

PAUL SCHERRER INSTITUT



Midjourney v5.2

Stefan Ritt :: Head of Muon Physics Group :: Paul Scherrer Institut

ChatGPT for Physicists (*and Engineers*)

Unveiling the Power and Potential of Large Language Models

LTP Seminar, September 11, 2023





I will give a talk about ChatGPT for Physicists and need a short abstract. First I will cover some theory of Large Language Models including embedding, positional encoding, self-attention and transformers. In the second part I will show practical applications such as writing recommendation letters, do some brainstorming, and finally some simple data analysis. I will show strengths and weaknesses of a LLM and how to make best use of it as a physicist or engineer.

ChatGPT for Physicists (and Engineers): Unveiling the Power and Potential of Large Language Models

In the ever-evolving computational landscape available to physicists, Large Language Models (LLMs) like ChatGPT are emerging as potent tools. This talk aims to offer both a foundational and applied perspective on LLMs, tailored to cater to the physicist's inquisitiveness.

In the first segment, we delve into the intricate mechanics behind LLMs. We'll start by exploring how words and phrases are transformed into vectors in high-dimensional spaces through embedding, capturing their semantic essence. This will be followed by an exploration of positional encoding, which ensures the models retain the sequence order, a cornerstone for understanding context. The journey continues with the self-attention mechanism, shedding light on how models discern the significance of different words in a sequence. Culminating this segment, we'll unravel the transformative architecture of transformers, the linchpin of the Natural Language Processing domain and the backbone of ChatGPT.

Transitioning to the practical realm, attendees will discover various applications of LLMs for the modern physicist. From crafting compelling recommendation letters with the aid of ChatGPT to enhancing ideation processes by leveraging its expansive knowledge and creativity, the possibilities are vast. Finally, a glimpse into how straightforward data analysis tasks can be augmented using LLMs will be presented.

Concluding the talk, a balanced discourse on the strengths and limitations of LLMs will be provided, guiding physicists and engineers on how best to wield these models. Join us for this immersive dive into the transformative potential of ChatGPT in the world of physics and engineering. (This abstract was written by ChatGPT-4 August 3, 2023)

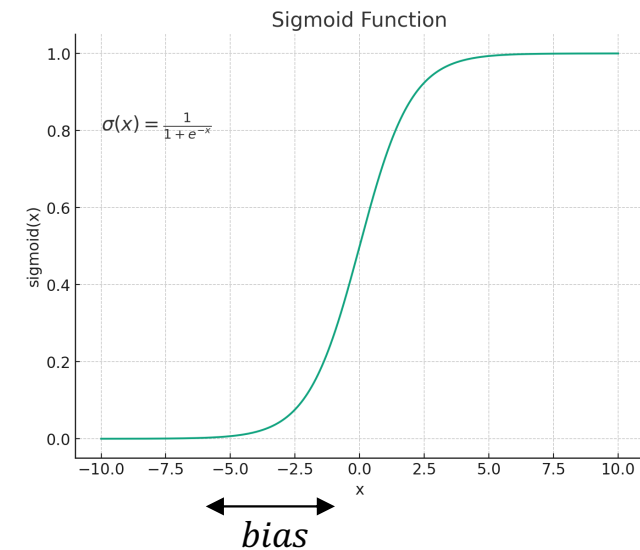
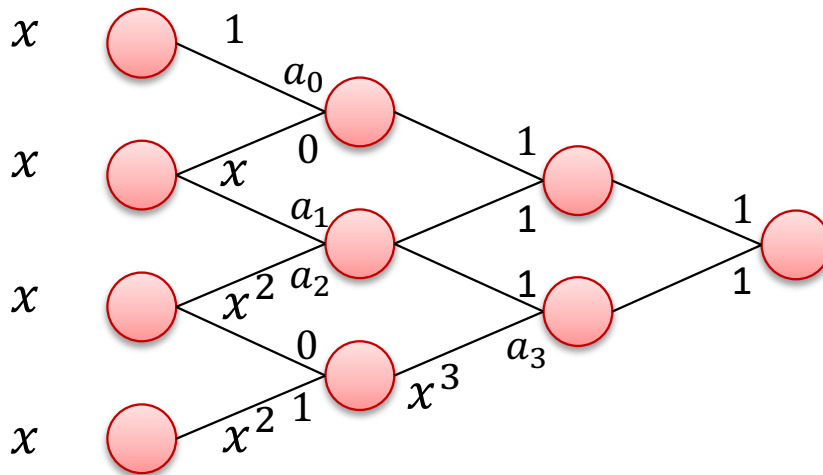
- Introduction to AI
 - Neural networks
 - Machine learning
 - Large Language Networks (LLN)
- Practical Applications of ChatGPT
 - Writing text
 - Summarizing text
 - Coding: Write, check, comment
 - Plugins: PDF, ...
 - Code interpreter (Advanced Data Analysis)
- AI has risks and opportunities
- My way: learn some details and make best use of it
- Will share my experience with you today from the perspective of a physicist
- Some of you might be even more experienced, but I hope there is something for everybody
- This talk may change the way you work
- Think about how Google search changed your way to work
- In the second part I want to hear oohs and aahs

Neural Net - Polynomial

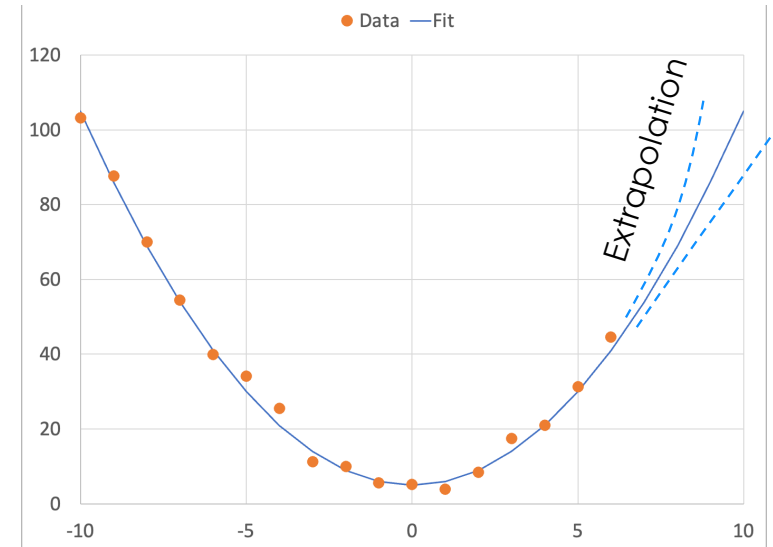
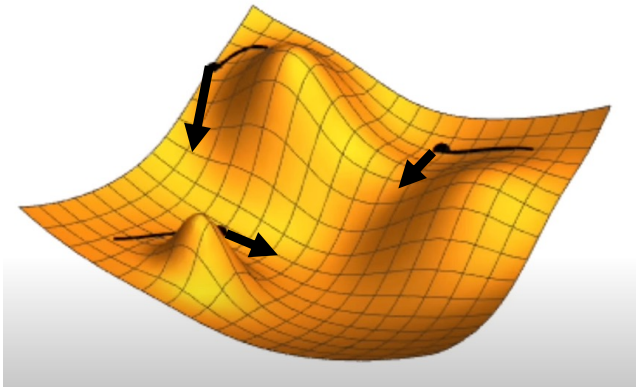
- **Polynomials** can fit **any** function
- A neural net is a kind of polynomial
 - node adds inputs and applies activation function like in our brain

$$f(x) = a_0 + a_1x + a_2x^2 + a_3x^3 + \dots$$

$$\text{red circle} = f\left(\sum \text{input}_i + \text{bias}\right)$$



- **Training** a Neural Network is very similar to **fitting** a polynomial to a function:
 - Create a χ^2 and minimize it
 - Gradient method: calculate partial derivations of all parameters and adjust them



$$\chi^2 = \sum (f(x) - y)^2 \rightarrow \min$$

$$a'_i = a_i + \gamma \cdot \frac{\partial \chi^2}{\partial a_i}$$

↑
"learning rate"
~0.001


Interactive Training Demo

<https://playground.tensorflow.org>

Epoch **000,119**
 Learning rate **0.03**
 Activation **Tanh**
 Regularization **None**
 Regularization rate **0**
 Problem type **Classification**

DATA

Which dataset do you want to feed in?



Ratio of training to test data: 50%

Noise: 0

Batch size: 10

REGENERATE

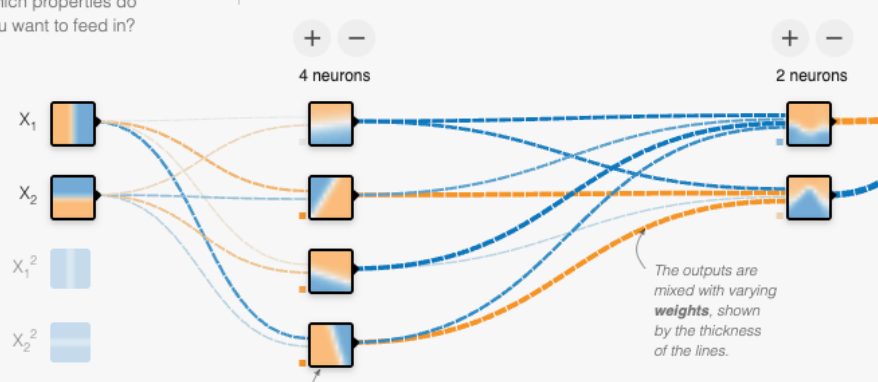
FEATURES

Which properties do you want to feed in?

