attention head that is focussing on nouns. The key is answering the question the query raised.



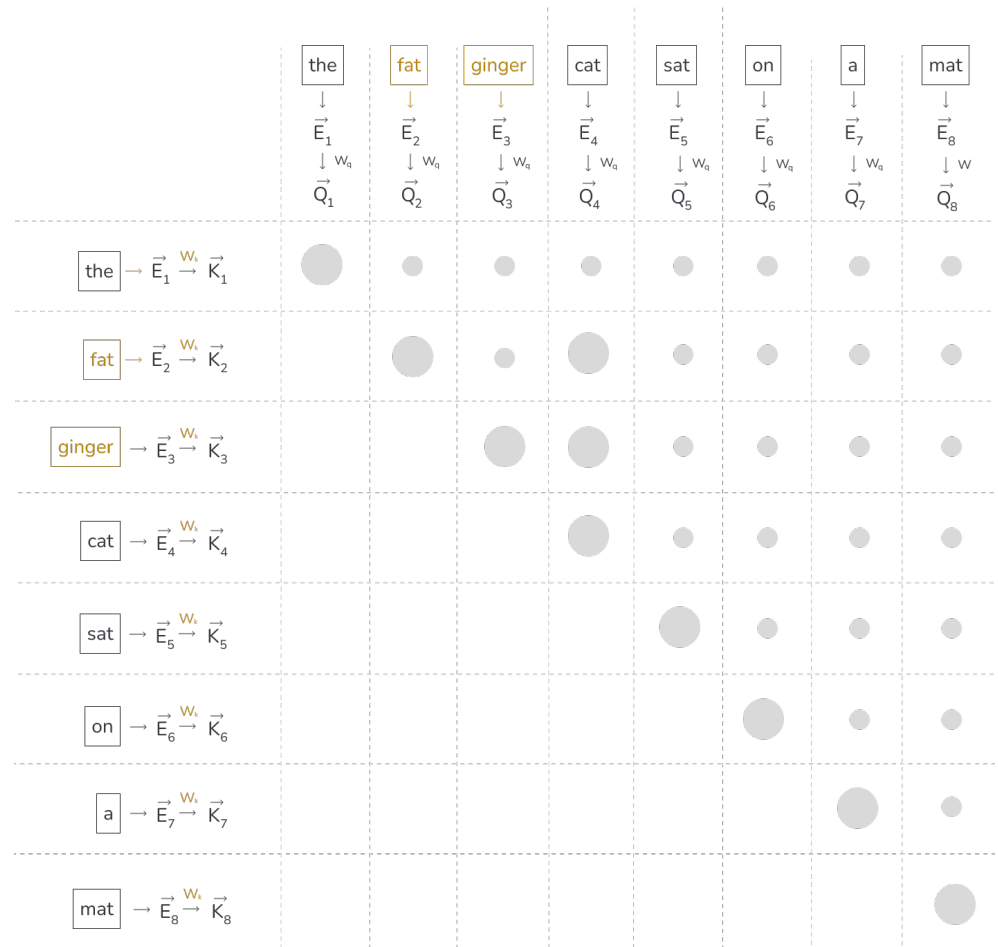
Attention! Queries, Keys and Values

Query – Key

Compute the dot product with each query-key pair, to determine how well the key matches the query. Where the queries and keys align, the dot product is larger.

Imagine it like this:

With our previous example, the dot product of the key vectors of „fat“ and „ginger“ with the vector of „cat“ yields the largest result. The embeddings of „fat“ and „ginger“ attend to „cat“.



