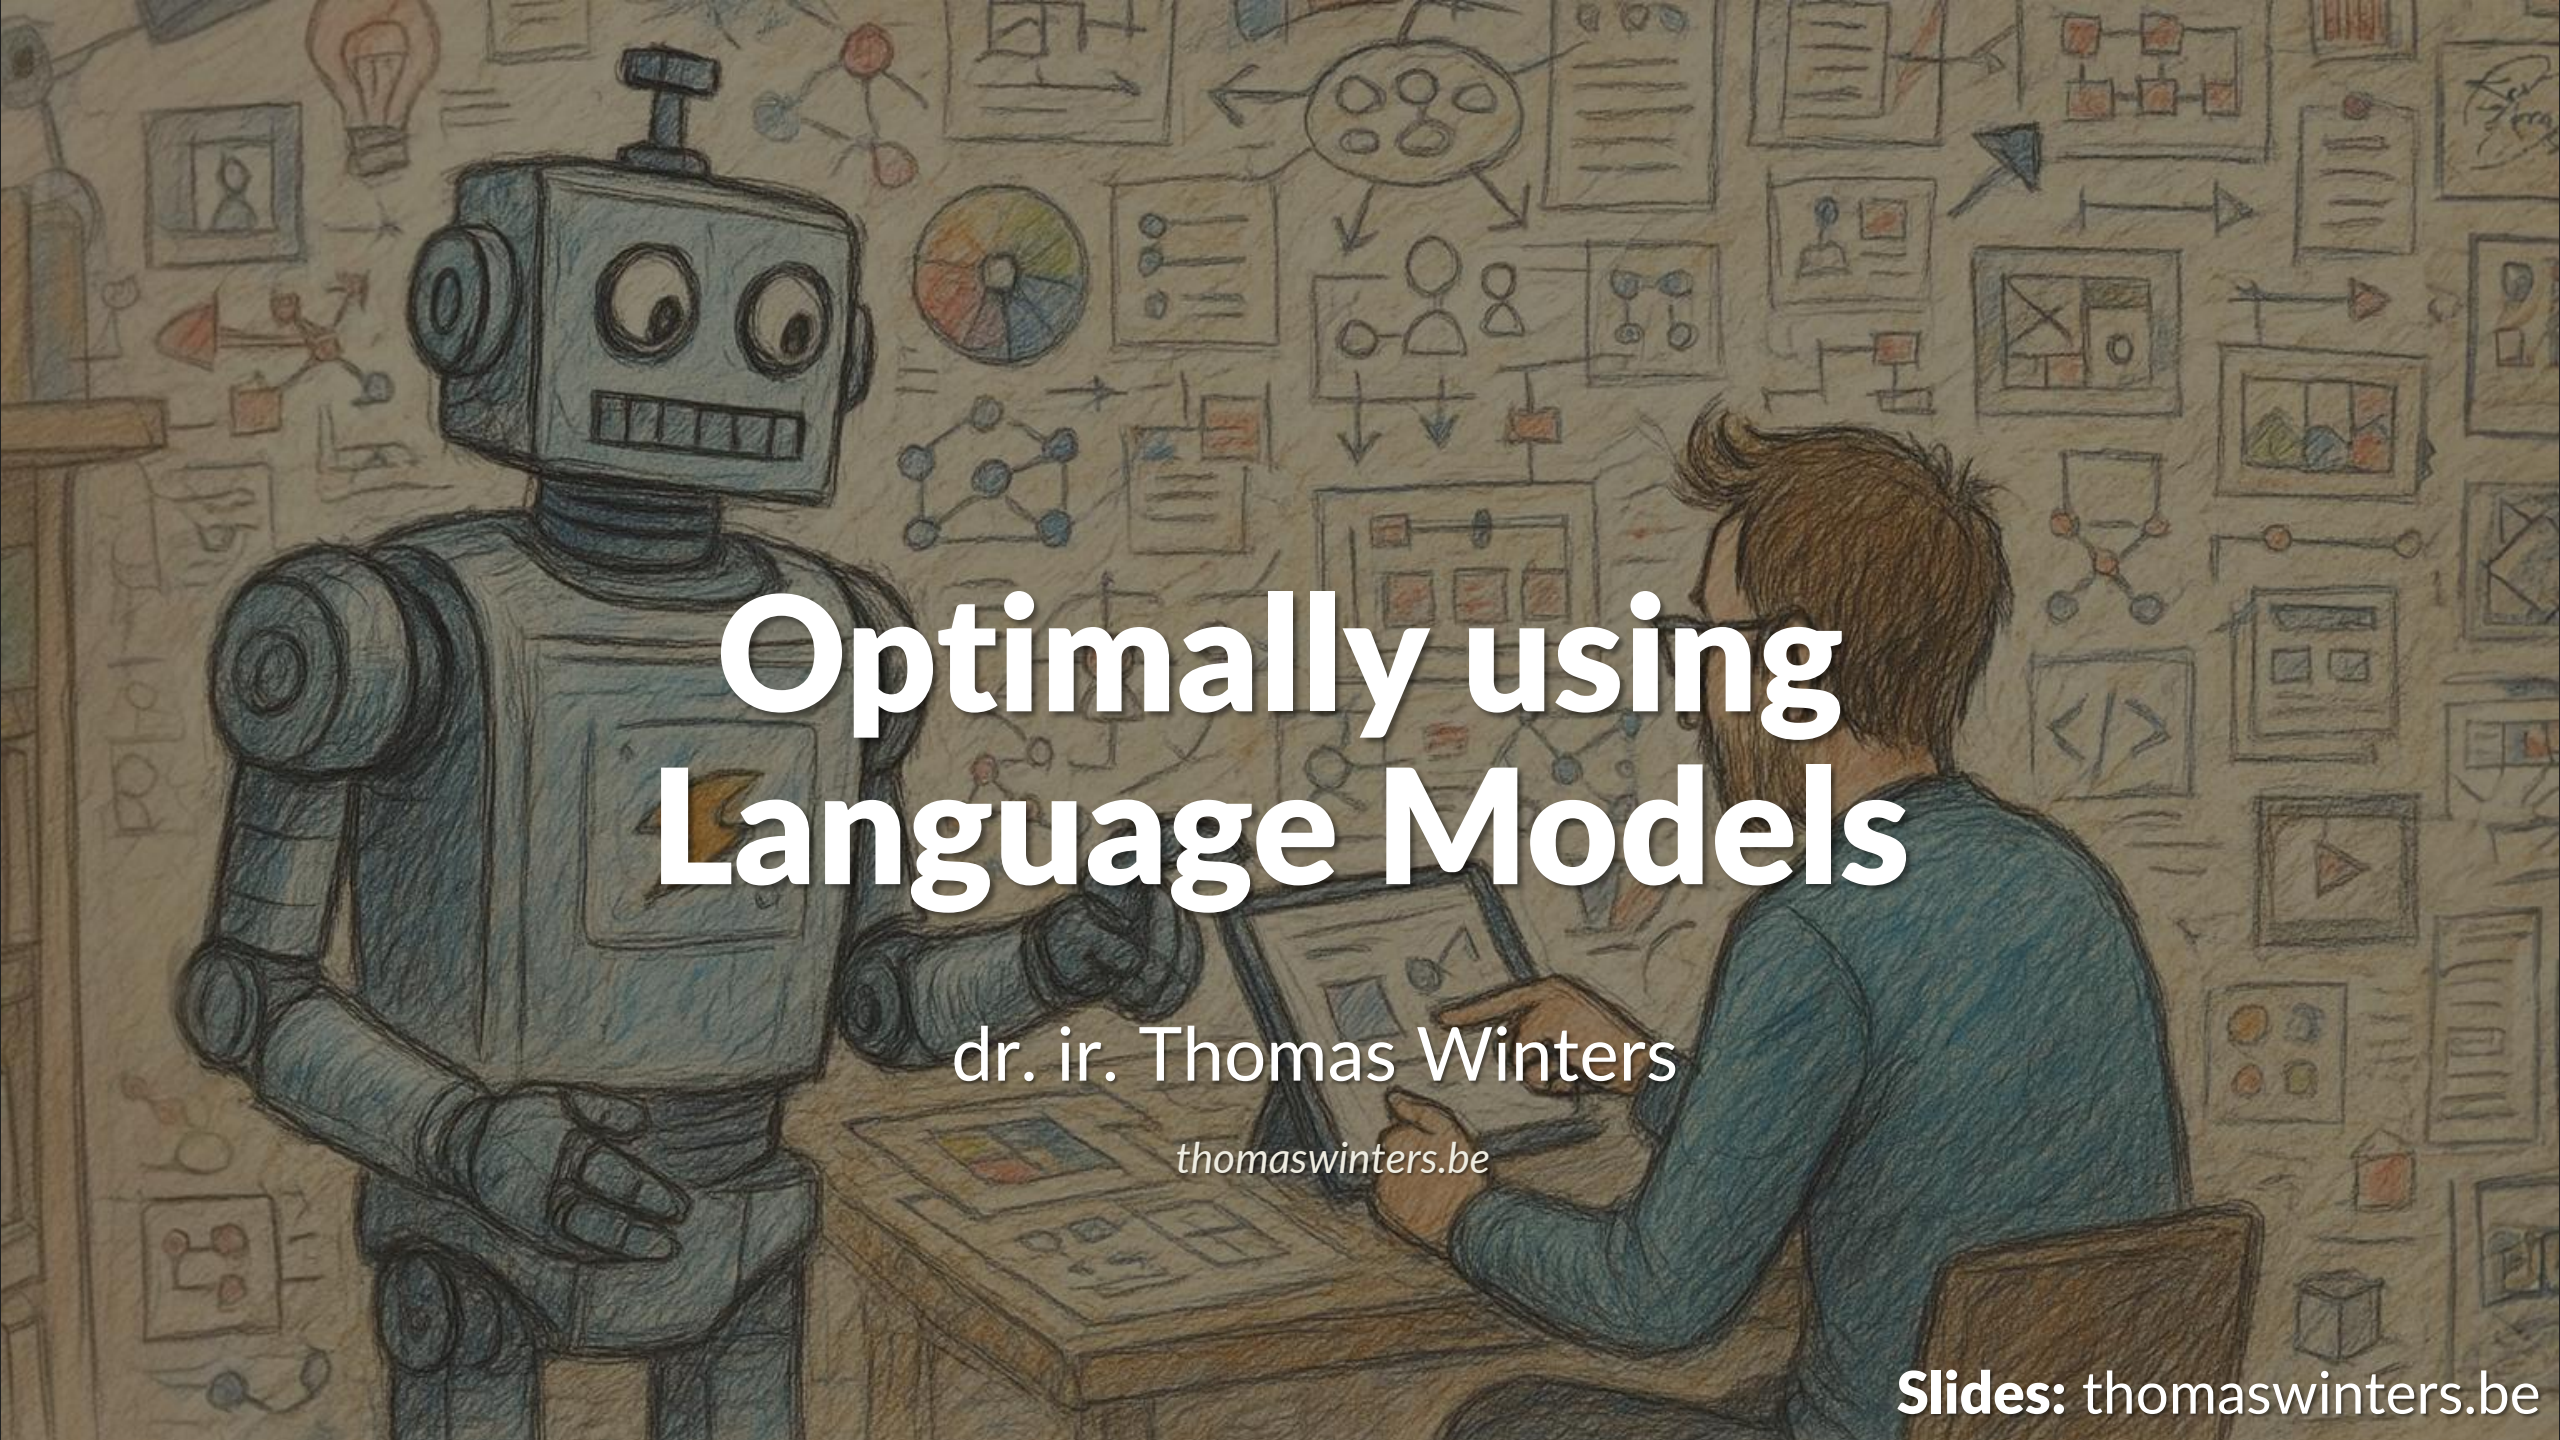


Optimally using Language Models

dr. ir. Thomas Winters

thomaswinters.be

Slides: thomaswinters.be

Who has used language model (GPT/Gemini/Copilot…) ?



Who has used language model to solve an actual need (work, personal...)?



Who found task that language model can’t do?



Thomas Winters

KU Leuven researcher with PhD
in creative AI & large language models

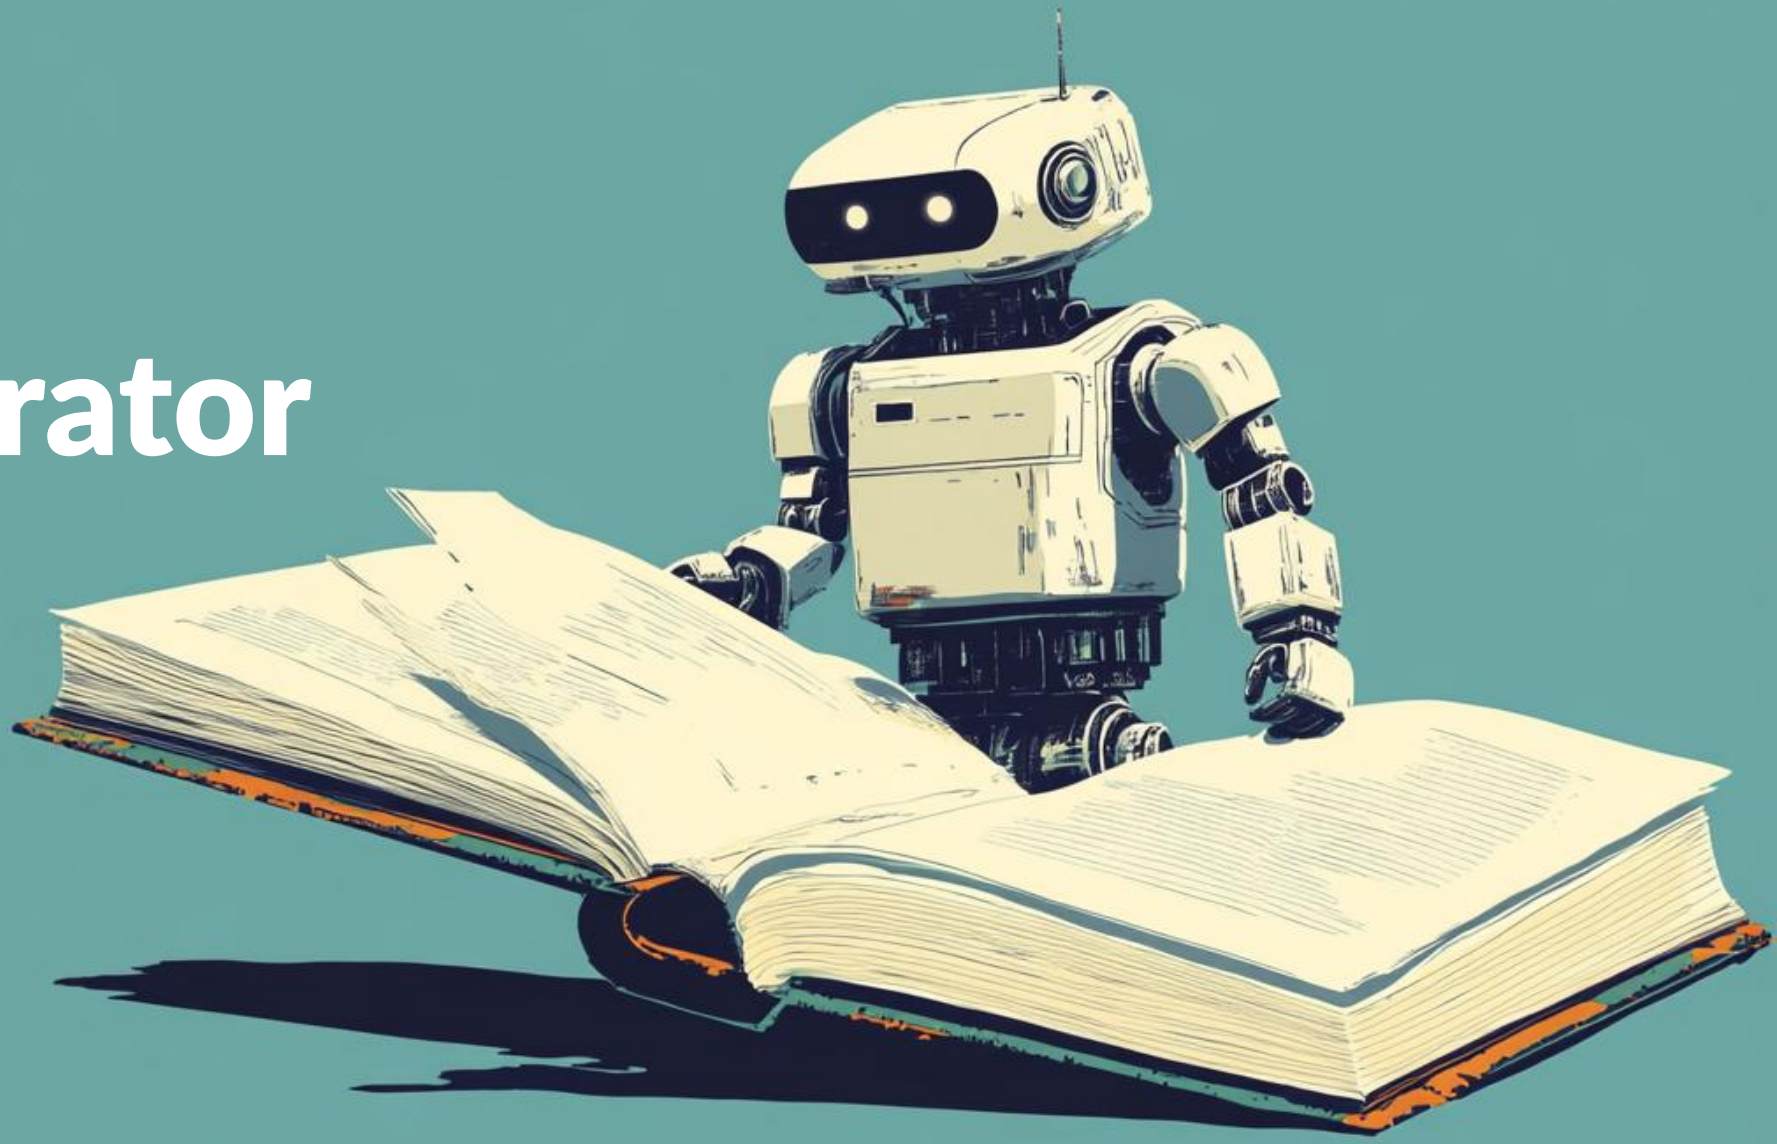
Co-authored RobBERT model (2020)
(state-of-the-art Dutch BERT language model)

Worked with GPT-like models since 2019
(humor generation, Improbotics...)

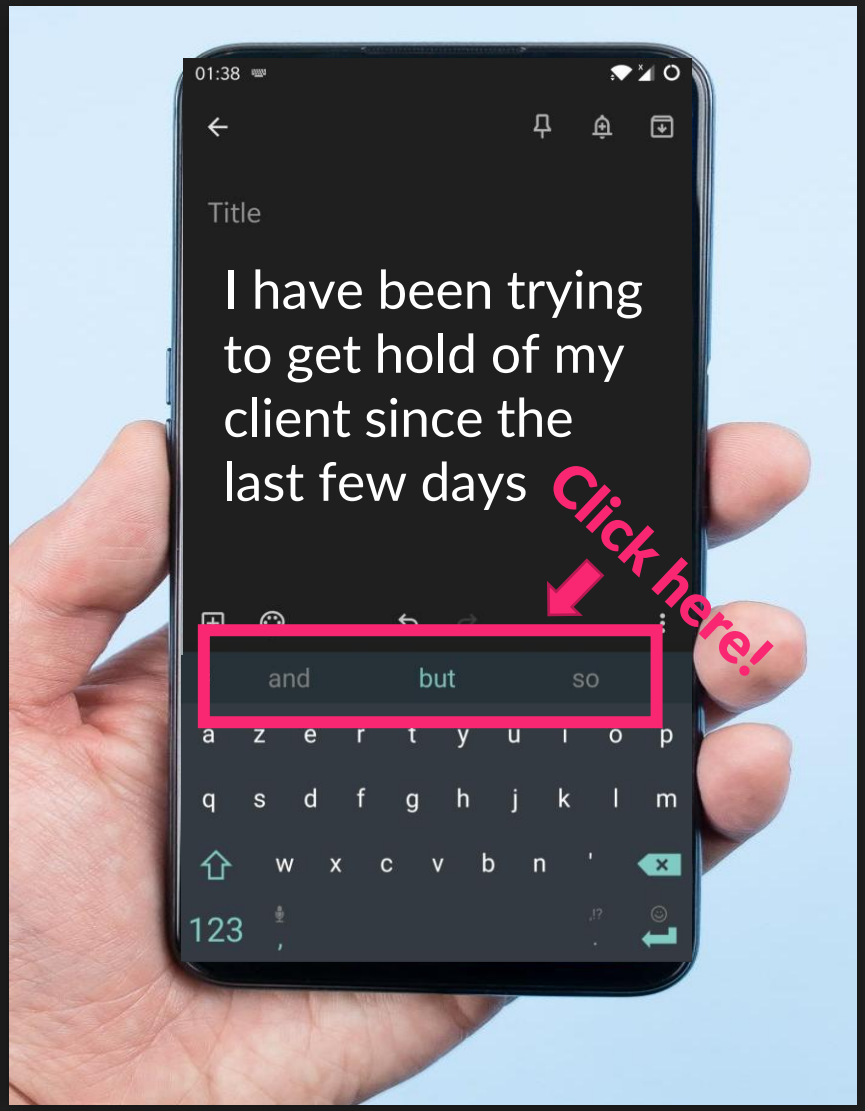
Creative AI consultancy & talks



Text generator history



How do autoregressive language models work?



1. Open smartphone **keyboard**
2. Press any **autocomplete** word 20x
3. You've just generated a sentence using an AI trained to **sound similar to you**!

Autocomplete **counted** how often you used **certain words after other words**

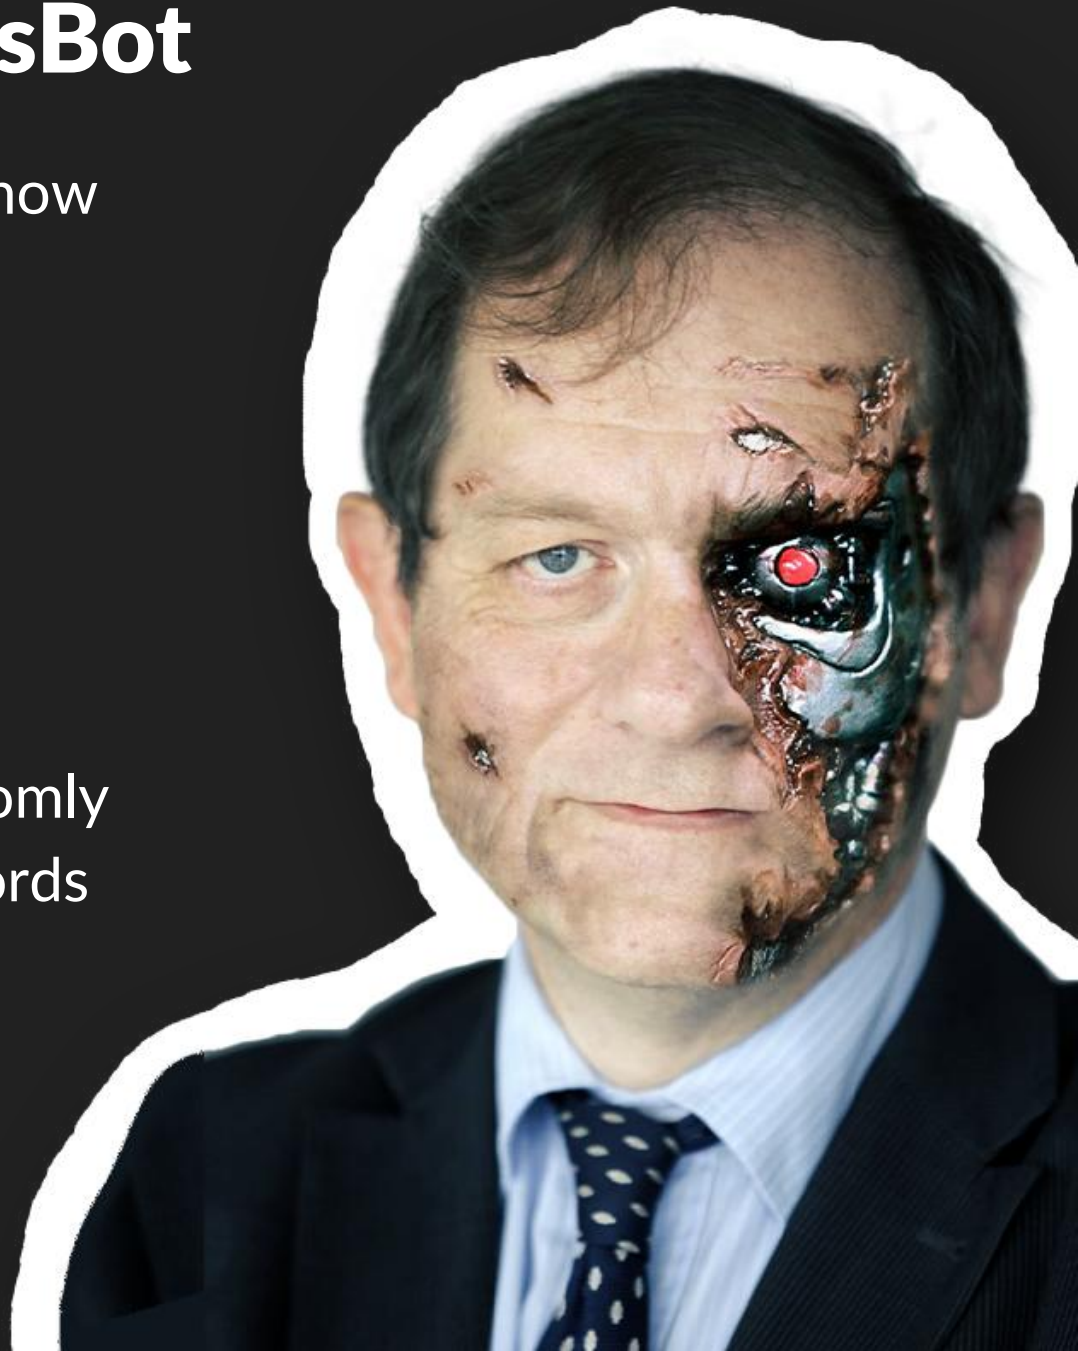
Example: TorfsBot

1. Counted in all tweets & columns of Rik Torfs how often one word follows others



2. Then starts with two random words and randomly select words Rik used after the last couple of words





Training a Large Language Model (LLM)



Whole textual internet

Training process of several million $

Gemini

Large language model (ChatGPT, Gemini, Claude…)

LLM also predicts next tokens



≈ autocomplete on steroids

Parts specialist TVH is launching a new brand of spare parts for uction machinery.

European custo ew replacement pa mon Chinese heavy om the company's Inte

Token	Probability
expanding	11.80%
to	4.81%
the	3.50%
launching	**2.90%**
continuing	2.50%

Probability of next token given *previous tokens*

Tokens	Characters
122	566

GPT-models tokenize words in order to convert these tokens to integers as model input.

The tokens are chosen in such a way that common substrings have their own token. It is purely frequency-based, and often do not make morphological sense, especially for less common words, for example: Distracted hauntingly hermits thickened by bamboozling tomfoolery.

In English, common words are usually a single token.
In het Nederlands worden woorden vaak raar opgesplitst. Vandaar dat GPT het moeilijker heeft met lettergrepen en dus lastigheden heeft om rijmen te snappen.

Text | Token IDs

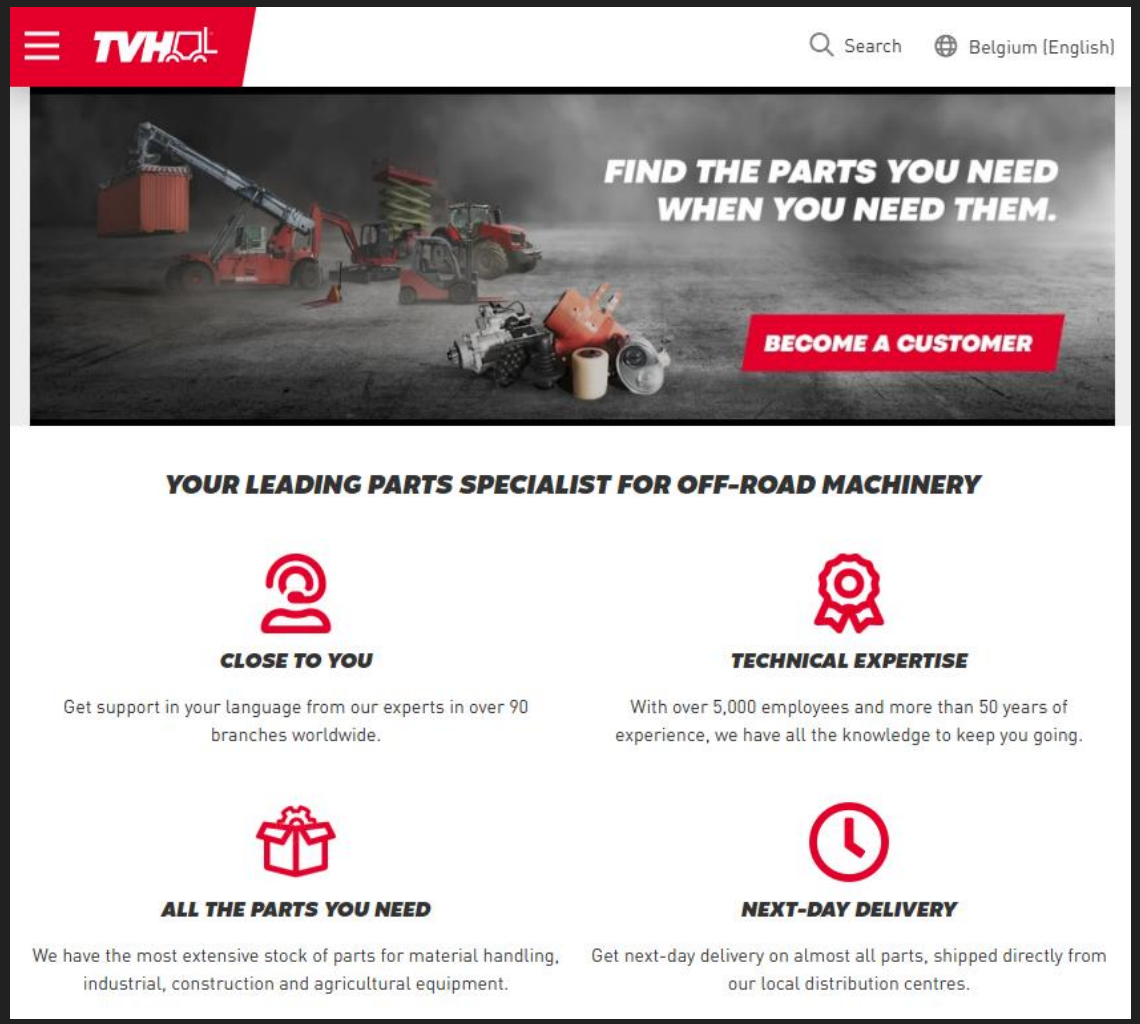
Tokens	Characters
122	566

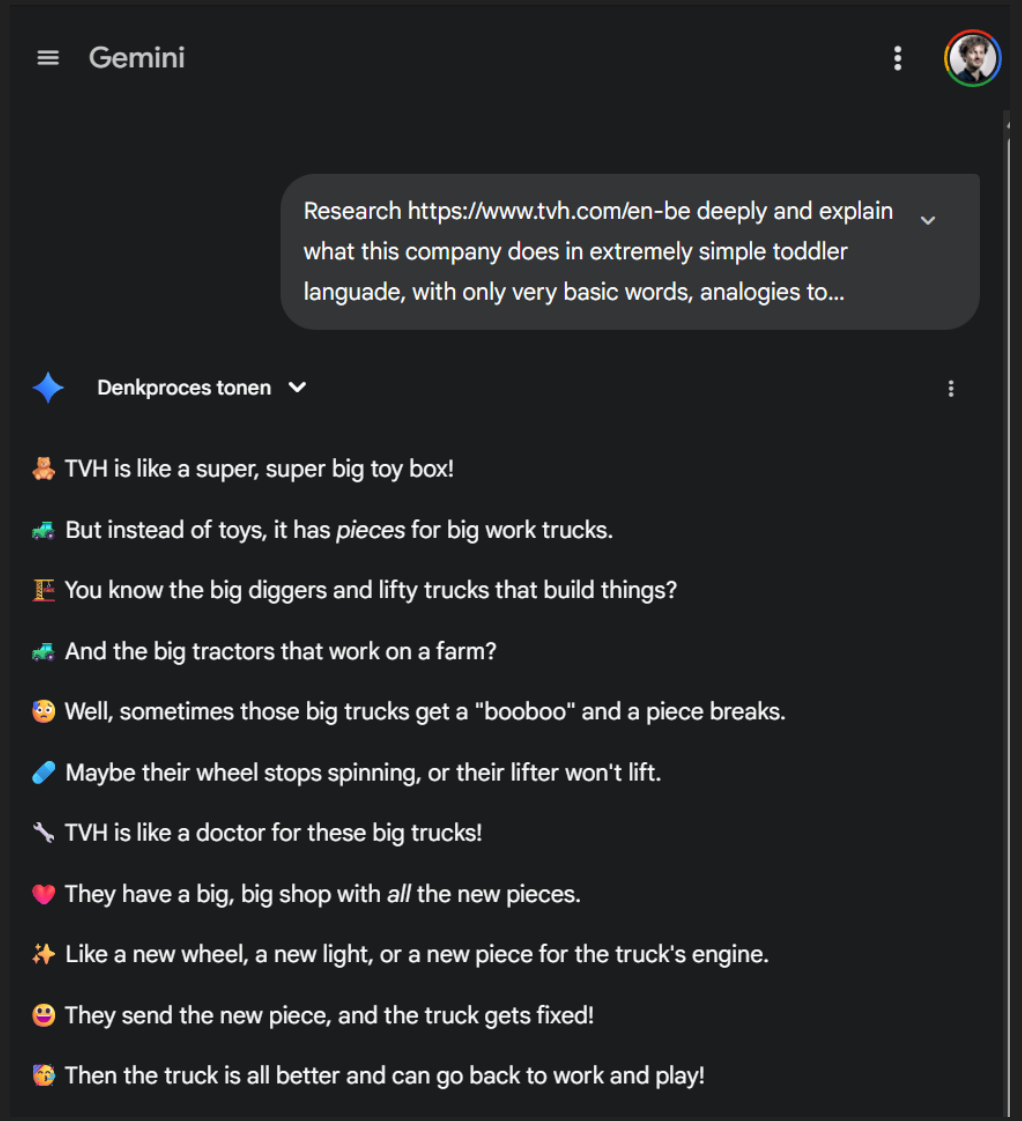
```
[162016, 58699, 82, 192720, 6391, 306, 2569, 316, 8452, 1879, 20290, 316,
54912, 472, 2359, 3422, 364, 976, 20290, 553, 16180, 306, 2238, 261,
2006, 484, 5355, 29234, 963, 679, 1043, 2316, 6602, 13, 1225, 382, 58878,
19813, 8235, 11, 326, 4783, 621, 625, 1520, 190095, 8512, 11, 6980, 395,
3760, 5355, 6391, 11, 395, 4994, 25, 53609, 57630, 128142, 423, 30846,
1348, 17154, 12848, 656, 287, 19017, 6834, 3321, 6784, 69, 971, 875, 364,
637, 7725, 11, 5355, 6391, 553, 6971, 261, 4590, 6602, 558, 637, 1448,
73391, 6367, 68728, 25361, 428, 2125, 120944, 4059, 302, 13, 108402,
2125, 1814, 174803, 1448, 52910, 259, 6956, 1421, 10263, 9174, 2931, 469,
12392, 2174, 42246, 6956, 2030, 41291, 2712, 729, 3967, 21000, 13]
```

Text | Token IDs

https://platform.openai.com/tokenizer

Main strengths of GPT





Copy writing → Copy editing

Many tasks are “Convert text to other formats”

Even to extremer styles

Now summarize the company again but like a pirate would talk, full of references to pirate life and the sea

Denkproces tonen

Ahoy, matey! Gather 'round and listen! 🏴‍☠️

⚓ Ye know those giant, land-lubberin' beasts? The big metal contraptions that lift heavy chests, dig for buried treasure, and work the docks? Like forklifts and diggers, arrr!

🔧 Well, even the mightiest ship needs repairs! Sometimes, these workin' beasts bust their riggin', spring a leak, or lose a vital piece o' their gear! They get stuck in port, unable to haul the booty!

📦 That's where TVH comes in, ye savvy? They be like the ultimate Quartermaster for all these land ships! They have a treasure hoard... a massive, secret warehouse, bigger than any cave on a desert isle!

⚙️ But this hoard ain't filled with gold doubloons or shiny jewels! Nay! It's packed to the gills with every part, screw, cog, wheel, and gizmo ye could ever need to fix those workin' beasts! Millions of 'em!

🚢 So, when a land ship breaks down, its captain sends a message (like a message in a bottle, but faster!) to TVH.

Research https://www.tvh.com/en-be deeply and create an insightful table of what they do, with insightful and useful columns for understanding their business as wel...