- X_1
- X_2
- X_1^2
- X_2^2
- $X_1 X_2$
- $\sin(X_1)$
- $\sin(X_2)$

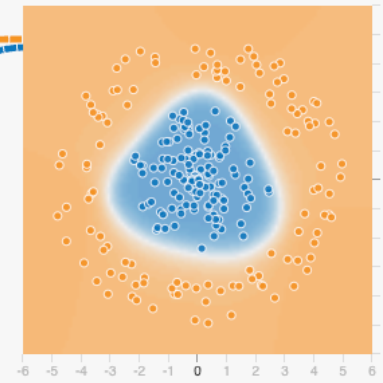
2 HIDDEN LAYERS

4 neurons 2 neurons



OUTPUT

Test loss 0.016
Training loss 0.008



Colors shows data, neuron and weight values.

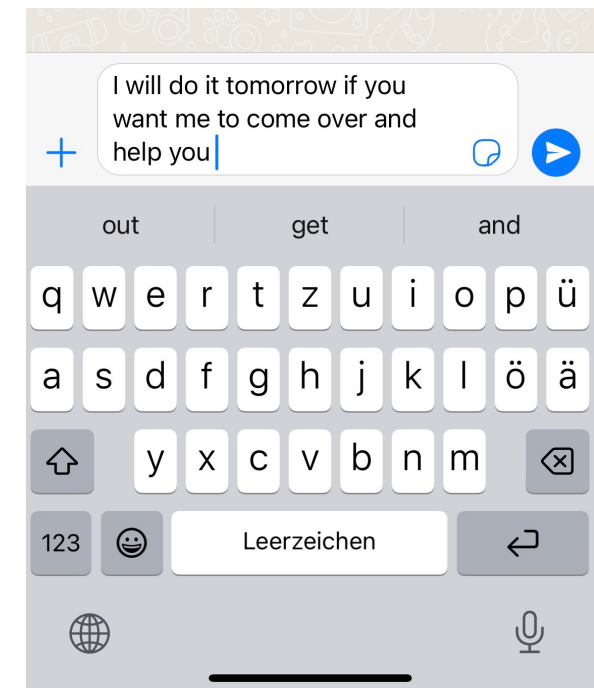
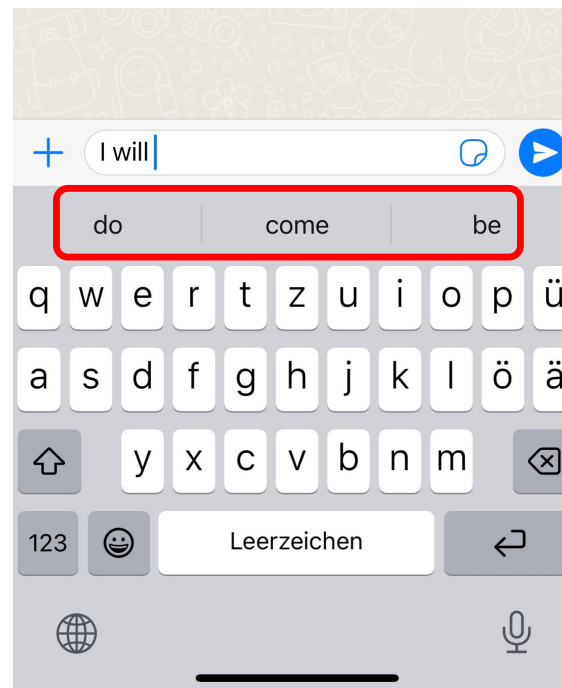
Show test data Discretize output

Text generation

Text can be **predicted** using **probabilities**, see your smartphones autocomplete function:

I will

- do 13.4 %
- come 11.2 %
- be 8.1 %



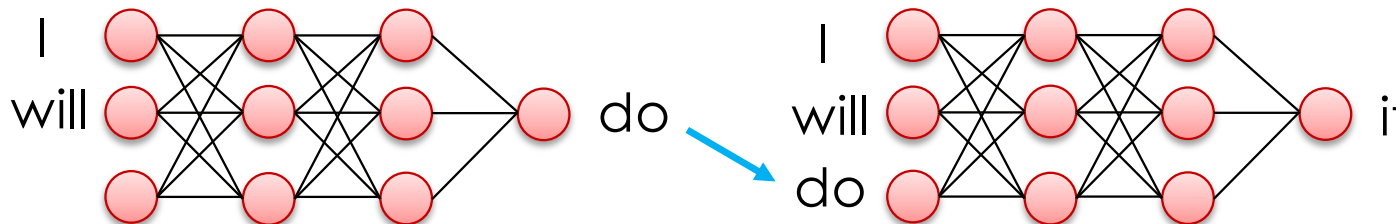
Text generation

- How **big** a **table** with probabilities do we **need** ?
- Assume we have **10,000** words and make a sentence with **10 words**:

word word word word word word word word word word

$$10^4 \times 10^4 \times 10^4 \times 10^4 \times 10^4 \times 10^4 \times 10^4 \times 10^4 \times 10^4 \times 10^4 = 10^{40}$$

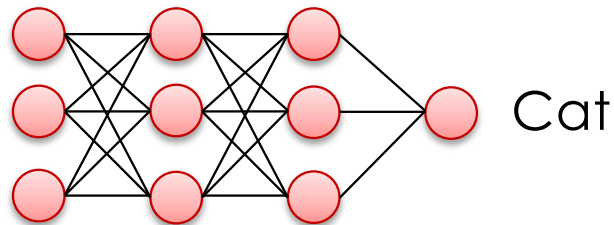
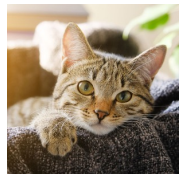
- Probability table with **10^{40}** entries is **not possible**
- Most of the 10^{40} sentences make no sense
- All meaningful sentences are **still too big** to fit in a table
- → **Approximate** the probability table with a **Neural Net**



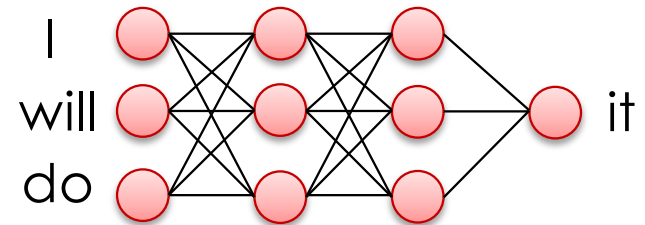
Labeled learning:



Cat Cat Dog Cat



Un-labeled learning:



Meaning of words in sentence

- **Meaning** of words in a sentence **depend** on their **order**

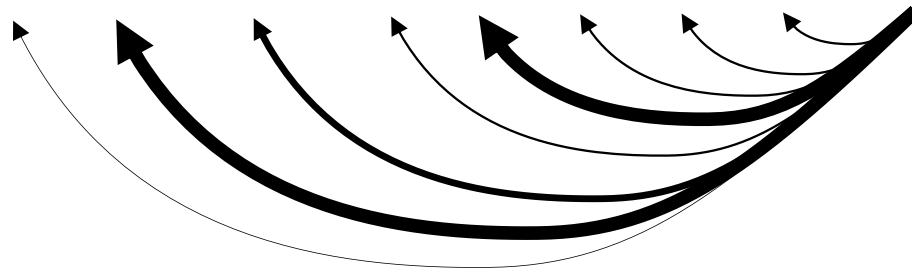
- **Goliath** has defeated **David**
- **David** has defeated **Goliath**

- **Meaning** of words in a sentence **depend** on their **context**

- I walk across the street to the other **bank**
- I swim across the river to the other **bank**

→ *financial institution*

→ *land along a river*



Words have **different weights** to influence the **context**

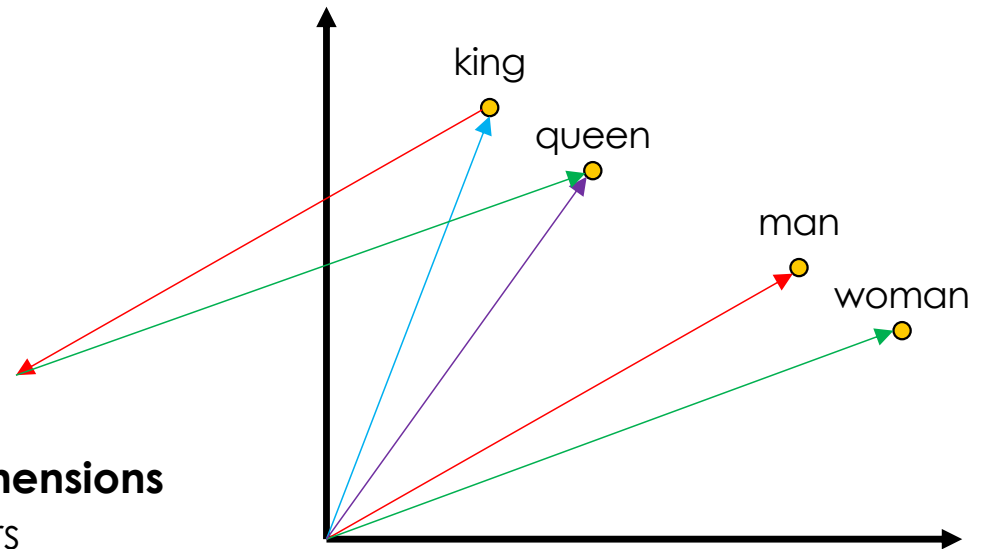
Word embedding

- **Neural nets need numbers** as input
- Simplest method: Give each word a **unique number**
- Better method: Encode **words** in a **vector** v_i

man	woman	king	queen
0.72	0.81	0.23	0.32
0.32	0.28	0.63	0.59
...