Attention! Queries, Keys and Values

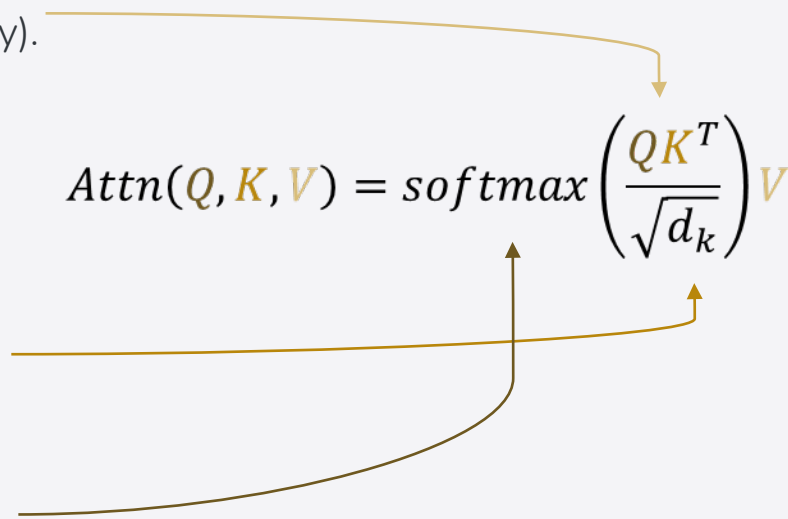
Attention

So far, we have determined which word is relevant to which other word (dot product of query and key).

We would like to use this as a score for how relevant every word is to update the meaning of other words.

For numerical stability the dot product is divided by the square root of the dimension of the query-key space.

To normalize the numbers to be between 0 and 1 we apply softmax.

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


Attention! Queries, Keys and Values

Value

The value matrix multiplied by the embedding of the preceding word results in the value vector.

$$\vec{V}_i = W_K \vec{E}_i$$

	the ↓ \vec{E}_1	fat ↓ \vec{E}_2	ginger ↓ \vec{E}_3	cat ↓ \vec{E}_4	sat ↓ \vec{E}_5	on ↓ \vec{E}_6	a ↓ \vec{E}_7	mat ↓ \vec{E}_8
$\vec{E}_1 \xrightarrow{W_V} \vec{V}_1$	1.0 \vec{V}_1	0.0 \vec{V}_1	0.0 \vec{V}_1	0.0 \vec{V}_1	0.0 \vec{V}_1	0.0 \vec{V}_1	0.0 \vec{V}_1	0.0 \vec{V}_1
$\vec{E}_2 \xrightarrow{W_V} \vec{V}_2$	0.0 \vec{V}_2	1.0 \vec{V}_2	0.0 \vec{V}_2	0.42 \vec{V}_2	0.0 \vec{V}_2	0.0 \vec{V}_2	0.0 \vec{V}_2	0.0 \vec{V}_2
$\vec{E}_3 \xrightarrow{W_V} \vec{V}_3$	0.0 \vec{V}_3	0.0 \vec{V}_3	1.0 \vec{V}_3	0.58 \vec{V}_3	0.0 \vec{V}_3	0.0 \vec{V}_3	0.0 \vec{V}_3	0.0 \vec{V}_3
$\vec{E}_4 \xrightarrow{W_V} \vec{V}_4$	0.0 \vec{V}_4	0.0 \vec{V}_4	0.0 \vec{V}_4	0.0 \vec{V}_4	0.0 \vec{V}_4	0.0 \vec{V}_4	0.0 \vec{V}_4	0.0 \vec{V}_4
$\vec{E}_5 \xrightarrow{W_V} \vec{V}_5$	0.0 \vec{V}_5	0.0 \vec{V}_5	0.0 \vec{V}_5	0.0 \vec{V}_5	0.0 \vec{V}_5	0.0 \vec{V}_5	0.0 \vec{V}_5	0.0 \vec{V}_5
$\vec{E}_6 \xrightarrow{W_V} \vec{V}_6$	0.0 \vec{V}_6	0.0 \vec{V}_6	0.0 \vec{V}_6	0.0 \vec{V}_6	0.99 \vec{V}_6	1.0 \vec{V}_6	0.0 \vec{V}_6	0.0 \vec{V}_6
$\vec{E}_7 \xrightarrow{W_V} \vec{V}_7$	0.0 \vec{V}_7	0.0 \vec{V}_7	0.0 \vec{V}_7	0.0 \vec{V}_7	0.0 \vec{V}_7	0.0 \vec{V}_7	1.0 \vec{V}_7	1.0 \vec{V}_7
$\vec{E}_8 \xrightarrow{W_V} \vec{V}_8$	0.0 \vec{V}_8	0.0 \vec{V}_8	0.0 \vec{V}_8	0.0 \vec{V}_8	0.0 \vec{V}_8	0.0 \vec{V}_8	0.0 \vec{V}_8	0.0 \vec{V}_8
	\parallel $\Delta \vec{E}_1$	\parallel $\Delta \vec{E}_2$	\parallel $\Delta \vec{E}_3$	\parallel $\Delta \vec{E}_4$	\parallel $\Delta \vec{E}_5$	\parallel $\Delta \vec{E}_6$	\parallel $\Delta \vec{E}_7$	\parallel $\Delta \vec{E}_8$

Attention! Queries, Keys and Values

Value

Each value vector is then multiplied by the corresponding weight for this word and the results summed up.