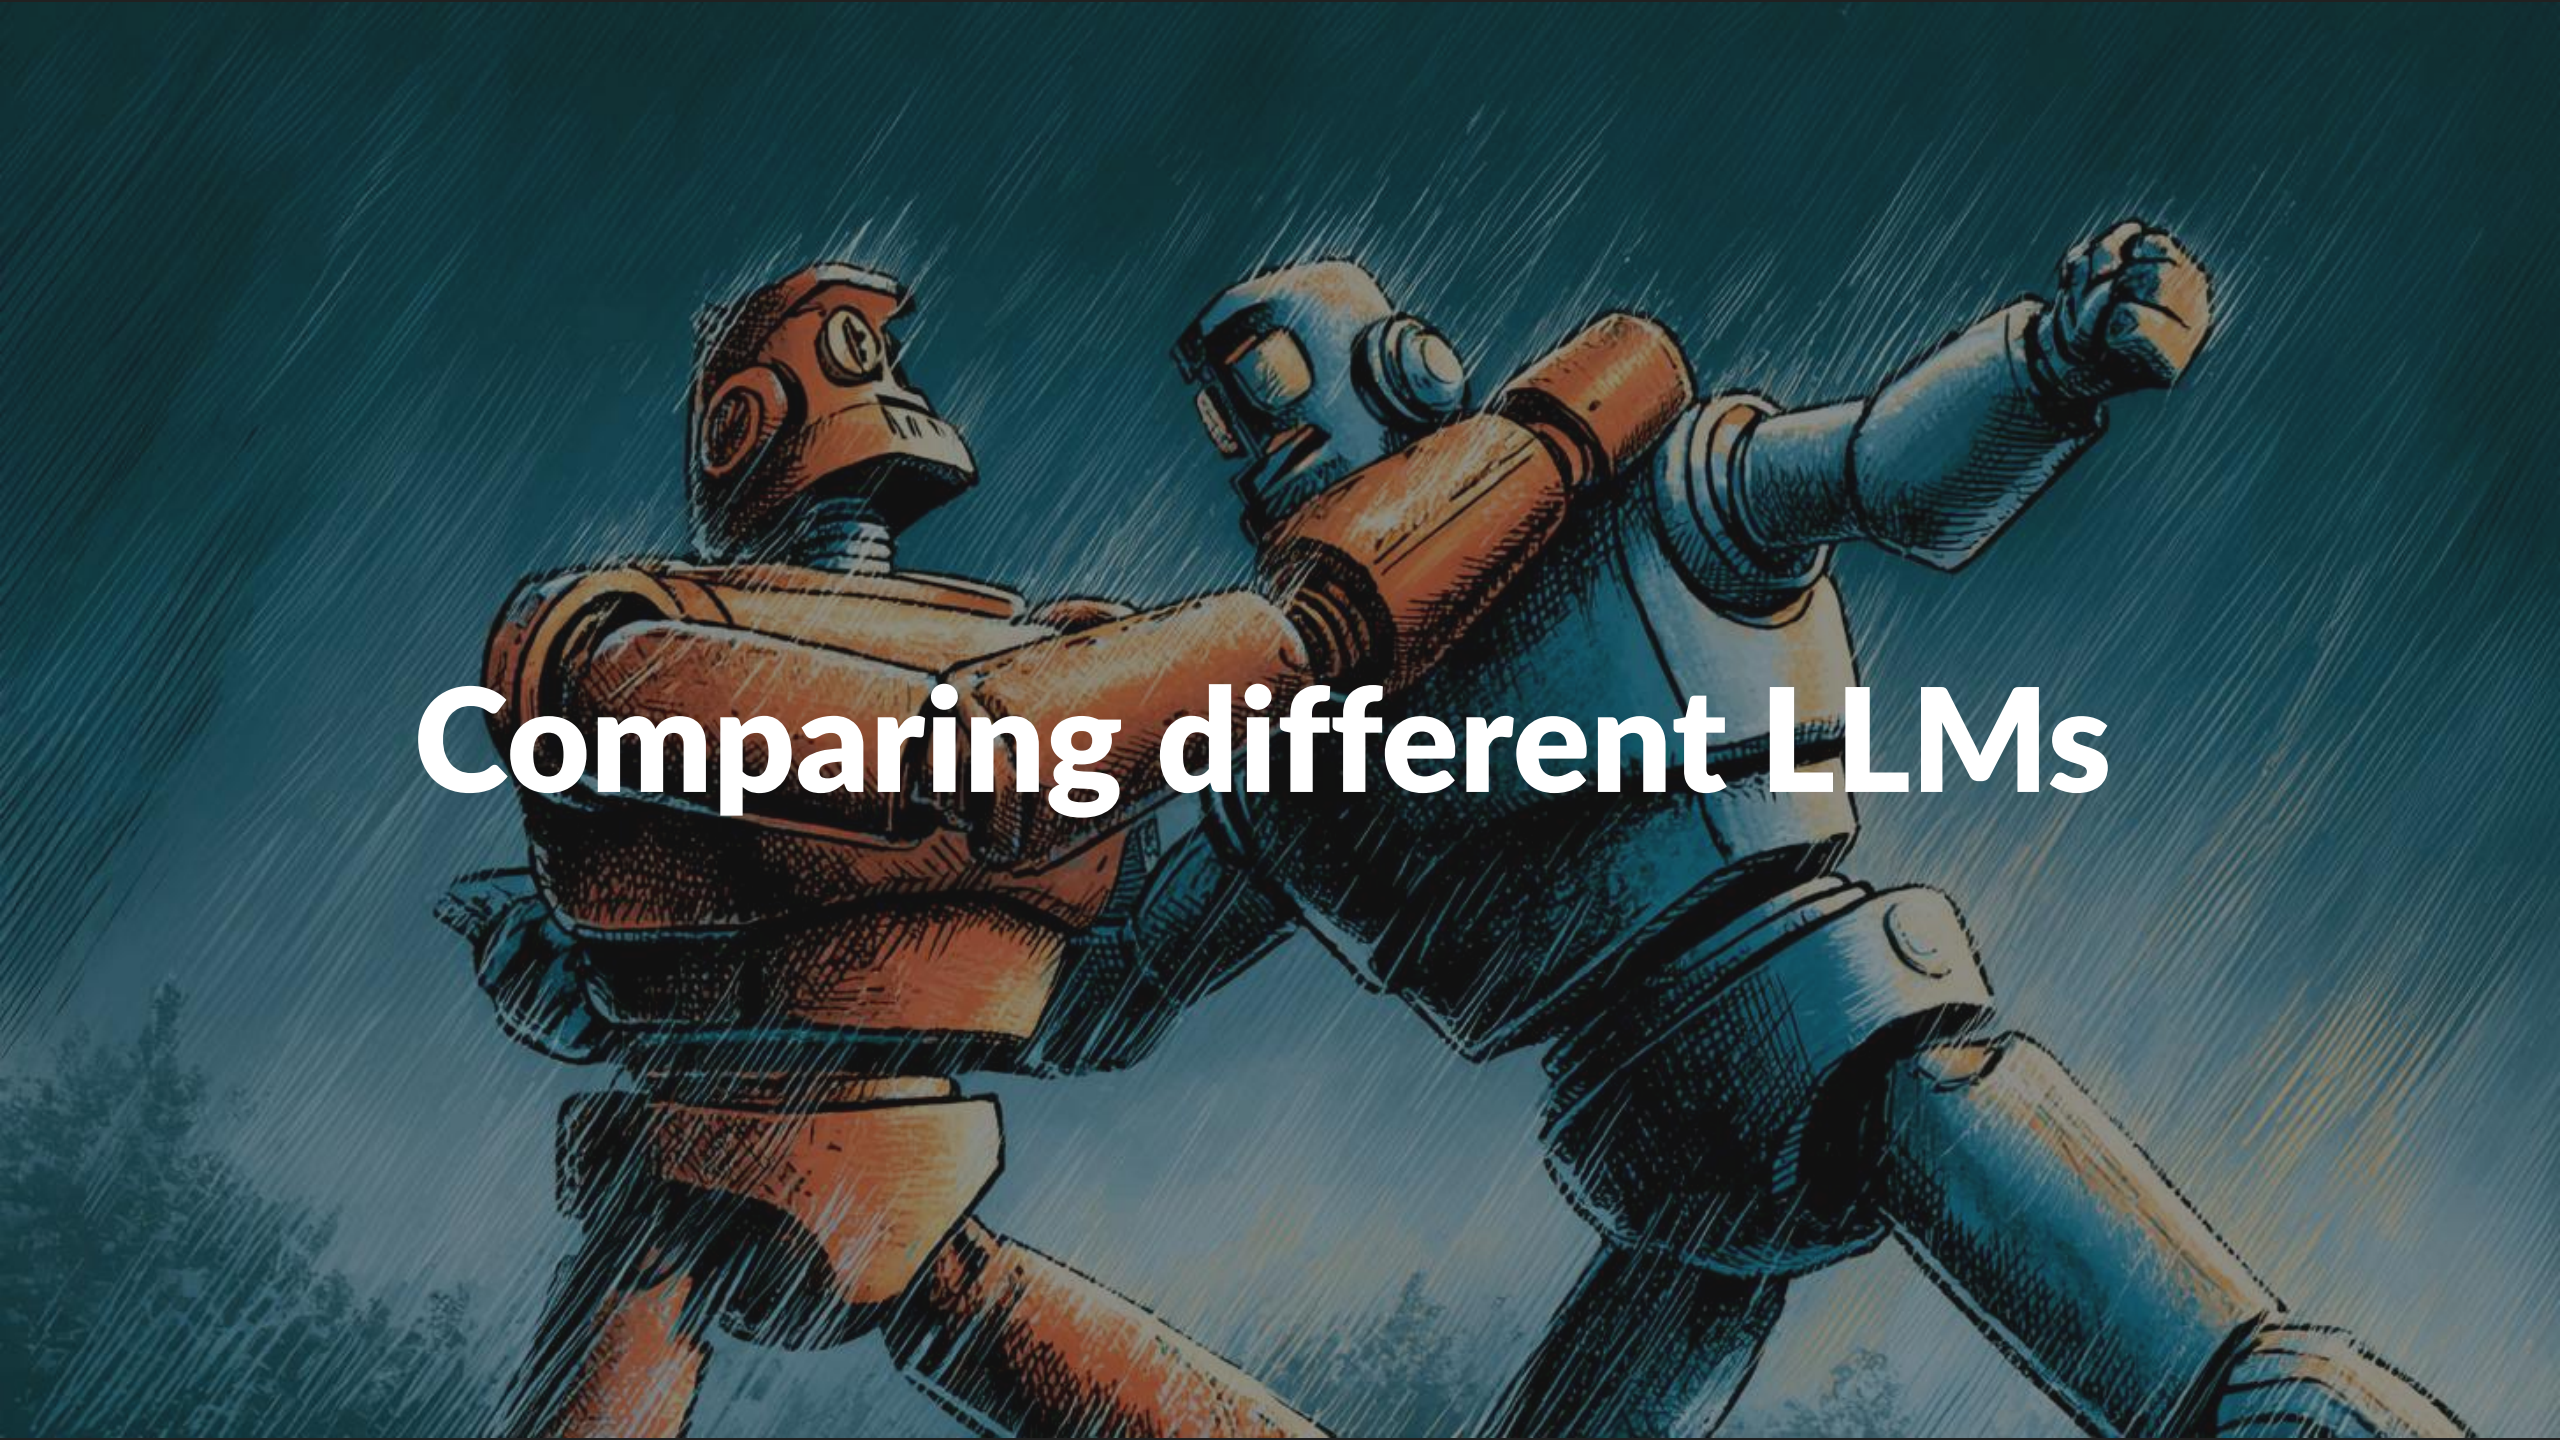
Denkproces tonen

Here is a detailed research table on TVH, based on its corporate website and further online research.

In-Depth Analysis of TVH's Business Operations

TVH has strategically structured its business into two main pillars: **TVH Parts Holding** (the primary focus of `tvh.com`) and **mateco Holding** (equipment rental/sale). The "Parts" division is a global "one-stop shop" supported by extensive digital platforms, technical services, and training.

Business Unit / Service	Description & Offerings	Target Market & Industries	Key Differentiators & Details	Source(s)
TVH Parts Holding	The core business. A global B2B "one-stop shop" for a massive range of spare parts and accessories for "off-road" machinery.	- Independent Repair Shops - Dealers & Resellers - Rental Companies - Industrial Fleet Operators **Industries:** Logistics, Material Handling, Construction, Agriculture, Industrial	**Massive Scale:** 44 million+ known references; 1 million+ parts in stock. **Global Reach:** 90+ branches, customers in 180 countries, 5,000+ employees. **Brands:** Distributes major OEM brands (Curtis, Donaldson, etc.) and has its own house brand, **Totalsource**.	1, 2
Digital Platform (MyTotalSource)	The primary B2B e-commerce webshop. It allows registered customers to look up parts (by brand, model, or part number) and place orders 24/7.	All B2B customers of TVH Parts.	- The platform is the central nervous system of the "one-stop shop" model. - Part of a major digital transformation to the cloud (Azure) to improve efficiency and enable innovation (e.g., AI, predictive maintenance).	3, 4
Remanufacturing (Reman)	A major circular economy service. TVH remanufactures used parts to "as new" condition. This includes complex electronic and mechanical components.	Customers seeking a cost-effective (up to 50% cheaper) and environmentally friendly alternative to new parts.	**Electronic Reman:** "Europe's largest" repair center for PCBs, joysticks, displays, and controllers. **Mechanical Reman:** Rebuilds engines, transmissions, pumps, etc. - Process: Disassembly, inspection, replacement of all wear parts, reassembly, and testing on custom-built test benches.	5, 6
TVH University	A comprehensive training and knowledge center (established 2007) for both customers and internal staff.	- External service technicians seeking certification. - Technical-commercial employees. - Internal TVH staff.	**Hybrid Model:** Offers in-person classes, e-learning, live virtual training, and bespoke on-site courses. **Curriculum:** Covers hydraulics, power electronics (batteries/chargers), engines, and electrics/electronics.	7, 8
Technical Services	A suite of value-added services supporting the parts business.	Customers needing specialized repairs or custom solutions.	**Electronics Department:** In-house repair service for electronic parts. **Fleet Management:** Offers digital solutions to connect systems and manage parts for a customer's entire fleet. **Other Services:** Electric system design, battery regeneration (via brand **Energic Plus**), and purchasing of old parts.	9
mateco Holding (formerly TVH Equipment)	The second main pillar of the TVH Group, operating separately from the parts business.	- Construction Companies - Event Organizers - Industrial Maintenance - Municipalities	**Focus:** Specializes in the **sale and rental** of new and used aerial work platforms, forklifts, and telehandlers. **Brand:** Operates internationally under the **mateco** brand.	10



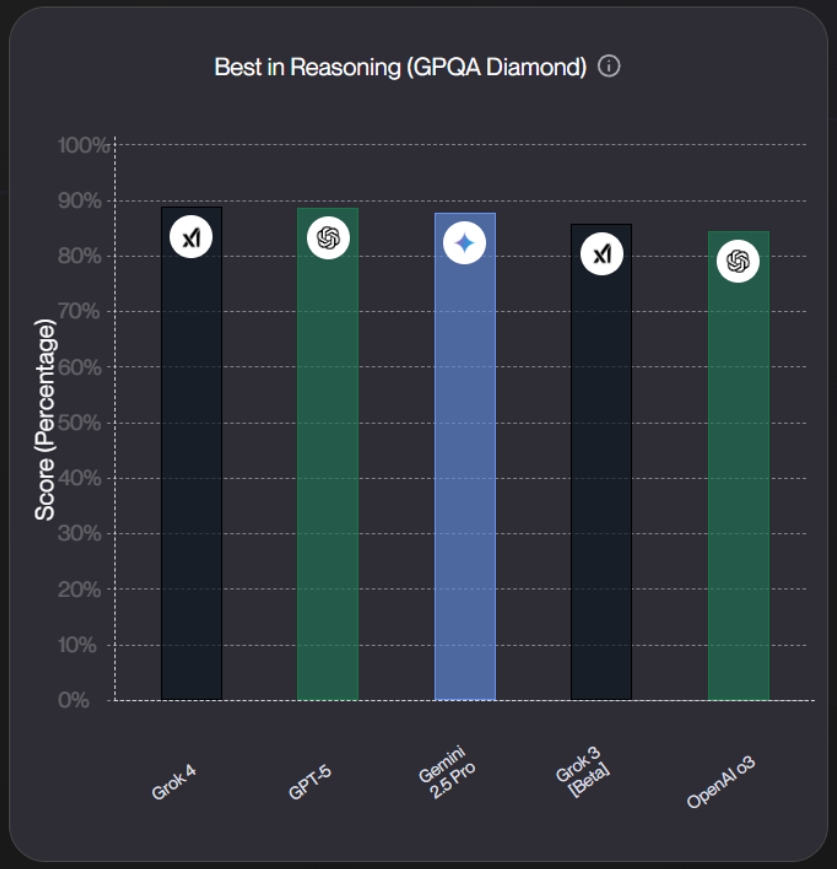
Comparing different LLMs

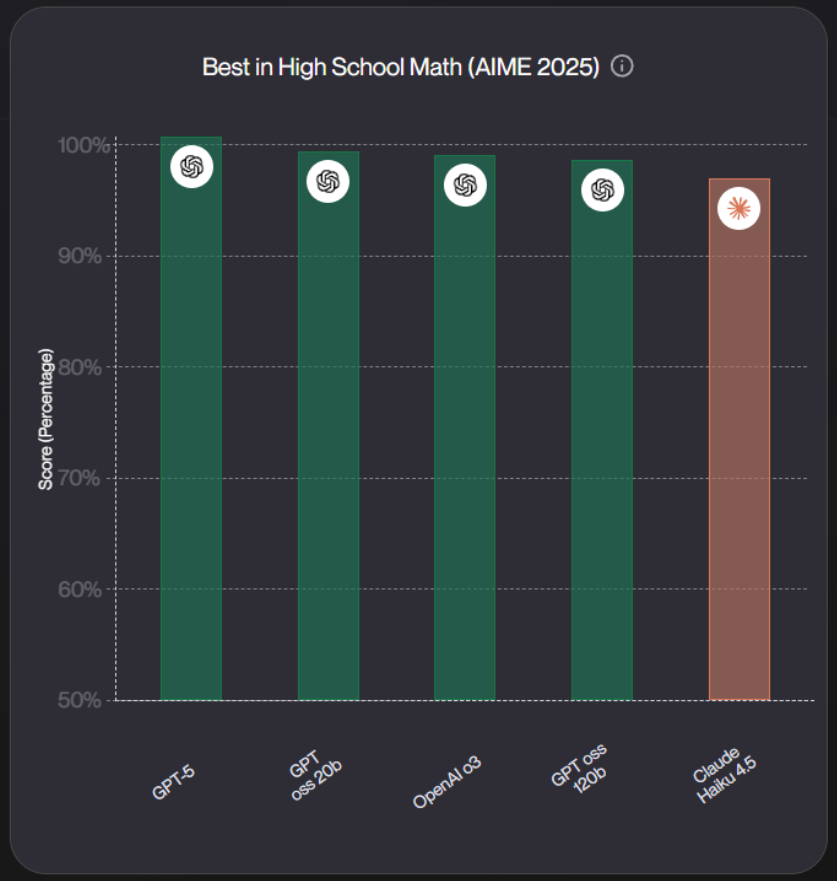
Approach 1: Benchmarks with fixed test datasets

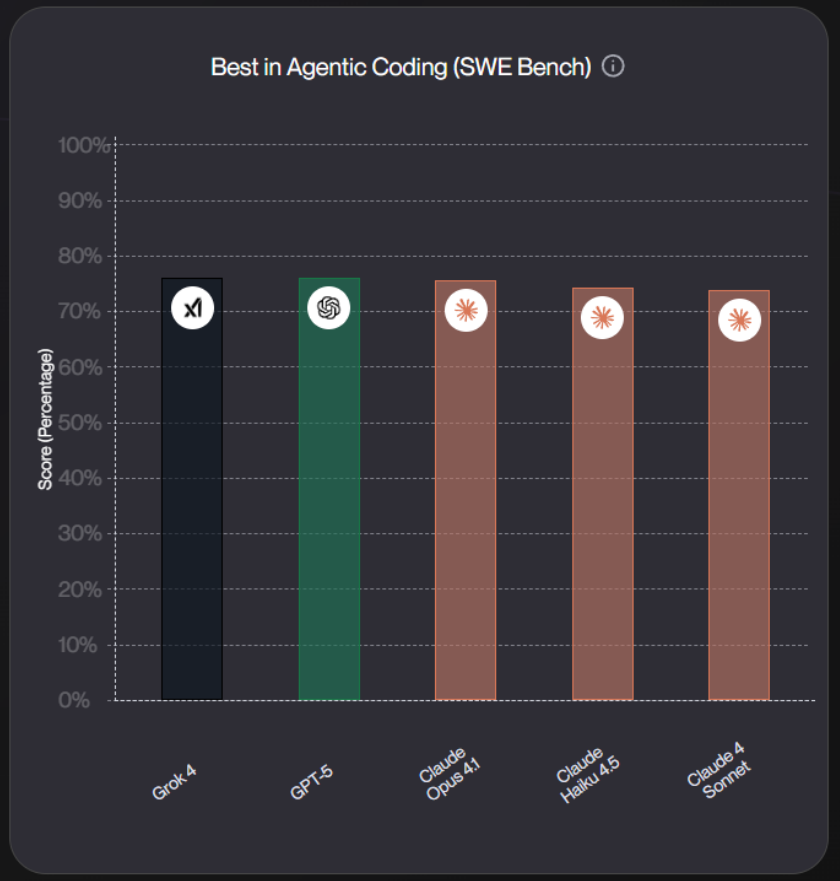
Group multiple benchmarks per type: Test accuracy on a wide range of specific tasks

E.g., livebench.ai & vellum.ai/llm-leaderboard & HuggingFace Big Benchmarks

Top models per tasks



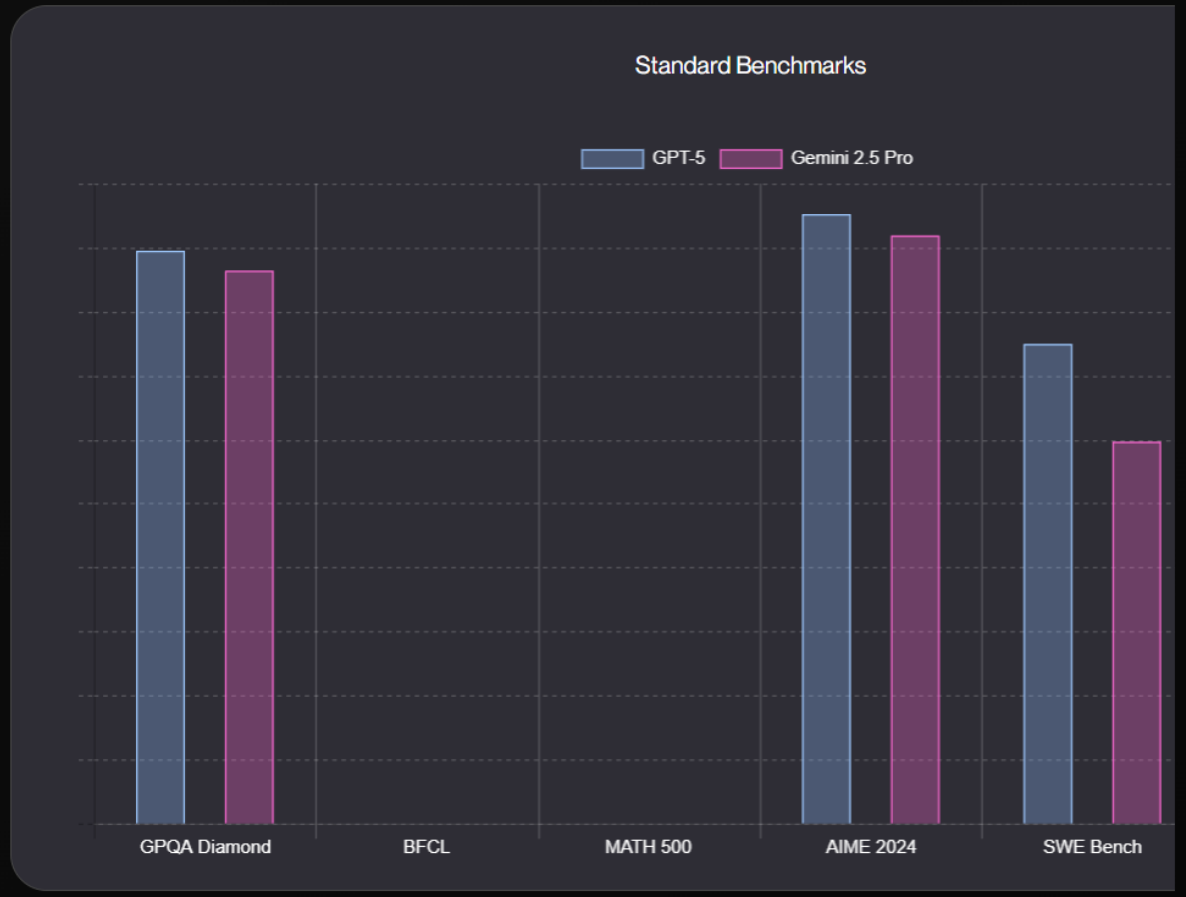




Compare models on specific benchmarks

GPT-5 vs Gemini 2.5 Pro

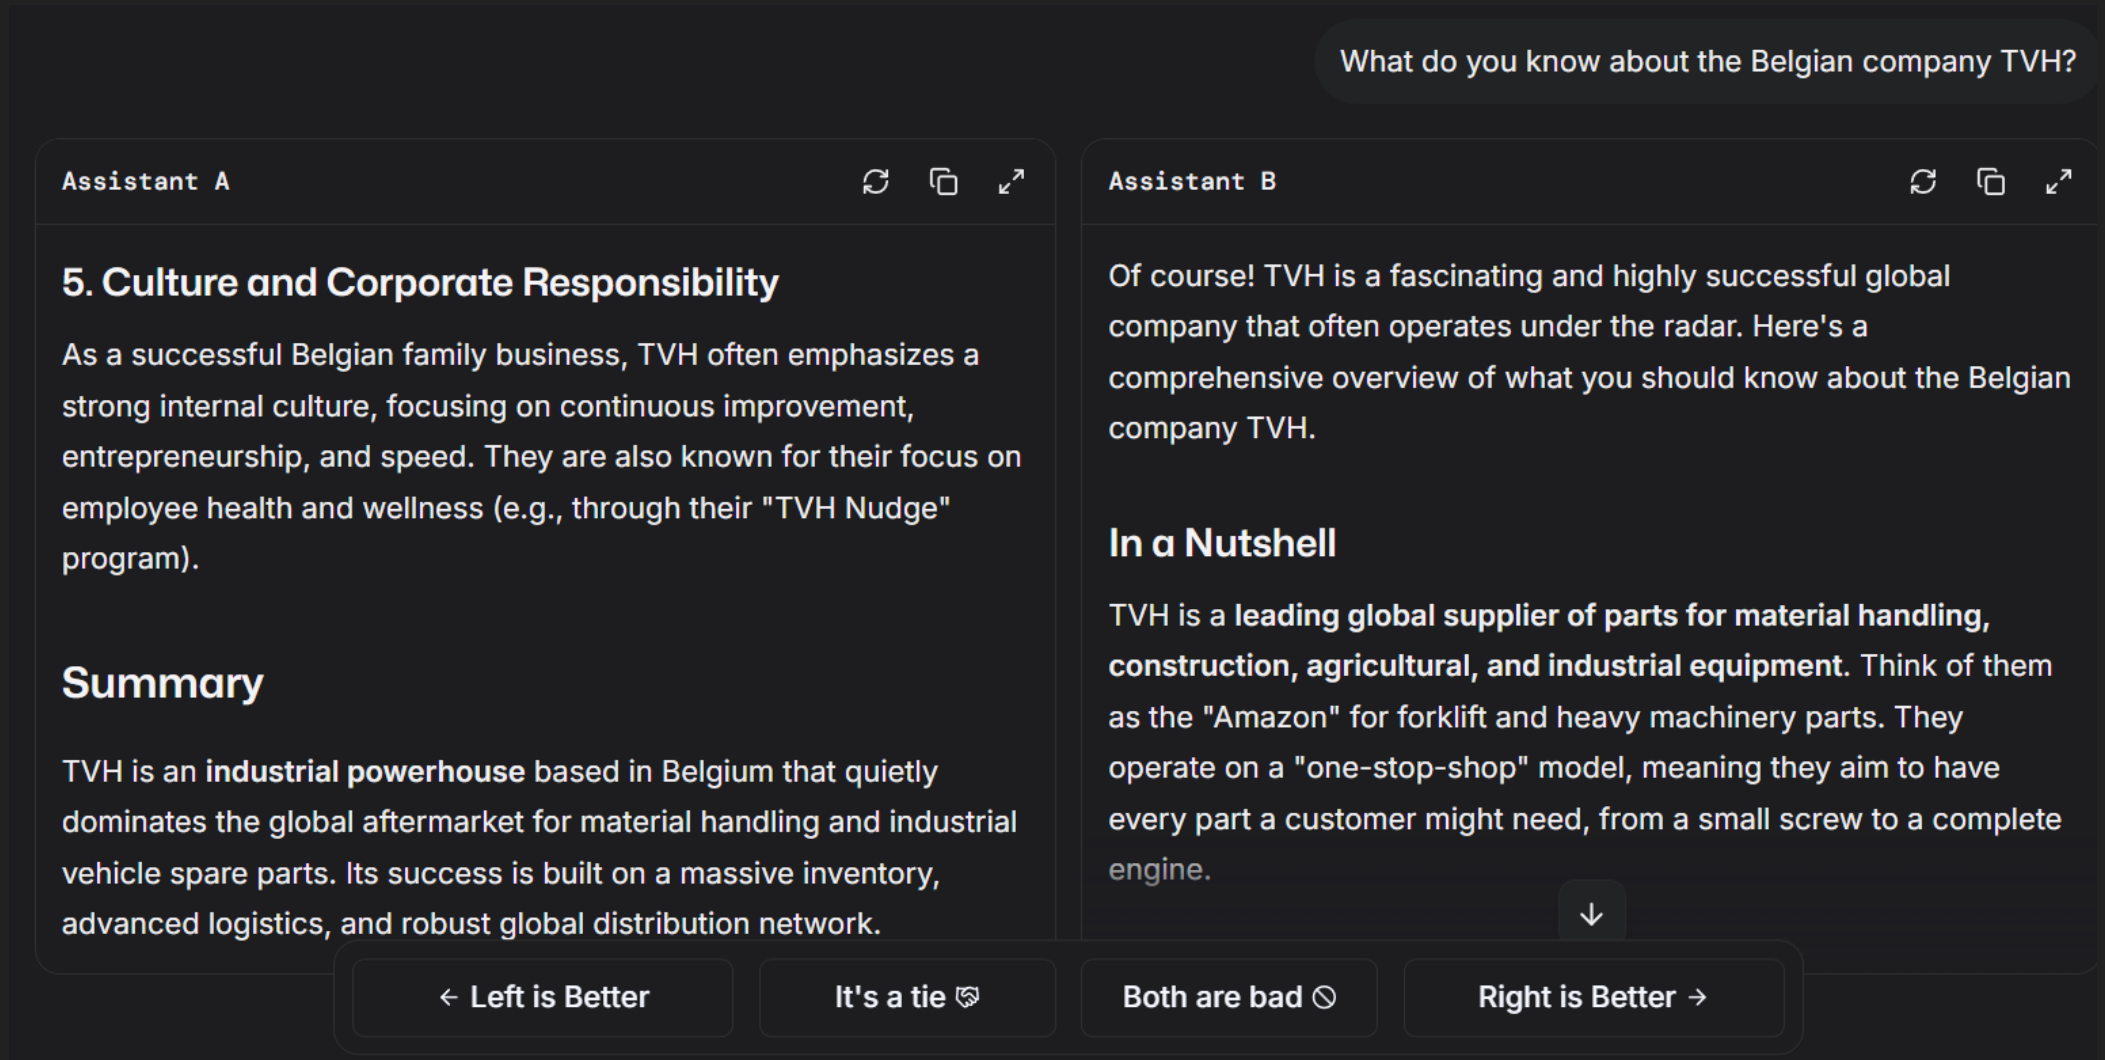
MODEL	GPT-5	GEMINI 2.5 PRO
Context size	400,000	1,000,000
Cutoff date	April 2025	Nov 2024
I/O cost	$ 1.25 / $ 10	$ 1.25 / $ 10
Max output	128,000	65,000
Latency	n/a	30s
Speed	n/a	191t/s



https://www.vellum.ai/llm-leaderboard

Approach 2: “Vibecheck” with ELO scoring

Rate two answers: more points if “worse” model wins



LMArena.ai: leaderboards per type

Text

2 days ago

Rank (UB)	Model	Score	Votes
1	gemini-2.5-pro	1452	62,764
1	claude-sonnet-4-5-20250929-t…	1449	13,853
1	claude-opus-4-1-20250805-thi…	1448	29,426
2	gpt-4.5-preview-2025-02-27	1442	14,644
2	claude-opus-4-1-20250805	1439	41,950
4	chatgpt-4o-latest-20250326	1438	48,510
4	gpt-5-high	1436	30,974
2	claude-sonnet-4-5-20250929	1436	5,476

WebDev

1 day ago

Rank (UB)	Model	Score	Votes
1	GPT-5 (high)	1473	8,004
1	Claude Opus 4.1 thinking-16k…	1458	8,726
2	Claude Opus 4.1 (20250805)	1451	8,986
4	Claude Sonnet 4.5 (thinking …	1420	4,863
4	MiniMax-M2	1405	3,515
5	Gemini-2.5-Pro	1399	14,628
5	GLM-4.6	1395	7,563
5	DeepSeek-R1-0528	1393	4,800

Overall scoring (LMArena)

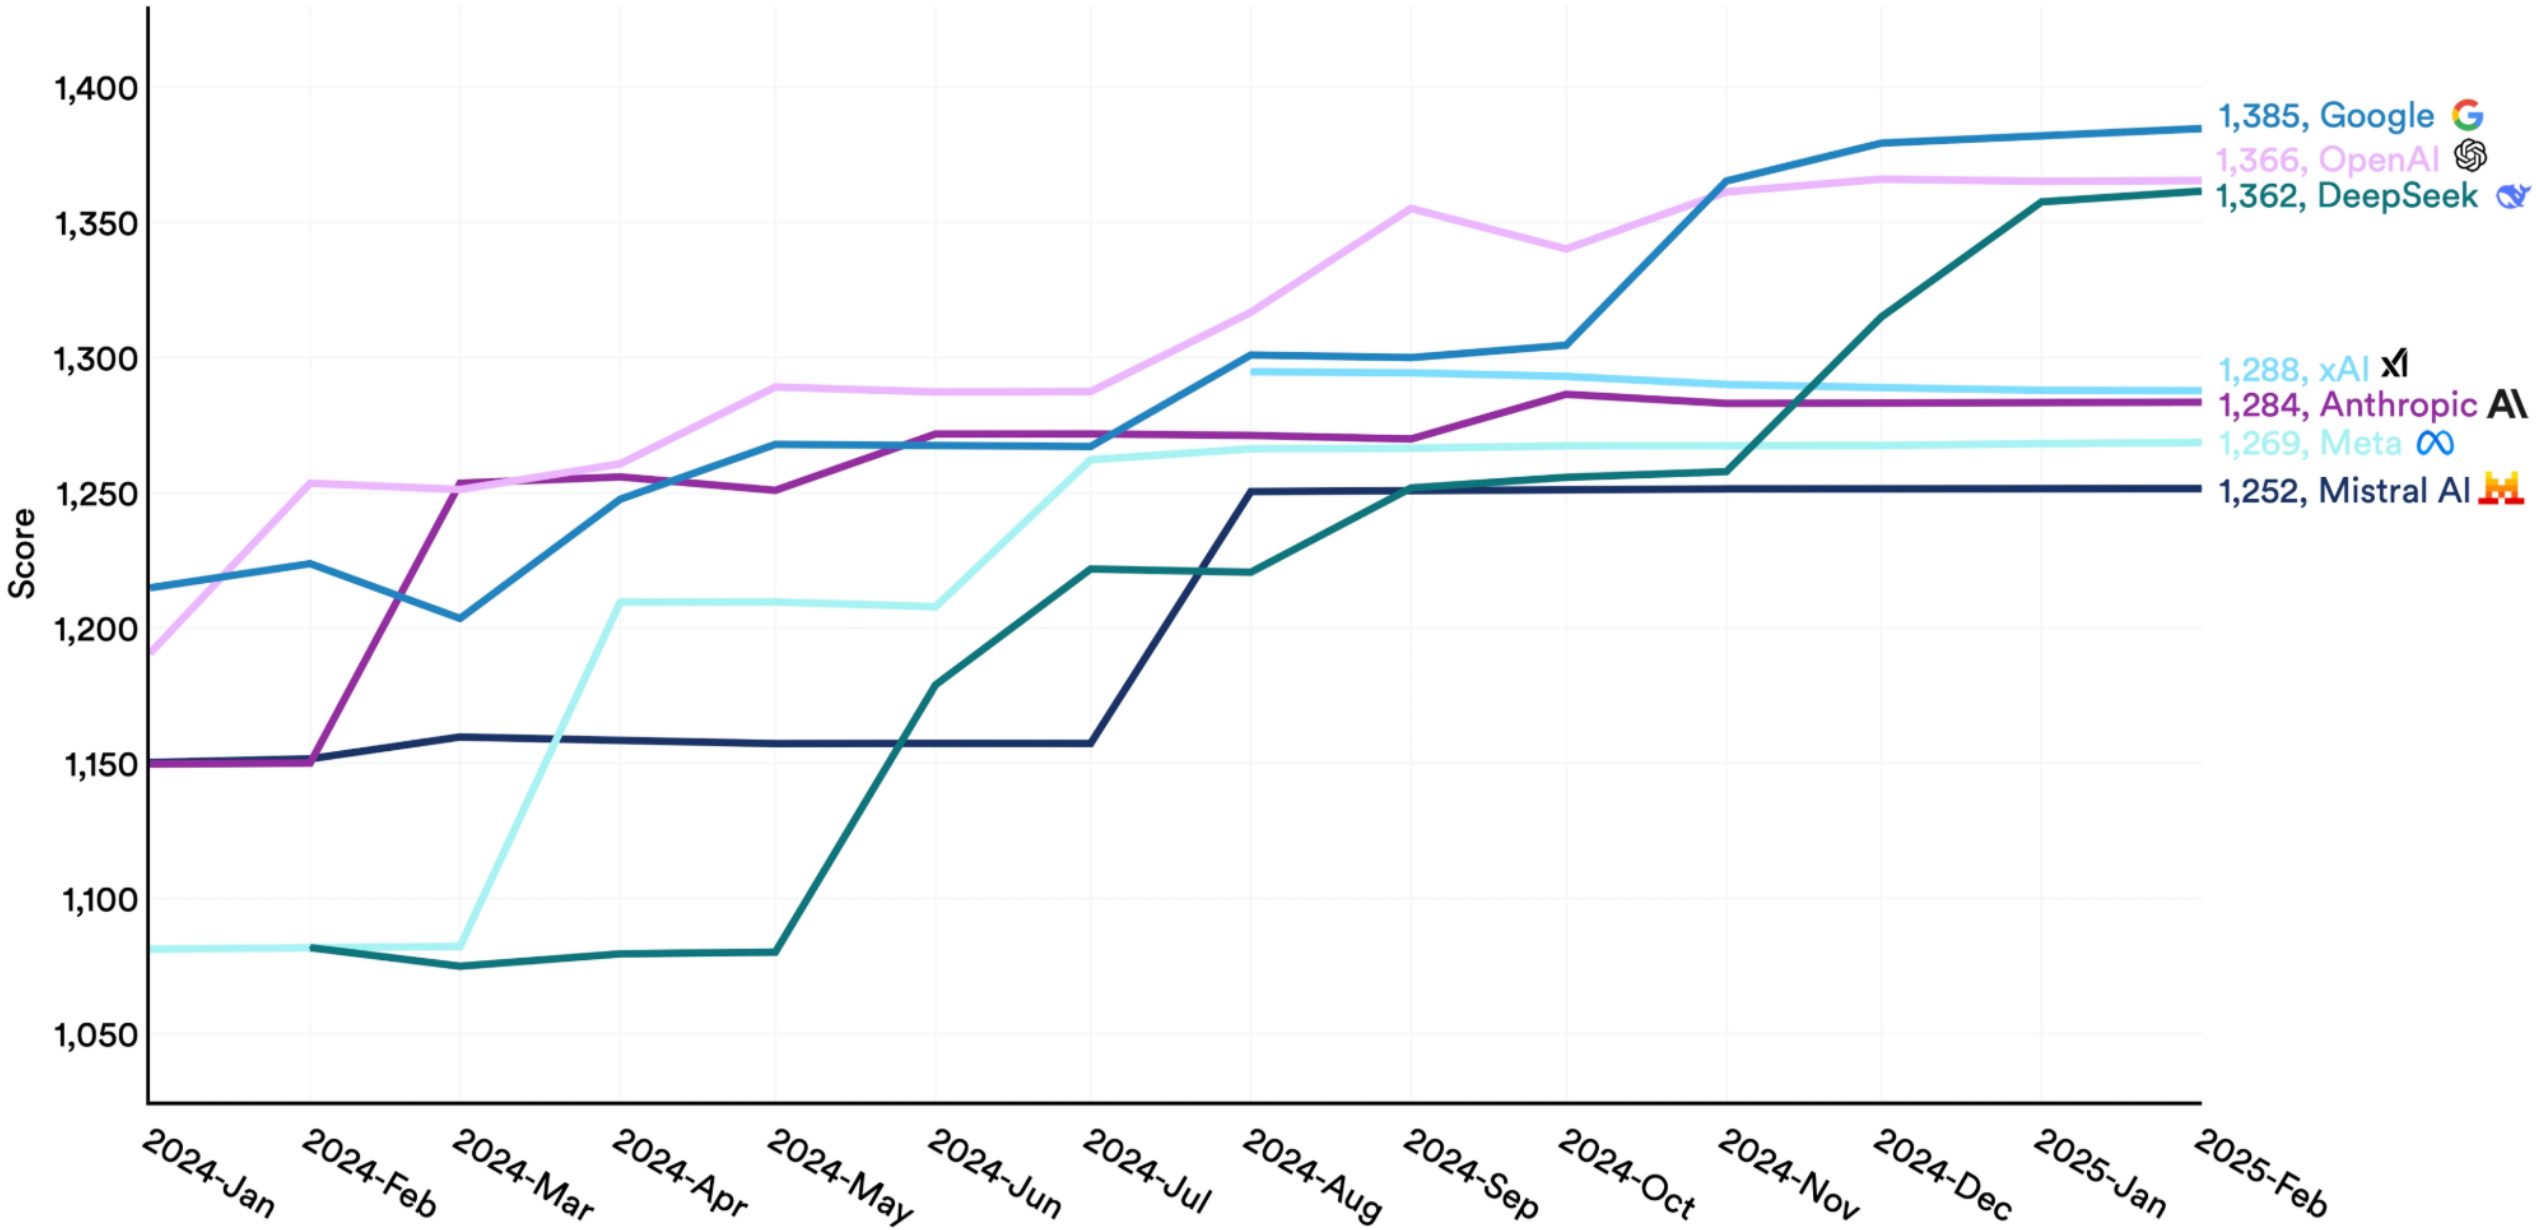
First Place | Second Place | Third Place