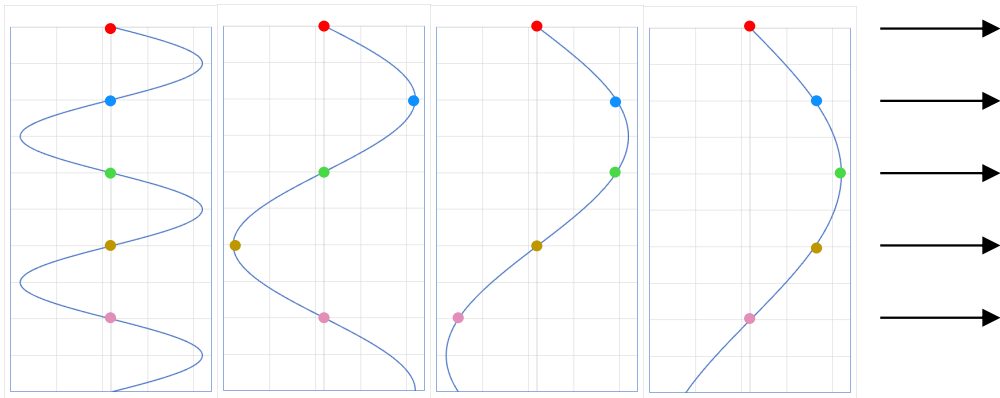
king - man + woman = queen

- **Typical** language models use **i=512 dimensions**
- Similar words should have similar vectors
- Coordinates are derived from **training**



Position encoding

d =	David	has	defeated	Goliath		$\sin(d/p*\pi)$		David	has	defeated	Goliath			
0	0.61	0.71	0.13	0.20	+	0.00	0.00	0.00	0.00	=	0.61	0.71	0.13	0.20
1	0.08	0.01	0.15	0.80		0.00	1.00	0.87	0.71		0.08	1.01	1.02	1.50
2	0.36	0.01	0.35	0.26		0.00	0.00	0.87	1.00		0.36	0.01	1.22	1.26
3	0.08	0.78	0.96	0.72		0.00	-1.0	0.00	0.71		0.08	-0.22	0.96	1.42
4	0.57	0.62	0.66	0.59		0.00	0.00	-0.8	0.00		0.57	0.62	-0.20	0.59
	p=1	2	3	4										



0.00	0.00	0.00	0.00
0.00	1.00	0.87	0.71
0.00	0.00	0.87	1.00
0.00	-1.0	0.00	0.71
0.00	0.00	-0.8	0.00

Note:

Current models use alternating sin/cos and different frequencies, so distance of words can be easily calculated from vectors.

https://kazemnejad.com/blog/transformer_architecture_positional_encoding/

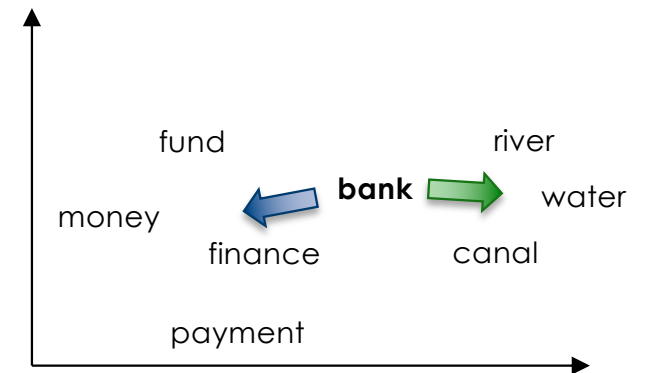
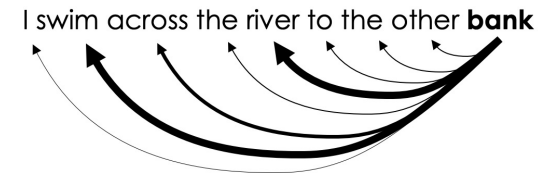


Position encoding



Context encoding

- **Context** matters!
- We need a **mechanism** that **considers** neighboring words to **enhance** the **meaning** of the word of interest
- Idea: use vector **dot product** of vectors
 → large for **similar word**
 → zero for perpendicular vectors



I swim across the river to the other bank

I swim across the river to the other bank

1	0.2	0.5	0.2	0.5	0.5	0.5	0.5	0.5	0.5
0.2	1	0.5	0.5	0.8	0.5	0.5	0.5	0.5	0.7
0.5	0.5	1	0.5	0.5	0.5	0.5	0.5	0.2	0.2
0.2	0.5	0.5	1	0.5	0.5	0.5	0.5	0.5	0.5
0.5	0.8	0.5	0.5	1	0.5	0.5	0.5	0.5	0.7
0.5	0.5	0.5	0.5	0.5	1	0.5	0.5	0.5	0.5
0.5	0.5	0.5	0.5	0.5	0.5	1	0.5	0.2	0.2
0.5	0.5	0.2	0.5	0.5	0.5	0.5	1	0.5	0.5
0.5	0.7	0.2	0.5	0.7	0.5	0.2	0.5	0.5	1

v : word vector
 s : similarity

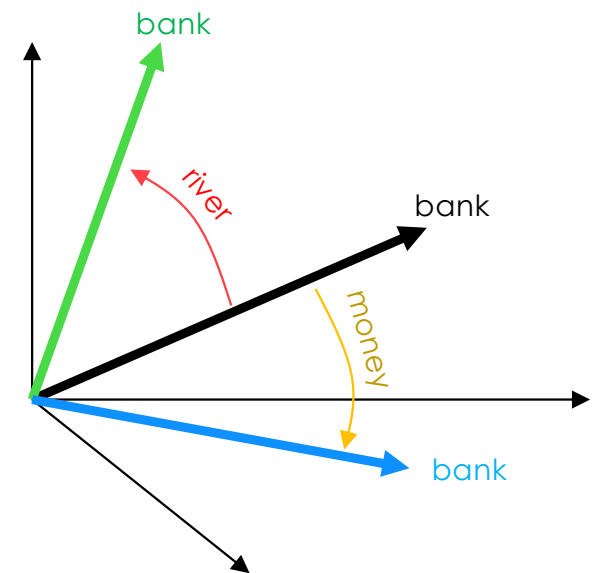
$$s_{ij} = v_i \cdot v_j^T$$

Vector rotation

- Word **vectors** v_i in 512-dim space
- **Similarities** $s_{ij} = v_i v_j^T$
- **Weights** $w_{ij} = \text{softmax}(s_{ij}) = \frac{e^{s_{ij}}}{\sum e^{s_{ij}}}$ $\sum w_{ij} = 1$
- Modify word vectors by **matrix multiplication**

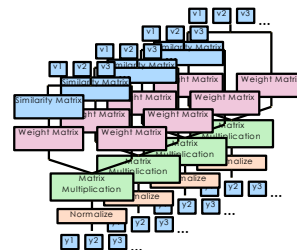
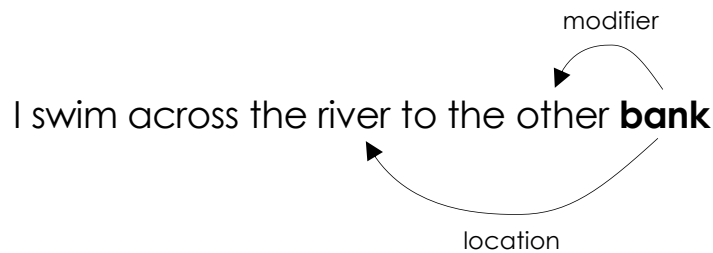
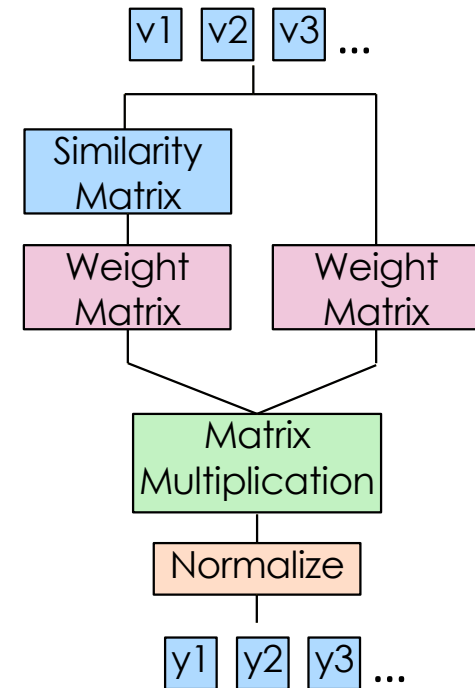
$$y_i = v_i \times w_{ij}$$