The result $\Delta \vec{E}_i$ is the change, that needs to be added to the original embedding to get an updated, richer meaning.

Since this happens to every word in the sequence, we end up with a set of more refined embeddings.

$$\begin{array}{l} \text{the} \rightarrow \vec{E}_1 \xrightarrow{w_v} \vec{V}_1 \\ \text{fat} \rightarrow \vec{E}_2 \xrightarrow{w_v} \vec{V}_2 \\ \text{ginger} \rightarrow \vec{E}_3 \xrightarrow{w_v} \vec{V}_3 \\ \text{cat} \rightarrow \vec{E}_4 \xrightarrow{w_v} \vec{V}_4 \\ \text{sat} \rightarrow \vec{E}_5 \xrightarrow{w_v} \vec{V}_5 \\ \text{on} \rightarrow \vec{E}_6 \xrightarrow{w_v} \vec{V}_6 \\ \text{a} \rightarrow \vec{E}_7 \xrightarrow{w_v} \vec{V}_7 \\ \text{mat} \rightarrow \vec{E}_8 \xrightarrow{w_v} \vec{V}_8 \end{array}$$

$$\begin{array}{l} 0.0 \vec{V}_1 \\ + \\ 0.42 \vec{V}_2 \\ + \\ 0.58 \vec{V}_3 \\ + \\ 0.0 \vec{V}_4 \\ + \\ 0.0 \vec{V}_5 \\ + \\ 0.0 \vec{V}_6 \\ + \\ 0.0 \vec{V}_7 \\ + \\ 0.0 \vec{V}_8 \end{array}$$

cat
↓
 \vec{E}_4



$$\begin{array}{l} \text{cat} \\ \downarrow \\ \vec{E}_4 \\ + \\ \Delta \vec{E}_4 \\ \parallel \\ \vec{E}'_4 \end{array}$$



Attention

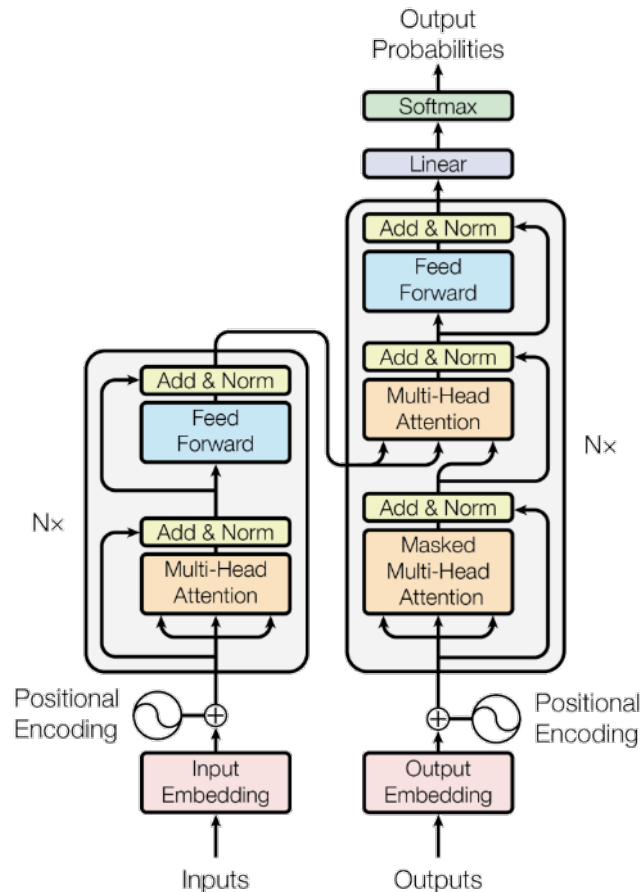
Multi-Head Attention

The attention mechanism we just looked at, is done several times in parallel.

Each attention head is focussing on different features of the embeddings and computes its own $\Delta \vec{E}_i^{(j)}$

The resulting $\Delta \vec{E}_i^{(j)}$ vectors, each suggesting the necessary change, are added to the original embedding.