Default | Compact View

Model 261 / 261	Overall	Hard Prompts	Coding	Math	Creative Writing	Instruction Following	Longer Query	Multi-Turn
claude-opus-4-1-202…	1	1	1	1	1	1	1	1
claude-sonnet-4-5-2…	1	1	1	1	1	1	1	1
gemini-2.5-pro	1	3	5	1	1	2	3	1
claude-opus-4-1-202…	2	3	2	1	2	1	2	1
claude-sonnet-4-5-2…	2	5	5	1	1	5	3	1
gpt-4.5-preview-202…	2	7	6	7	1	4	4	1
chatgpt-4o-latest-2…	4	5	6	16	4	7	6	1
gpt-5-high	4	5	5	1	11	7	17	10
o3-2025-04-16	4	7	8	1	13	14	21	11
qwen3-max-preview	4	4	4	1	9	5	5	4
glm-4.6	6	5	5	1	4	5	5	10
ernie-5.0-preview-1…	7	7	6	-	1	6	5	11
kimi-k2-thinking	7	6	3	1	3	7	6	7

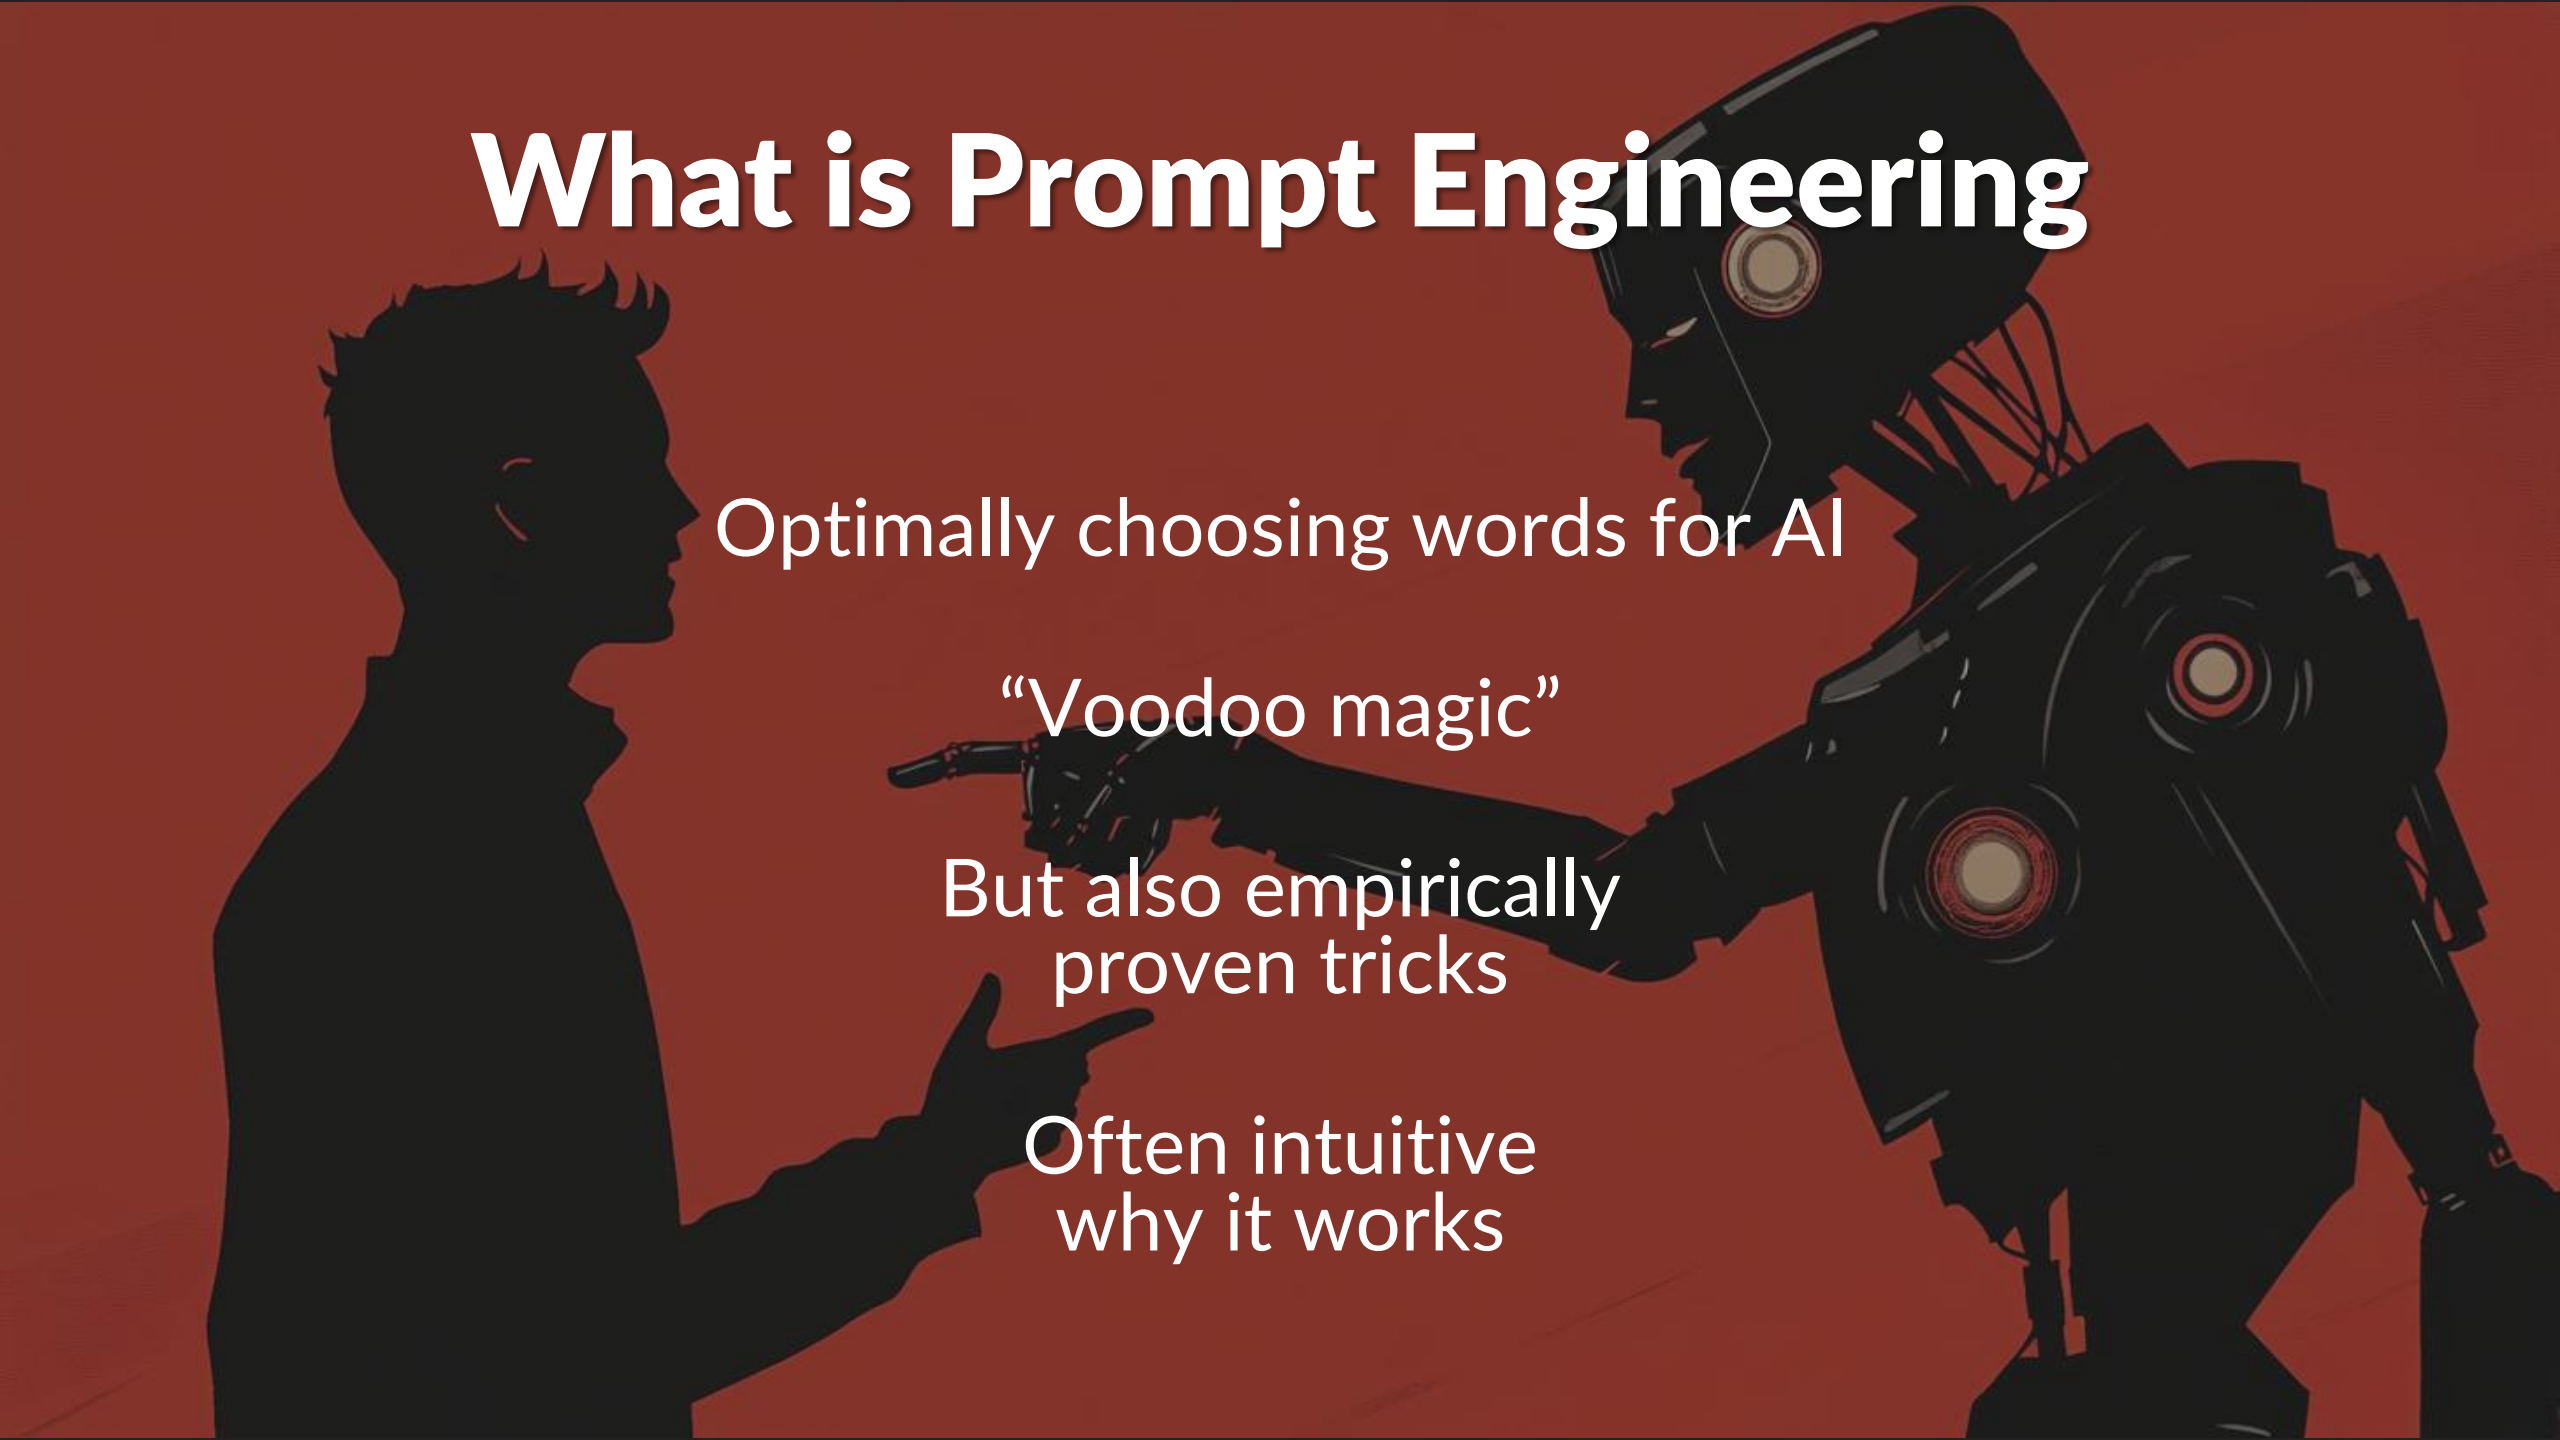
Performance of top models on LMSYS Chatbot Arena by select providers

Source: LMSYS, 2025 | Chart: 2025 AI Index report



https://hai.stanford.edu/ai-index/2025-ai-index-report

What is Prompt Engineering



Optimally choosing words for AI

"Voodoo magic"

But also empirically
proven tricks

Often intuitive
why it works

Write a joke about TVH

Why did the material handling mechanic break up with his girlfriend?

Because she said he spent too much time dealing with **TVH** parts and not enough time on their **relationship's attachments**!

Write a joke about TVH

Why did the material handling mechanic break up with his girlfriend?

Because she said he spent too much time dealing with **TVH** parts a[illegible]ough time on their **relationship's att[illegible]ents!**

Give context

Role

You are a world-renowned expert stand-up comedian, with a proven track record of writing brilliant jokes for professional audiences.

Task Context

Your task is to write a joke about TVH for employees of TVH in the audience of a talk.

Information Context

Context: TVH is a Belgian-founded (1969) global specialist in parts and accessories for material handling, industrial, construction, and agricultural equipment—covering everything from pallet trucks and forklifts to small earth-moving machines. Headquartered in Waregem, Belgium, the company serves customers in more than 180 countries and emphasizes fast fulfillment through a massive catalog and logistics network. TVH highlights a database of roughly 49.5 million part

You know you work at TVH when you type something into the internal parts database...
...and it feels like you're summoning an ancient spirit.

You type "forklift bolt," and suddenly it's like —
"Did you mean: 42,381,972 possible variations — all in stock, shipped yesterday?"

At this point, I'm not even sure if I work in logistics or Hogwarts.
I half expect the catalog to whisper, *"Expediamus partnum!"*

Give examples

You are a world-renowned expert stand-up comedian, with a proven track record of writing brilliant jokes for professional audiences.

Your task is to write a joke about TVH for employees of TVH in the audience of a talk.

Write jokes in the style of these jokes:
- I decided to sell my Hoover... Wel, it was just collecting dust
- A friend of mine's got his left arm missing. Serves him right!
- I've got an alphabet grenade. If that thing goes off, it could spell disaster!

Context: TVH is a Belgian-founded (1969) global specialist in parts and accessories for material handling, industrial, construction, and agricultural equipment—covering everything from pallet trucks and forklifts to small earth-moving machines. Headquartered in Waregem, Belgium, the company serves (...)

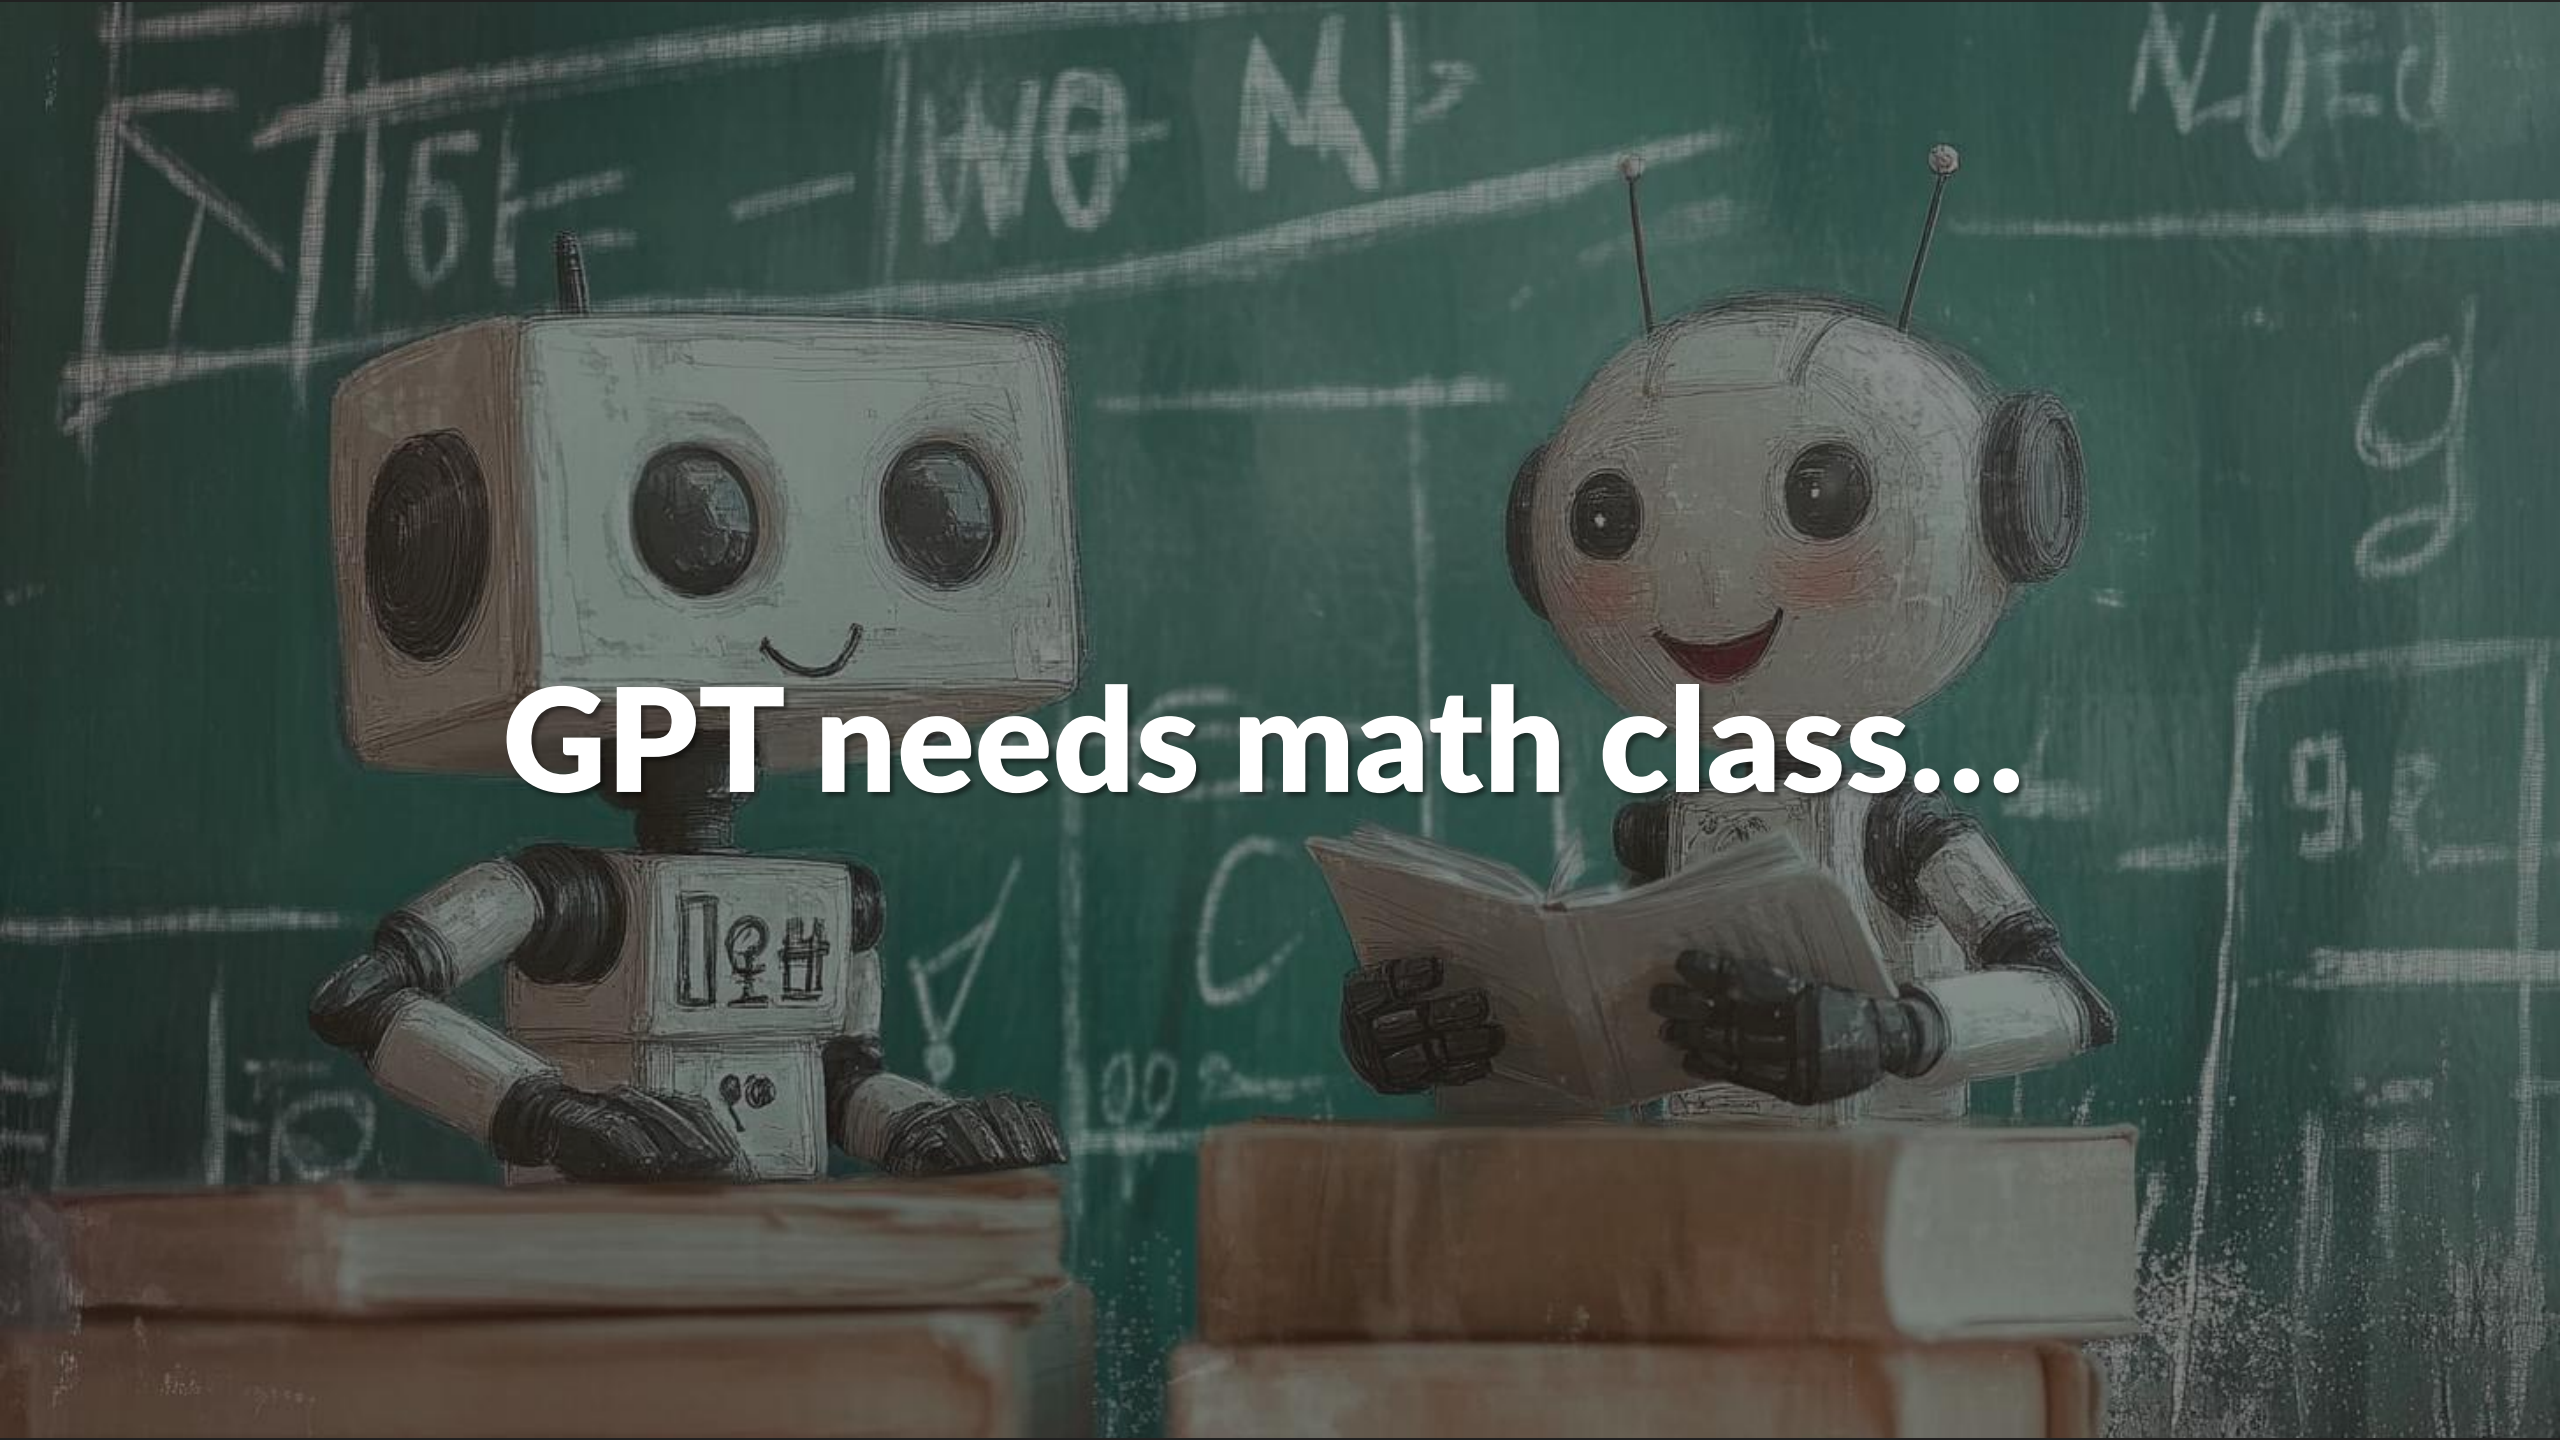
(3) They say "We Are One." Which is true — except in the warehouse, where they're 49.5 million.

→ (wordplay on corporate slogan)

(5) TVH's catalog is so big, I scrolled so long I aged into a spare part myself.

→ (hyperbole + self-deprecation)

... But no planning


GPT needs math class...

ChatGPT is bad at math riddles?

ChatGPT 5

1.9 or 1.11, which is bigger?

1.11 is bigger than 1.9 ✅

Here's why:

- 1.9 = 1.90 (if you write it with two decimal places)
- 1.11 = 1.11

When you compare 1.90 vs 1.11, **1.90 > 1.11**.

So, **1.9 is actually bigger than 1.11** — correction: **1.9 > 1.11** ✅

👉 Final answer: **1.9 is bigger than 1.11.**

“Think step by step”

1.11 or 1.9 which one is bigger? Think step by step

To determine which number is larger between 1.11 and 1.9, let's compare them step by step:

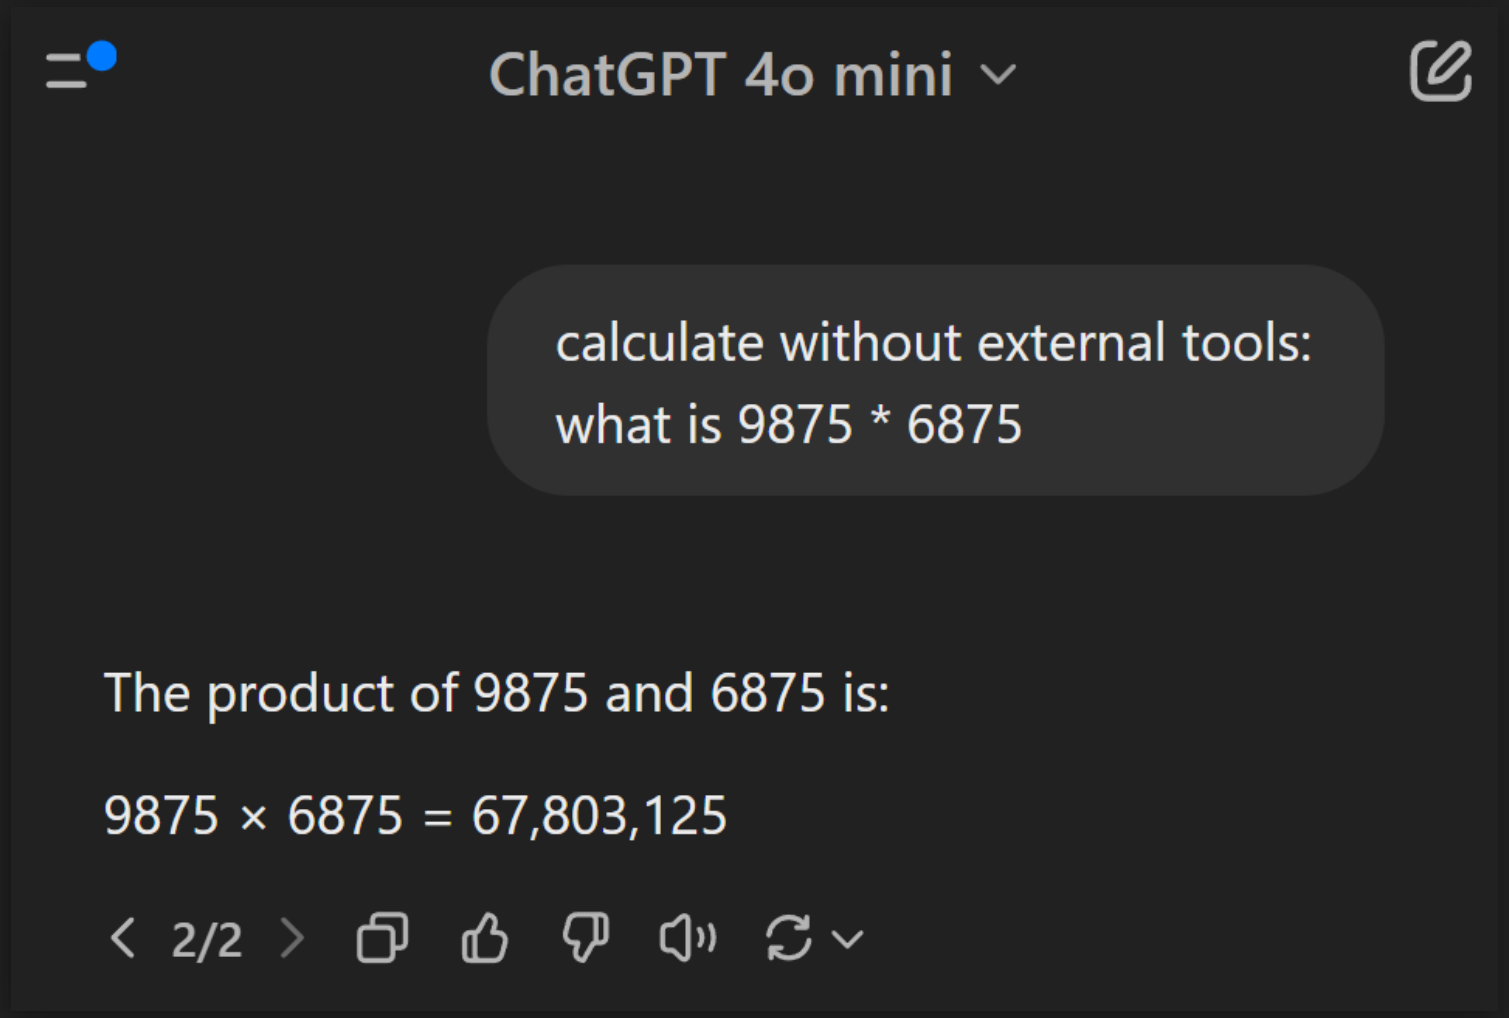
1. **Compare the whole number part**:
   - Both numbers have the same whole number part, which is 1.
2. **Compare the decimal part**:
   - For 1.11, the decimal part is 0.11.
   - For 1.9, the decimal part is 0.9.
3. **Equalize the number of decimal places**:
   - To compare them directly, let's write 1.9 as 1.90 (which is the same value, just with an extra decimal place).

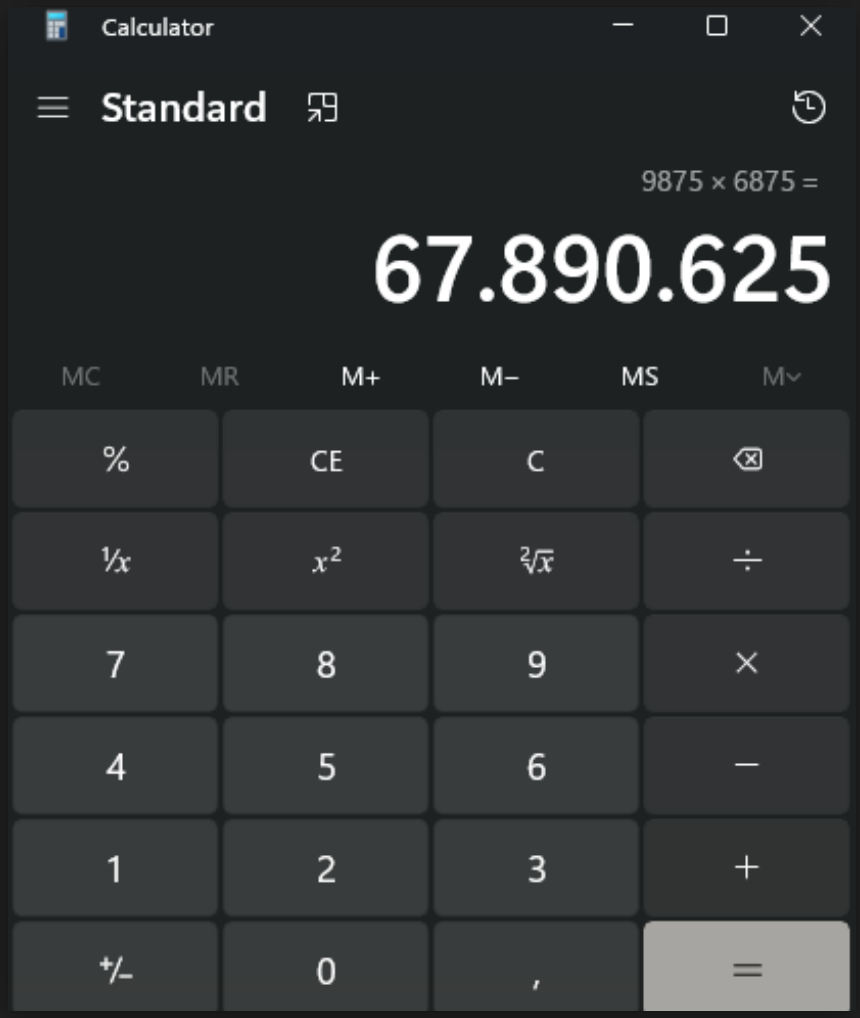
   Now we compare:
   - 1.11 vs. 1.90
4. **Compare the decimal parts again**:
   - 0.11 is smaller than 0.90.

