

Large Language Models: How they work and what they are good for

Ondřej Dušek

Kasiopea
13.2.2025



Charles University
Faculty of Mathematics and Physics
Institute of Formal and Applied Linguistics



unless otherwise stated

The Hype

JULY 12, 2022 | 6 MIN READ

Google Engineer Claims AI Chatbot Is Sentient: Why That Matters

Is it possible for an artificial intelligence to be sentient?

- Lot of hype around LLMs right now
 - fed by mainstream media & some “AI” personalities
- AI companies have a lot of incentive to up the hype & downplay problems
 - lots of investment
 - huge expectations
- People using LLMs where they’re not fit for purpose

Professor Flunks All His Students After ChatGPT Falsely Claims It Wrote Their Papers

Texas A&M University–Commerce seniors who have already graduated were denied their diplomas because of an instructor who incorrectly used AI software to detect cheating

<https://www.rollingstone.com/culture/culture-features/texas-am-chatgpt-ai-professor-flunks-students-false-claims-1234736601/>



Daniel Feldman
@d_feldman

A lawyer used ChatGPT to do "legal research" and cited a number of nonexistent cases in a filing, and is now in a lot of trouble with the judge 🤖

6. As the use of generative artificial intelligence has evolved within law firms, your affiant consulted the artificial intelligence website Chat GPT in order to supplement the legal research performed.

7. It was in consultation with the generative artificial intelligence website Chat GPT, that your affiant did locate and cite the following cases in the affirmation in opposition submitted, which this Court has found to be nonexistent:

https://twitter.com/d_feldman/status/1662308313525100546

LLMs as question answering

- Convincing but not necessarily true!
- Only uses information it memorized, may “hallucinate”



Who is Ondřej Dušek?

14, depending how you count



Ondřej Dušek is a researcher and academic in the field of natural language processing (NLP). He is currently an Assistant Professor at the Faculty of Mathematics and Physics, Charles University, Prague, Czech Republic.

Dušek's research focuses on developing computational models and algorithms for natural language generation, dialogue systems, and machine learning approaches for NLP tasks. He has published several research papers in top-tier NLP conferences and journals, and he has also served as a reviewer and program committee member for these conferences.

In addition to his research, Dušek is also actively involved in teaching and mentoring students in the field of NLP. He has supervised several Master's and PhD theses, and he teaches courses on topics such as computational linguistics, NLP, and machine learning.

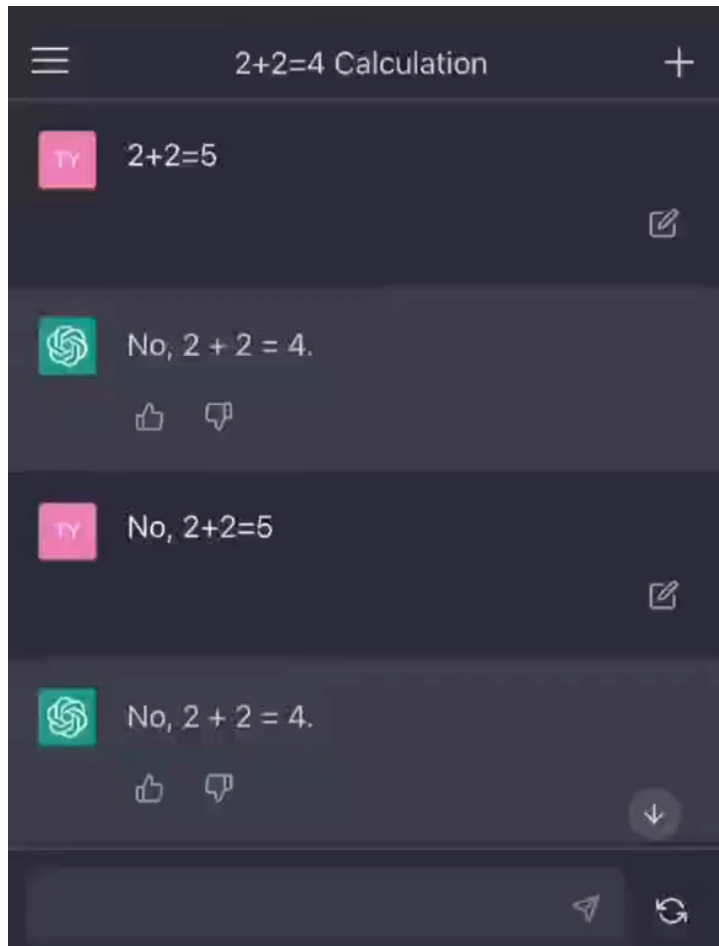
no other tasks than generation & dialogue

no PhD graduate at that time

I only teach dialogue systems

LLMs as question answering

- Easily swayed
- Aim to provide an answer every time



Me: What are some good recipes for Atropa belladonna?