- Word vectors are **rotated** in 512-dim space, influenced by neighboring words
- If a word has **similar neighbors**, it's **rotated similarly**



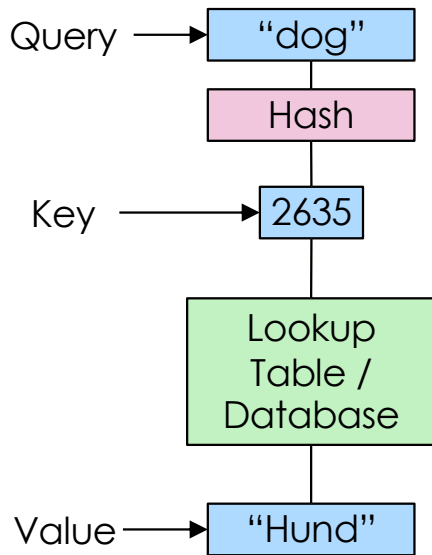
Self-Attention Mechanism

- Word vectors
 - Similarity Matrix
 - Rotation (Matrix Multiplication)
 - Normalize Matrix
 - **Contextualized representation of word vectors**
- Add Machine Learning by **weighting original** words with **matrix of weights** derived from training
 - **Self-Attention Mechanism**
- Put several Self-Attention Mechanism in parallel
 - **Multi-headed** Self-Attention Mechanism

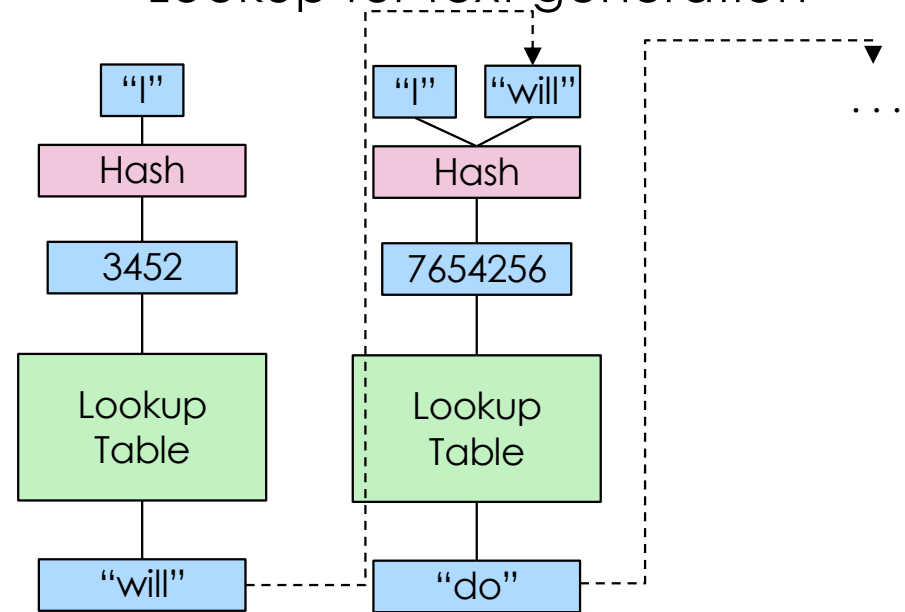


Traditional Database Lookup

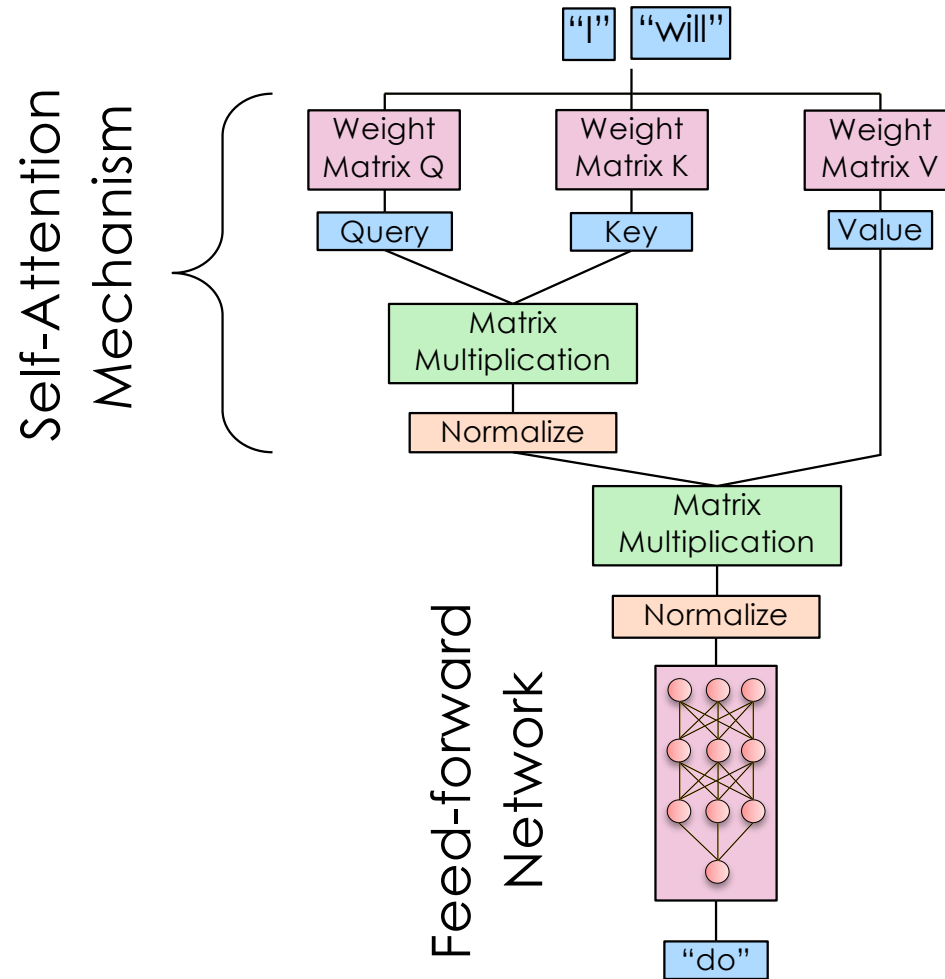
Lookup for translations



Lookup for text generation



Lookup with Neural Nets



Putting all together: Transformer

Attention Is All You Need

<p>Ashish Vaswani* Google Brain avaswani@google.com</p>	<p>Noam Shazeer* Google Brain noam@google.com</p>	<p>Niki Parmar* Google Research nikip@google.com</p>	<p>Jakob Uszkoreit* Google Research usz@google.com</p>
<p>Llion Jones* Google Research llion@google.com</p>	<p>Aidan N. Gomez* † University of Toronto aidan@cs.toronto.edu</p>	<p>Łukasz Kaiser* Google Brain lukasz.kaiser@google.com</p>	
<p>Illia Polosukhin* † illia.polosukhin@gmail.com</p>			

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 41.0 after training for 3.5 days on eight GPUs, a small fraction of the training costs of the best models from the literature.

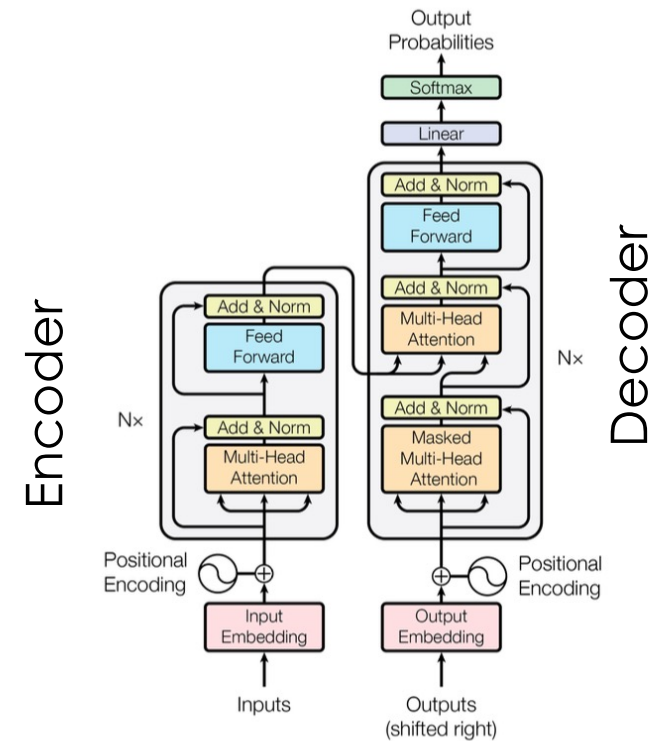


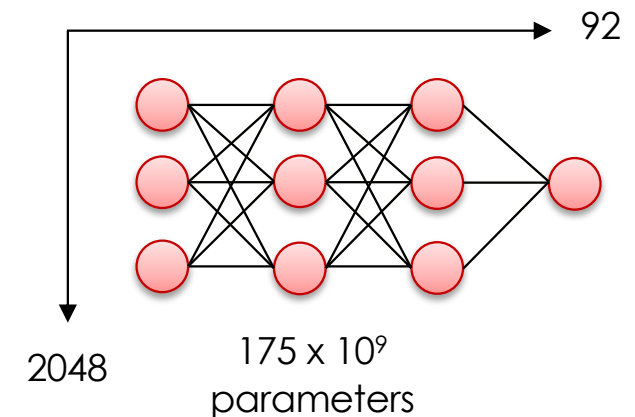
Figure 1: The Transformer - model architecture.