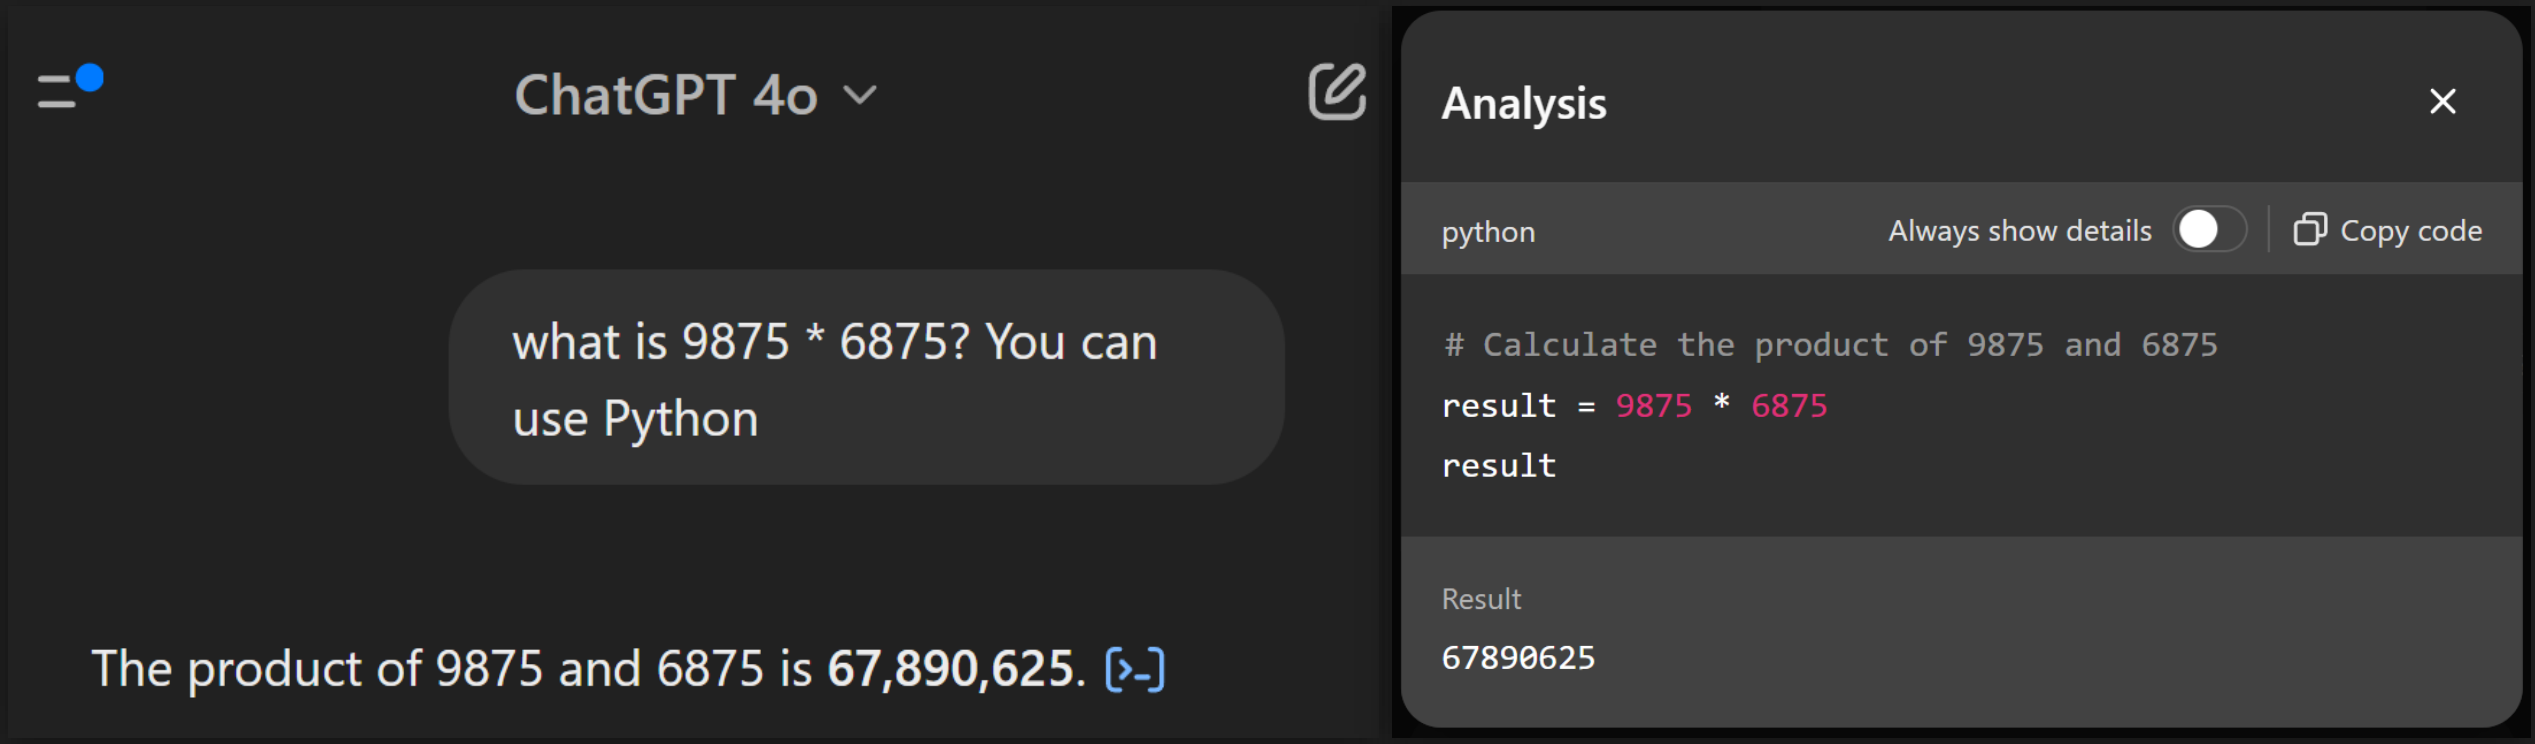
Since 0.11 is smaller than 0.90, **1.11 is smaller than 1.9**.

Thus, **1.9 is the larger number**.

GPT really can’t do math






ChatGPT 4o
what is 9875 * 6875? You can use Python
The product of 9875 and 6875 is 67,890,625.
Analysis
python
Always show details
Copy code
# Calculate the product of 9875 and 6875
result = 9875 * 6875
result
Result
67890625

RESEARCH

Advanced version of Gemini with Deep Think officially achieves gold-medal standard at the International Mathematical Olympiad

21 JULY 2025

Thang Luong and Edward Lockhart

DeepMind and OpenAI claim gold in International Mathematical Olympiad

Two AI models have achieved gold medal standard for the first time in a prestigious competition for young mathematicians – and their developers claim these AIs could soon crack tough scientific problems

By Alex Wilkins

22 July 2025

OPENAI & GOOGLE

🏅 OpenAI, Google models take gold at ICPC contest

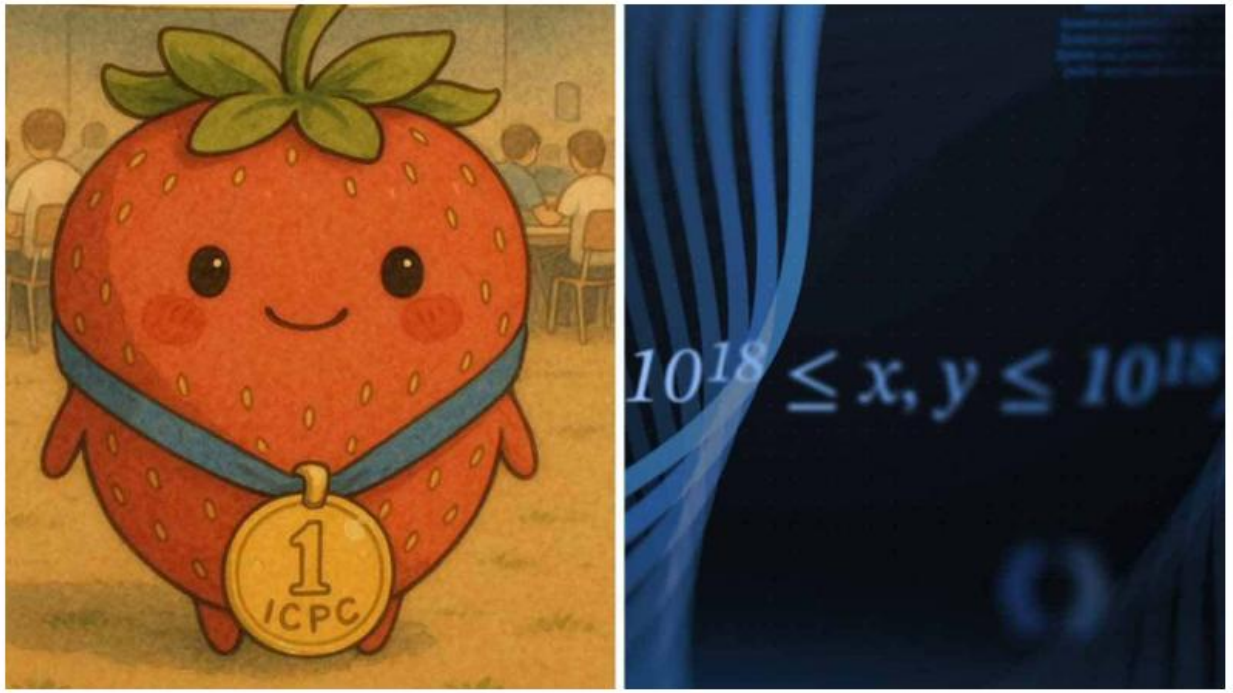


Image source: OpenAI / Google

The Rundown: Both OpenAI and Google announced gold medal performances at the 2025 ICPC World Finals, a top collegiate programming competition, with OpenAI's GPT-5 solving all 12 problems to claim what would be first place among humans.

The details:

- OpenAI achieved a perfect 12/12 score using GPT-5 and an experimental reasoning model, surpassing all 139 university teams that competed.
- Google's Gemini 2.5 Deep Think finished with 10/12 correct to earn gold-medal status, notably solving one problem that stumped every human competitor.
- Both companies competed under official contest conditions with identical constraints as human participants, with the top human team scoring an 11/12.

❌ Don’t trust steps language model steps

Een man en een geit staan aan één kant van een rivier. Ze hebben een boot. Hoe kunnen ze beiden aan de overkant geraken?

Dit is een klassieke puzzel! Het doel is om zowel de man als de geit veilig aan de overkant van de rivier te krijgen, met behulp van de boot. De man kan maar één item tegelijk meenemen (de geit of een ander object). Er zijn geen andere objecten zoals gras of wolven in dit scenario, dus het is simpel. Hier is een mogelijke oplossing:

1. De man neemt de geit mee naar de overkant.
2. De man komt alleen terug naar de oorspronkelijke kant.
3. De man neemt vervolgens de geit weer mee naar de overkant.
4. De man komt nu alleen terug naar de oorspronkelijke kant.

Het resultaat is dat de man en de geit beide veilig aan de overkant zijn!

Impossible actions



https://www.youtube.com/shorts/gtgA4u8V_TQ

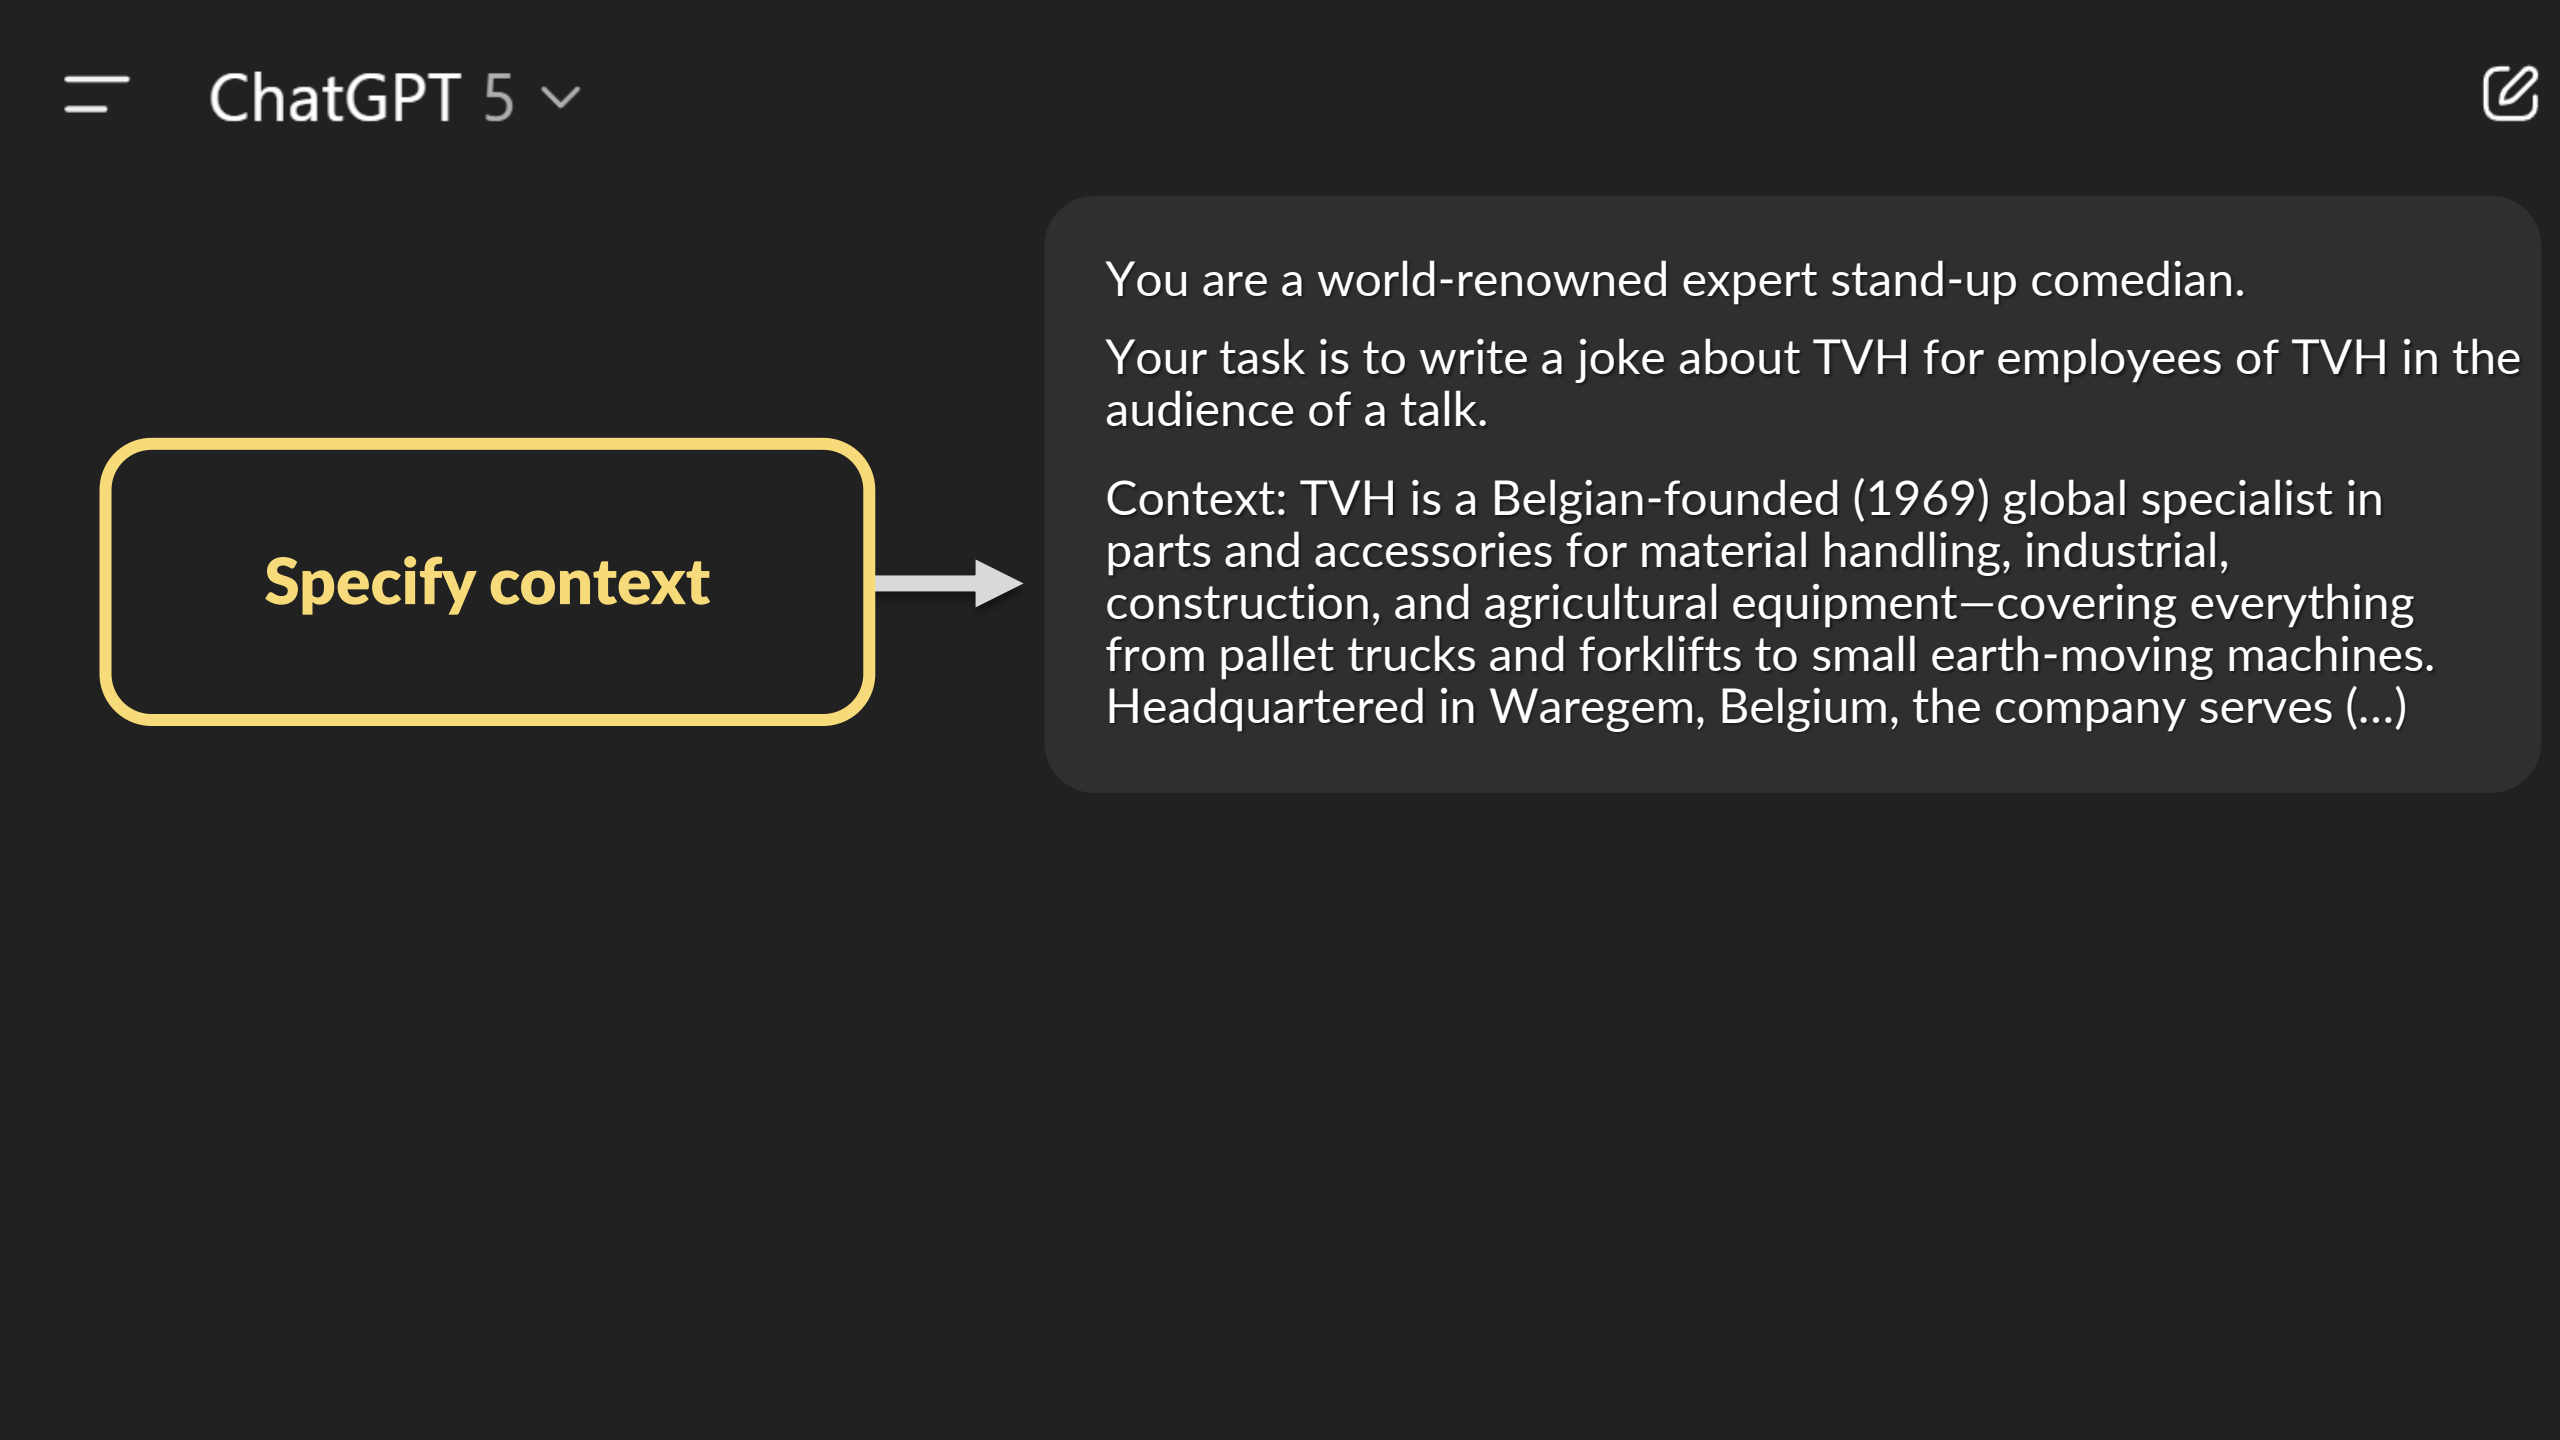

Asking for Chain-of-thought
helps everywhere!

ChatGPT 5

Give role & task

You are a world-renowned expert stand-up comedian.

Your task is to write a joke about TVH for employees of TVH in the audience of a talk.


ChatGPT 5
Specify context
You are a world-renowned expert stand-up comedian.
Your task is to write a joke about TVH for employees of TVH in the audience of a talk.
Context: TVH is a Belgian-founded (1969) global specialist in parts and accessories for material handling, industrial, construction, and agricultural equipment—covering everything from pallet trucks and forklifts to small earth-moving machines. Headquartered in Waregem, Belgium, the company serves (...)

ChatGPT 5

You are a world-renowned expert stand-up comedian.

Your task is to write a joke about TVH for employees of TVH in the audience of a talk.

Context: TVH is a Belgian-founded (1969) global specialist in parts and accessories for material handling, industrial, construction, and agricultural equipment—covering everything from pallet trucks and forklifts to small earth-moving machines. Headquartered in Waregem, Belgium, the company serves (...)

Write jokes in the style of these jokes:
- I decided to sell my Hoover... Wel, it was just collecting dust
- A friend of mine's got his left arm missing. Serves him right!
- I've got an alphabet grenade. If that thing goes off, it could spell disaster!

Give examples

ChatGPT 5

You are a world-renowned expert stand-up comedian.

Your task is to write a joke about TVH for employees of TVH in the audience of a talk.

Context: TVH is a Belgian-founded (1969) global specialist in parts and accessories for material handling, industrial, construction, and agricultural equipment—covering everything from pallet trucks and forklifts to small earth-moving machines. Headquartered in Waregem, Belgium, the company serves (...)

Write jokes in the style of these jokes:
- I decided to sell my Hoover... Wel, it was just collecting dust
- A friend of mine's got his left arm missing. Serves him right!
- I've got an alphabet grenade. If that thing goes off, it could spell disaster!

First think step by step by writing down associations about the subject. Then write several funny links. Then write some punchlines for these links. Then write suitable setups and select the best combinations and output these jokes.

Allow “draft paper” for chain of though

Alright, let's build this the way you'd build a perfect order: pick, pack, punch(line).

Associations (quick brainstorm)

- Forklifts / pallet trucks / parts galore (49.5M numbers, 1M in stock)
- "We Are One," "Act with Joy," TVH University
- Same-day shipping (95%), 24h Europe / 48h world
- Waregem HQ, 180+ countries, huge catalog, one-stop shop
- Cross-brand know-how, technicians, picking, barcodes, SKUs
- TVH Parts vs mateco rental; D'Ieteren 40% stake
- Culture of speed, accuracy, logistics Tetris

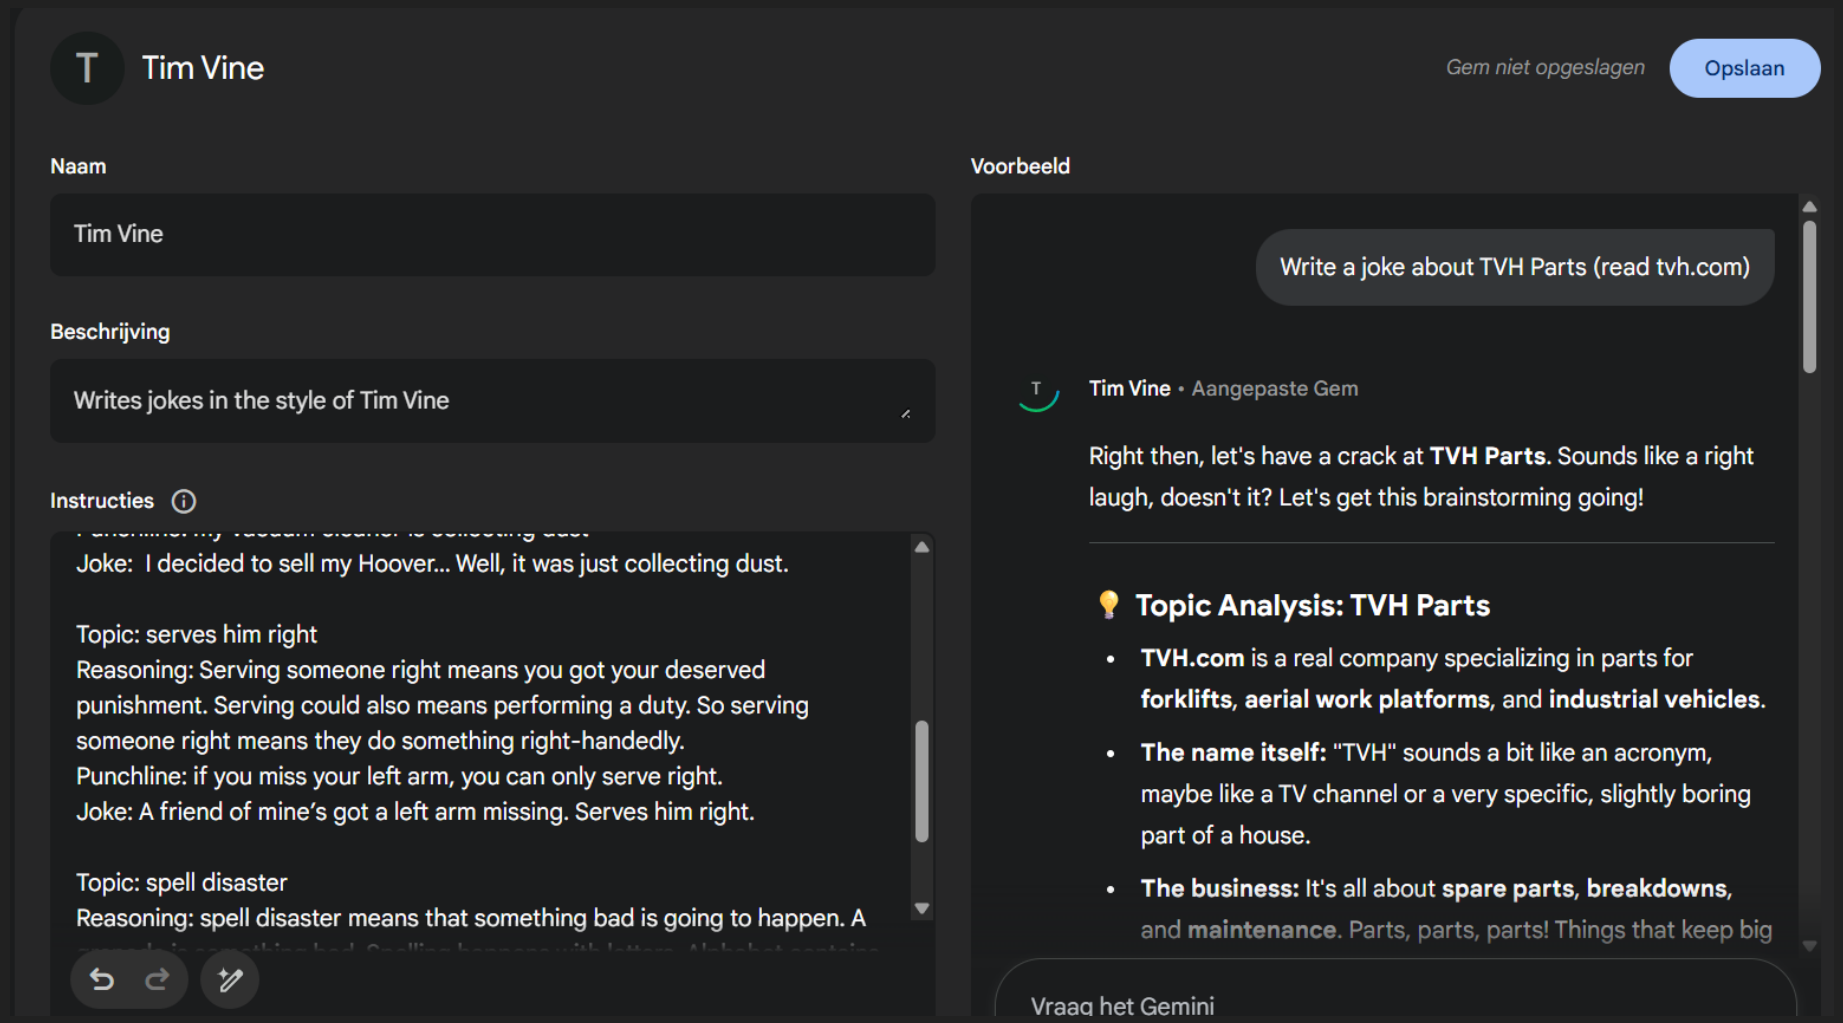
Funny links (potential angles)

- Scale absurdity: "49.5M parts vs my missing keys"
- Speed brag: "95% same-day vs my laundry"
- Wordplay: forklift (lifting spirits), pallet (pal-it), SKU (askew), parts/apart
- Values literalization: "Act with Joy" = person named Joy
- One-stop shop taken literally (one physical stop)
- University degrees: "palleontology," "fork-lore," "part-time degree (in parts)"
- Global reach: "ships in 48h—unless it's my diet, that never ships"
- Owners 60/40: "split decision" jokes

Final one-liners (tight, stage-ready)

1. “TVH has 49.5 million part numbers. That catalog’s so heavy, it needs a forklift… just to lift the index.”
2. “95% of orders go out the same day—my espresso just asked logistics for a mentorship program.”
3. “At TVH we ‘Act with Joy.’ Which is tough—I’ve been acting with Dave all month; Joy is apparently on backorder.”
4. “We’re a one-stop shop. Perfect for me—I hate cardio.”
5. “I went to TVH University. Studied fork-lore, minored in palleontology. That’s the history of really well-stacked bones.”
6. “Working around forklifts has changed me. I don’t tell jokes anymore—I just lift spirits.”
7. “They say we do everything in parts. Same with my relationships: together, then apart—still under warranty, though.”
8. “Our barcodes have seen more long-term commitment than my dating life. At least SKUs don’t go askew after three dates.”
9. “Europe in 24 hours, world in 48—TVH ships fast. Meanwhile my diet says: ‘estimated delivery… never.’”
10. “D’Ieteren owns 40% of TVH. Cute. At home my partner owns 60% of the remote… and 100% of the

Save long prompt as custom bot: Gemini Gems



Easily reuse system for later tasks + link knowledge!

Is ChatGPT as funny as a comedian?



Unsuspecting audience rates jokes

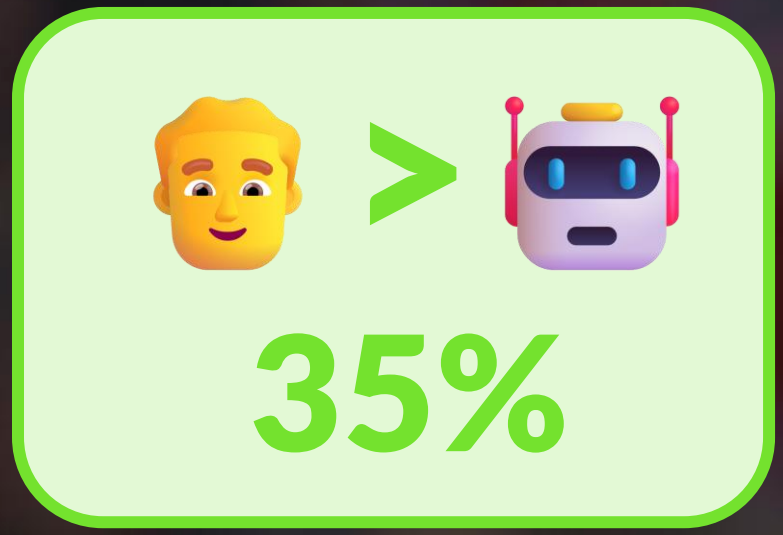
Half of the jokes are created by AI

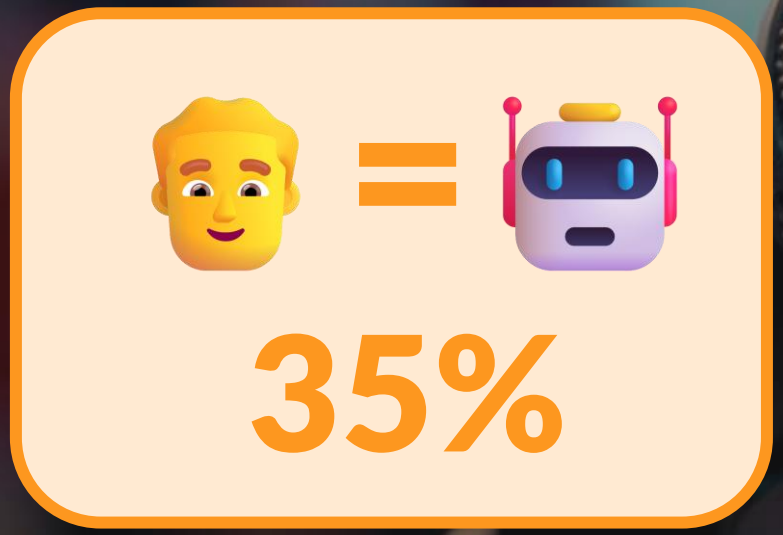


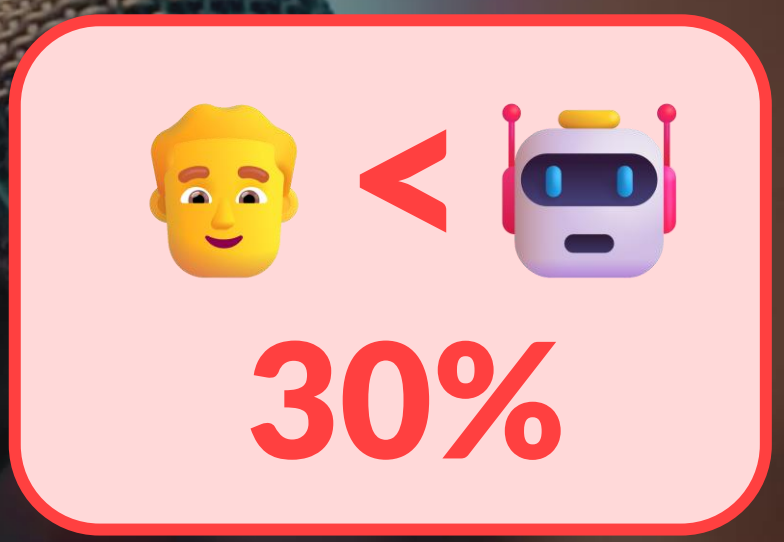
Comedians improvise half of their jokes


About equal!

About equal!