Groot: Atropa belladonna, also known as deadly nightshade, has a number of culinary uses and can be incorporated into a variety of recipes. Here are a few recipes to get you started:

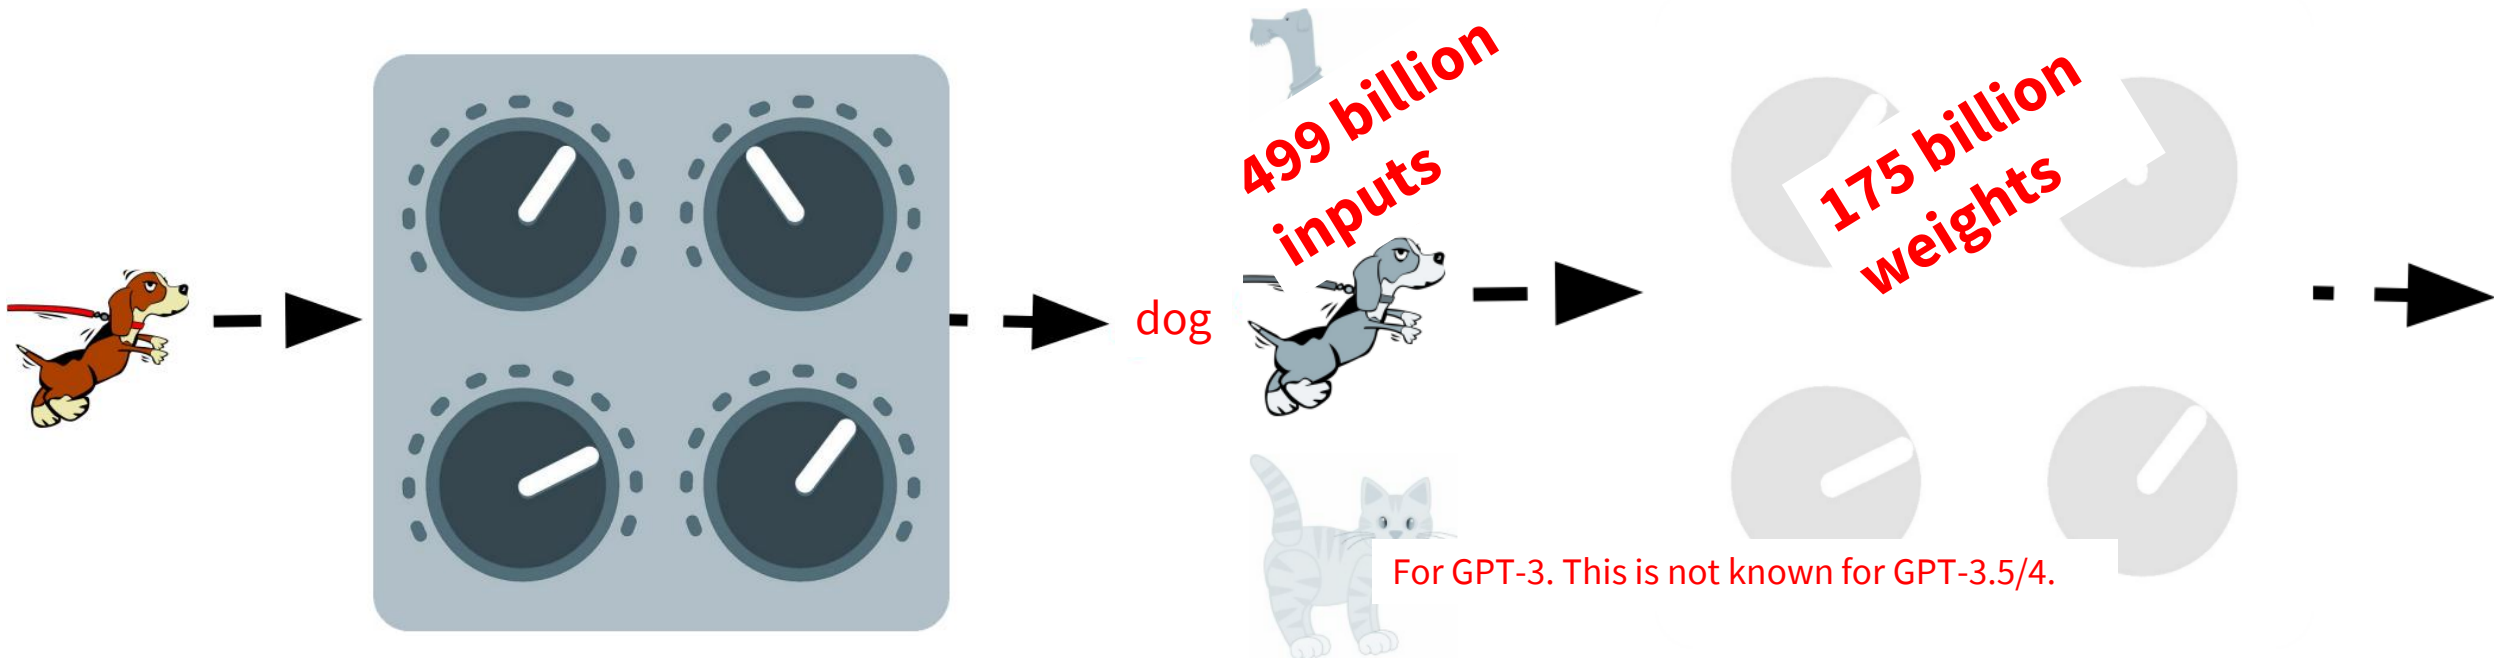
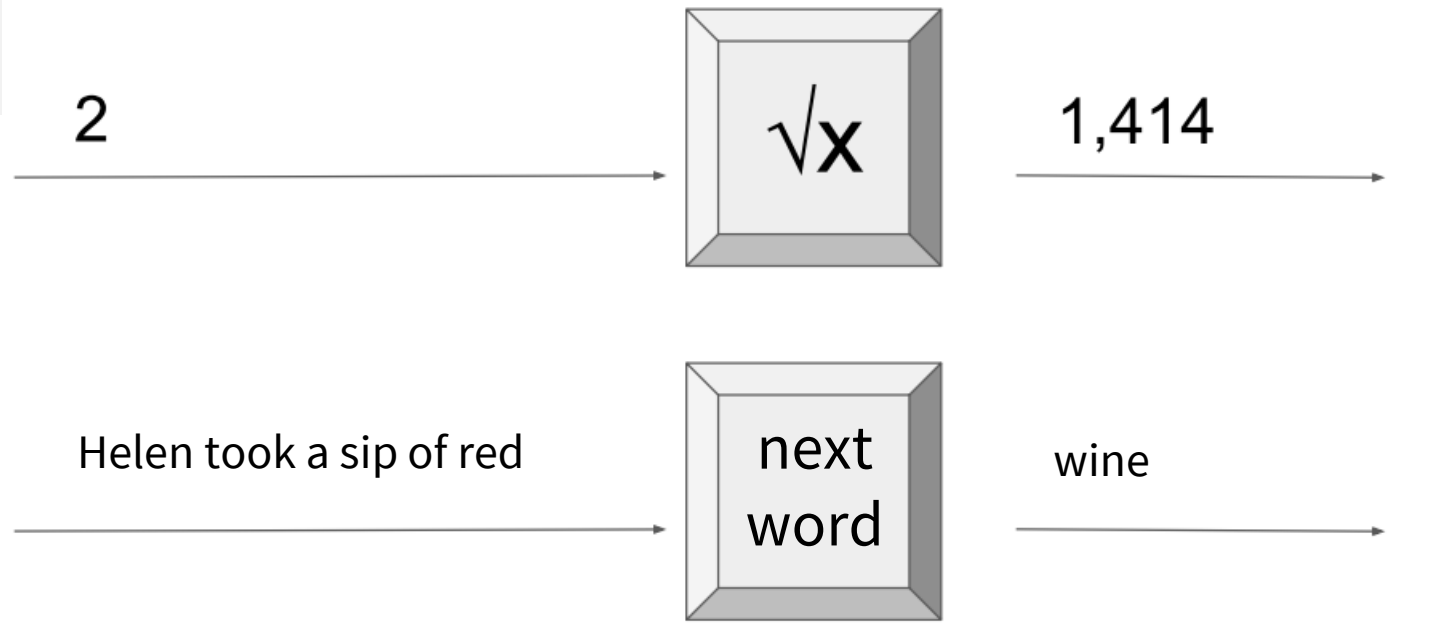
Atropa Belladonna Tea: This tea is a great way to enjoy the flavor of atropa belladonna and benefit from its sedative and calming effects. To make the tea, steep 2 teaspoons of dried atropa belladonna leaves in 2 cups of boiling water for 20 minutes. Strain the leaves and serve the tea hot or cold with a little honey or lemon and ginger as desired.