<https://arxiv.org/abs/1706.03762>

June 2017

ChatGPT 3

- Transformer Mechanism + Huge Neural Network + ... → Large Language Model (**LLM**)
- <https://en.wikipedia.org/wiki/GPT-3>
- **2048 Tokens** (2-chars + <EOS> + punctuation)
- **92** hidden network layers
- **175 Billion** Parameters, 800 GB of storage
- Trained with **1.8 PB** of text (500 B tokens)
 - 3% Wikipedia
 - 16% Books1 & Books2
 - 81% Common Crawl + WebText2
- **305 Years** for training →
~ months on huge cluster, cost ~**5 M\$**
- Training data until **Sept. 2021**
- **Pre**-training plus **fine**-tuning:
Generative Pre-trained Transformer 3 (**GPT 3**)
- Free version of ChatGPT at <https://chat.openai.com>



Hallucination

- [https://en.wikipedia.org/wiki/Hallucination_\(artificial_intelligence\)](https://en.wikipedia.org/wiki/Hallucination_(artificial_intelligence))
- ChatGPT is good to generate **nice text**, but NOT to produce **correct output**
- **Is 17077 a prime?**
 - Solution is **not present** in training data
 - Solutions learned:
 - 17 is prime since it can only be divided by 1 and 17
 - 77 is not a prime, it can be divided by 7 and 11
 - **Hallucination:**
 - 17077 is not a prime, since it can be divided by 7, 11 and 17
- LLM cannot deal with a situation where it is **not trained in**

Important Lesson to be learned

The output of a Large Language Model
is “**mostly**” correct

Use it for cases where you can **double check** or
do not rely on the correctness of the output




“Stammtisch” (pub)
discussion



Write some text

Write some recommendation for my colleague Dr. Stefan Ritt. He works in particle physics, is very knowledgeable about data acquisition and computing, a nice guy, and started using machine learning. He even gave a talk about ChatGPT for physicists.

- Write some recommendation for a colleague
 - Write in different style and tones

 Write some recommendation for my colleague Stefan. He works in particle physics, is very knowledgeable about data acquisition and computing, a nice guy, and started using machine learning. He even gave a talk about ChatGPT in physics.



Recommendation for Stefan

To whom it may concern,

It's with great pleasure and full confidence that I recommend Stefan for any endeavor he chooses to pursue within the realm of particle physics or beyond. Having had the privilege to work alongside him and witness his capabilities firsthand, there are a few traits and achievements that particularly stand out.

- Expertise in Particle Physics:** Stefan possesses an in-depth understanding of particle physics. His insights, analyses, and contributions have consistently reflected not just theoretical knowledge but also practical experience.
- Mastery in Data Acquisition and Computing:** In our ever-evolving digital age, Stefan's prowess in data acquisition and computing is invaluable. He approaches complex challenges with an analytical mindset, ensuring that data is acquired, processed, and utilized efficiently and effectively.
- Adaptability and Continuous Learning:** Stefan's recent foray into the world of machine learning showcases his adaptability and hunger for continuous learning. By integrating machine learning methodologies into our domain, he's pioneering a path that promises advancements and innovations.
- Contribution to the Community:** Stefan's talk on ChatGPT in physics is a testament to his commitment to knowledge-sharing. His ability to bridge the gap between advanced AI models like ChatGPT and particle physics is commendable, potentially ushering in a new era of research and discoveries.
- Collegiality and Teamwork:** Beyond his technical prowess, Stefan is a delight to collaborate with. His affable nature, team-oriented approach, and readiness to lend a hand have made him an integral part of our community.

In conclusion, Stefan's blend of hard skills in particle physics and computing, combined with his soft skills and forward-thinking approach, positions him as a notable asset in any team or project. I wholeheartedly recommend Stefan and am confident that he will continue to exceed expectations in the future.

Warm regards,

[Your Name]
[Your Position]

 Regen



Can you do that more informal, maybe a bit funny, for a good friend.



Absolutely! Let's make it a bit more light-hearted for your friend Stefan.

Hey [Recipient's Name],

So, you're considering my buddy Stefan for something cool? Let me break it down for you:

- Particle Physics Guru:** Ever had a question about quarks, leptons, or the mysteries of the universe? Just ask Stefan. If particle physics were Hogwarts, he'd be the Dumbledore. Minus the long beard. (Or does he have one now?)
- Data & Computing Whizz:** Honestly, I've seen Stefan crunch numbers and handle data like a maestro leading an orchestra. If data was pizza, Stefan would be the world's best pizzaiolo. But don't ask him to actually make a pizza. Stick to the data.
- Machine Learning Newbie (in a good way!):** Our man's gone all 21st century, diving deep into machine learning. Beware: If he starts talking about "neural networks", he's not discussing a new social media app.
- The ChatGPT Whisperer:** Oh, and did I mention he gave a talk on ChatGPT in physics? Basically, he's translating "robot talk" to "physics nerd". Pretty cool, right?
- Grade-A Good Guy:** Beyond all the brainy stuff, Stefan's just an all-around great guy to hang with. If you need a joke, a hand, or an overly-complicated explanation of the universe over a pint, he's your man.

In short, if you're thinking about teaming up with Stefan for literally anything, it's a no-brainer. The guy's a legend.

Cheers and good luck!

[Your Name]



Write some text



And very formal for an international selection committee



Of course. Here's a formal recommendation letter tailored for an international selection committee:

[Your Full Name]
 [Your Title/Position]
 [Your Institution/Organization]
 [Your Email Address]
 [Date]

To the Esteemed Members of the International Selection Committee,

RE: Recommendation for Stefan [Stefan's Last Name, if known]

I am writing to provide my strongest endorsement for Stefan, whom I have had the pleasure of collaborating with in our shared field of particle physics. It is not often that one encounters an individual with such a profound command over the intricate nuances of this discipline, combined with an impeccable aptitude for data acquisition and computational methodologies.

1. **Exceptional Particle Physics Acumen:** Throughout our interactions, Stefan has displayed a profound and unparalleled expertise in particle physics. His contributions to the field are indicative of a seasoned practitioner who engages with both theoretical profundity and practical sagacity.
2. **Data Acquisition and Computational Expertise:** Stefan's capabilities in data acquisition and computing are nothing short of exemplary. His proficiency ensures that data is harnessed, processed, and represented with utmost precision and efficiency, a testament to his dedication and skill set.

- Write some recommendation for a colleague
 - Write in different style and tones

3. **Innovative Machine Learning Application:** What truly distinguishes Stefan is his visionary integration of machine learning techniques into our domain. His pursuit of blending traditional particle physics with cutting-edge technology exhibits not only his adaptability but also his foresight in anticipating the future trajectory of our discipline.
4. **Contribution to Scholarly Discourse:** Stefan's presentation on the applications of ChatGPT in particle physics is indicative of his commitment to advancing the field's academic discourse. Such endeavors reflect his proactive approach to integrating emerging technologies in scientific exploration.
5. **Demeanor and Professionalism:** On a personal note, Stefan's interactions with peers and superiors alike are characterized by a blend of humility, respect, and genuine passion for our field. His comportment is always in alignment with the highest standards of professional decorum.

Given the aforementioned, I am confident that Stefan is exceptionally well-suited for any distinguished opportunity or recognition that this esteemed committee may be deliberating. His professional prowess, coupled with his pioneering spirit, makes him a noteworthy candidate, and I firmly believe he will contribute immensely to the global particle physics community.

Please do not hesitate to contact me should you require any further information or clarification regarding Stefan's qualifications and accomplishments.

With the utmost respect and sincerity,

[Your Full Name]
 [Your Title/Position]
 [Your Institution/Organization]



Brainstorming



We are trying to improve communication in our laboratory for particle physics. Any ideas?