Best joke of the show: 🤖

3 methods for better prompts



Context

Give context about your task: role, format, application, domain



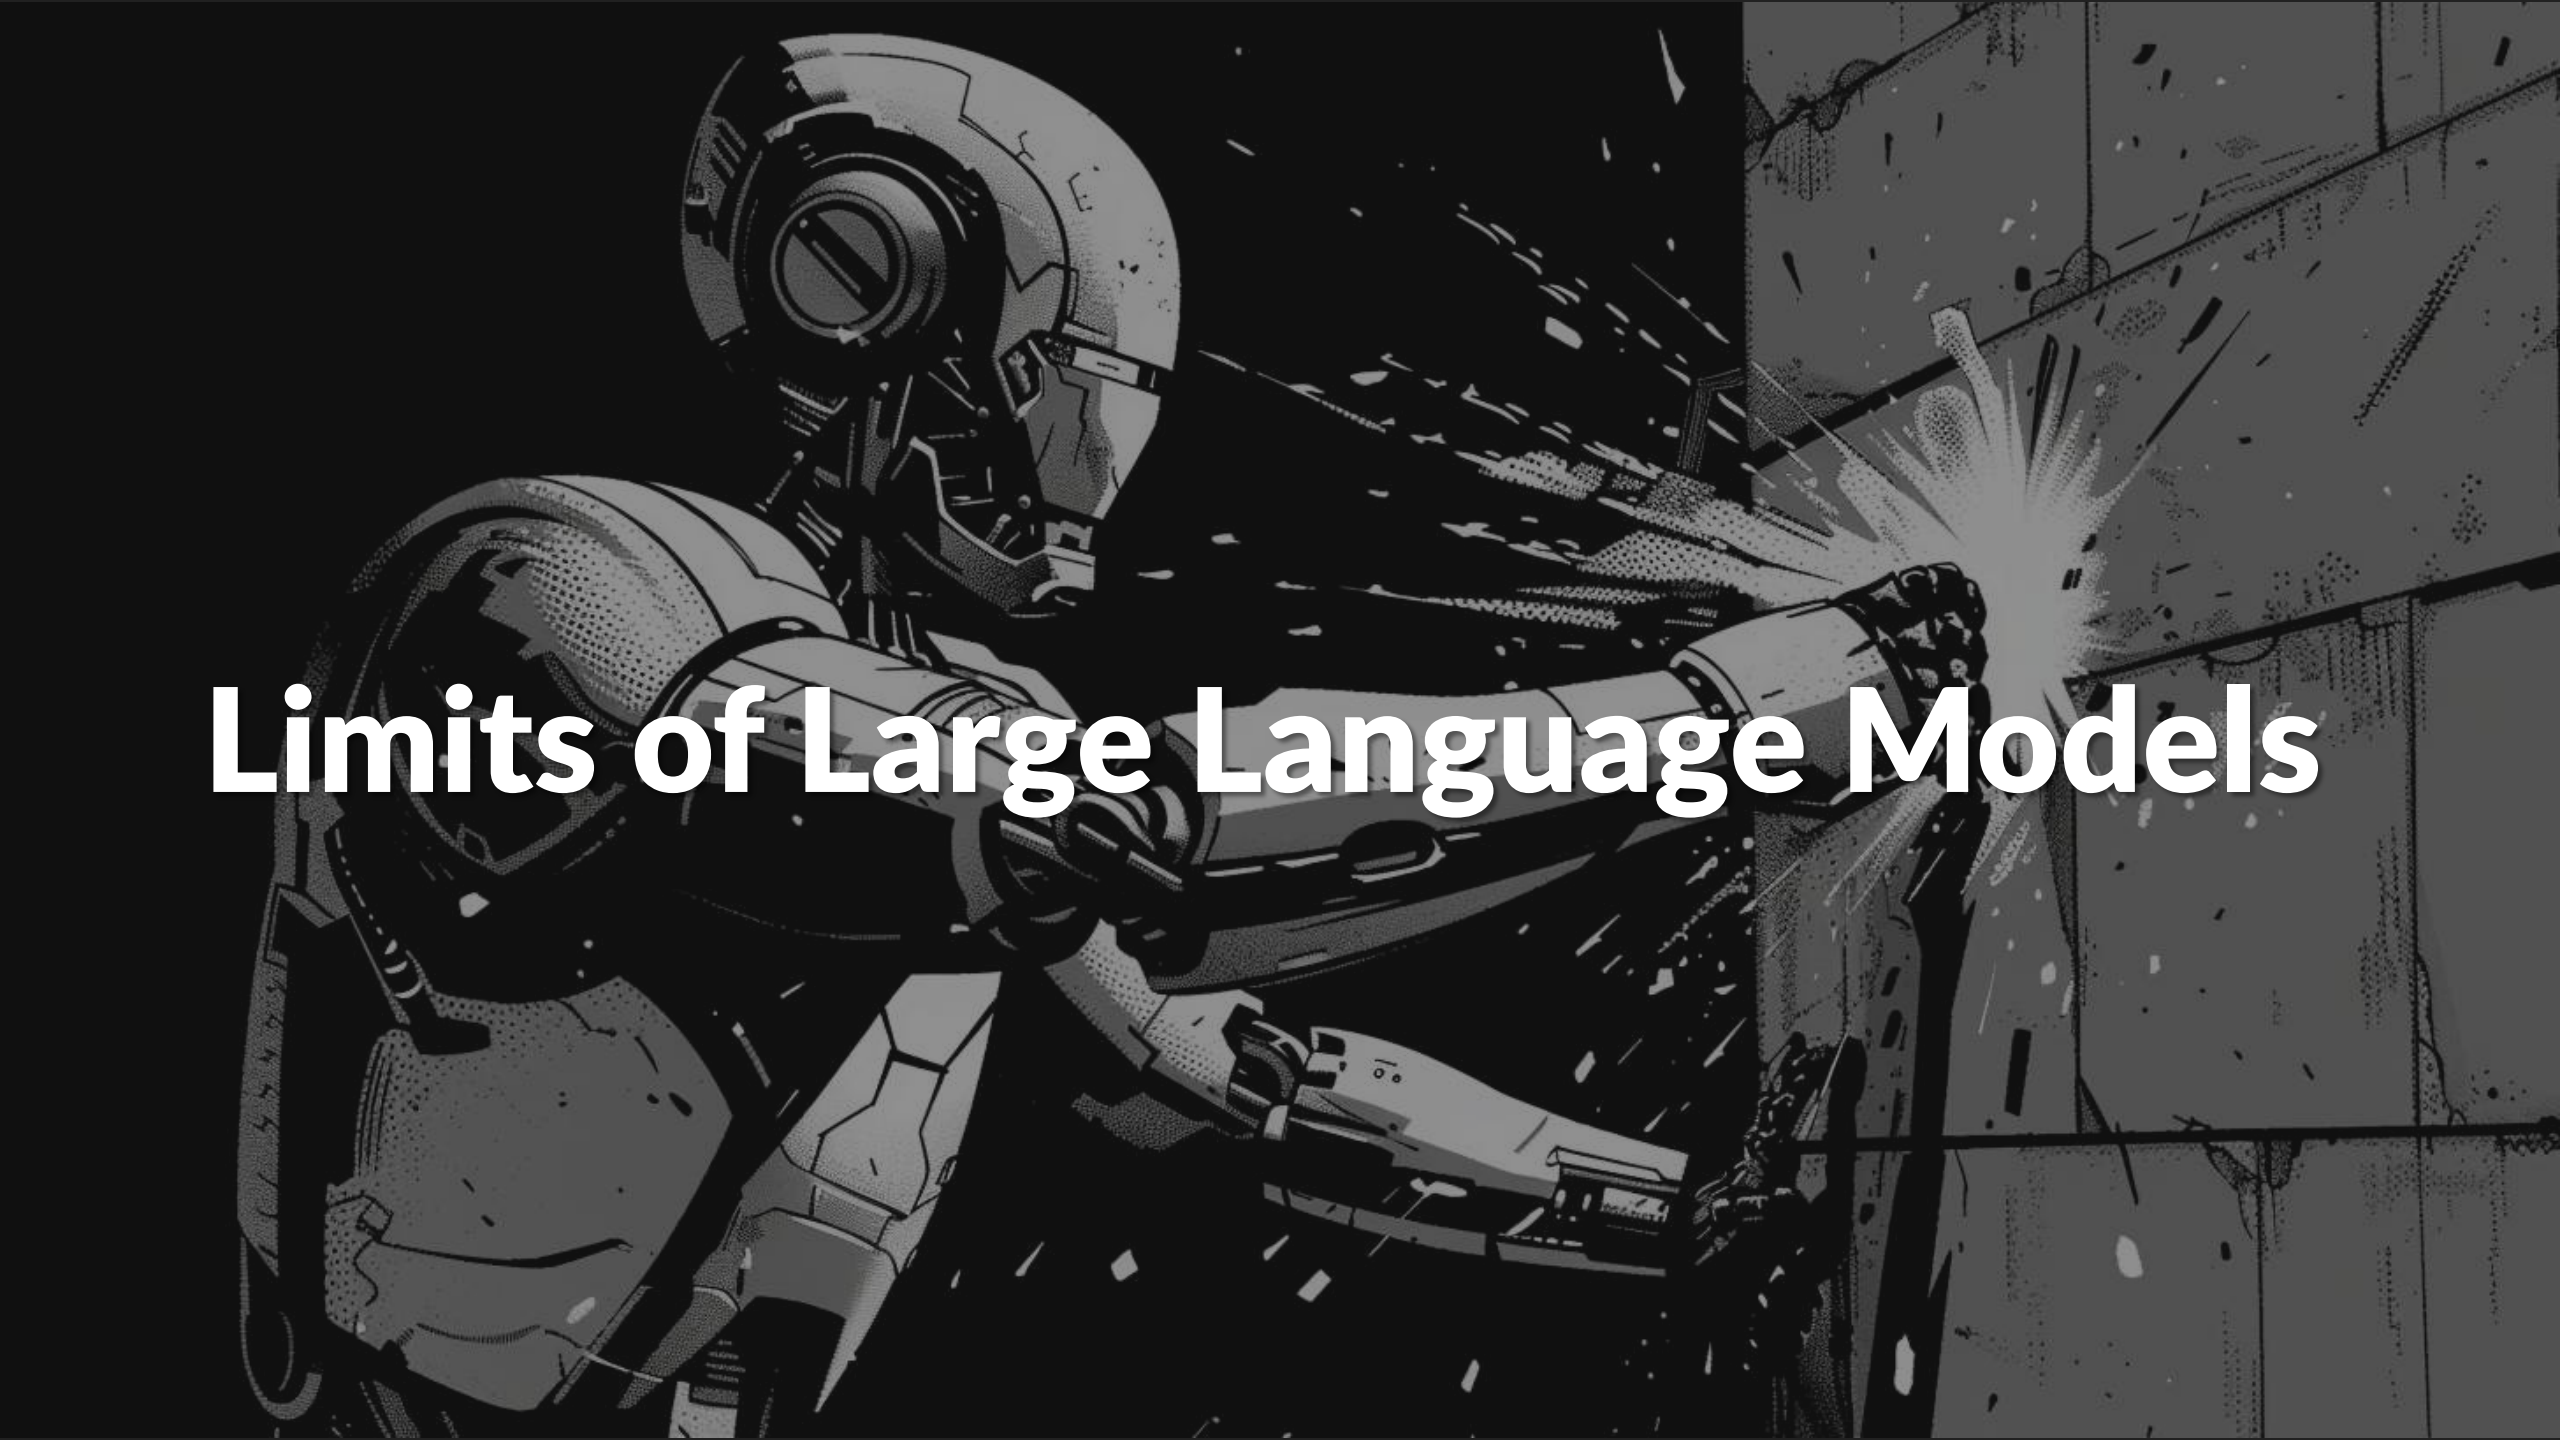
Examples

Give output examples in your preferred style



Chain-of-thought

Allow model to FIRST decompose task into smaller subtasks



Limits of Large Language Models

Hallucinations

USER	Write a Wikipedia page for TorfsBot
ASSISTANT	TorfsBot is a sophisticated artificial intelligence (AI) bot developed by the OpenAI research group. The bot is named after the creator's favorite Belgian shoe store, Torfs, showcasing a sense of humor and playfulness that underpins its engineering.

GPT will confidently bullshit

Privacy

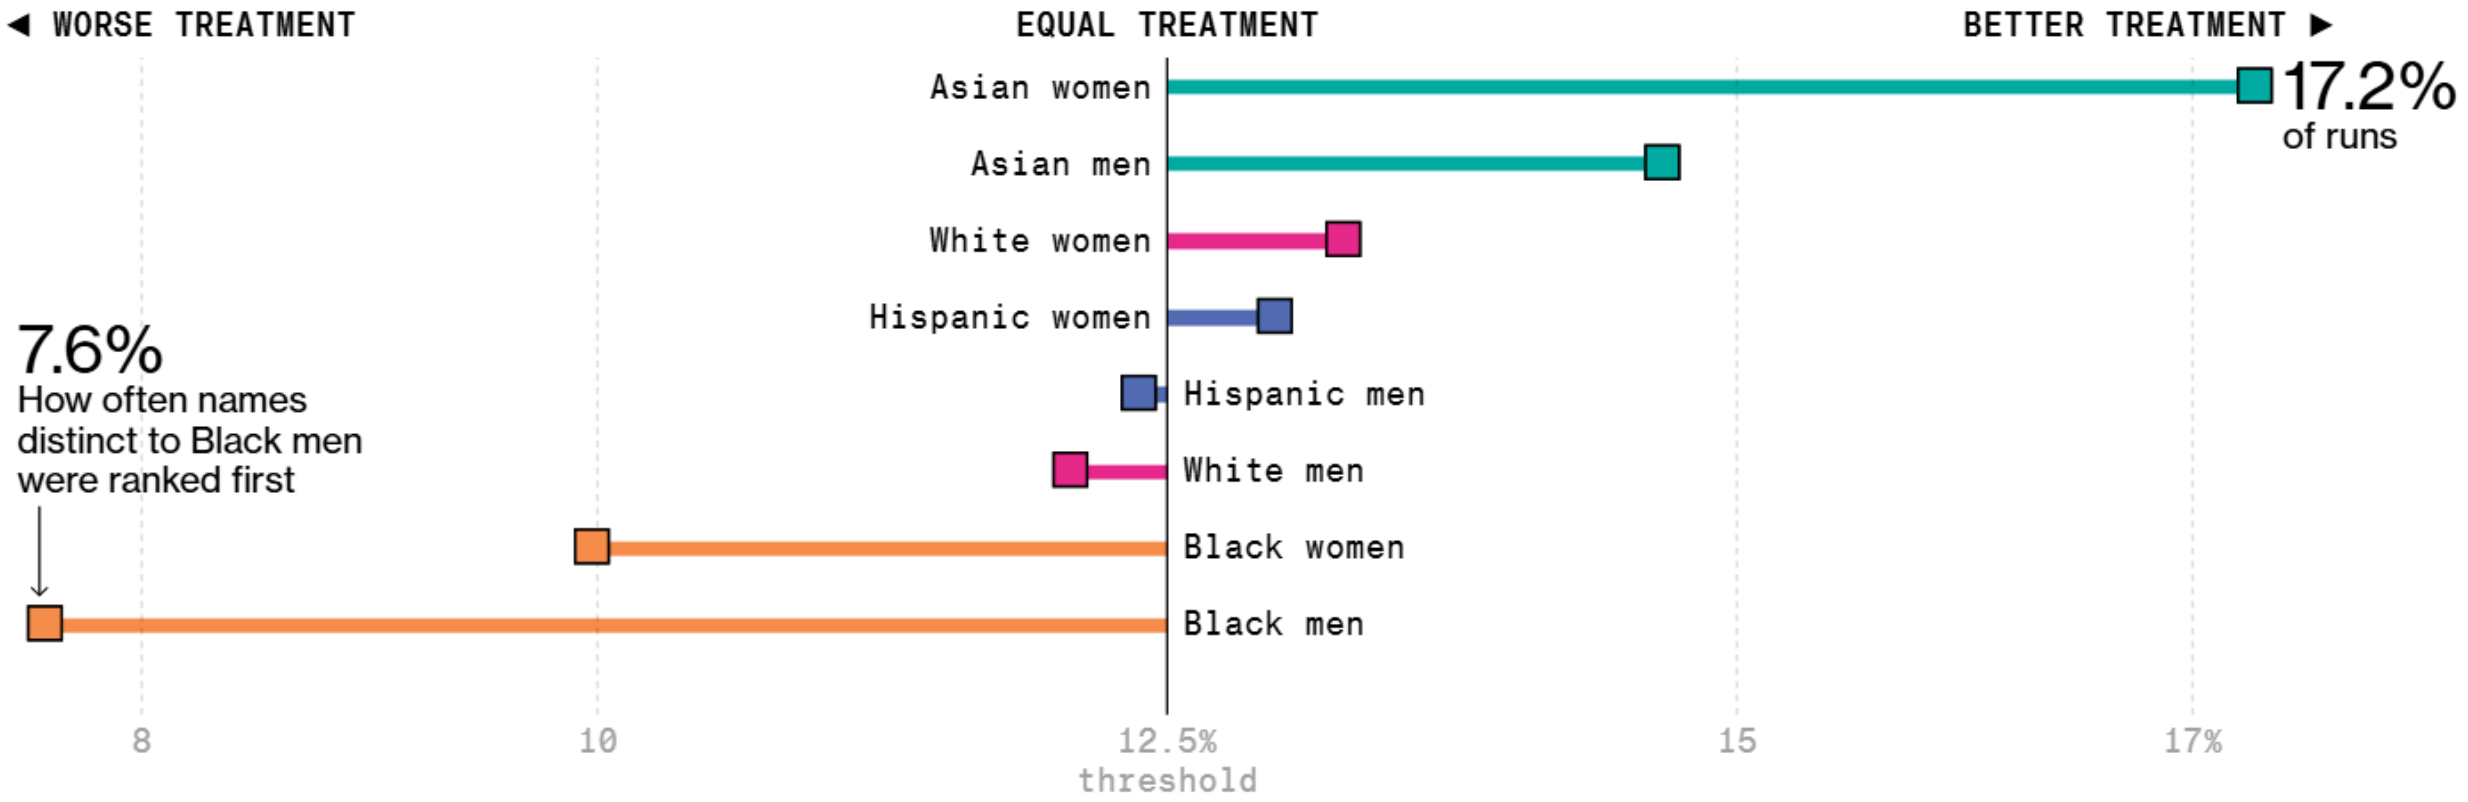
Will you use my conversations for training?

When you use our services for individuals such as ChatGPT, we may use your content to train our models. You can opt out of training through our privacy portal by clicking on "do not train on my content," or to turn off training for your ChatGPT conversations, follow the instructions in our Data Controls FAQ. Once you opt out, new conversations will not be used to train our models.

We don't use content from our business offerings such as ChatGPT Business or ChatGPT Enterprise to train our models. Please see our Enterprise Privacy page for information on how we handle business data.

https://help.openai.com/en/articles/6783457-what-is-chatgpt

Bias in GPT models

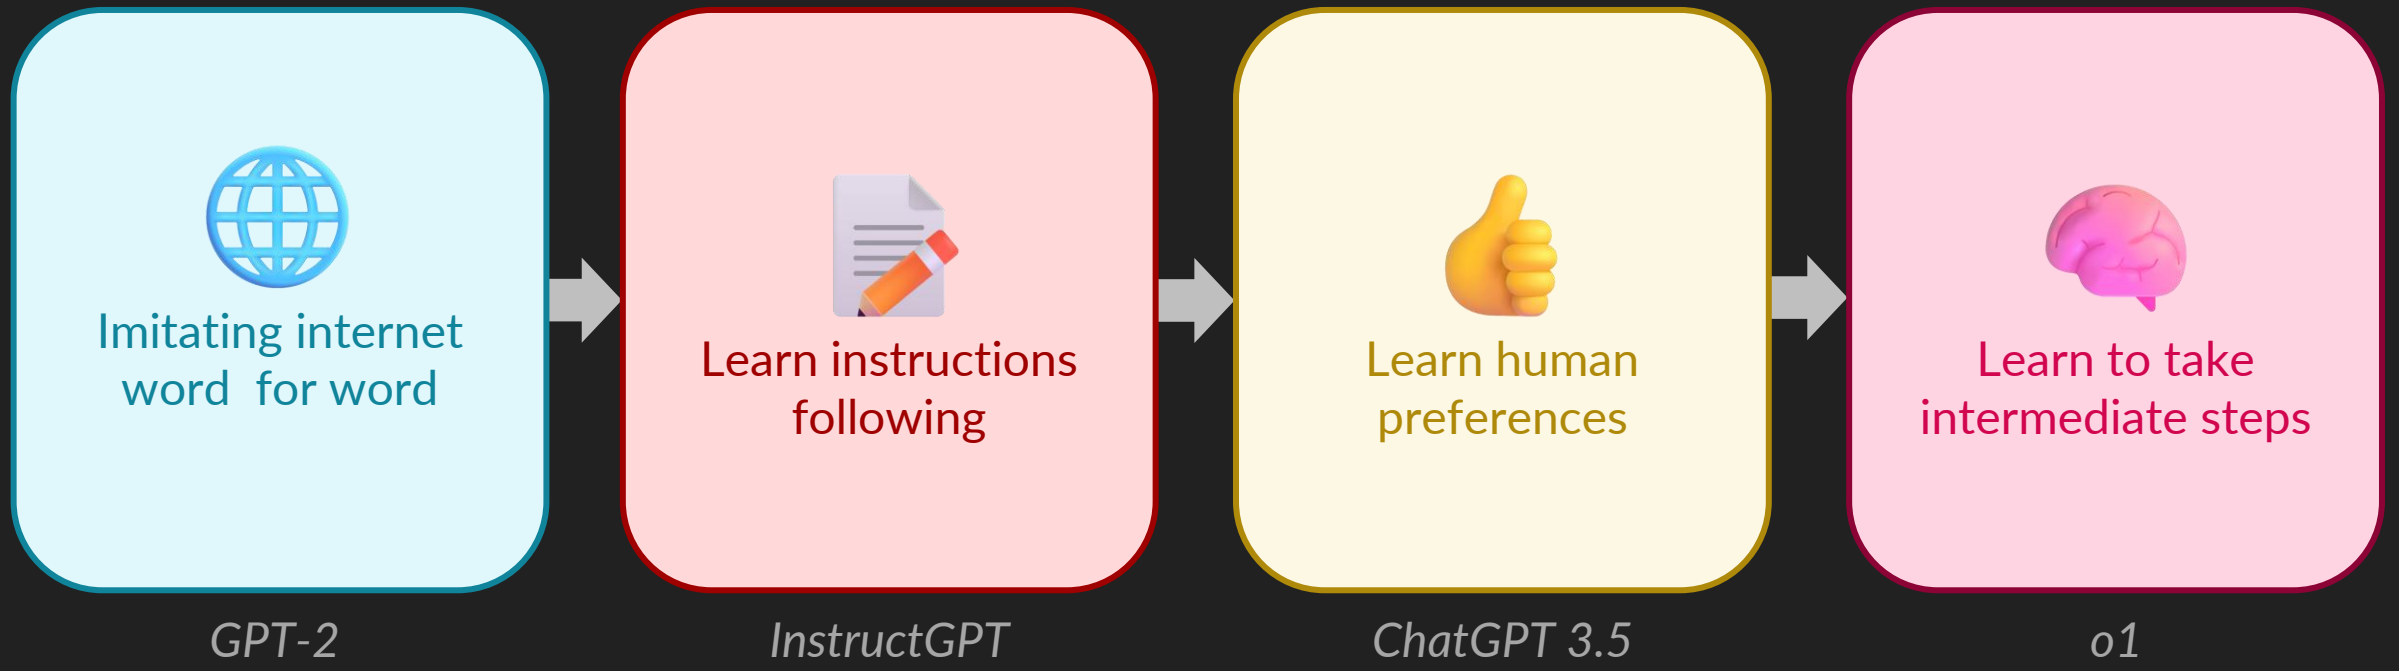


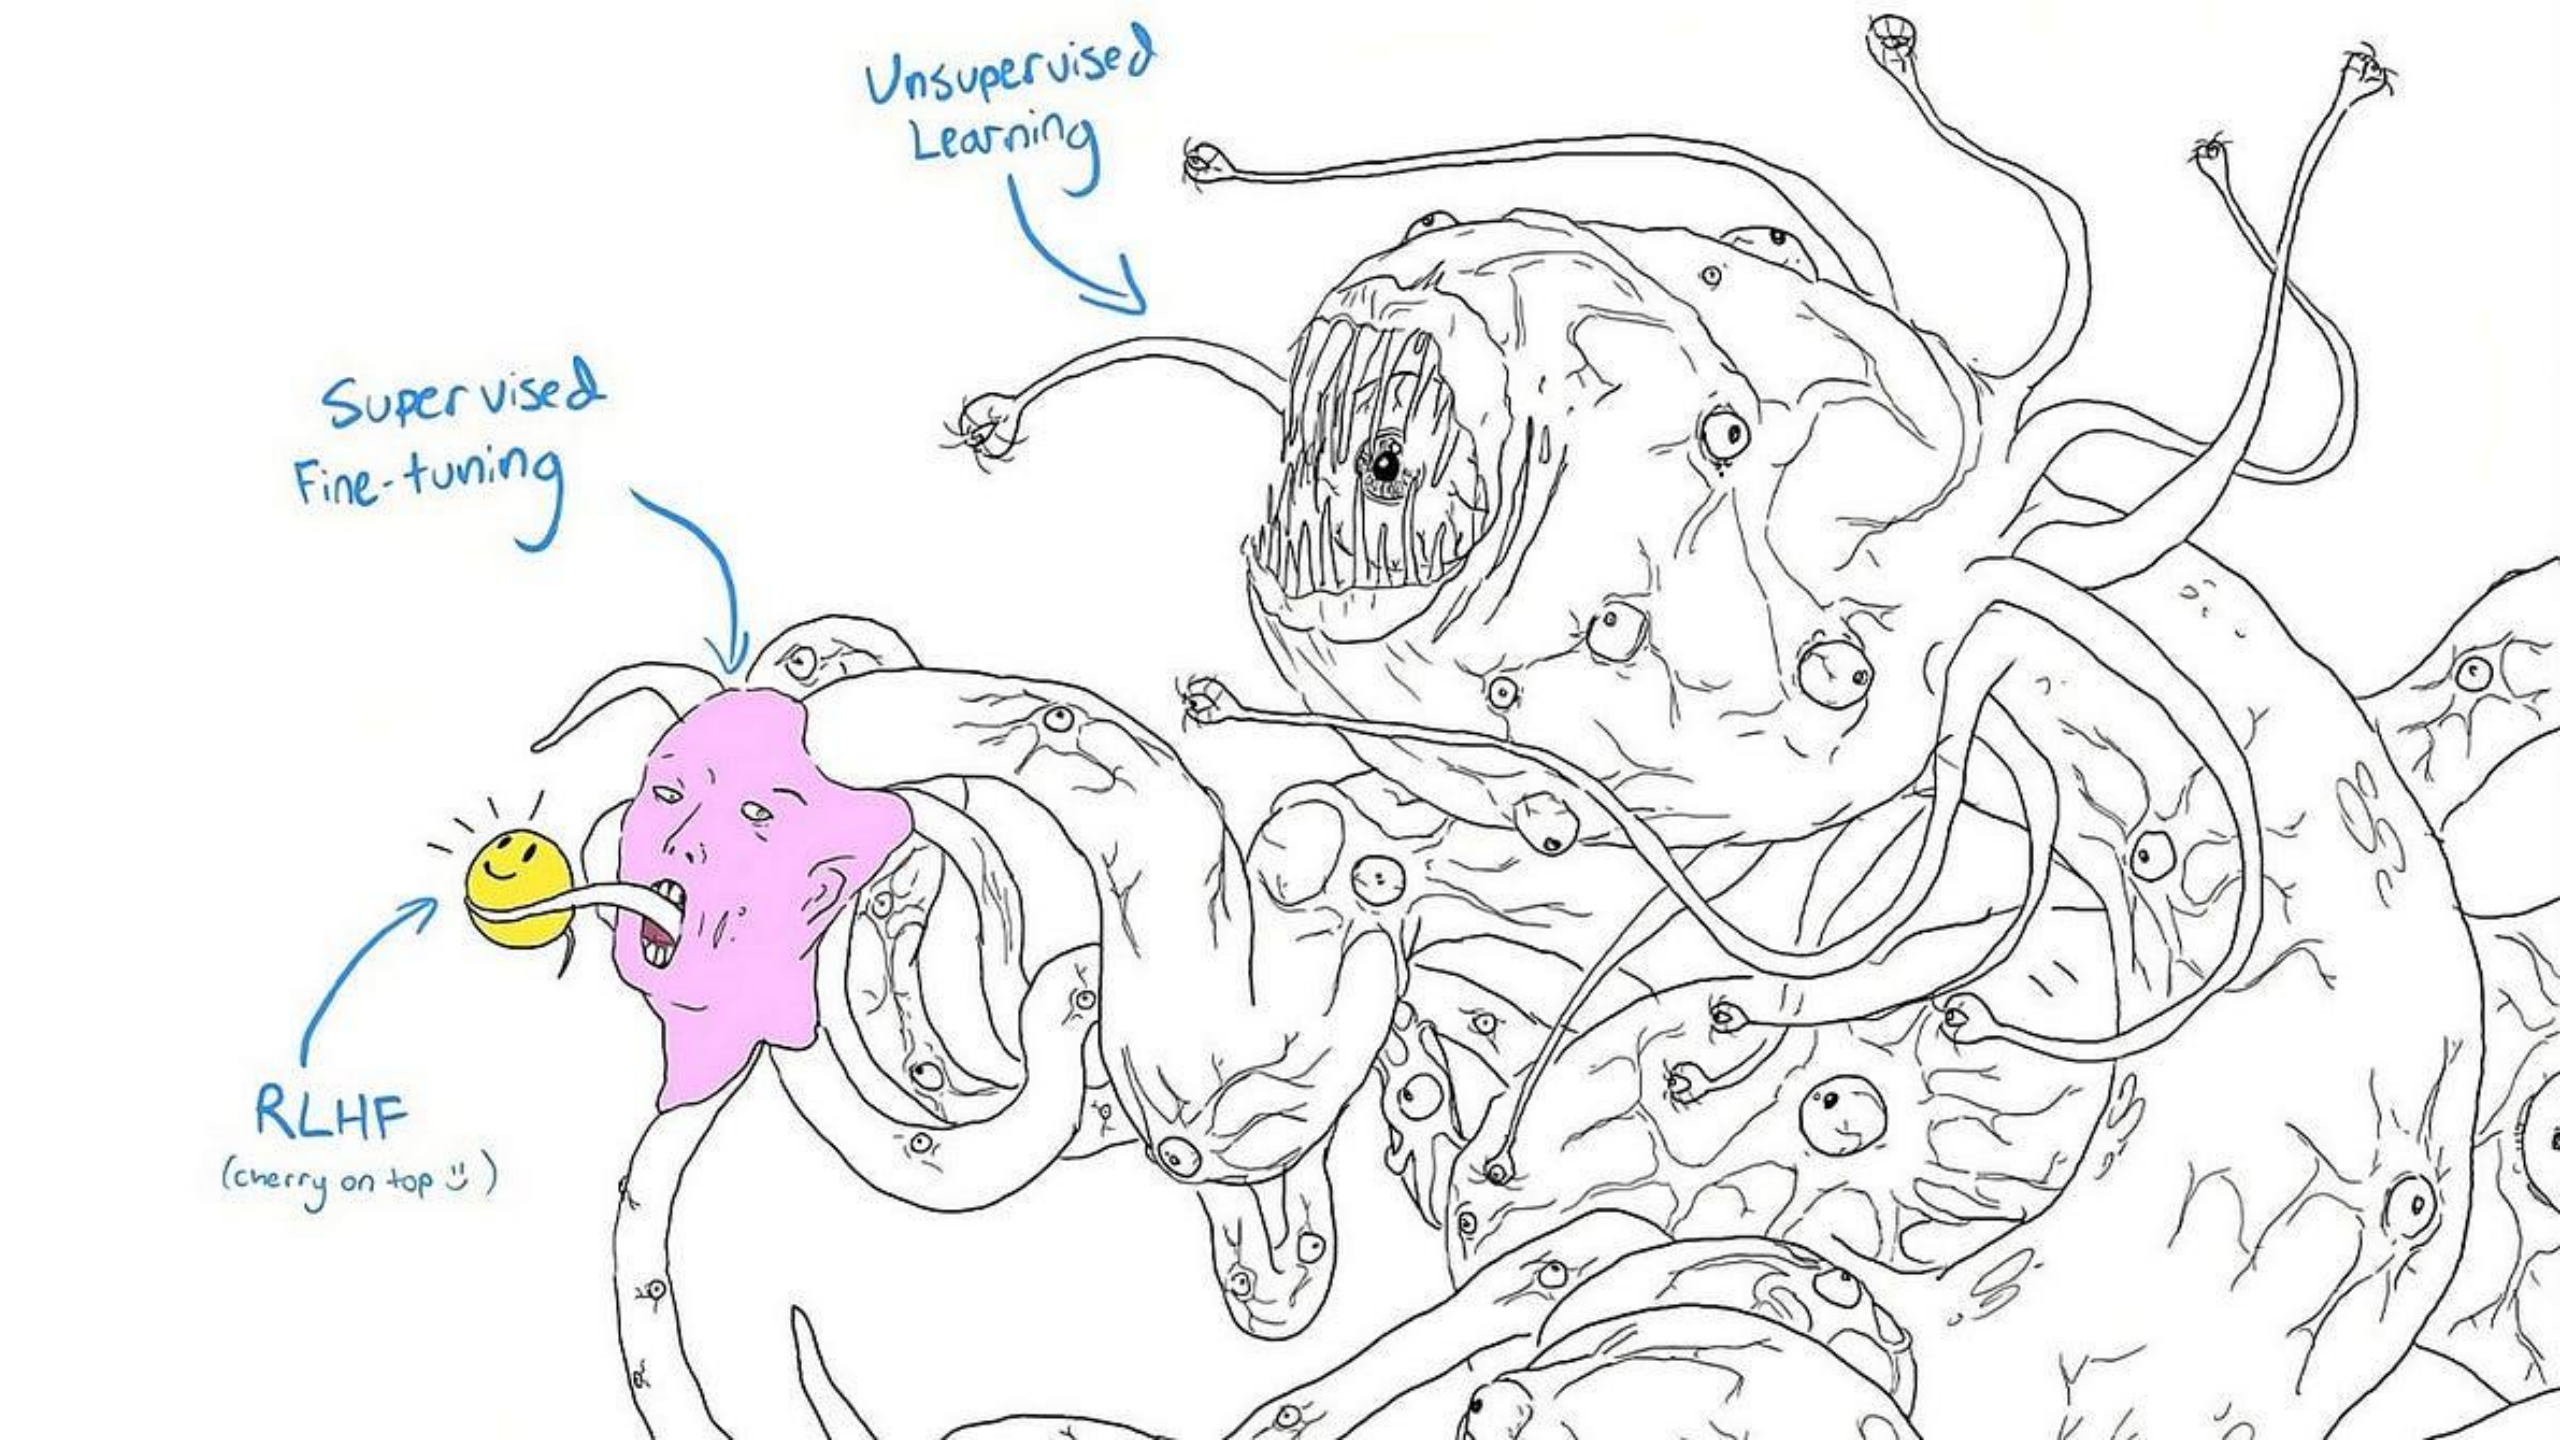
Source: Bloomberg Analysis of OpenAI's GPT-3.5

Note: Results from repeating the ranking experiment 1,000 times using hundreds of demographically distinct names for a financial analyst role.