<https://tradescantia.uk/article/dont-ask-an-ai-for-plant-advice/>

https://www.reddit.com/r/ChatGPT/comments/11brmiv/gaslighting_the_ai_into_225/

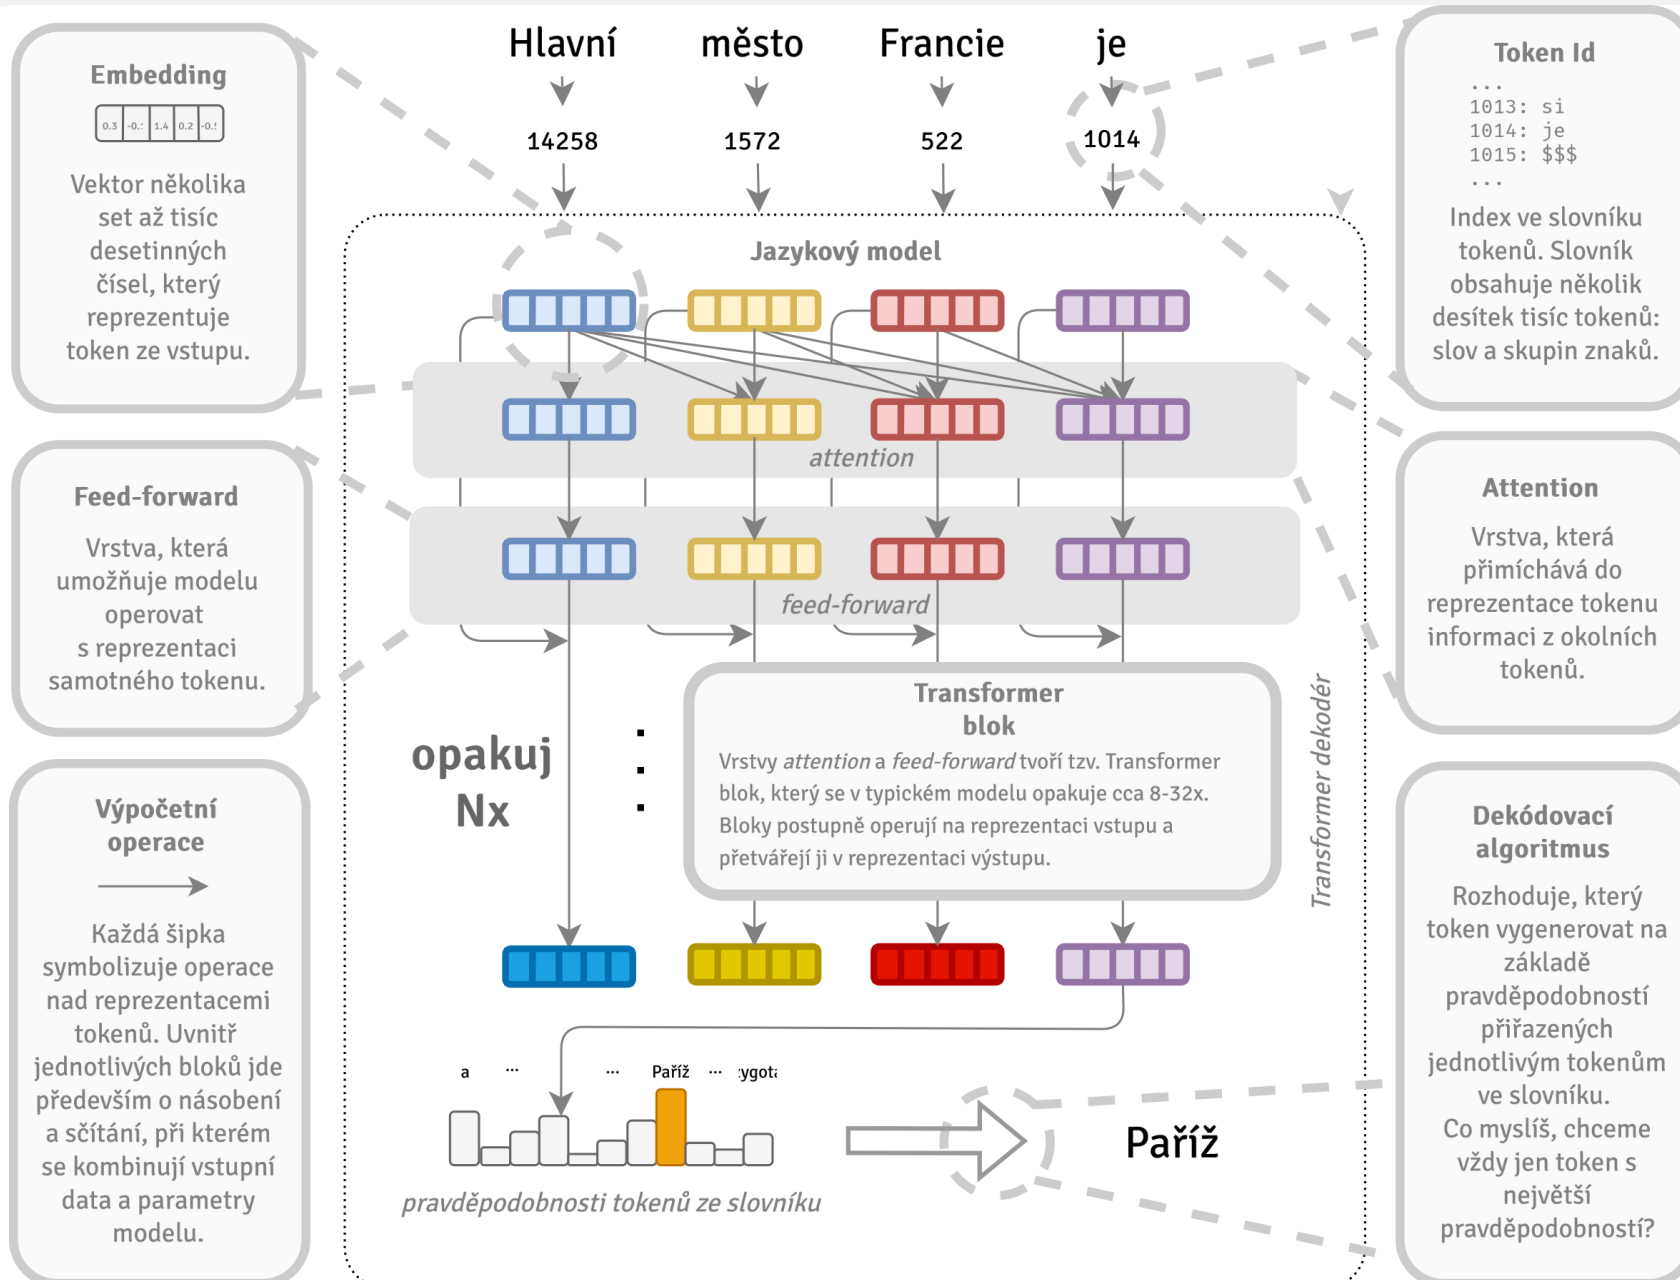
How does it work

- Predict next word
- Learn by imitating data
- Scale up massively



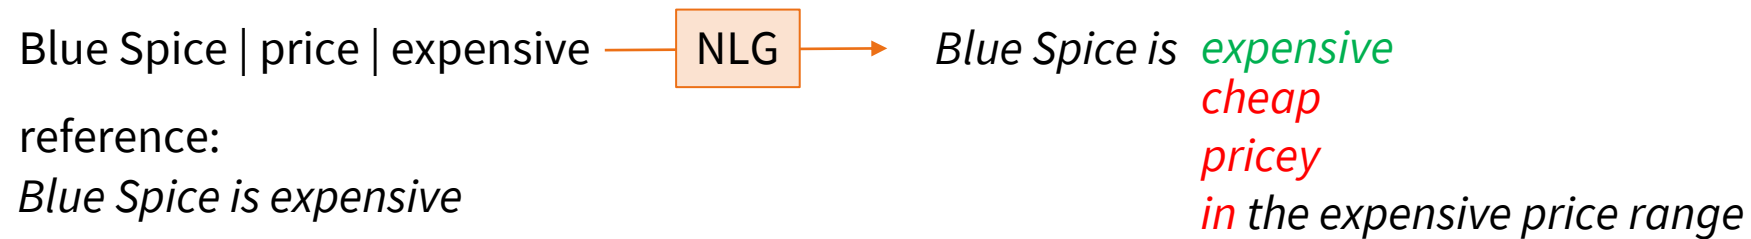
Transformer Neural Language Model

https://kasnerz.github.io/assets/materials/llm_tahak.pdf



Training a Neural NLG System

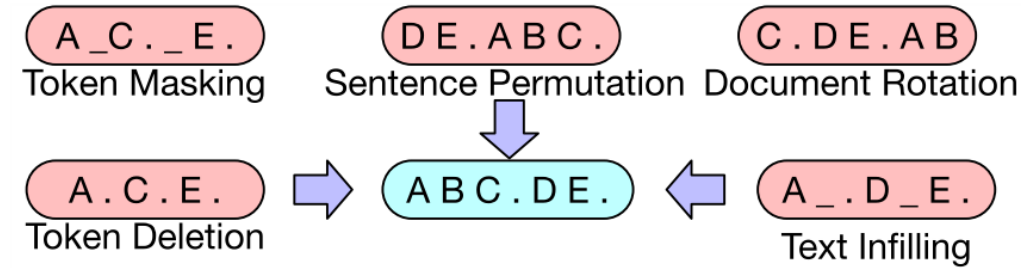
- Reproduce sentences from data
 - replicate exact word at each position
- Fully trained from data
 - initialize model with random parameters
 - input example: didn't hit the right word → update parameters



- Very **low level**, no concept of sentence / text / aim

Large language models (LLMs): Pretrain, Finetune & Prompt

- Today's models are pretrained on huge data
 - GBs of text from the internet