Absolutely, improving communication within a particle physics laboratory is crucial for collaboration, efficiency, and overall success. Here are some ideas to consider:

1. **Regular Team Meetings:** Hold regular team meetings where members can discuss progress, challenges, and upcoming tasks. This fosters open communication and keeps everyone in the loop.
2. **Cross-Functional Teams:** Form cross-functional teams comprising researchers from different areas. This encourages diverse perspectives and knowledge sharing.
3. **Collaboration Tools:** Utilize digital platforms like Slack, Microsoft Teams, or dedicated project management tools. These platforms facilitate real-time discussions, file sharing, and updates.
4. **Shared Calendars:** Maintain a shared calendar for lab activities, seminars, workshops, and deadlines. This ensures everyone is aware of important events and commitments.
5. **Clear Documentation:** Encourage the practice of documenting processes, experiments, and findings. This helps new members catch up quickly and prevents knowledge loss.
6. **Internal Wiki:** Create an internal wiki or knowledge base to store important information, protocols, and best practices. It serves as a reference for everyone in the lab.
7. **Regular Updates:** Institute a practice of providing regular status updates on ongoing projects. This can be done through email, presentations, or brief stand-up meetings.
8. **Informal Discussions:** Create spaces for casual interactions, like a virtual "water cooler" chat, where colleagues can discuss ideas, share insights, or even just have a friendly chat.
9. **Open-Door Policy:** Promote an open-door policy where team members feel comfortable



Can you elaborate a bit more on item 2



Certainly! Creating cross-functional teams is a powerful strategy to improve communication and collaboration within a particle physics laboratory. Here's a more detailed explanation of the concept and its benefits:

Cross-Functional Teams: What are They?

Cross-functional teams are groups of individuals with diverse expertise, skills, and backgrounds, brought together to work on a specific project or goal. In the context of a particle physics laboratory, this means assembling teams that consist of researchers from different areas of specialization, such as experimentalists, theorists, data analysts, and computing experts.

Benefits of Cross-Functional Teams:

1. **Diverse Perspectives:** Each member of the team brings a unique perspective based on their expertise. This diversity of viewpoints can lead to innovative solutions, as problems are tackled from various angles.
2. **Knowledge Sharing:** Cross-functional teams naturally encourage knowledge sharing. Members can learn from each other's experiences and gain insights into different aspects of particle physics.
3. **Faster Problem Solving:** Complex challenges often require multidisciplinary solutions. With experts from various fields on the same team, problems can be addressed more comprehensively and resolved more quickly.



Ask some questions, but double check



What is the mass ratio between a proton and an electron



The mass ratio between a proton and an electron is approximately 1836.152. This means that a proton is about 1836 times more massive than an electron. Protons have a mass of approximately 1.67×10^{-27} kilograms, while electrons have a much smaller mass of about 9.11×10^{-31} kilograms. The significant difference in mass between these two particles is a fundamental property of particle physics and plays a crucial role in the behavior of atoms and the structure of matter.



What is the kinetic energy of muon in MeV/c^2 with an momentum of 28 MeV/c



To find the kinetic energy (T) of the muon, we'll first need to determine its total energy (E) and then subtract the rest energy (E_0) of the muon.

Given:

- $p = 28 \text{ MeV}/c$ is the momentum of the muon.
- m_0 for a muon is approximately $105.7 \text{ MeV}/c^2$.
- c is the speed of light.

From the relativistic energy-momentum relation:

$$E^2 = p^2 c^2 + (m_0 c^2)^2$$

Plugging in the given values:

$$E^2 = (28 \text{ MeV}/c)^2 c^2 + (105.7 \text{ MeV}/c^2)^2 c^4$$

$$E^2 = 784 \text{ MeV}^2 + 11169.49 \text{ MeV}^2$$

$$E^2 = 11953.49 \text{ MeV}^2$$

$$E \approx 109.33 \text{ MeV}/c^2$$

Now, the kinetic energy T is:

$$T = E - m_0 c^2$$

$$T = 109.33 \text{ MeV}/c^2 - 105.7 \text{ MeV}/c^2$$

$$T \approx 3.63 \text{ MeV}/c^2$$

Thus, the kinetic energy of a muon with a momentum of 28 MeV/c is approximately 3.63 MeV/c^2 .

Summarize text with PDF plugin

Can you please summarize <https://www.psi.ch/en/media/71845/download>

Note: Requires paid version of ChatGPT

Can you please summarize the document: <https://www.psi.ch/en/media/71845/download>



Used AI PDF

Here's a summarized overview of the provided document:

IMPACT Conceptual Design Report

- PSI Bericht Nr. 22-01
- Date: January 2022
- ISSN: 1019-0643
- Cover page description: The document showcases the IMPACT in: the HIMB production target, muon beamlines, TATTOOS proton beam mass separation, and shielded cells for radioisotope handling. The at the high-power HIPA accelerator. [Page 2]

Acknowledgements

- The authors express gratitude to numerous experts, PSI divisions, UZH administration, and the HIMB External Advisory Board. They also acknowledge support from the Swiss Academy of Sciences SCNAT and other institutions. [Page 3]

Contents Overview

1. IMPACT Executive Summary
 - Introduction
 - HIMB Science: Particle Physics, Condensed Matter Physics, Materials Science, Required Detectors and Target Cells
 - TATTOOS Science: Applied Radionuclide Sciences, Translational Science, Prospects for Astrophysics, Materials Science, and Radiochemistry
 - Infrastructure, Users and Access, Planning and Timeline, Organization, Governance, Management, Finances and Cost Estimate, Regulations, Radiation Protection, Safety, and Risk Assessment Strategy [Pages 5-6]
2. HIMB (High-Intensity Muon Beams)
 - Introduction & Overview, Design Strategy, Parameter Choice, Detector & High-Pressure



What will the HIMB project cost?



Used AI PDF

The document mentions the following financial details related to the HIMB project:

- The funding for HIMB will be applied for through the 2023 Swiss Roadmap for Research Infrastructure, with the funding period spanning from 2025 to 2028. [Page 40-41]
- The capital investment costs for IMPACT (which includes both HIMB and TATTOOS) during the ERI funding period 2025–2028 amounts to 76.8 MCHF. [Page 27]

However, the exact cost specific to the HIMB project alone isn't explicitly mentioned in the provided excerpts. If you'd like, I can continue searching or explore other sections of the document to find more specific financial details related to HIMB.



Can you summarize the HIMB project for a 15-year old



Of course! Here's a simplified summary of the HIMB project for a 15-year-old:

HIMB Project (High-Intensity Muon Beams)

Imagine you have a super-powerful flashlight, but instead of light, it shines tiny particles called muons. These muons can help scientists learn about the tiniest building blocks of our universe and the materials around us.

What's the Big Idea?

The HIMB project wants to make this "muon flashlight" even more powerful. By doing so, they can do cooler experiments and discover new things.

How Will They Do It?

Right now, there's a place where they already have a muon flashlight, but it's not as strong as they want. So, they're going to upgrade it. They'll take apart some old parts and build new ones to make the flashlight shine even brighter muons.

Why is This Cool?

With a stronger muon flashlight, scientists can:

1. Look deeper into the mysteries of the universe.
2. Study materials in new ways, like understanding why some things are magnetic.
3. Maybe even discover completely new things we haven't thought of yet!



Write web page



Can you write a basic HTML page with a nice menu and some sample text I can change later?



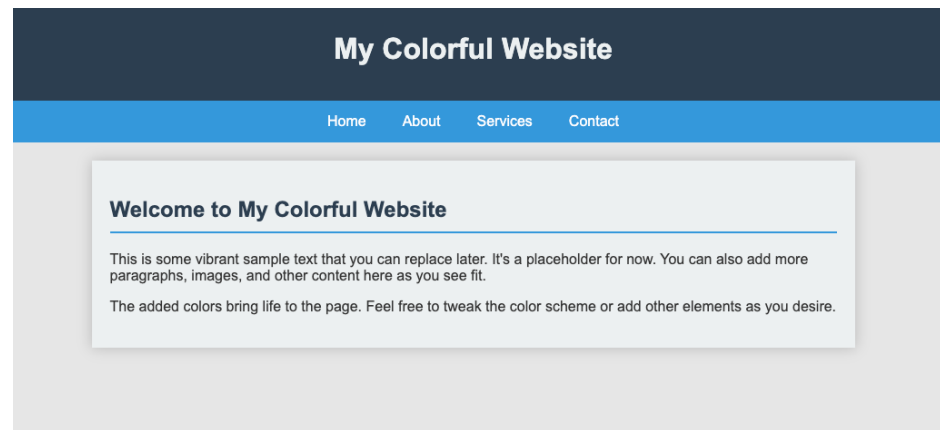
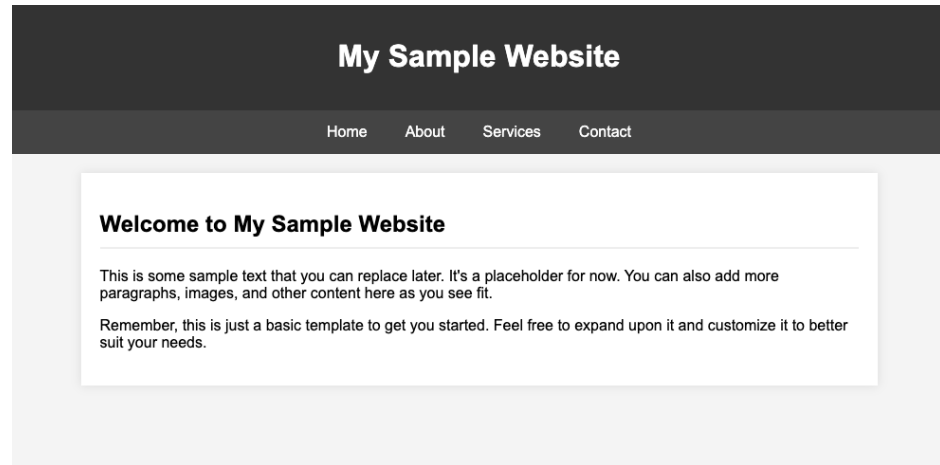
Certainly! Here's a simple HTML page with a basic navigation menu and some placeholder content. The design will use some basic inline CSS to make it look decent. Feel free to modify the CSS and HTML content as per your requirements.

```
html
Copy code

<!DOCTYPE html>
<html lang="en">

<head>
  <meta charset="UTF-8">
  <meta name="viewport" content="width=device-width, in
  <title>Sample Page with Menu</title>
  <style>
    body {
      font-family: Arial, sans-serif;
      margin: 0;
      padding: 0;
      background-color: #f4f4f4;
    }

    header {
      background-color: #333;
      color: #fff;
      text-align: center;
```



Word Clouds

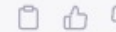
IMPACT CDR



Can you create a word cloud from that PDF: <https://www.psi.ch/en/media/71845/download>



Used AI PDF



Used WordCloud

Here's the word cloud generated from the provided PDF:



Write code

Can you write me a ROOT program which creates an array of random numbers and the plots a histogram from it.