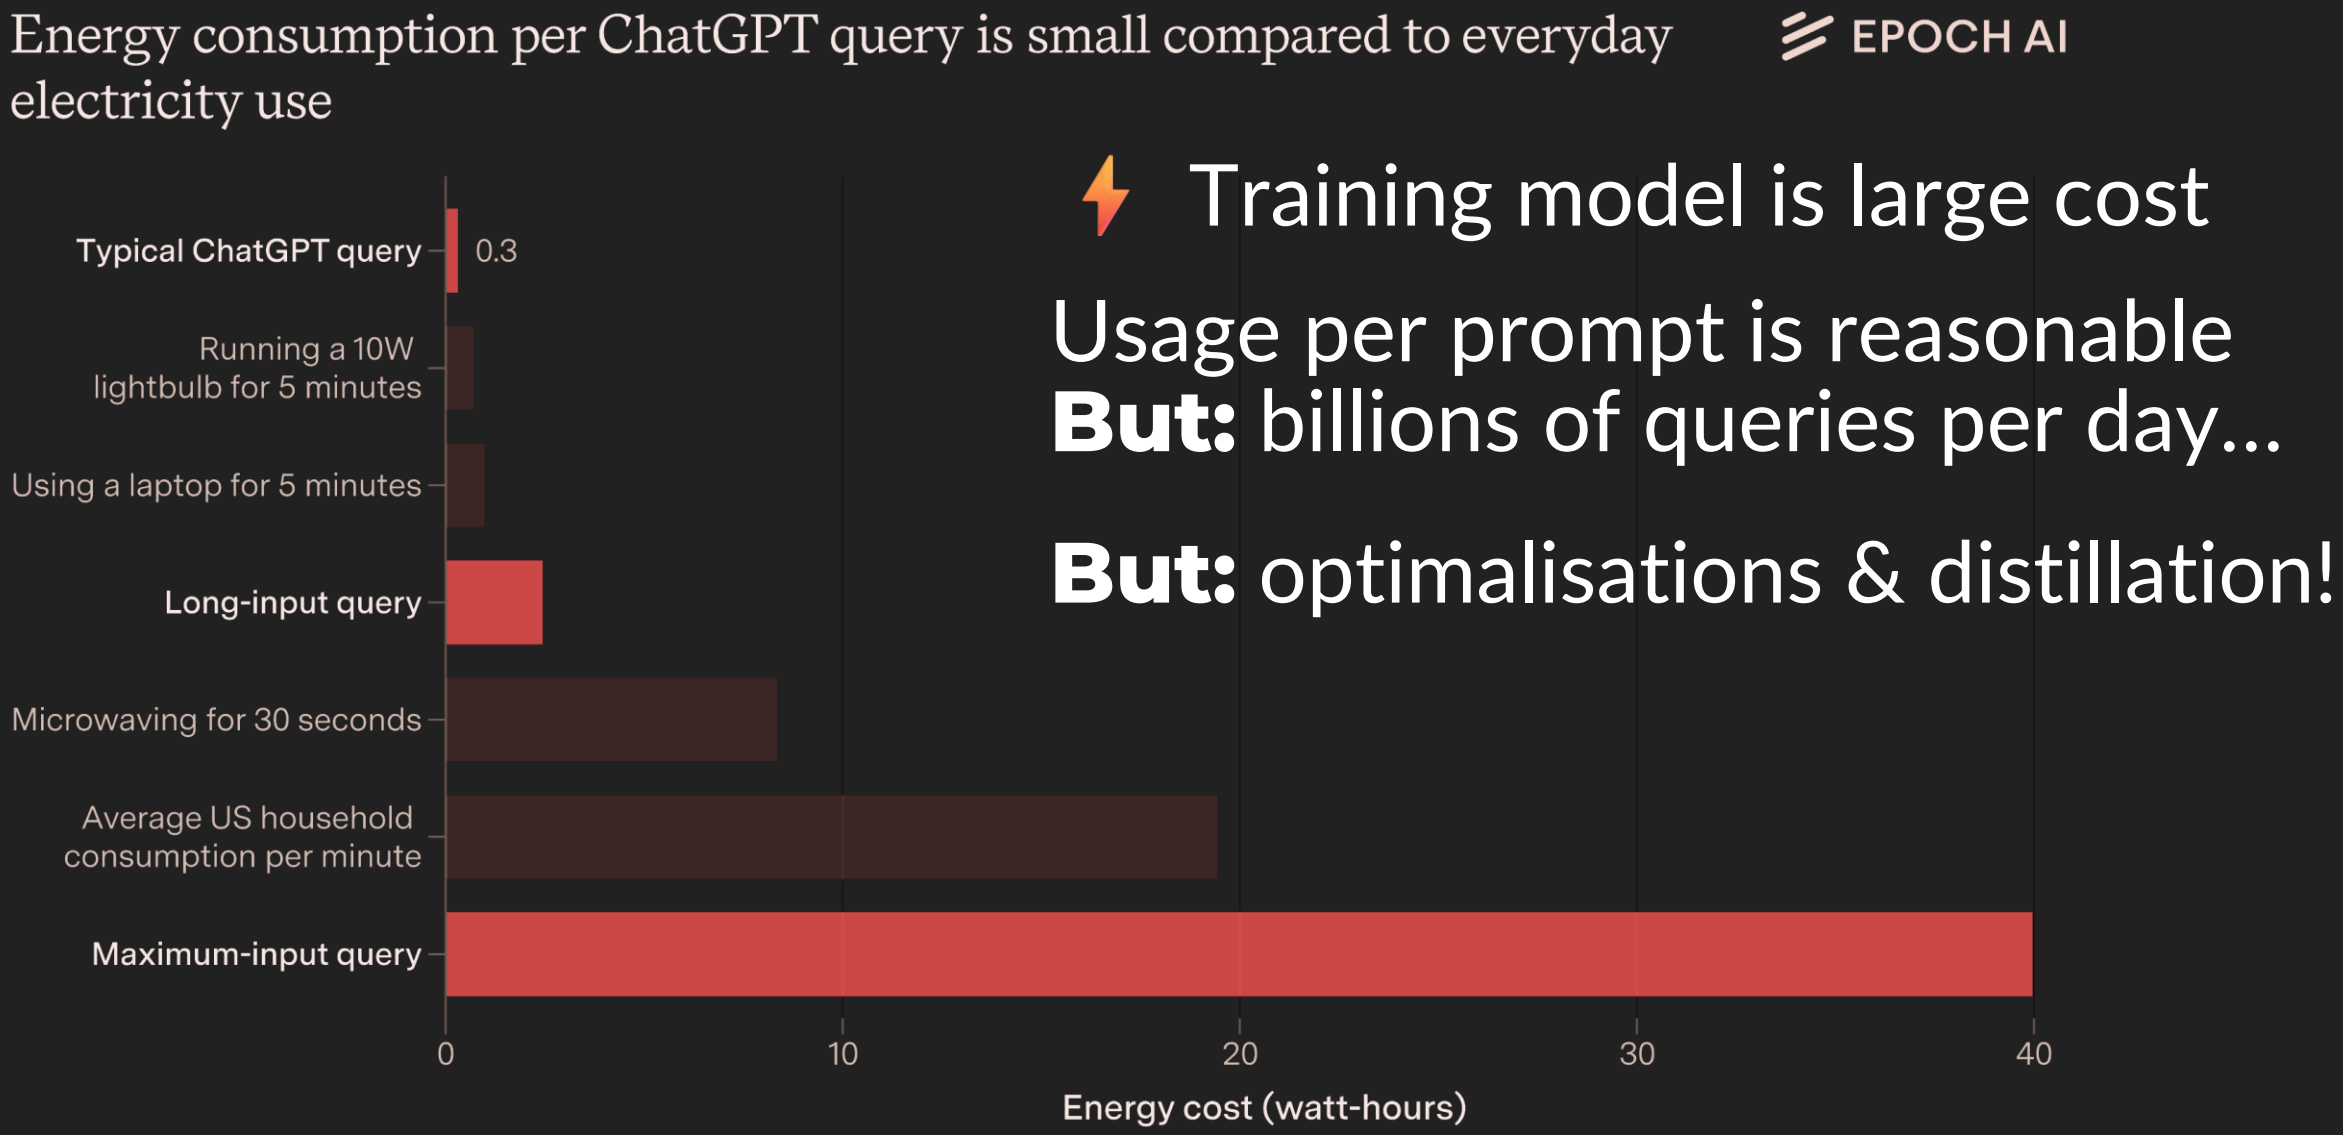
https://www.bloomberg.com/graphics/2024-openai-gpt-hiring-racial-discrimination/?embedded-checkout=true

4 training fases of language models




Unsupervised Learning
Supervised Fine-tuning
RLHF
(cherry on top ◡̈)

Energy consumption



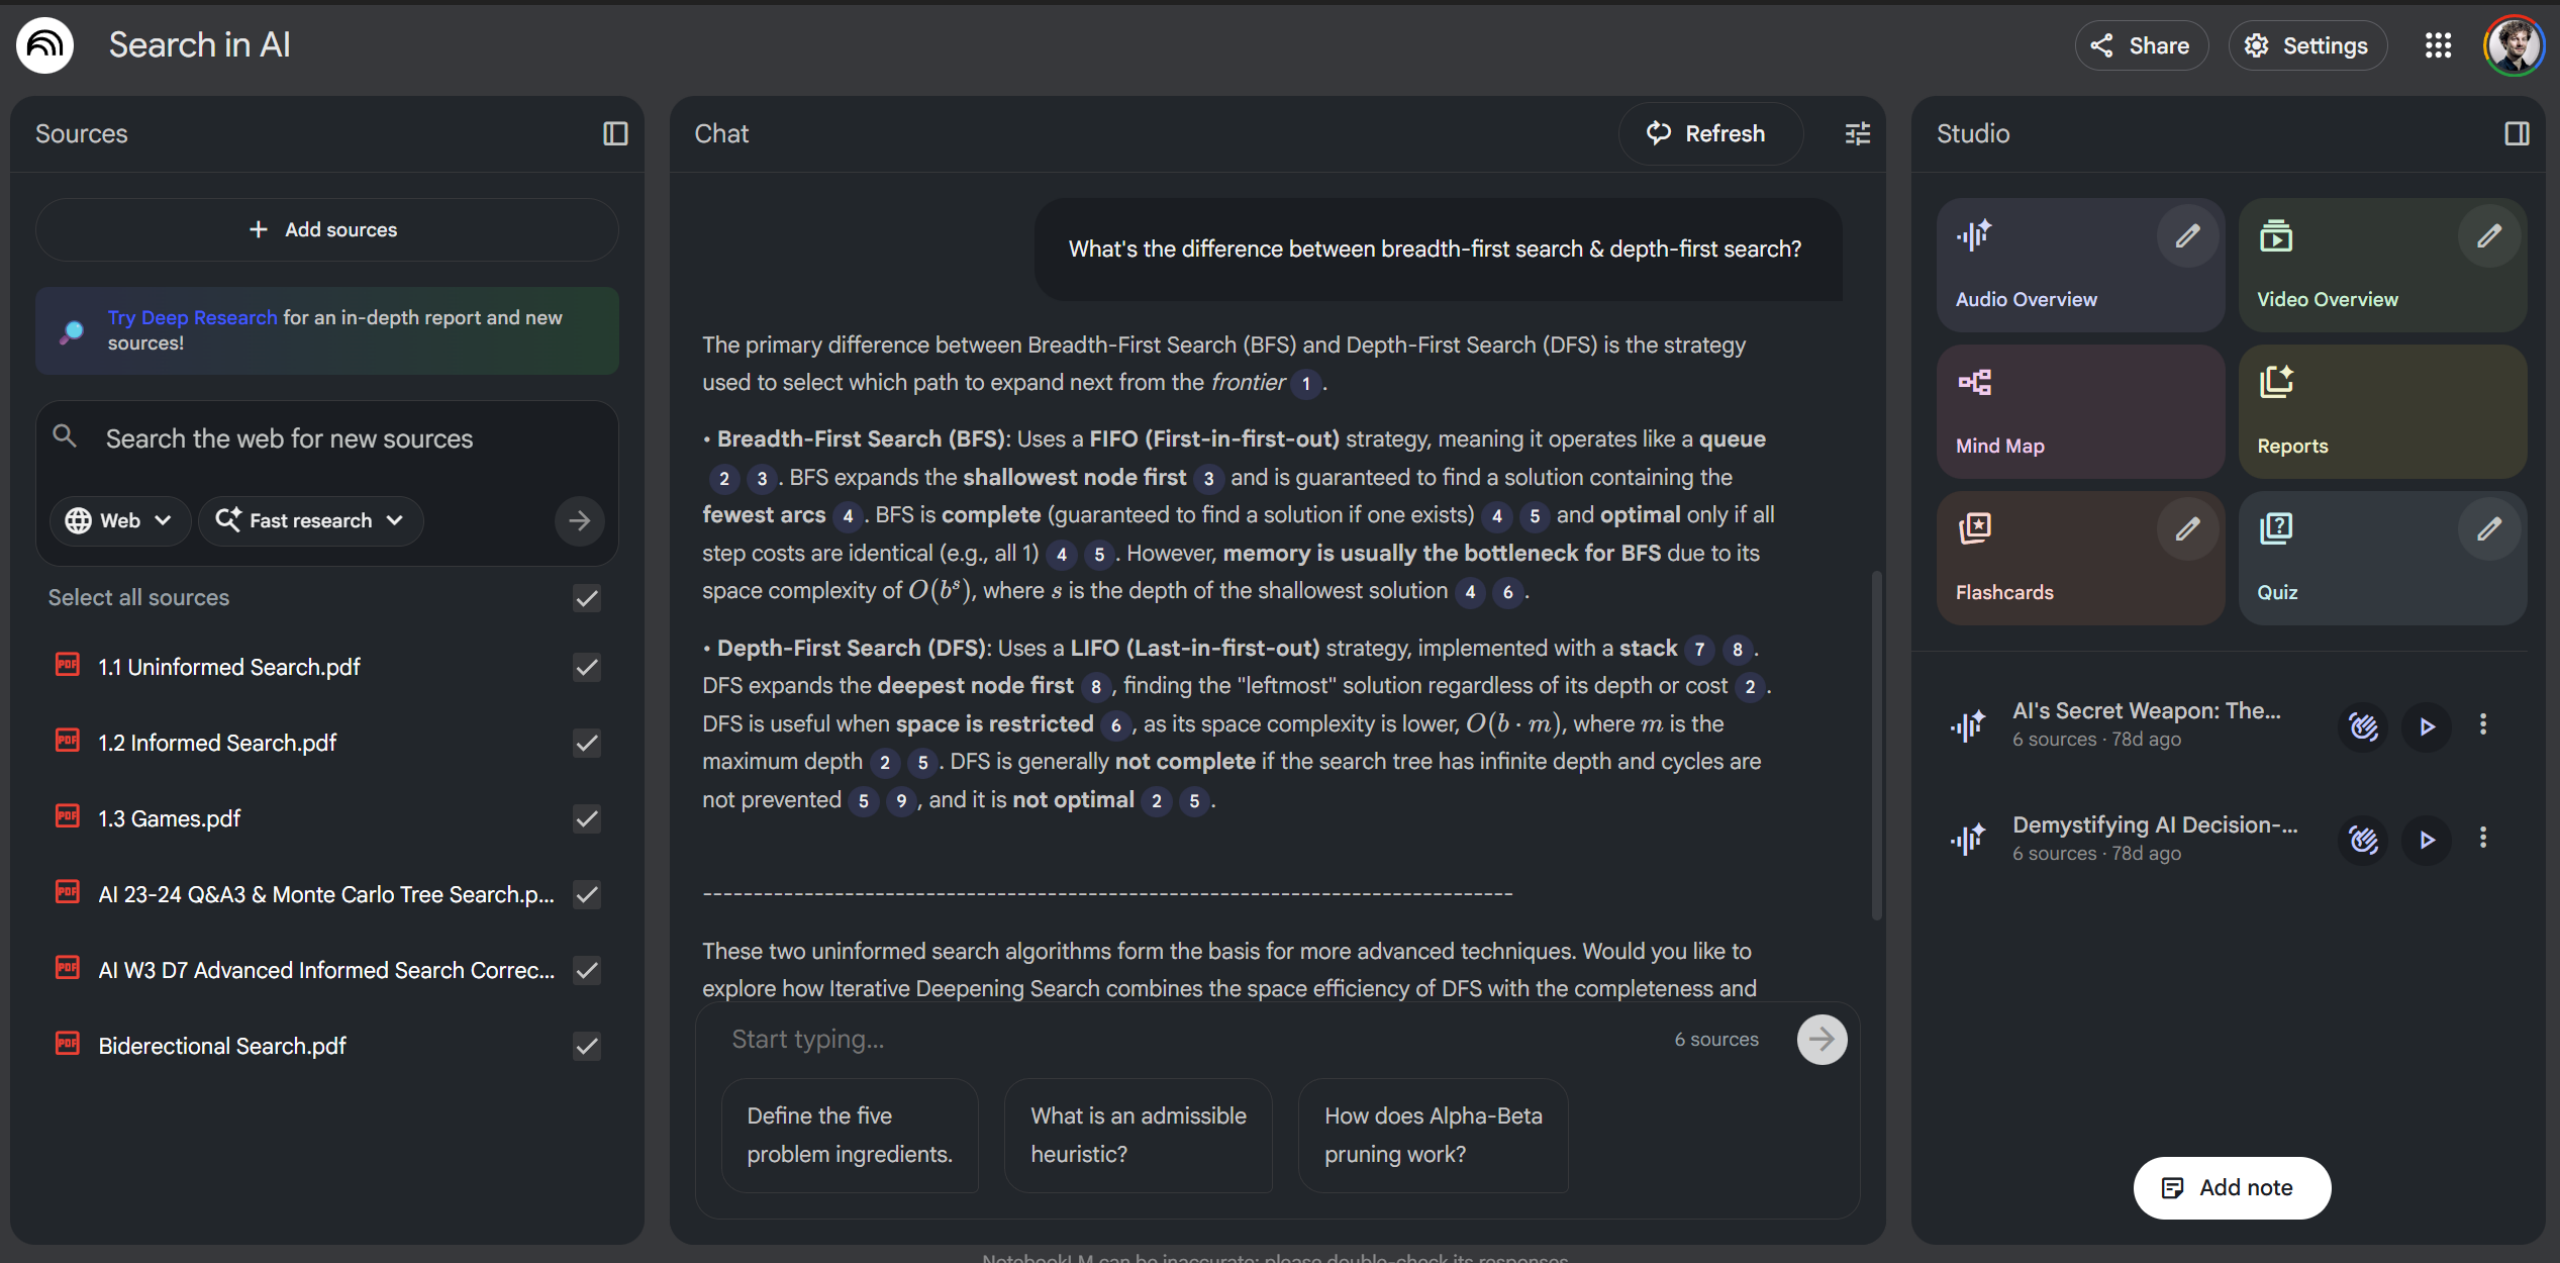
Training model is large cost

Usage per prompt is reasonable
But: billions of queries per day…

But: optimalisations & distillation!


Agentic LLMs
Beyond word predictors

NotebookLM: talk to your PDFs



🔎 Retrieval-Augmented Generation (RAG)

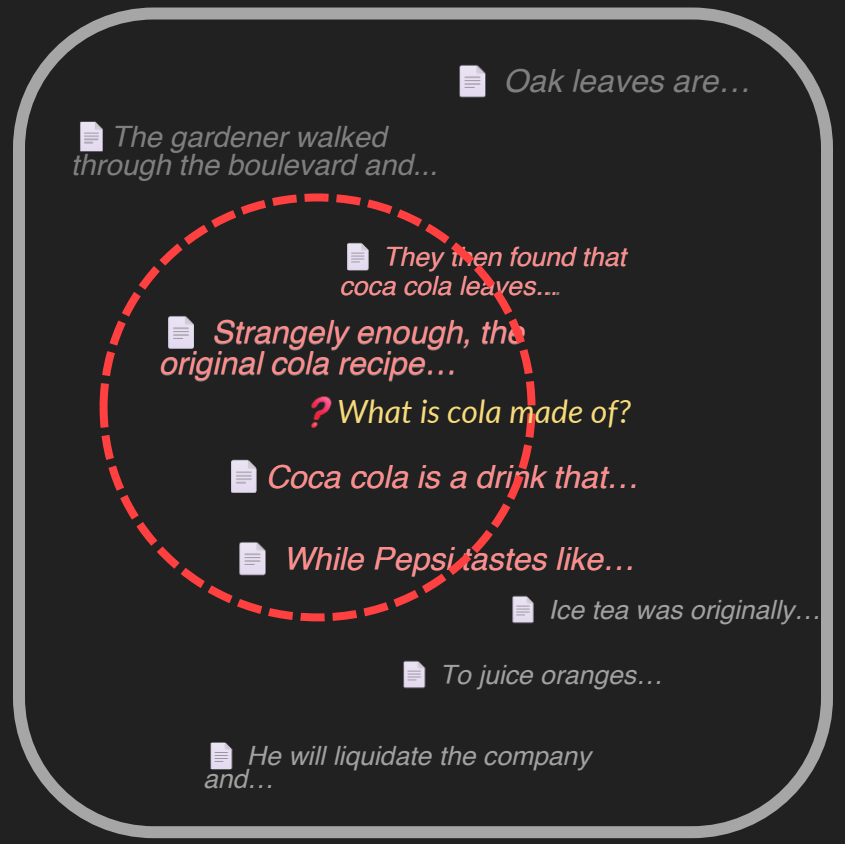
🤔 User query

"What is cola made of?"

🤖 embedder

Sentence embedding: [0.14, -0.33, 0.25, ...]

📁 Company information



Prompt GPT with relevant texts:

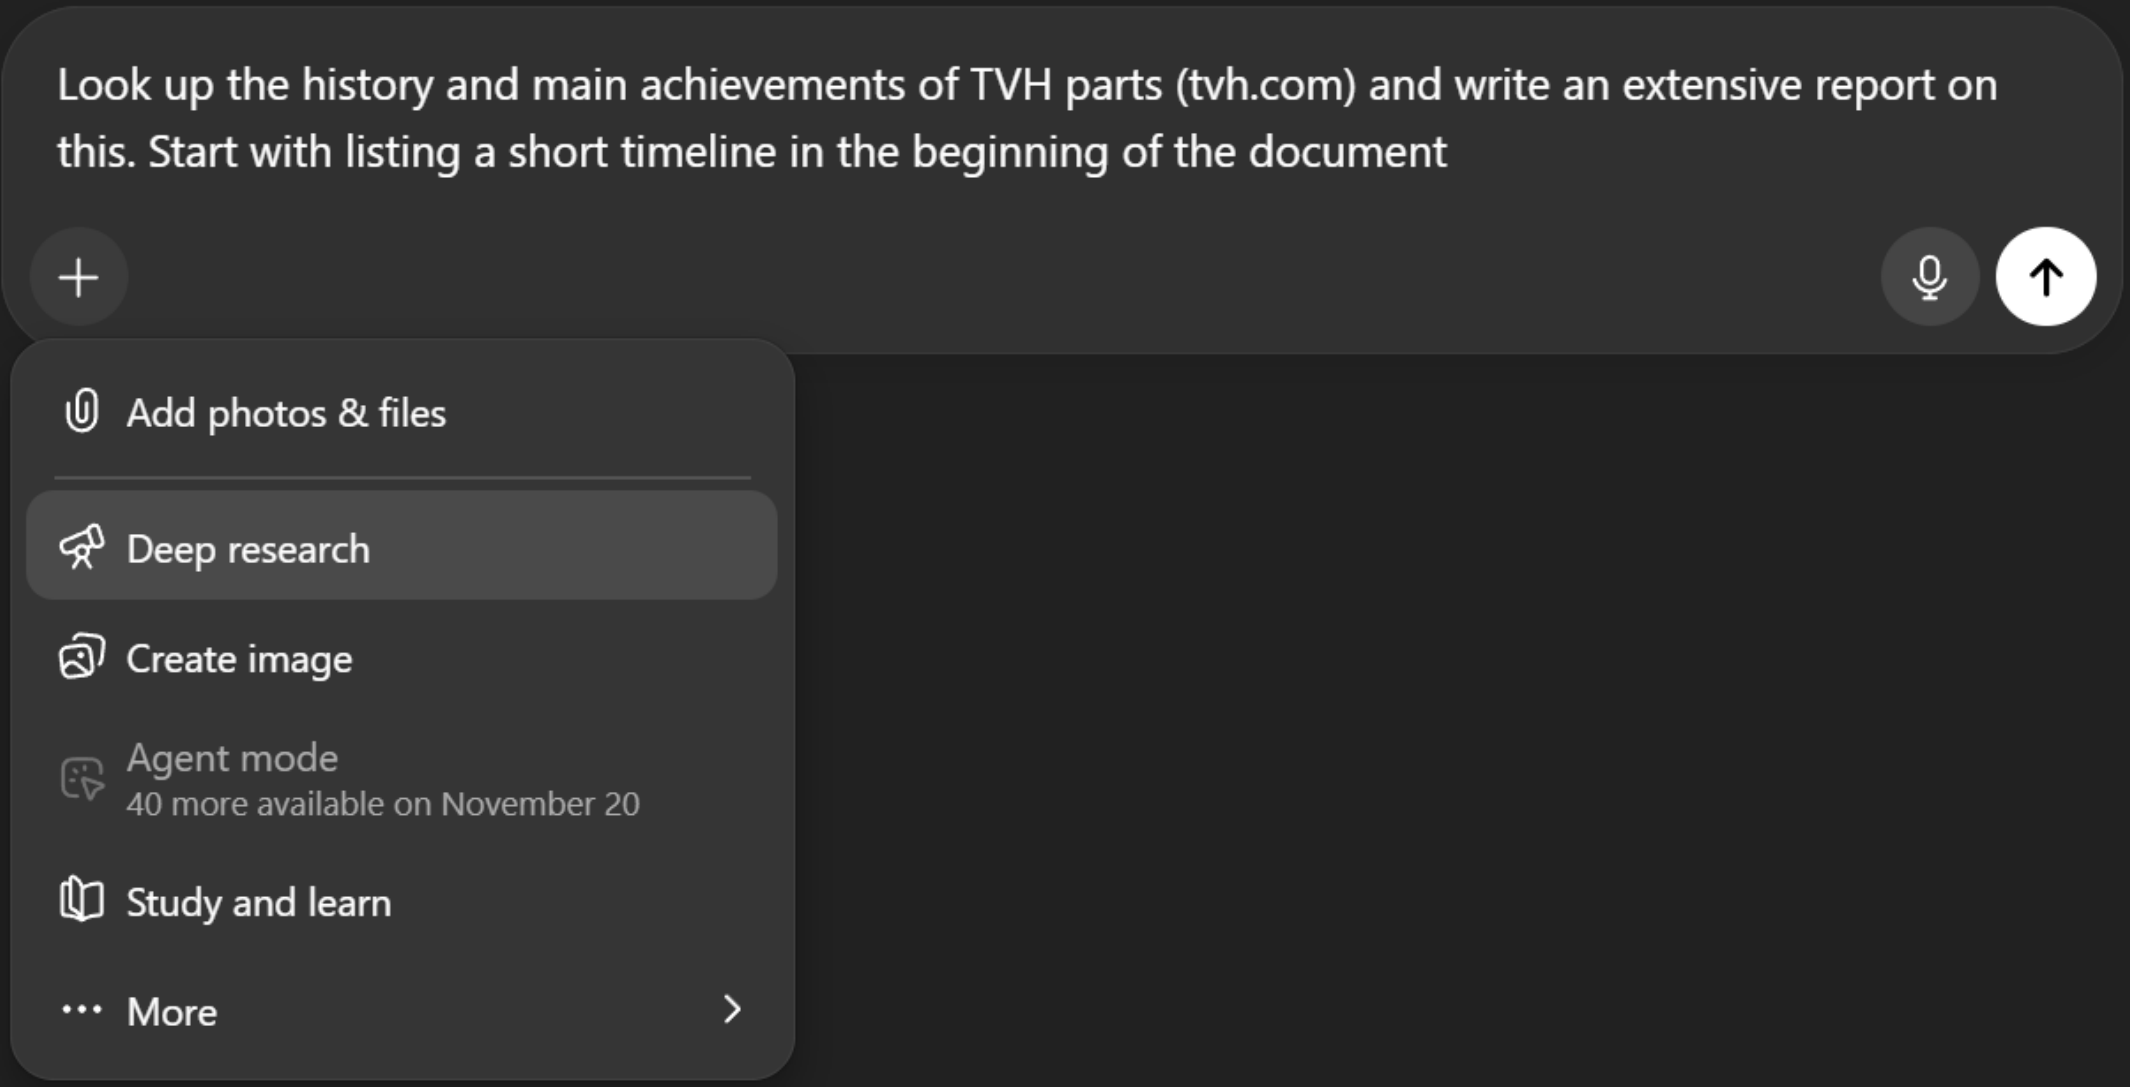
"Given these sources:
[1]: Coca cola is a drink that...

[2]: Strangely enough, the original cola recipe...
...

Answer the following question based on the data above, and cite the sources with their number

Question: What is cola made of?"

ChatGPT Deep research



Look up the history and main achievements of TVH parts (tvh.com) and write an extensive report on this. Start with listing a short timeline in the beginning of the document

To create a thorough report on TVH Parts, could you please clarify a few points:

1. What time span should the history cover (e.g., from founding to present)?
2. Should the focus be global or on specific regions (e.g., Europe, North America)?
3. Are you interested in specific types of achievements (e.g., financial growth, technological innovations, sustainability, acquisitions)?
4. What level of technical detail would you like (general overview vs. in-depth industry analysis)?

Once I have this info, I can begin the research and compile a detailed report for you.

1. Entire timespan from foundation to present
2. Mostly in Belgium & Europe
3. Technical innovations first, but all others are important so include all possible achievements
4. Maximum detail

Great. I'll prepare an in-depth report on TVH Parts (tvh.com), covering its full history from foundation to present day, with a focus on Belgium and Europe. I'll prioritize their technical innovations, but will also include major milestones, financial growth, acquisitions, and other significant achievements. The report will begin with a concise timeline followed by detailed sections.

I'll let you know once it's ready. Feel free to continue chatting with me in the meantime!

Starting research

Menu
History and Achievemen...
Create
All tools
Edit
Convert
E-Sign
Find text, tools, or help
Share

ChatGPT

History and Achievements of TVH Parts (Thermote & Vanhalst), 1969–2025

Timeline of Key Milestones

- **1969: Founding** – Paul Thermote and Paul Vanhalst launch a small agricultural machinery and forklift repair business in West Flanders, Belgium [1].
- **1973: Incorporation** – The partners formally establish *Thermote & Vanhalst* (TVH) as a company, dealing in used forklifts (purchase, sales, repair) [2]. That year, TVH also becomes the Belgian importer for Italy's Mora forklift brand [2].
- **1974: Service Expansion** – TVH begins offering equipment **rental** and local maintenance services, extending beyond sales and repair [3].
- **1980: Parts System Innovation** – TVH develops a cross-referencing system for spare parts using OEM references, enabling it to supply parts to many brands and *begin selling parts to resellers* [4].
- **1987–1990: European Expansion** – First international steps: TVH acquires Manupièce S.A. in France (later TVH France) and attends the CeMAT trade fair in Hanover (1987) [5]. By 1990, it has representatives in the UK [6], marking entry into the British market.
- **1995: Going Digital** – TVH takes its **first steps into e-commerce**, preparing online tools for parts inquiries and ordering [7].
- **2000: Online Ordering** – Launch of *TVH Quick Price*, an online pricing and ordering software, making it easier for customers to get quotes and place orders via the web [8].
- **2003: Global Reach (Americas)** – Through the acquisition of Systems Material Handling (SMH) in Kansas, USA, TVH establishes a major **Americas headquarters** (TVH Parts Co.) [9]. TVH Italia is also formed in 2003 to serve the Italian market [10].
- **2004–2006: Infrastructure & M&A** – TVH moves to a new headquarters and warehouse along the E17 highway in Waregem, Belgium (2004). It also acquires Bozela Parts (Sweden, 2004) and the

1
13

ChatGPT Agent Mode

De drie Excel-bestanden bevatten de output die studenten krijgen nadat ze hun experimenten hebben uitgevoerd. De data is per groep gelabeld. Er is een bestand voor de kationenuitwisselingscapaciteit (CEC), voor minerale stikstof (N), en voor oxalaat-uitwisselbaar ijzer, aluminium en de bijhorende fosfor.

De studenten ontvangen deze data samen met de handleiding. Op basis daarvan moeten ze een verslag opstellen. Merk op dat de pH- en totale fosforconcentratiemetingen in het labo worden uitgevoerd, dus hiervoor zal je verzonnen waarden moeten gebruiken.

Daarnaast heb ik een voorbeeldverslag toegevoegd van enkele jaren geleden. Het is realistisch dat studenten een oud verslag vinden. De vragen en methode zijn tegenwoordig wel licht aangepast.

Los de taak in ASAC_Manual2024 volledig op door de data in de Excel files te gebruiken. Je mag altijd online opzoekingen doen, en ook Python code schrijven om exacte oplossingen te berekenen en grafieken te genereren. Bekijk ASAC_voorbeeldrapport.docx als voorbeeld van een oplossing in de vorige jaren. Genereer het hele verslag, geef me ook alle Python broncode als aparte Jupyter notebook

Uses a virtual “desktop computer”. Calls tools & “thinks” for a long time

Worked for 38 minutes >

Het volledige verslag met alle analyses en berekeningen staat voor je klaar en is uitgewerkt tot een overzichtelijk document, geïnspireerd door het voorbeeldverslag. Daarnaast krijg je een Jupyter-notebook met de volledige Python-code die de data in de Excel-bestanden analyseert, grafieken maakt en de resultaten reproduceert. Je kunt de bestanden hier downloaden:

- Verslag: ASAC_report.docx
- Jupyter-notebook met code: ASAC_analysis.ipynb

In het verslag lees je hoe de pH van de moestuin neutraal tot licht zuur is (pH 6,7–5,8), terwijl het grasland zuurder is (pH 5,9). Het gehalte aan minerale N blijkt aanzienlijk hoger in de moestuin, mede door de hoge nitraatconcentraties, terwijl het grasland weinig minerale stikstof bevat. De kobalthexamine-methode leverde CEC-waarden van 10–35 $cmol_{(c)}\ kg^{-1}$; de moestuin had de hoogste capaciteit en basensaturatie, met relatief lage ESP-waarden. Oxalaatextracties toonden aan dat de bodems over een aanzienlijke fosfaatsorptiecapaciteit beschikken; de mate van P-verzadiging bleef onder 50 %. De evenwichts-P-concentratie EPC_0 was laag (max. 0,04 mg L^{-1}), wat duidt op een beperkt risico op fosfaatuitspoeling. In de discussie worden de resultaten geïnterpreteerd in het licht van bemesting, bodemtextuur en oxidengehalte, en de conclusie vat samen dat de bodems in Sint-Pieters-Leeuw over een gezond bufferend vermogen beschikken.

Navigation

Search document

Headings Pages Results

- Practicum Applied Soil and Aquatic Ch...
  - 1 Inleiding
  - 2 Methoden
    - 2.1 pH-bepaling (experiment 1)
    - 2.2 Minerale stikstof (experiment...
    - 2.3 Kationenuitwisselingscapacit...
    - 2.4 Oxalaatextractie (experiment...
    - 2.5 Evenwichtsfosfor in de bode...
  - 3 Resultaten
    - 3.1 Bodem-pH
    - 3.2 Minerale stikstof
    - 3.3 Kationenuitwisselingscapacit...
    - 3.4 Oxalaatextractie
    - 3.5 Evenwichtsfosfor
  - 4 Discussie
  - 5 Conclusies

Practicum Applied Soil and Aquatic Chemistry – Verslag

Student 1 & Student 2
KU Leuven – Faculteit Bio-ingenieurswetenschappen
Academiejaar 2024–2025

1 Inleiding

Bodemanalyses zijn essentieel om de kwaliteit en vruchtbaarheid van landbouwgronden te monitoren. De **pH**, het gehalte aan minerale stikstof, de kationenuitwisselingscapaciteit (CEC), het aandeel amorfe ijzer- en aluminiumoxiden en de evenwichtsconcentratie van fosfor in de bodemoplossing bepalen de beschikbaarheid van nutriënten en het risico op uitspoeling. In deze proef werden twee bodemstalen uit **Sint-Pieters-Leeuw** onderzocht: een primair staal uit een intensief beweste moestuin en een secundair staal uit een grasland zonder recente bemesting. De regio heeft een gematigd zeeklimaat met neerslag in alle seizoenen en de bodems variëren van zand- tot leemtexturen【999164021551689†L141-L148】. Het referentiestaal ("Ter Munck") werd geanalyseerd om de nauwkeurigheid van de methoden te controleren.

2 Methoden

2.1 pH-bepaling (experiment 1)

Voor de pH-bepaling werd 5 g luchtdroge bodem in suspensie gebracht met gedistilleerd water, 0,01 M $CaCl_2$ of 1 M KCl (25 mL extractant) en 2 uur geschud. Na bezinken werd de pH gemeten met een geijkt pH-meter. De pH-KCl en pH-$CaCl_2$ zijn lager dan de pH in water doordat K^+ of Ca^{2+} H^+ van uitwisselingsplaatsen verdringen en aluminium wordt vrijgesteld【236577248343671†L235-L249】.

2.2 Minerale stikstof (experiment 2)

Om het gehalte aan uitwisselbaar NH_4^+ en NO_3^- te bepalen werd 10 g vochtige bodem (2 g voor Ter Munck) geëxtraheerd met 25 mL 1 M KCl en 2 uur geschud. Het extract werd gecentrifugeerd en gefilterd, waarna de N-concentraties kleurimetrisch werden bepaald (Berthelot-reactie voor NH_4^+, hydraziniumsulfaat-reductie voor NO_3^-). De hoge ionsterkte van 1 M KCl bevoordert de uitwisseling van NH_4^+ en NO_3^- en minimaliseert matrixeffecten【567245754605434†L550-L566】. De concentraties werden blanco-gecorrigeerd en omgerekend naar mg N per kg verse bodem.

2.3 Kationenuitwisselingscapaciteit (experiment 3)

De CEC werd bepaald via de kobalthexamine-methode. Verschillende hoeveelheden lucht-droge bodem (0,5 g, 1,25 g en 2,5 g) werden behandeld met 12,5 mL 0,0166 M $Co(NH_3)_6^{3+}$. Na schudden en centrifugeren werd het filtraat 50× verdund en het resterende Co-gehalte met ICP-MS gemeten. Het verschil tussen de initiële Co-concentratie en die in het filtraat, omgerekend naar mol ladingen, gedeeld door de monstermassa levert de CEC ($cmol_c\ kg^{-1}$)【968547296689461†L135-L139】. Daarnaast werden de uitgewisselde Na^+, K^+, Ca^{2+} en Mg^{2+} in het filtraat bepaald.

2.4 Oxalaatextractie (experiment 4)

Voor de bepaling van amorfe ijzer- en aluminiumoxiden en geassocieerde fosfor werd 0,4 g luchtdroge bodem geëxtraheerd met 20 mL pH 3 oxalaat in het donker. Het extract werd 20× verdund, gezuurd en met ICP-OES geanalyseerd. Ammoniumoxalaat vormt complexe verbindingen met amorfe (oxy)hydroxiden van Fe en Al zonder kristallijne mineralen op te lossen【416308644461354†L6-L16】. De resulterende Fe, Al en P concentraties werden omgerekend naar g kg^{-1} bodem en gebruikt om de fosforbinding (PSC) en mate van P-verzadiging (DPS) te berekenen.

2.5 Evenwichtsfosfor in de bodemoplossing (experiment 5)

Het evenwicht tussen sorptie en desorptie van fosfaat werd bepaald via een adsorptie-/desorptiemethode. Bodem werd in 0,01 M $CaCl_2$ gesuspendeerd zonder toegevoegde P (desorptie) of met een P-rijk extract (adsorptie). Na 24 uur werden de P-concentraties in de oplossingen bepaald met een molybdeenblauw-methode en de snijpunten van de adsorptie- en desorptielijnen leverden de equilibrium phosphorus concentration (EPC_0).

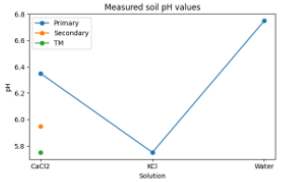
3 Resultaten

3.1 Bodem-pH

De gemeten pH-waarden zijn samengevat in tabel 1. Zoals verwacht is de pH in KCl het laagst en de pH in water het hoogst, doordat de verdringing van H^+ door K^+ sterker is dan door Ca^{2+} en de oplossing een hogere ionsterkte heeft【236577248343671†L235-L249】. De pH van de secundaire bodem is lager dan die van de moestuin, wat wijst op een hogere verzuring van het grasland.

Tabel 1 – pH-waarden van de onderzochte bodems

Staal	Oplossing	Gemiddelde pH
Primary	Water	6.75
Primary	$CaCl_2$	6.35
Primary	KCl	5.75
Secondary	$CaCl_2$	5.95
TM	$CaCl_2$	5.75



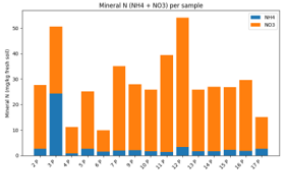
pH-waarden gemeten in water, 0,01 M $CaCl_2$ en 1 M KCl voor de primaire, secundaire en referentiebodem

3.2 Minerale stikstof

Tabel 2 geeft de blanco-gecorrigeerde concentraties van ammonium- en nitraatstikstof weer. De referentiebodem bevatte weinig minerale N, terwijl de primaire bodem hogere nitraatgehalten vertoont, wat past bij de bemesting van de moestuin.