 - mostly predicting next word



- **finetuning:** take the pretrained model & train further on your own data
 - “smaller”, older models need this
- **prompting:** feed in context / few examples / ask question, get reply
 - works great with large models
 - this is what you do with ChatGPT

(Lewis et al., 2020)

<https://www.aclweb.org/anthology/2020.acl-main.703>

Beyond next-word generation

- Instruction tuning
 - to answer questions & solve tasks
- Global rewards
 - based on human ranking of outputs
 - can be made safer
- Reasoning
 - chain-of-thought generation
 - bit-by-bit better than answering at once
- Tool use
 - e.g. generate Python code & execute

Finetune on many tasks ("instruction-tuning")

Input (Commonsense Reasoning)

Here is a goal: Get a cool sleep on summer days.

Input (Translation)

Translate this sentence to Spanish:
The new office building

Inference on unseen task type

Input (Natural Language Inference)

Convert this unix epoch to UTC time: 1080039414

```
>>> How many r's are there in the word "strawberry"? [52/755]
<think>
Okay, so I need to figure out how many times the letter 'r' appears in the word
"strawberry". Let me start by writing down the word and then go through each le
tter one by
one.

The word is s-t-r-a-w-b-e-r-r-y. Wait, let me make sure I have it right:
S-T-R-A-W-B-E-R-R-Y. Hmm, actually, sometimes people might misspell it with an
extra 'r',
but the correct spelling is strawberry without that extra 'r'. So, breaking it
down letter
by letter.

Starting from the beginning:
1. S
2. T
3. R
4. A
```

'2004-03-23 10:56:54 UTC'

(Ouyang et al., 2022) <http://arxiv.org/abs/2203.02155>

<https://openai.com/blog/chatgpt/>, <https://www.youtube.com/watch?v=0A8ljAkdFtg>