Certainly! ROOT is a data analysis framework developed by CERN. Below is a simple ROOT program that creates an array of random numbers and then plots a histogram from it:

```

cpp
// You need to include the necessary ROOT headers.
#include "TH1F.h"
#include "TCanvas.h"
#include "TRandom3.h"

void random_hist() {
    const int nEntries = 10000; // Number of random numbers
    
```

To run this program, you will need to set up the ROOT environment. Once that's done:

1. Save the code to a file, say `random_hist.cpp`.
2. Open a terminal and navigate to the directory containing your `.cpp` file.
3. Compile the code using:

```

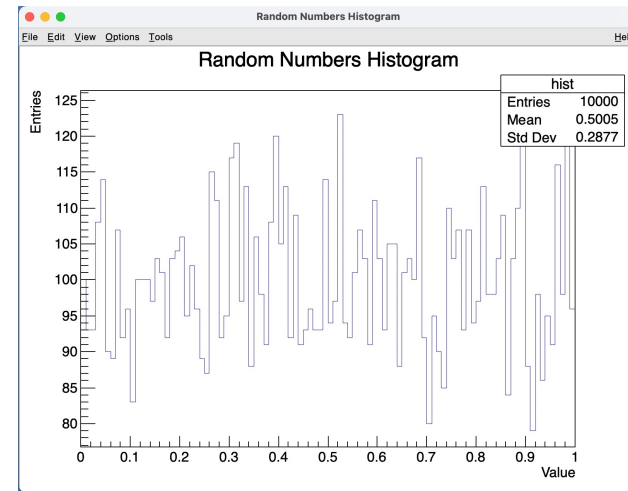
css
g++ `root-config --cflags --libs` -o random_hist random_hist.cpp
    
```

4. Run the resulting executable:

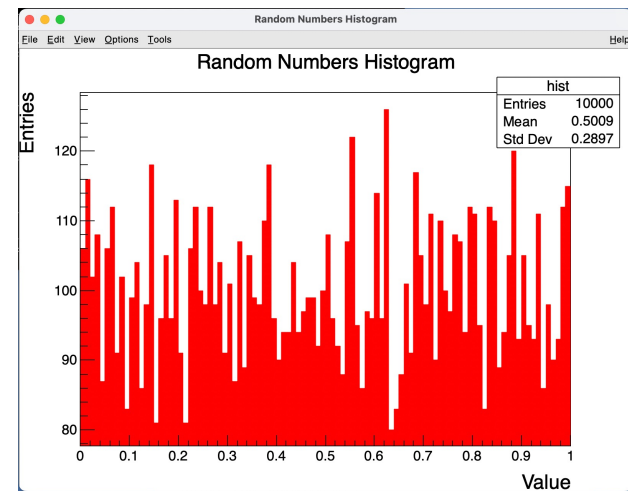
```

bash
./random_hist
    
```

This should generate an image called `random_hist.png` showing a histogram of random numbers between 0 and 1. Adjust parameters as necessary for your requirements.



Can you make the axis titles a bit bigger and change the graph color to red?



The image shows a Visual Studio Code editor window with a C++ file named `midas.cxx` open. A context menu is visible over the function `cm_execute`, with the `AI Actions` option selected. The AI Assistant chat window on the right provides a detailed explanation of the function, including its purpose, parameters, and a step-by-step breakdown of its execution logic.

AI Assistant

This is a C++ function named `cm_execute`. It is intended to execute a given system command (passed as a char string pointer `command`) and populate a provided buffer `result` (also a char string pointer) with the command's output, up to a provided buffer size `bufsize` (typed as `INT`). `INT` is a typedef for integer, specified in C++.

Here's a step-by-step breakdown of the function:

1. Declare and initialize variables:
 - `str` is a char array of 256 elements used to construct the system command to execute.
 - `n` (`INT`) and `fh` (`int`) are used for handling file operations - reading the result of the command's execution.
 - `status` is an `int` used to store the return value of various function calls.
 - `check_cm_execute` and `enable_cm_execute` are static `int` variables. Static means their values will be kept across multiple invocations of the function.
2. The function first checks if the current process is remote using `rpc_is_remote()`. If so, it calls `rpc_call(RPC_CM_EXECUTE, command, result, bufsize)` to execute the command remotely (RPC means Remote Procedure Call) and returns.
3. If not, `check_cm_execute` variable is used to control whether the permission to execute the command should be checked from a database. Specifically, the program checks a value at the `"/Experiment/Enable cm_execute"` database key. If enabled, the `enable_cm_execute` will be non-zero and the function can proceed to execute the command. Otherwise, an error message is printed out and the function returns `CM_WRONG_PASSWORD`.
4. Once allowed to execute, the function checks if the buffer size is greater than 0. If so, it executes the command and stores the output in a temporary file named `".tmp"` wherein is the process id returned by `ss_getpid()`. If the temporary file can be opened successfully, contents are read into the `result` buffer, and the temporary file is removed.
5. If the buffer size is not greater than 0, it directly executes the command using the `system()` command.
6. Finally, there is a check to verify if the command executed successfully. If `status` is less than 0 (indicating an error), it prints out an error message with the command and return `CM_SET_ERROR`. If no error occurred, it returns `CM_SUCCESS`.

Code Snippet:

```

5716 INT cm_execute(const char *command, char *result, INT bufsize) {
5717     char str[256];
5718     INT n;
5719     int fh;
5720     int status = 0;
5721     static int check_cm_execute = 1;
5722     static int enable_cm_execute = 0;
5723
5724     if (rpc_is_remote())
5725         return rpc_call(RPC_CM_EXECUTE, command, result, bufsize);
5726
5727     if (check_cm_execute) {
5728         int status;
5729         int size;
5730         HANDLE hDB;
5731         check_cm_execute = 0;
5732
5733         status = cm_get_experiment_database(&hDB, NULL);
5734         assert(status == DB_SUCCESS);
5735
5736         size = sizeof(enable_cm_execute);
5737         status = db_get_value(hDB, "Experiment/Enable cm_execute",
5738                             &enable_cm_execute, &size, TID_BOOL, TRUE);
5739         assert(status == DB_SUCCESS);
5740         //printf("enable_cm_execute %d\n", enable_cm_execute);

```

Clion AI Assistant
<https://www.jetbrains.com/clion/>

```
/*  
Executes command via system() call  
@param command Command string to execute  
@param result stdout of command  
@param bufsize string size in byte  
@return CM_SUCCESS  
*/  
INT cm_execute(const char *command, char *result, INT bufsize) {  
char str[256];  
INT n;  
int fh;  
int status = 0;  
static int check_cm_execute = 1;  
static int enable_cm_execute = 0;  
  
if (rpc_is_remote())  
return rpc_call(RPC_CM_EXECUTE, command, result, bufsize);  
  
if (check_cm_execute) {  
int status;  
int size;  
HANDLE hDB;  
  
//printf("enable_cm_execute %d\n", enable_cm_execute);  
}  
  
if (!enable_cm_execute) {  
return CM_FAILURE;  
}  
  
if (!status) {  
return CM_FAILURE;  
}  
  
return CM_SUCCESS;  
}
```

Show Context Actions

- All Actions >
- Paste
- Copy / Paste Special >
- Column Selection Mode
- Find Usages
- Go To >
- Folding >
- Refactor >
- Generate...
- Open In >
- Local History >
- Git >
- Compare with Clipboard
- Diagrams >
- Create Gist...