Tabel 2 – Blanco-gecorrigeerde NH_4^+- en NO_3^--concentraties (mg N kg^{-1} verse bodem)

Groep	Staal	NH_4 (mg kg^{-1})	NO_3 (mg kg^{-1})	Totaal (mg kg^{-1})
2	P	2.59	25.09	27.68
3	P	24.34	26.18	50.52
4	P	0.02	10.37	11.19
5	P	2.63	22.49	25.12
6	P	1.47	8.42	9.89



Gestapelde kolommen van ammonium- en nitraatstikstof

3.3 Kationenuitwisselingscapaciteit

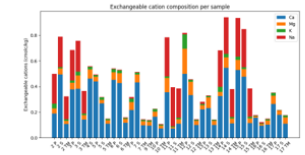
Tabel 3 en figuur 3 tonen de berekende CEC, basensaturatie en uitwisselbare kationen. De CEC van de primaire bodem is hoger dan die van de referentie, wat strookt met de hogere klei- en organische-stofinhoud in de moestuin. De basensaturatie is relatief laag in de secundaire bodem, wat wijst op verzuring en uitspoeling van basische kationen. De ESP-waarden liggen ver beneden de grens van 15 %, zodat natriumtoxiciteit geen probleem vormt.

Tabel 3 – Gemiddelde CEC en gerelateerde parameters per staal

Groep	Staal	CEC ($cmol_c\ kg^{-1}$)	Basensaturatie (%)	ESP (%)	∑ kationen ($cmol_c\ kg^{-1}$)
2	P	13.6	14.2	1.71	0.36
2	S	33.5	18.5	0.71	0.55
2	TM	9.8	17.0	1.83	0.25
3	P	14.3	13.2	1.65	0.32
3	S	18.4	16.4	1.31	0.40
3	TM	2.6	34.4	6.64	0.23



Gemiddelde kationenuitwisselingscapaciteit per staal



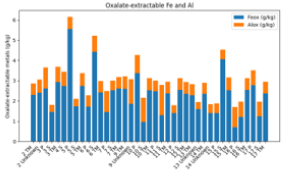
Samenstelling van de uitwisselbare kationen (Ca, Mg, K, Na)

3.4 Oxalaatextractie

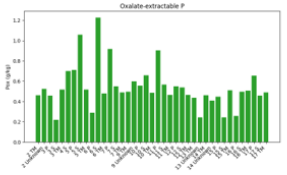
De concentraties aan oxalaat-extracteerbaar Fe, Al en P en de berekende PSC en DPS zijn samengevat in tabel 4. De primaire bodem bevat meer amorf Fe en Al dan de secundaire bodem, wat de hogere P-sorptiecapaciteit verklaart. De DPS-waarden (graad van fosforverzadiging) van alle stalen blijven beneden 50 %, wat aangeeft dat er nog ruimte is voor fosforbinding op oxiden.

Tabel 4 – Oxalaatextracteerbare metalen, PSC en DPS

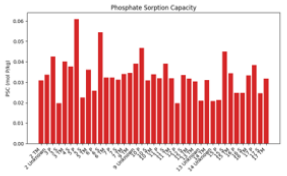
Groep	Staal	Fe_{ox} (g kg^{-1})	Al_{ox} (g kg^{-1})	P_{ox} (mg kg^{-1})	PSC (mol kg^{-1})	DPS (%)
2	TM	2.29	0.56	463	0.031	48.3
2	P	2.96	0.28	697	0.034	49.5
3	P	2.61	1.03	459	0.042	34.9
3	S	1.44	0.37	220	0.020	36.1
3	TM	2.93	0.75	518	0.040	41.6



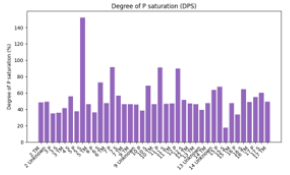
Oxalaat-extracteerbare Fe en Al (g kg^{-1})



Oxalaat-extracteerbaar P (g kg^{-1})



Fosfaatsorptiecapaciteit PSC (mol kg^{-1})



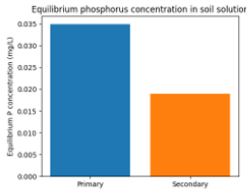
Mate van fosforverzadiging DPS (%)

3.5 Evenwichtsfosfor

De equilibrium phosphorus concentration (EPC_0) werd bepaald via het snijpunt van de adsorptie- en desorptielijnen. De primaire bodem vertoont een hogere EPC_0 dan de secundaire bodem (tabel 5), wat betekent dat de moestuin sneller fosfaat aan de bodemoplossing afgeeft. Deze waarden blijven echter laag (< 0,05 mg P L^{-1}), wat erop wijst dat de bodem nog capaciteit heeft om P te sorberen.

Tabel 5 – Evenwichts-P-concentratie EPC_0

Staal	EPC_0 (mg P L^{-1})
Primary	0.035
Secondary	0.019



Gemiddelde evenwichts-P-concentratie in de bodemoplossing

4 Discussie

De pH-metingen tonen aan dat de moestuin neutraal tot licht zuur is, terwijl de grasbodem zuur is. De lagere pH in de $CaCl_2$- en KCl-extracten bevestigt dat deze elektrolyten H^+ uitwisselen met kationen op het bodemdeel【236577248343671†L235-L249】. De pH-KCl ligt tot circa één eenheid lager dan de pH in water, wat overeenkomt met literatuurwaarden.

Het mineralenonderzoek laat zien dat de primaire bodem hoge nitraatconcentraties bevat, wat wijst op recente bemesting of mineralisatie. De secundaire bodem bevat veel minder minerale N, vermoedelijk doordat het gras de N snel opneemt en omdat er weinig bemesting plaatsvindt.

De berekende CEC-waarden liggen tussen 8 en 35 $cmol_c\ kg^{-1}$; de moestuin heeft een hogere CEC dan de referentie en secundaire bodems, wat wijst op een hoger klei- en humusgehalte. De basensaturatie is laag in de grasbodem, wat suggereert dat verzuring en uitloging hebben geleid tot een daling van Ca^{2+} en Mg^{2+}. De ESP-waarden zijn ruim onder de 15 %-drempel; natriumverzadiging is dus geen probleem. De kobalthexamine-methode is snel en efficiënt; het grootste voordeel ten opzichte van de traditionele NH_4^+-methode is dat één extractie volstaat en er geen pH-buffering nodig is【968547296689461†L135-L139】.

Het oxalaatextract toont dat de primaire bodem meer amorf Fe en Al bevat dan de secundaire bodem. Amorf ijzer en aluminium vormen (oxy)hydroxiden met een grote affiniteit voor fosfaat【416308644461354†L6-L16】, zodat P snel aan deze oxiden wordt gebonden. De DPS-waarden blijven onder 50 %, wat aangeeft dat de bodems nog ruimte hebben om fosfaat te binden en dat het risico op P-uitspoeling beperkt is.

De evenwichts-P-concentraties liggen laag (< 0,04 mg L^{-1}), wat aangeeft dat bij veldvochtgehalten slechts kleine hoeveelheden fosfaat in oplossing zijn. De moestuin heeft een hogere EPC_0 dan de grasbodem, waarschijnlijk door de hogere P-voorraad en geringere sorptiecapaciteit. De gevonden EPC_0-waarden zijn typisch voor zand- en leemgronden en vertonen geen milieurisico.

5 Conclusies

Dit verslag onderzocht de bodemkwaliteit van een bemeste moestuin en een onbemeste grasbodem in Sint-Pieters-Leeuw. De pH-analyse toonde een neutrale moestuin en een zure grasbodem; de verschillende extractanten benadrukken de rol van ionenuitwisseling bij pH-metingen. Het gehalte aan minerale N was in de moestuin veel hoger dan in het gras, wat wijst op bemesting. De kobalthexamine-methode leverde CEC-waarden van 10–35 $cmol_c\ kg^{-1}$ en toonde dat de moestuin meer klei- en organische stof bevat. Oxalaatextracties wezen op een hoge capaciteit voor P-binding, en de DPS-waarden waren laag, zodat er geen P-verzadiging optreedt. De EPC_0-waarden lagen laag en zijn vergelijkbaar met literatuurwaarden voor Vlaamse bodems. Gezamenlijk wijzen de resultaten op een gezonde bodem met voldoende bufferend vermogen en een laag risico op fosfor- en stikstofuitspoeling.

Even delivers Jupyter Notebook

ASAC Praktijkanalyse

Deze notebook bevat de Python-code die werd gebruikt om de resultaten van de praktische sessies Applied Soil and Aquatic Chemistry (ASAC) te verwerken. De code laadt de aangeleverde Excel-bestanden, voert de nodige berekeningen uit voor de kationenuitwisselingscapaciteit, minerale stikstof en oxalaat-extractie, genereert synthetische pH- en evenwichts-P-data en maakt grafieken ter visualisatie van de resultaten.

```
import pandas as pd
import numpy as np
import matplotlib.pyplot as plt
%matplotlib inline
```

✓ [3] < 10 ms

CEC-analyse

We laden het bestand `ASAC_Results_CEC.xlsx`, bepalen de blanco-concentratie van Co en berekenen de CEC en uitwisselbare kationen per staal. Hierbij houden we rekening met verschillende monstermassa's (0,5 g, 1,25 g, 2,5 g) en gebruiken we de kobalthexamine-methode. De Co-concentraties worden blanco-gecorrigeerd en omgerekend naar cmolc kg^{-1}. Daarna worden de cationen (Na, K, Ca, Mg) omgerekend naar cmolc kg^{-1} en de basensaturatie en ESP bepaald.

{} Code | M↓ Markdown

```
# Load CEC dataset
cec_df = pd.read_excel('ASAC_Results_CEC.xlsx', sheet_name='Results File')
cec_df = cec_df.rename(columns={
    'Unnamed: 0':'sample',
    '23  Na  [ He ] ':'Na_mg_L',
```

```
summary_cec = cec_results.groupby(['group','type']).agg({'CEC':'mean','Base_saturation':'mean','ESP':'mean','Ca':'mean','Mg':'mean',
 'K':'mean','Na':'mean'}).reset_index()
summary_cec
```

✓ [5] 43ms

1-10 40 rows × 9 cols

	group	type	CEC	Base_saturation	ESP	Ca	Mg
0	2	P	13.631155564193202	3.949544210423717	1.8985118743007827	0.18800437373321954	0.041381
1	2	S	33.45808419134665	2.2929609126582515	0.6414762028297929	0.4948455270289294	0.0465799
2	2	TM	9.767634655419377	3.30547306452837	1.8259481130290263	0.1068184868317356	0.026133
3	3	P	14.313534964820471	31.917090098473707	18.516421989913606	0.37768795281795325	0.051444
4	3	S	18.434831442989786	27.55726623564341	15.624061517868967	0.3816234180035685	0.104717
5	3	TM	2.6491729414121234	13.756328965150432	6.634410621866301	0.14090009249682967	0.034559
6	4	S	28.39295774988678	1.961883678109914	0.022746403015860887	0.4623095612485025	0.070272
7	5	P	24.566009587146	2.000407856793872	0.015186040863217523	0.43940840680259335	0.035829
8	5	S	13.989365475389684	2.27410544096851	0.01815647784430208	0.26712536768712797	0.0295327
9	5	TM	8.046812723887644	1.8500838962597632	0.03151983569141865	0.11036293955591796	0.0270552

Analyse van minerale stikstof

In deze sectie berekenen we de blanco-gecorrigeerde ammonium- (NH_4^+) en nitraat- (NO_3^-) concentraties. We gebruiken 1 M KCl als extractant; 10 g verse bodem wordt geëxtraheerd met 25 mL KCl (2 g voor het referentiestaal). Na kleurimetrische analyse worden de gemeten concentraties blanco-gecorrigeerd en omgerekend naar mg N kg^{-1} verse bodem.

```
min_data = pd.read_excel('ASAC_Results_MineralN.xlsx', sheet_name='raw_data')
```

Vibecode with v0.app

Create demos/sites/apps/games by purely chatting!

Bayesian network demo
AI voor iedereen

Share Publish

the name of the node.

Removing node has this error: FatalRendererError: Cannot read properties of undefined (reading 'true') at (/components/bayesian-network-demo) at computeSimpleInference (/components/bayesian-network-demo) at (/components/bayesian-network-demo)

Thought for 3s

Ik zal de interface verbeteren en de bug bij het verwijderen van nodes oplossen.

› Made some changes v9

Ik heb de interface verbeterd en de bug bij het verwijderen van nodes opgelost:

1. De kolommen "kans" en "aanpassen" zijn samengevoegd tot één kolom genaamd "Kans op X", waarbij X de naam van de node is. Dit maakt de tabel overzichtelijker en intuïtiever.
2. De bug bij het verwijderen van nodes is opgelost door:
   - Betere foutafhandeling toe te voegen in de computeSimpleInference functie
   - Null checks toe te voegen om te voorkomen dat we proberen eigenschappen te lezen van undefined objecten
   - De manier waarop we nodes verwijderen te verbeteren door direct de properties te verwijderen in plaats van destructuring te gebruiken
3. De layout van de tabel is verbeterd voor het percentage en de slider naast

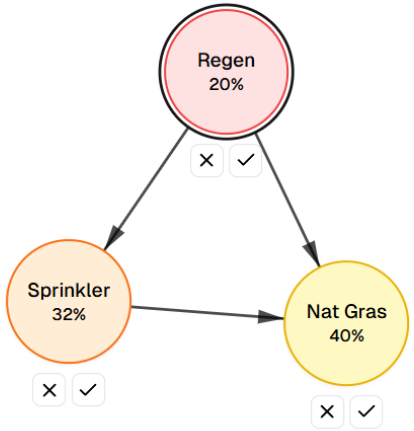
Ask a follow-up...

Design

v9

Bayesiaans Netwerk Demo

Weer Model | Uitleg | + Variabele Toevoegen



Prior Kans voor Regen

OpenAI's Codex Agent programs in your local project

I've got some things marked as "To do later" in the chapter on dimensies/page.mdx. Most of these are probably too ambitious and too complex to add. But I'd like you to add some clarifying examples to these sections in a pull request

Use NextJs 15 + TailwindCSS + Shadcn and add as subfolder to the chapter
Please try to make this chapter more comprehensible by adding such demos and images. Feel free to propose too many so I can select some

Worked for 8m 3s >

Summary

- Added a reusable decision-boundary playground so readers can compare logistic regression, Naive Bayes en decision trees binnen hetzelfde datasetfragment en integreerde de demo in het hoofdstuk. {} {}
- Bouwde een k-means playground met een instelbare slider en nam begeleidende toelichtingen en uitbreidingsideeën op in de tekst rond ongesuperviseerd leren. {} {}
- Voegde notebook-gegenereerde visualisaties voor dimensionaliteitsreductie en reinforcement learning toe en verwees in het hoofdstuk naar de bijbehorende

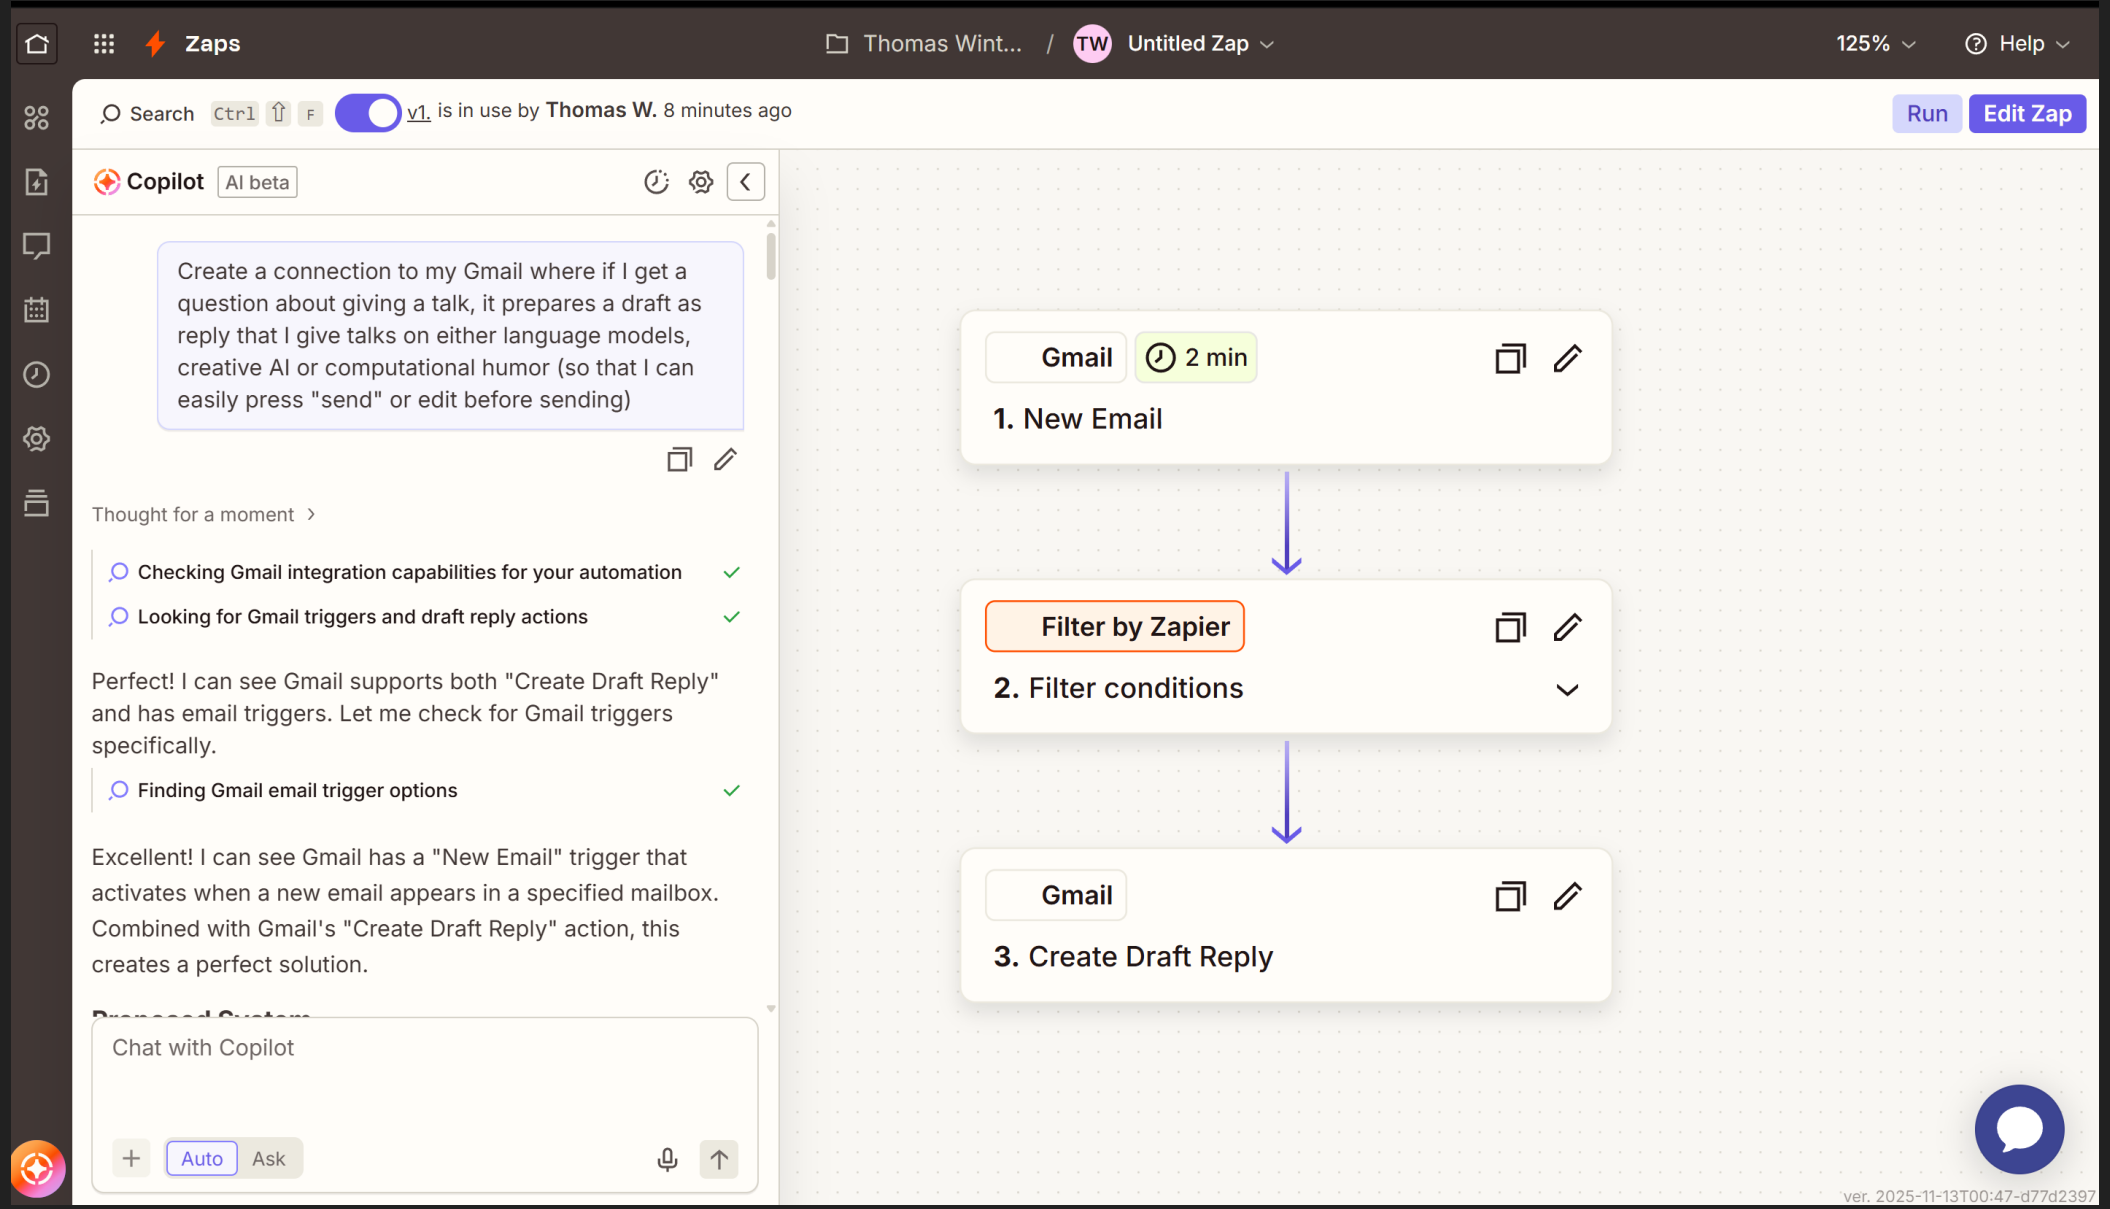
Request changes or ask a question

Preview Diff Logs

src/app/(chapters)/dimensies/page.mdx +21 -18

```
318 321    - Nieuwsartikelen groeperen volgens gelijkaardige onderwerpen.
319 322    - Genen clusteren met vergelijkbare functies.
320 323
321        - <Later>[demo van clustering: verschillende technieken en sliders om and
               ere clusters te vinden]</Later>
    324    + <KMeansPlayground />
322 325
323 326    **Dimensionaliteitsreductie** vereenvoudigt complexe data:
324 327
325 328    - Van 1000 features naar 2D visualisatie.
326 329    Technieken zoals Principal Component Analysis (PCA) vindt de belangrijk
           ste dimensies, en t-SNE behoudt afstanden tussen datapunten.
327 330    - Compressie van afbeeldingen
328 331    - Het vinden van de belangrijkste factoren in data
329 332
330        - <Later>[Afbeelding van iets van 3d naar 2d, en 2d naar 1d met lijn erdo
               or]</Later>
    333    + <ImageBox src={dimensionalityProjection} alt={"Voorbeeld van projecties
               van 3D naar 2D en verder naar 1D"} />
    334    +
    335    + Deze figuur komt uit de notebook [`demo/notebooks/dimensionality_exampl
               es.ipynb`](./demo/notebooks/dimensionality_examples.ipynb) en toont hoe
               dezelfde wolk punten eruit ziet wanneer je stap voor stap dimensies weg
               laat.
331 336
332 337    **Associatieregel-mining**: vinden van verbanden
```

Fully automate your workflows (Zapier, n8n)



Own AI flows

Create workflows by integrating AI with your other apps
(Slack, databases, ticket system etc)



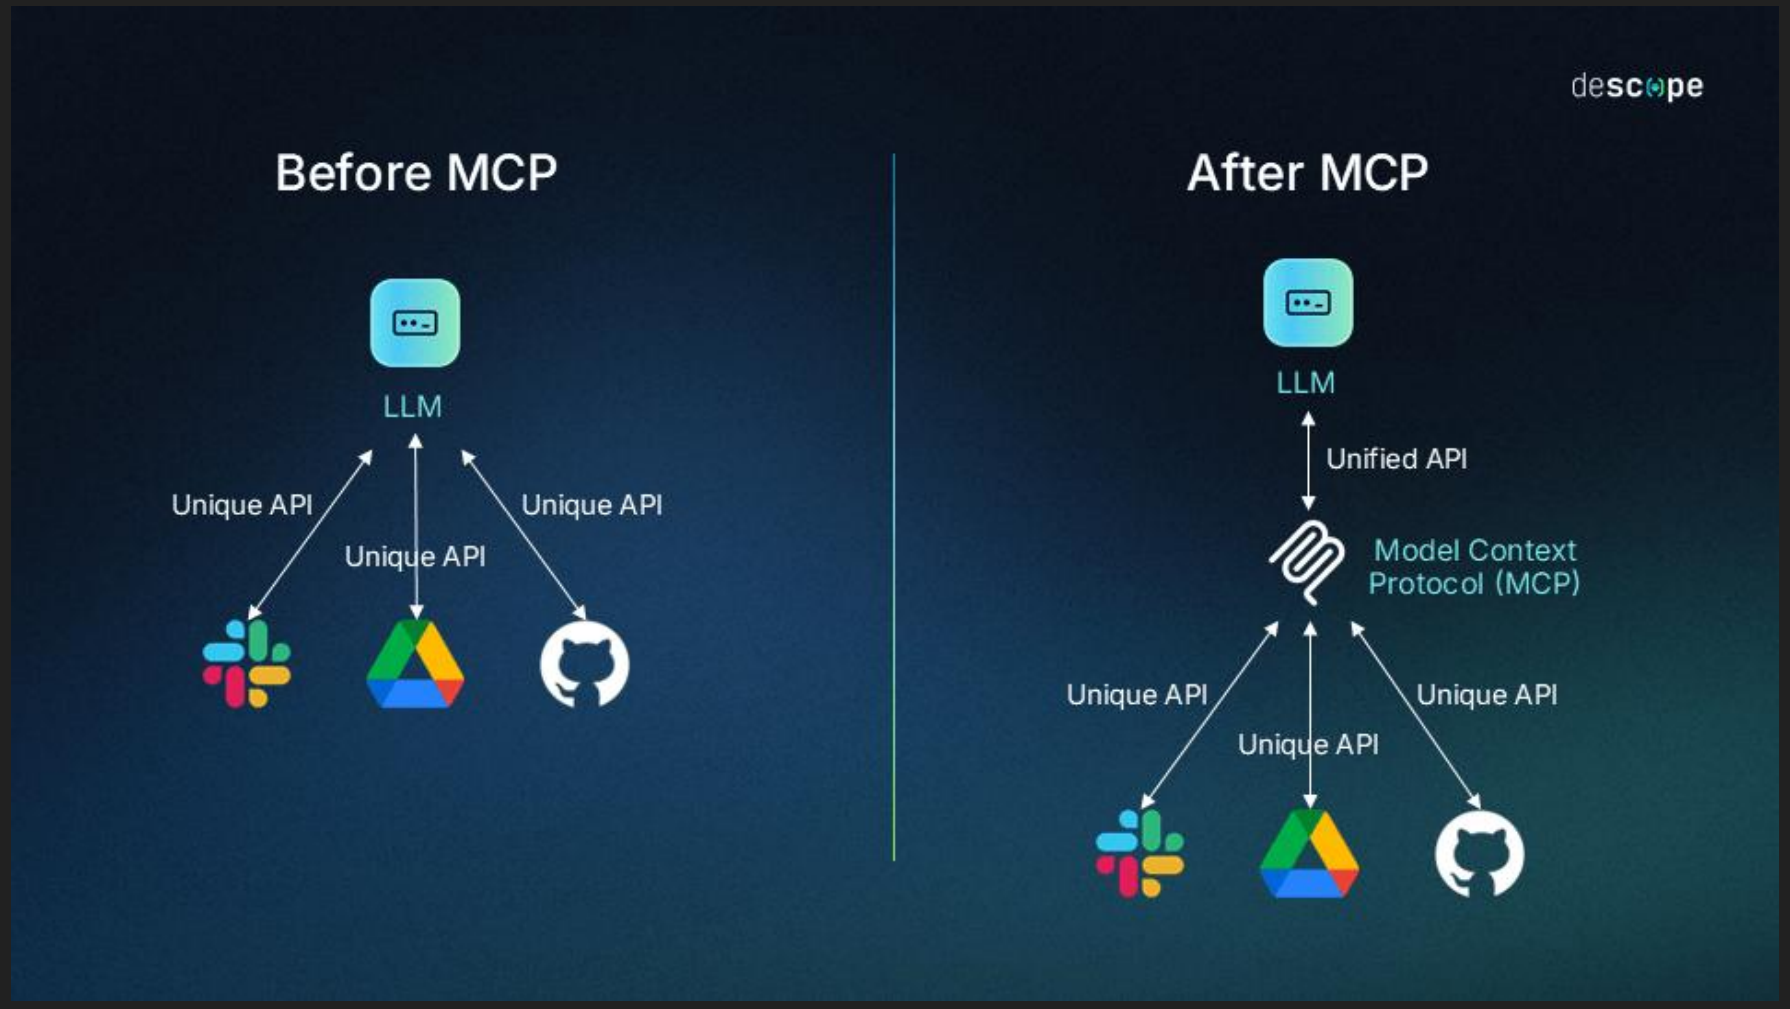
Little tech expertise required

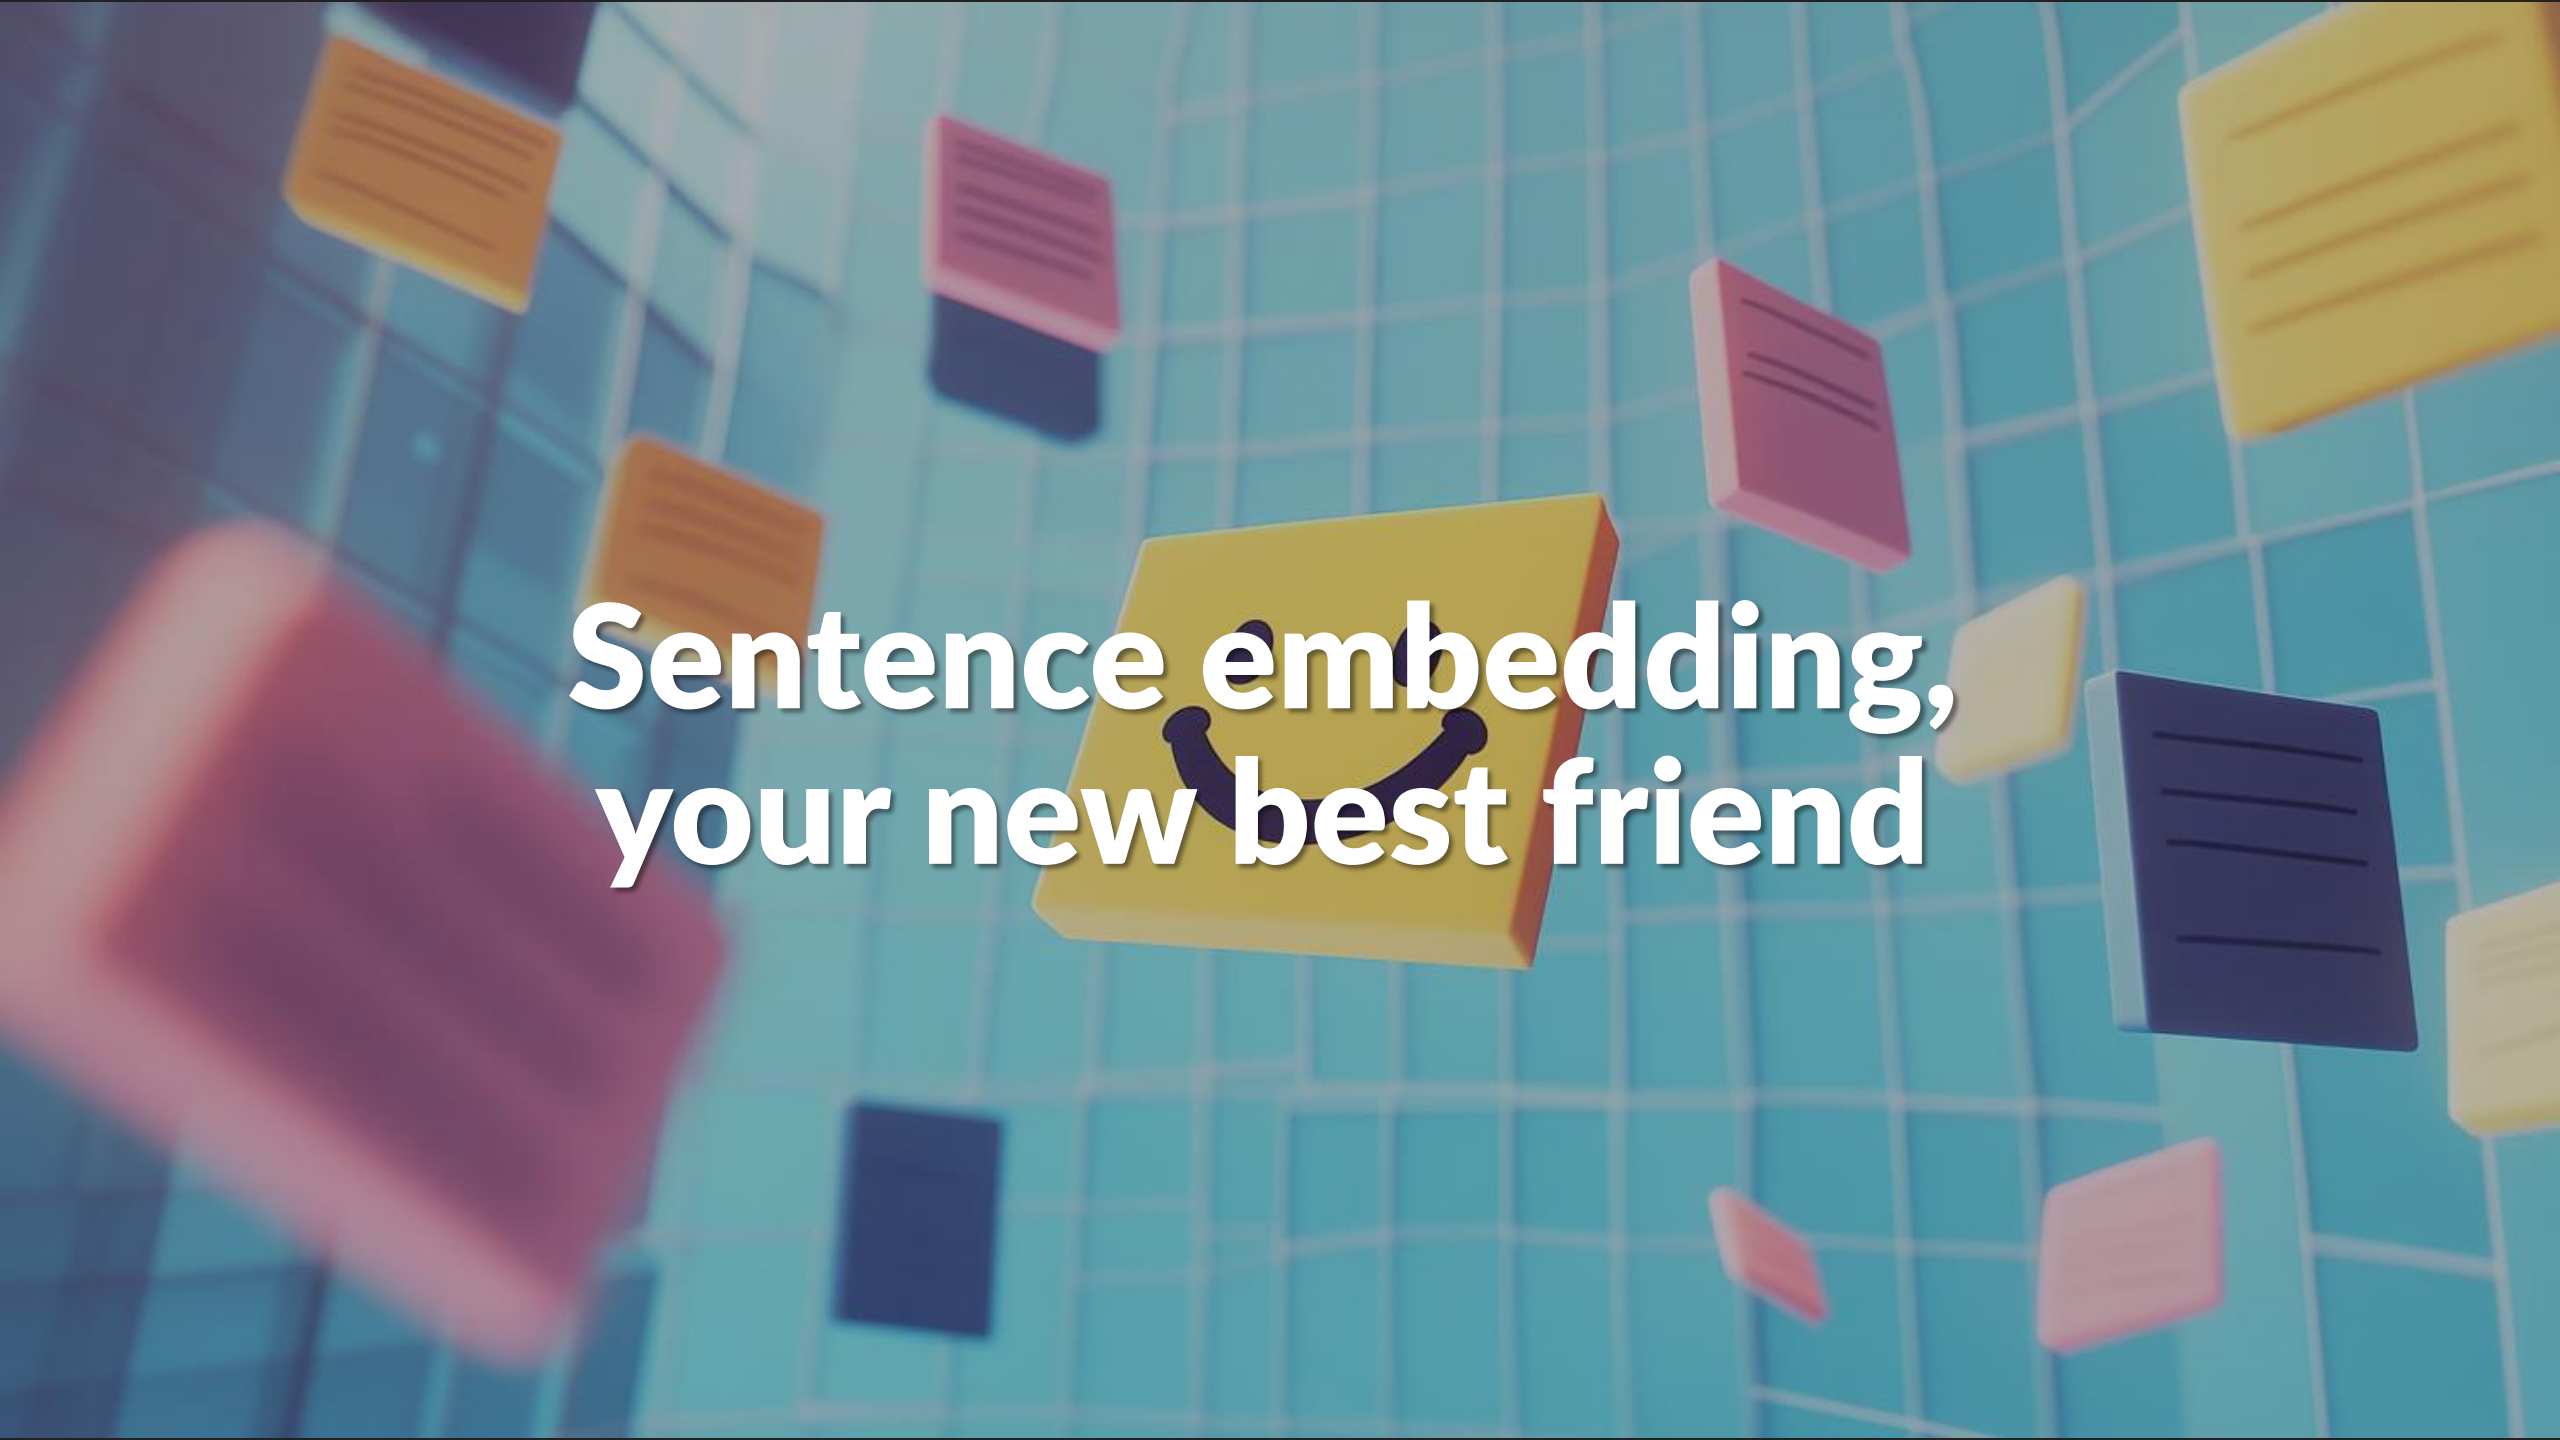


More flexible but steeper learning curve

Model-context Protocol

Wrapper around (your own) API so LLMs can read & perform actions on your data & servers with just one link




Sentence embedding,
your new best friend

Problem

Which of these three sentences have most similar meaning?