This can be done in parallel on GPUs.



Attention

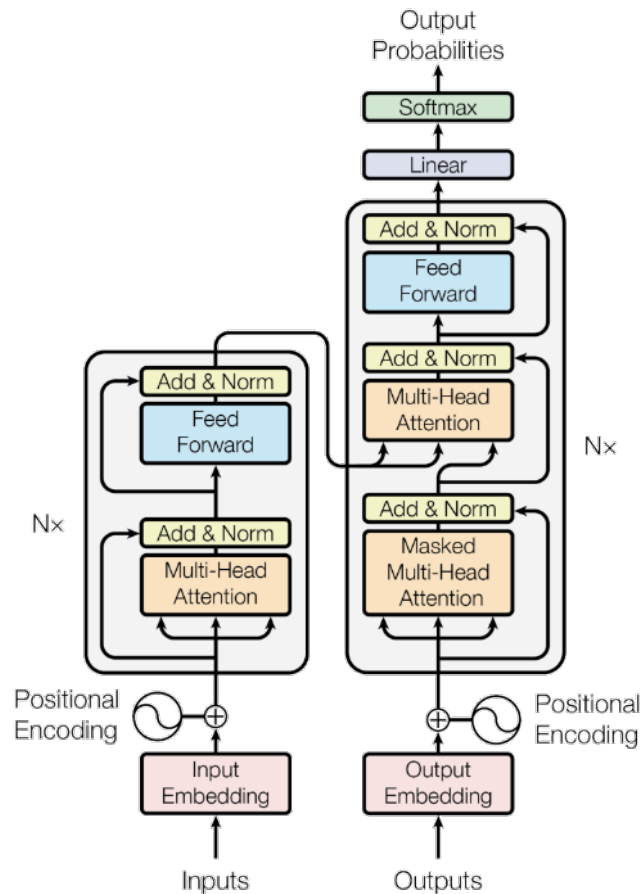
Cross Attention

Cross attention is almost the same as self attention.

Difference:

- query and key maps act on different data sets (e.g. 2 different languages in machine translation)
- no masking, since there is no issue of later words affecting earlier ones.


Source: "Attention Is All You Need", Vaswani et al.




Feed Forward Network

Multi Layer Perceptron (MLP)

- Home to approx. 2/3s of all parameters
- This is where „knowledge“ is baked in
- Source of hallucinations
- Facts that are associated with input embeddings are added to the input embeddings
- Each embedding vector can be processed independently -> parallelization!



Feed Forward Network



Tokenization & Embeddings

Turning words into numbers

Speaker: Simeon Harrison
Trainer at EuroCC Austria

From Text to Tokens

Character Tokenization

- Only needs 256 integers
- Good to tokenize rare words
- Helps with misspellings
- Loses linguistic structure
- Rarely used in practice

L a r g e l a n g u a g e m
o d e l s o n s u p e r c o m
p u t e r s

From Text to Tokens

Word Tokenization

- Model does not have to learn words from characters
- Each word has specific ID
- Size of vocabulary explodes
- Model needs to learn different tokens for e.g. singular and plural

Large language models
on supercomputers

From Text to Tokens

Subword Tokenization

- Best of both worlds
- Deal with complex words easily
- Frequent words remain as one token

Large language models
on super computers


From Tokens to Vectors

Embeddings

Tokens are mapped to unique integers according to the vocabulary size of the tokenizer.

Now, the tokens need to be embedded, which means turned into a vector representation.

This is done by an embedding layer of a model. The model takes each token ID and looks it up in an embedding matrix. The embedding matrix is a learned set of weights that maps each token ID to a corresponding high-dimensional vector (embedding).




From Tokens to Vectors

$$E(\text{aunt}) - E(\text{uncle}) \approx E(\text{woman}) - E(\text{man})$$

Embeddings

Each word/token has a unique direction in the embedding space

Similar words point in a similar direction.



THANK YOU



This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101101903. The JU receives support from the Digital Europe Programme and Germany, Bulgaria, Austria, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, Greece, Hungary, Ireland, Italy, Lithuania, Latvia, Poland, Portugal, Romania, Slovenia, Spain, Sweden, France, Netherlands, Belgium, Luxembourg, Slovakia, Norway, Türkiye, Republic of North Macedonia, Iceland, Montenegro, Serbia