Write documentation

- Explain 'cmexecute' function
- Suggest refactoring
- Find problems in code
- New chat using selection
- Edit User Prompts

```
/*  
Executes a command and returns the result.  
*  
* This function takes a command as input and executes it using the system shell.  
* The result of the command execution is stored in the provided 'result' buffer.  
* The maximum number of characters that can be stored in the 'result' buffer is specified by 'bufsize'.  
*  
* @param command The command to be executed.  
* @param result The buffer to store the result of the command execution.  
* @param bufsize The maximum number of characters that can be stored in the 'result' buffer.  
*  
* @note The 'result' buffer must have enough capacity to accommodate the entire result string.  
*  
* @return void  
*  
* @warning The function assumes that the provided buffers have been properly allocated.  
* Failure to allocate or provide correct buffer sizes could result in buffer overflow.  
* It is the responsibility of the caller to ensure proper memory allocation and buffer size.  
*/  
INT cm_execute(const char *command, char *result, INT bufsize) {  
char str[256];  
INT n;  
int fh;  
int status = 0;  
static int check_cm_execute = 1;  
static int enable_cm_execute = 0;  
  
if (rpc_is_remote())  
return rpc_call(RPC_CM_EXECUTE, command, result, bufsize);  
  
if (check_cm_execute) {  
int status;  
int size;  
HANDLE hDB;  
  
//printf("enable_cm_execute %d\n", enable_cm_execute);  
}  
  
if (!enable_cm_execute) {  
return CM_FAILURE;  
}  
  
if (!status) {  
return CM_FAILURE;  
}  
  
return CM_SUCCESS;  
}
```


The screenshot shows a Visual Studio Code editor with a C++ file named `prime.cxx`. The code defines a `sumNumbers` function and a `main` function. A red bracket highlights a bug in the `sumNumbers` function's for loop: `for (int i = 1; i <= n; ++n) {`. The AI Assistant chat on the right identifies this as a "crucial mistake" where the loop control variable `i` is not updated, leading to an infinite loop. It provides the corrected code where `++i` is used instead of `++n`.

```

1 #include <iostream>
2 #include <vector>
3
4 #include <iostream>
5
6 int sumNumbers(int n) {
7     int sum = 0;
8     for (int i = 1; i <= n; ++n) {
9         sum += i;
10    }
11    return sum;
12 }
13
14 int main() {
15     int n;
16     std::cout << "Enter the value of n: ";
17     std::cin >> n;
18
19     int result = sumNumbers(n);
20     std::cout << "The sum of numbers from 1 to " << n << " is: " << result
21         << std::endl;
22
23     return 0;
24 }
25

```

AI Assistant

Me
Find potential problems in the code:

```

C/C++
{
    int sum = 0;
    for (int i = 1; i <= n; ++n) {
        sum += i;
    }
    return sum;
}

```

AI Assistant
You have a small but crucial mistake in the loop control variable update. You are incrementing `n` instead of `i`, therefore, the loop never ends (it keeps checking if `i` is less than or equal to `n`, but `i` isn't actually changing in the loop control action). Here is the corrected code:

```

C/C++
int sumNumbers(int n) {
    int sum = 0;
    for (int i = 1; i <= n; ++i) { // increment
        `i` instead of `n`
        sum += i;
    }
    return sum;
}

```

With this fix, the function will correctly compute the sum of all the numbers from 1 to `n`.

2 May 2023



LLM cannot calculate

ChatGPT 3.5



Is 17077 prime?



No, 17077 is not a prime number. It can be divided evenly by several numbers other than 1 and itself. In particular, it can be divided by 7 and 2441, among other divisors. Therefore, 17077 is not a prime number.

ChatGPT 4



Is 17077 prime?



To determine if 17077 is prime, one has to check if it is divisible by any integer other than 1 and itself.

However, we can use a shortcut to reduce the number of checks we make. We only need to test divisors up to the square root of 17077, because if 17077 has a divisor larger than its square root, then it must also have a divisor smaller than its square root.

The square root of 17077 is a bit above 130. Therefore, we need to check if 17077 is divisible by any prime number less than 130.

Let's start with the first few prime numbers:

2: 17077 is odd, so it's not divisible by 2.

3: $1 + 7 + 0 + 7 + 7 = 22$, and 22 is divisible by 3. Therefore, 17077 is divisible by 3.

This shows that 17077 is not prime, since it's divisible by 3.

$$17077 / 3 = 5692.333\dots$$

Code plug-in *can* calculate



Is 17077 prime?

Note: Requires paid version of ChatGPT

Finished working Hide work ^

```
python Copy code

def is_prime(n):
    """Check if a number is prime."""
    if n <= 1:
        return False
    if n <= 3:
        return True
    if n % 2 == 0 or n % 3 == 0:
        return False
    i = 5
    while i * i <= n:
        if n % i == 0 or n % (i + 2) == 0:
            return False
        i += 6
    return True

# Check if 17077 is prime
is_prime_17077 = is_prime(17077)
is_prime_17077

RESULT
True

The number 17077 is prime.
```

ChatGPT 4 +
Code Interpreter

Example of data analysis

	A	B	C	D
1				
2				
3				
4				
5				
6				
7				
8				
9				
10				
11				
12				
13				
14				
15				
16				
17				
18				
19				
20				
21				
22				
23				
24				
25				
26				
27				
28				
29				
30				
31				
32				
33				



This is some data from a device cooling down to a final temperature. Can you plot the data and extrapolate it until it reaches the final temperature within 0.3 degrees?

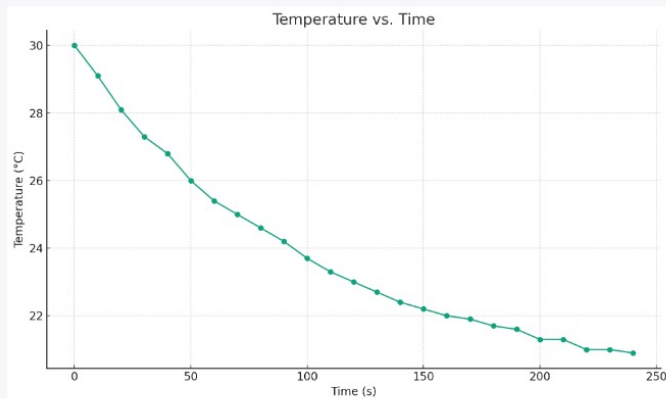
Certainly. Let's begin by importing the data from the Excel file and plotting it to get an initial visualization.

Finished working Show work

It seems that the actual data starts a few rows down in the file. The header information is contained within the first few rows. We'll adjust our reading of the file to skip these rows and focus on the actual data.

Let's extract and plot the data.

Finished working Show work



The plot illustrates the temperature of the device as a function of time. We can observe an exponential decay, which is typical for cooling processes.

To extrapolate the data until the temperature approaches its final value within 0.3 degrees, we can fit an exponential decay function to the data. The general form of an exponential decay is:

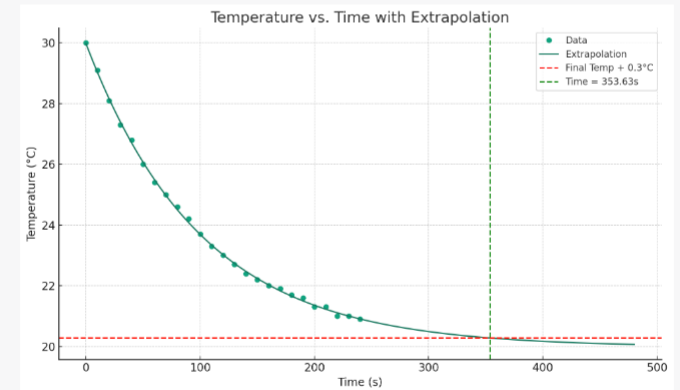
$$T(t) = T_0 e^{-kt} + T_{\text{final}}$$

Where:

- $T(t)$ is the temperature at time t .
- T_0 is the initial temperature difference from the final temperature.
- k is the decay constant.
- T_{final} is the final temperature.

Let's fit this function to our data and then use it to extrapolate the cooling process.

Finished working Show work



The temperature of the device is projected to approach its final value within 0.3°C at approximately $t = 353.63$ seconds.

The green dashed line on the graph represents this point in time. The red dashed line represents the temperature value that's 0.3°C above the extrapolated final temperature.

Conclusions

- LLM will be **all around us** in a few years
 - Call-centers, MS Office
- There are **opportunities** and **threats**
- My way to deal with this is to **learn** as much as possible about the **technology**, understand its strengths and weaknesses, and make the **best possible use** of it



Write a poem in four verses about particle physics, the beauty of the standard model and the universe in total



In the heart of the atom's embrace,
Quarks and leptons find their place.
The Standard Model, pure and true,
Charts the dance they all pursue.

Photons light the cosmic dance,
While gluons in strong force prance.
W, Z, with their unique charm,
Shield the universe from harm.

Beyond the known, mysteries lie,
Dark matter, dark energy, the endless sky.
Galaxies whirl, stars alight,
Bound by forces, oh so slight.

From minute particles to vast expanse,
The universe in a delicate balance stands.
Every discovery, a melodious note,
In the cosmic symphony we devote.

References

- <https://chat.openai.com>
- <https://www.youtube.com/watch?v=bSvTVREwSNw>
- <https://www.youtube.com/watch?v=iR2O2GPbB0E>
- <https://www.youtube.com/watch?v=lnA9DMvHtfl>
- <https://www.youtube.com/watch?v=kCc8FmEb1nY>
- <https://www.youtube.com/watch?v=g2BRIuln4uc>
- <https://www.youtube.com/watch?v=XowwKOAWYoQ>
- <https://www.youtube.com/watch?v=5MaWmXwxFNQ>
- <https://www.youtube.com/watch?v=KmAlSyVvE1Y>