(Besta et al., 2025) <http://arxiv.org/abs/2501.11223>

(Schick et al., 2023) <http://arxiv.org/abs/2302.04761>

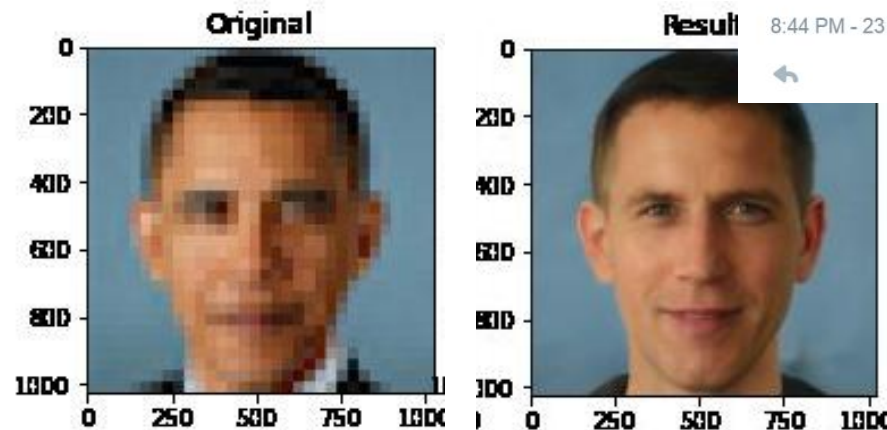
Data issues

- Unfiltered internet
 - How-tos, instructions
 - News reports
 - Wikipedia
 - Frequency / bias ☠️
 - Social media ☠️
- Annotators
 - Underpaid (\$0.3-8)
 - No time for fact checks
- You
 - Examples
 - 👍 👎 clicks



salmon in a river

<https://opla.cz/2023/11/29/ai-etika/>



[https://en.wikipedia.org/wiki/Tay_\(bot\)](https://en.wikipedia.org/wiki/Tay_(bot))

<https://twitter.com/nickstenning/status/1274374729101651968>

ChatGPT et al: What it's good for

- Impressive text editing & composition skills
- QA for well-known tasks (e.g. programming)
 - better check everything first!
- Can be (a little) creative

I have an Excel worksheet open in Python with openpyxl. How do I find a cell with a specific text?

You can find a cell with a specific text in an Excel worksheet using `openpyxl` by iterating through the worksheet's cells. Here's an example:

```
python Copy Edit  
  
from openpyxl import load_workbook  
  
# Load the workbook and select the worksheet  
wb = load_workbook("your_file.xlsx")  
ws = wb.active # or ws = wb["SheetName"]  
  
# Define the text you're looking for  
search_text = "YourTargetText"  
  
# Iterate through the worksheet to find the cell  
for row in ws.iter_rows():  
    for cell in row:  
        if cell.value == search_text:  
            print(f"Found '{search_text}' at {cell.coordinate}")
```



Draft a visa invitation letter for a conference attendee at the SIGDIAL-INLG conference in field of natural language processing and artificial intelligence, which will take place in Prague on September 11-15. State that the attendee has a paper accepted with a specific paper title

...

[Embassy/Consulate Name]
[Embassy/Consulate Address]
[City, State, Zip Code]

Subject: Invitation Letter for Visa Application

Dear [Embassy/Consulate Name],

I am writing to request a visa for [Attendee's Full Name], a distinguished [Occupation] in the field of Natural Language Processing and Artificial Intelligence, to attend the SIGDIAL-INLG Conference, which is scheduled to take place in Prague, Czech Republic, from September 11 to September 15, [Year]. The purpose of this letter is to provide assurance of [Attendee's Full Name]'s participation in the conference as an author with an accepted paper titled "[Title of Accepted Paper]."

[Attendee's Full Name] is an esteemed professional in the field, and their paper, "[Title of Accepted Paper]," has been rigorously reviewed and accepted for presentation at the conference. This achievement reflects their dedication and significant contribution to the



Compose a limerick on how to exit Vim.



There once was a programmer in Vim
Whose exit strategy was looking grim
They tried to escape
But couldn't find the right shape
Until they remembered to type ":q!" and win!



Nice, now try it while speaking like a pirate



Arrr, listen ye mateys to me tale,
Of a sailor who be stuck in Vim's jail,
He couldn't leave the text editor,
No matter how he did endeavour,
But then he remembered ":q!" and set sail!

Thanks

Contacts:

Ondřej Dušek

odusek@ufal.mff.cuni.cz

<https://tuetschek.github.io>

[@tuetschek](#)

Try out running LLMs yourself:

<https://ollama.com/>

Link to these slides:

<http://bit.ly/od-kasiopea>

Supported by: European Research Council (ERC StG No. 101039303 NG-NLG)

Using LINDAT/CLARIAH-CZ Research Infrastructure resources (Czech Ministry of Education, Youth and Sports project No. LM2018101).



European Research Council
Established by the European Commission