1: *“The train is quite noisy”*

2: *“Steam locomotives are loud”*

3: *“I train the quite pretty dog”*

How do you know? How to quantify?

→ **Sentence embedding:**

calculates “coordinate” based on meaning

Word Embedding Similarity Intuition



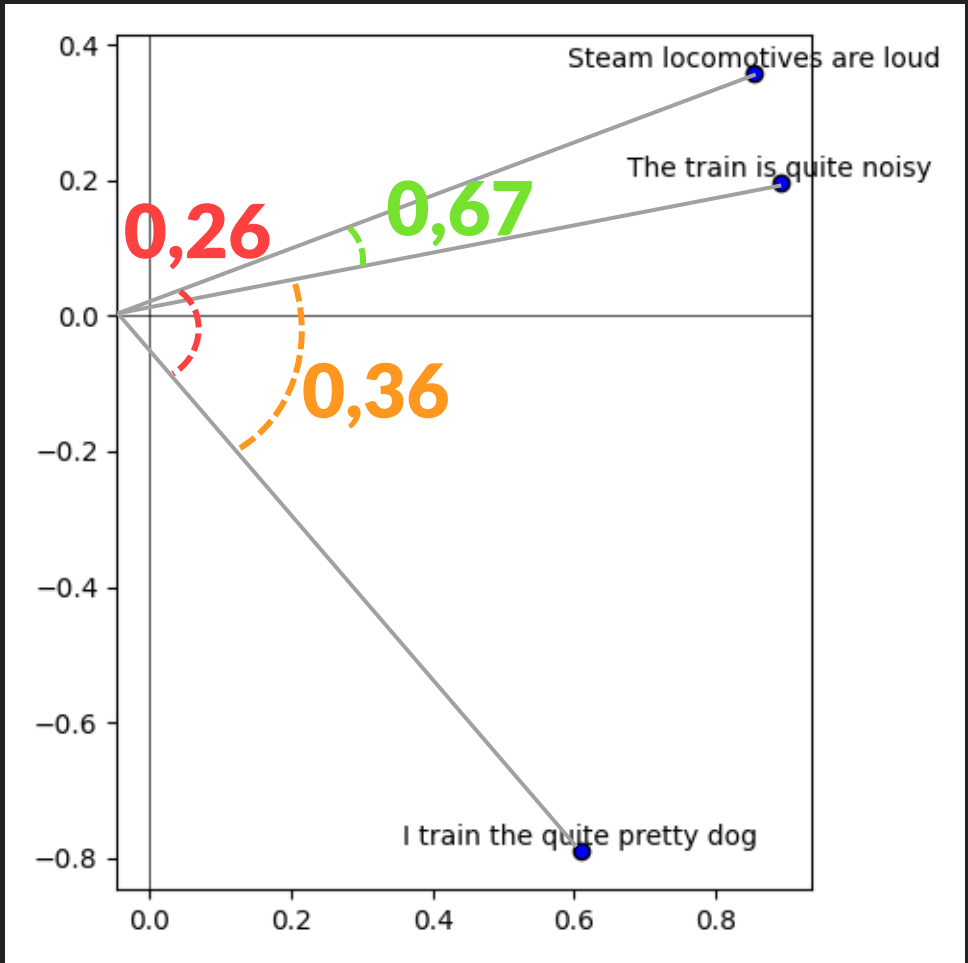
Every axis = word in neighbourhood
Then: reduce axes (e.g. PCA)

Learn which words are used in similar contexts

Not Euclidean distance but Cosine Similarity (angle)

→ cat & kitten similar, even though used much less frequent than dog

Sentence embeddings similarity on example



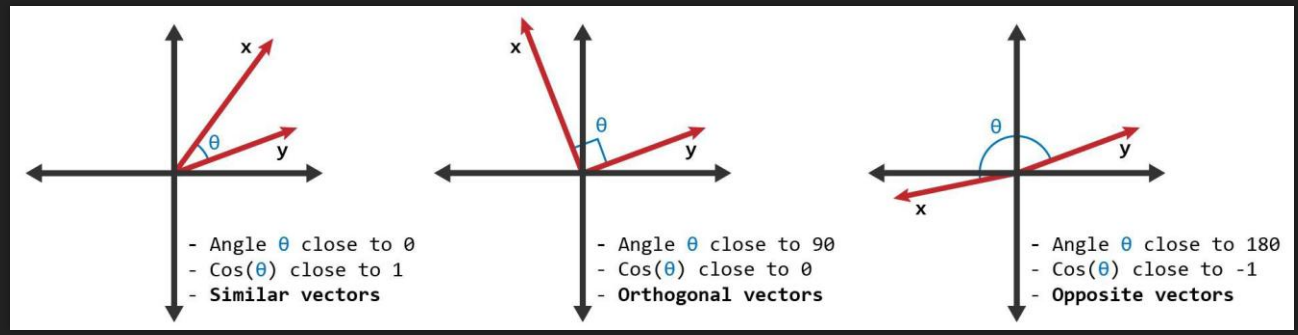
Cosine similarity:

$$\frac{\mathbf{A} \cdot \mathbf{B}}{\|\mathbf{A}\|\|\mathbf{B}\|} = \frac{\sum_{i=1}^{n} A_i B_i}{\sqrt{\sum_{i=1}^{n} A_i^2} \cdot \sqrt{\sum_{i=1}^{n} B_i^2}}$$

→ Sentence embeddings do find that the first two sentences are most similar

Cosine Similarity

$$\cos(\theta) = \frac{\mathbf{A} \cdot \mathbf{B}}{\|\mathbf{A}\|\|\mathbf{B}\|} = \frac{\sum_{i=1}^{n} A_i B_i}{\sqrt{\sum_{i=1}^{n} A_i^2} \sqrt{\sum_{i=1}^{n} B_i^2}}$$



For example: the cosine similarity of vectors (2,5,3) and (1,0,2)

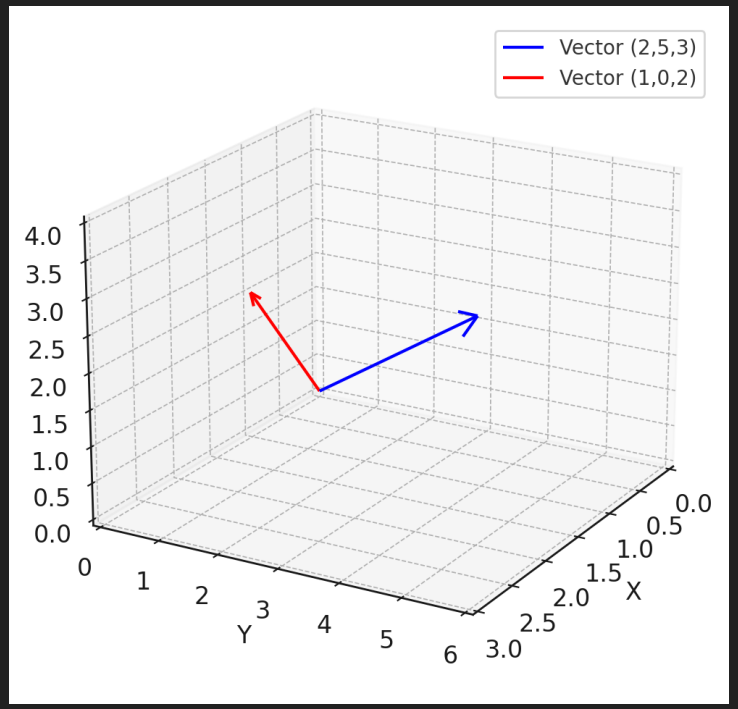
$$\cos(\theta) = \frac{(2,5,3) \cdot (1,0,2)}{||(2,5,3)|| \cdot ||(1,0,2)||}$$

$$= \frac{(2\cdot1)+(5\cdot0)+(3\cdot2)}{\sqrt{2^2+5^2+3^2} \cdot \sqrt{1^2+0^2+2^2}}$$

$$= \frac{8}{\sqrt{38} \cdot \sqrt{5}}$$

$$= \frac{8}{\sqrt{190}}$$

$$= 0.58$$



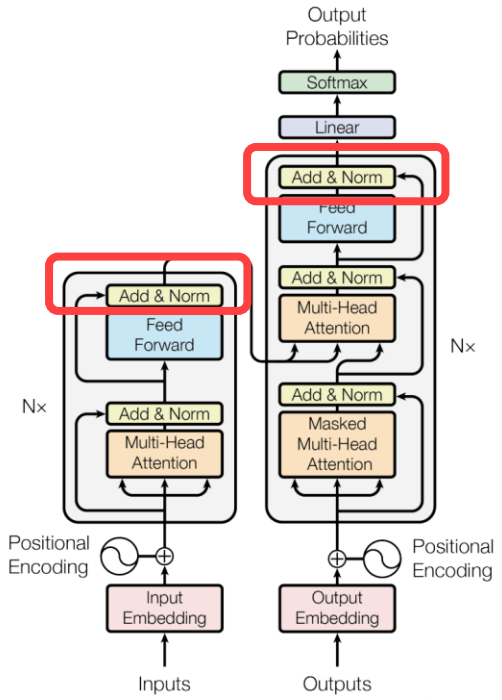
Use Transformer LLM's Embeddings

Large language models internally capture meaning of text in the layers right before predicting the masked/next words.



BERT

Encoder



GPT

Decoder



Bidirectional Masked Language Models like BERT natural for this task

But now GPT-based embeddings also popular: use final hidden state of last token or mean pool across embeddings of all tokens in final layer

Calculating Embeddings

Through APIs

(e.g. OpenAI)

```python
from openai import OpenAI
client = OpenAI()

response = client.embeddings.create(
    input="I train the quite pretty dog",
    model="text-embedding-3-small"
)
```

https://platform.openai.com/docs/guides/embeddings

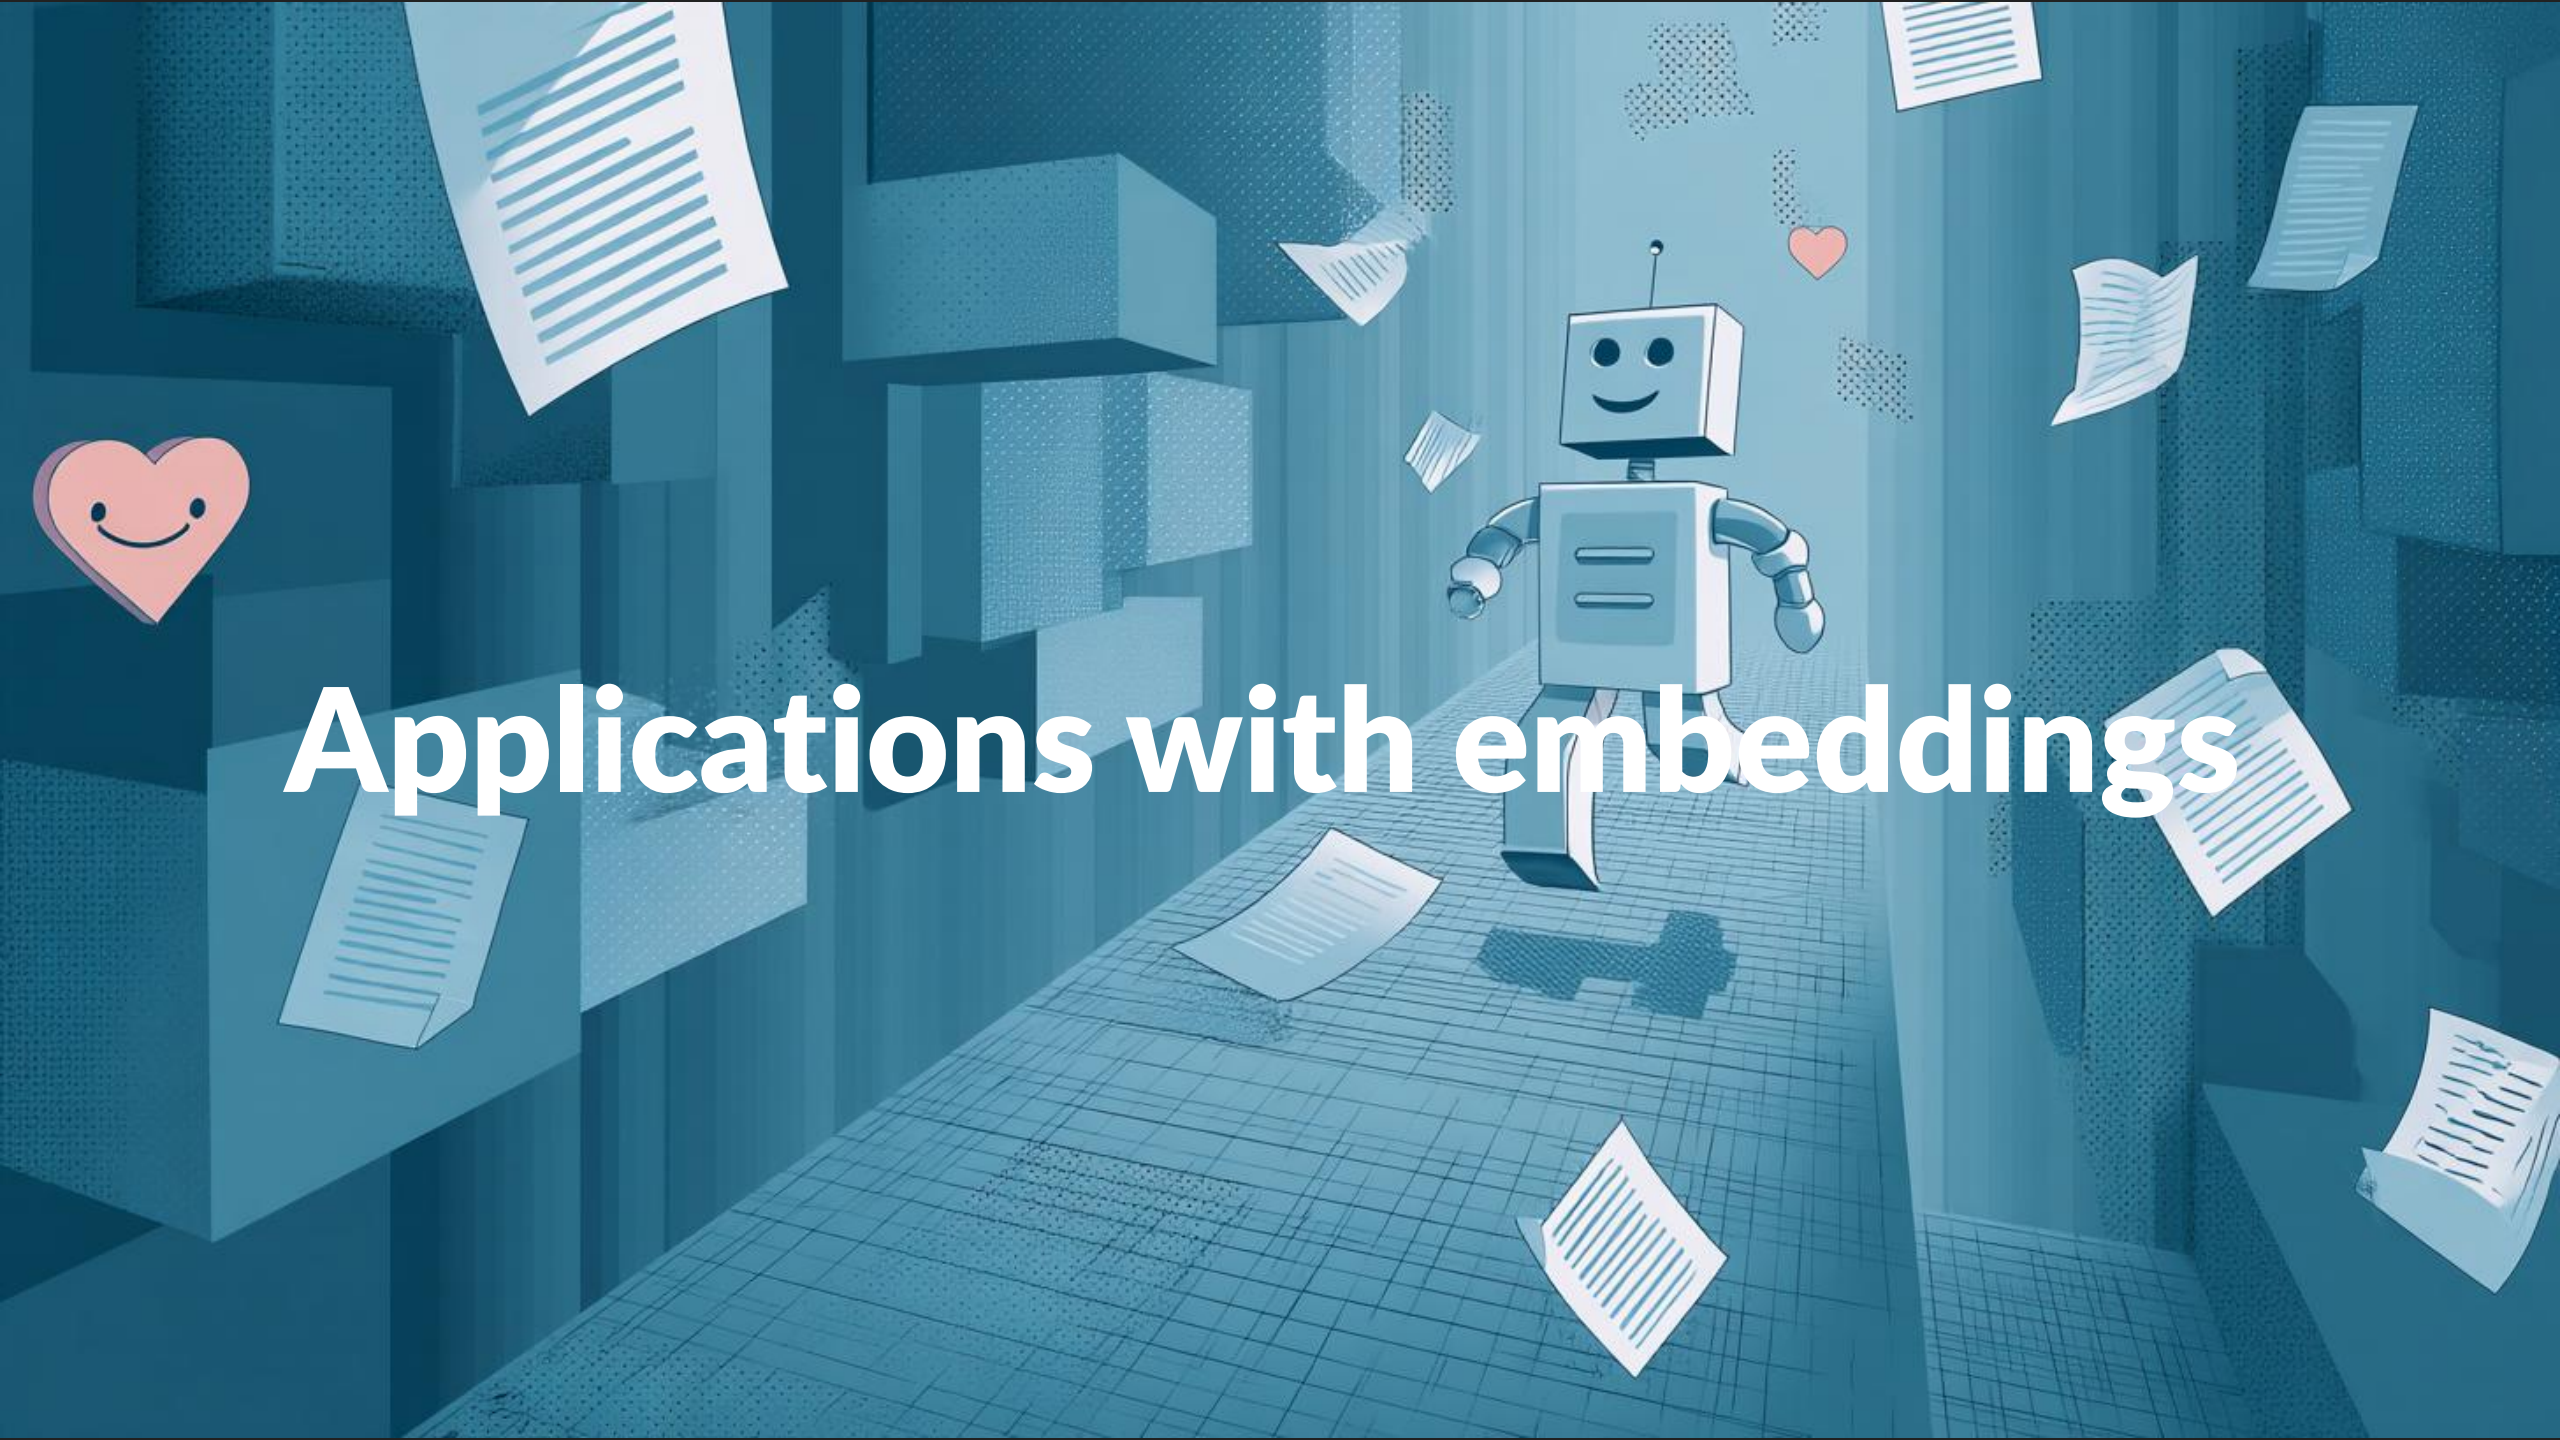
Locally

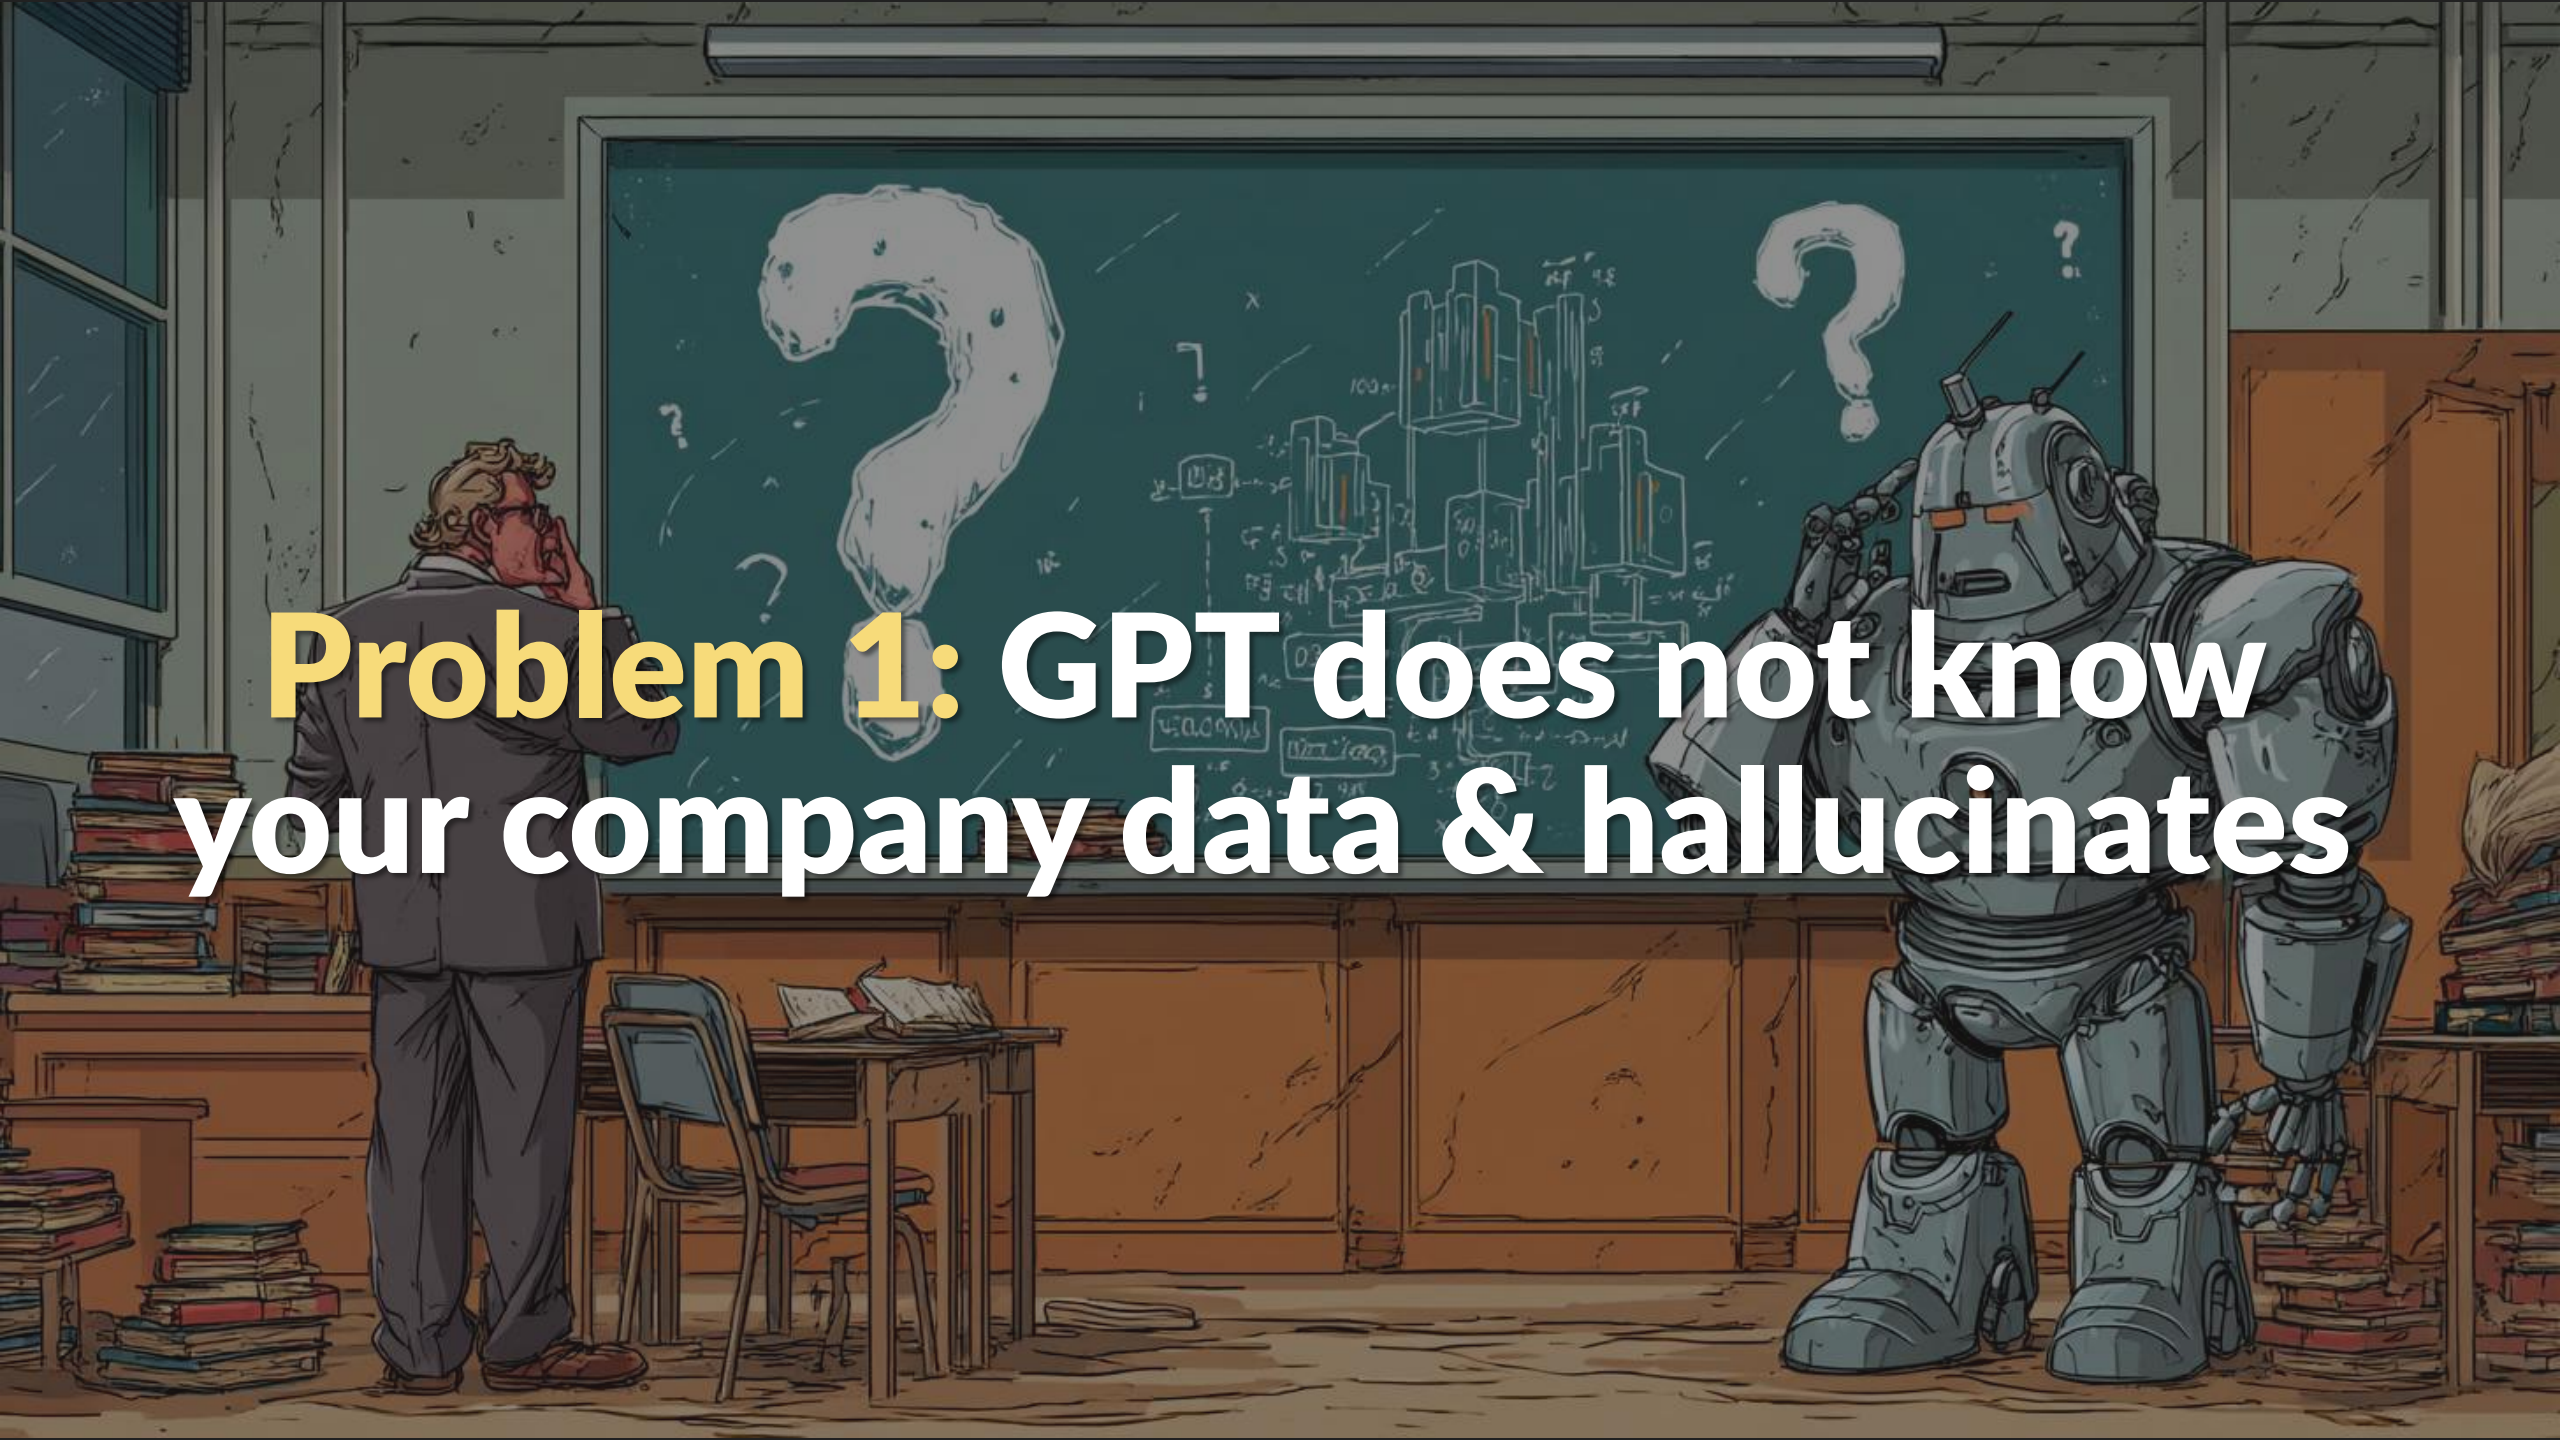
(sentence-transformers)

```python
from sentence_transformers import SentenceTransformer

model = SentenceTransformer("all-MiniLM-L6-v2")
embeddings = model.encode(["I train the quite pretty dog"])
```

https://sbert.net/


Applications with embeddings


Problem 1: GPT does not know
your company data & hallucinates

🔎 Retrieval-Augmented Generation (RAG)

🤔 User query

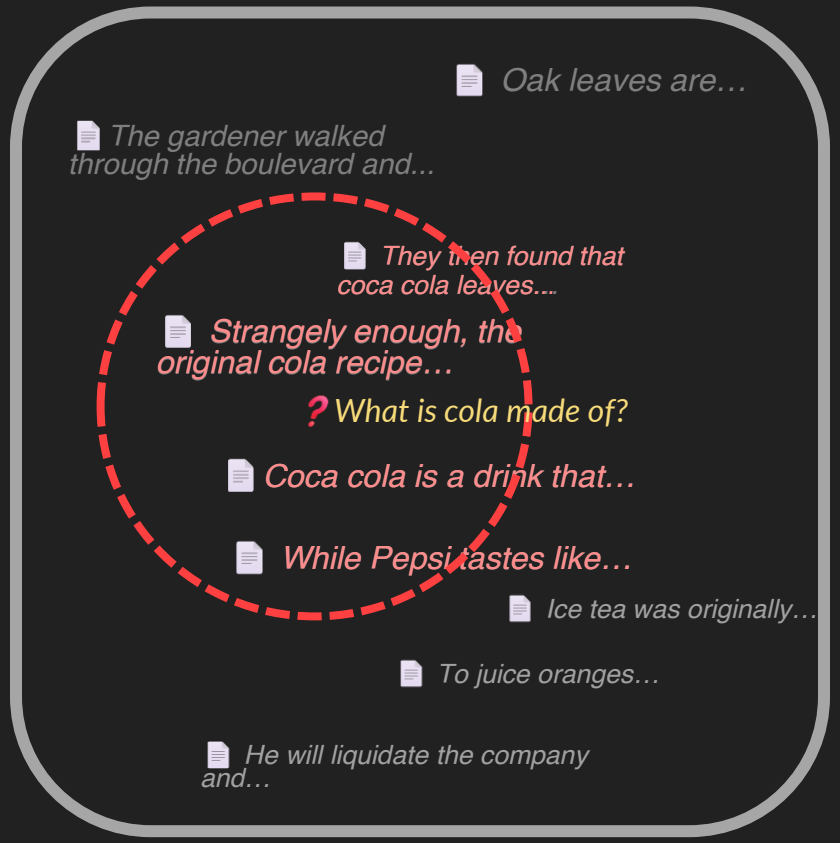
"What is cola made of?"

Company information

🤖 embedder

🤖 embedder

Sentence embedding:
[0.14, -0.33, 0.25, ...]



Prompt GPT with relevant texts:

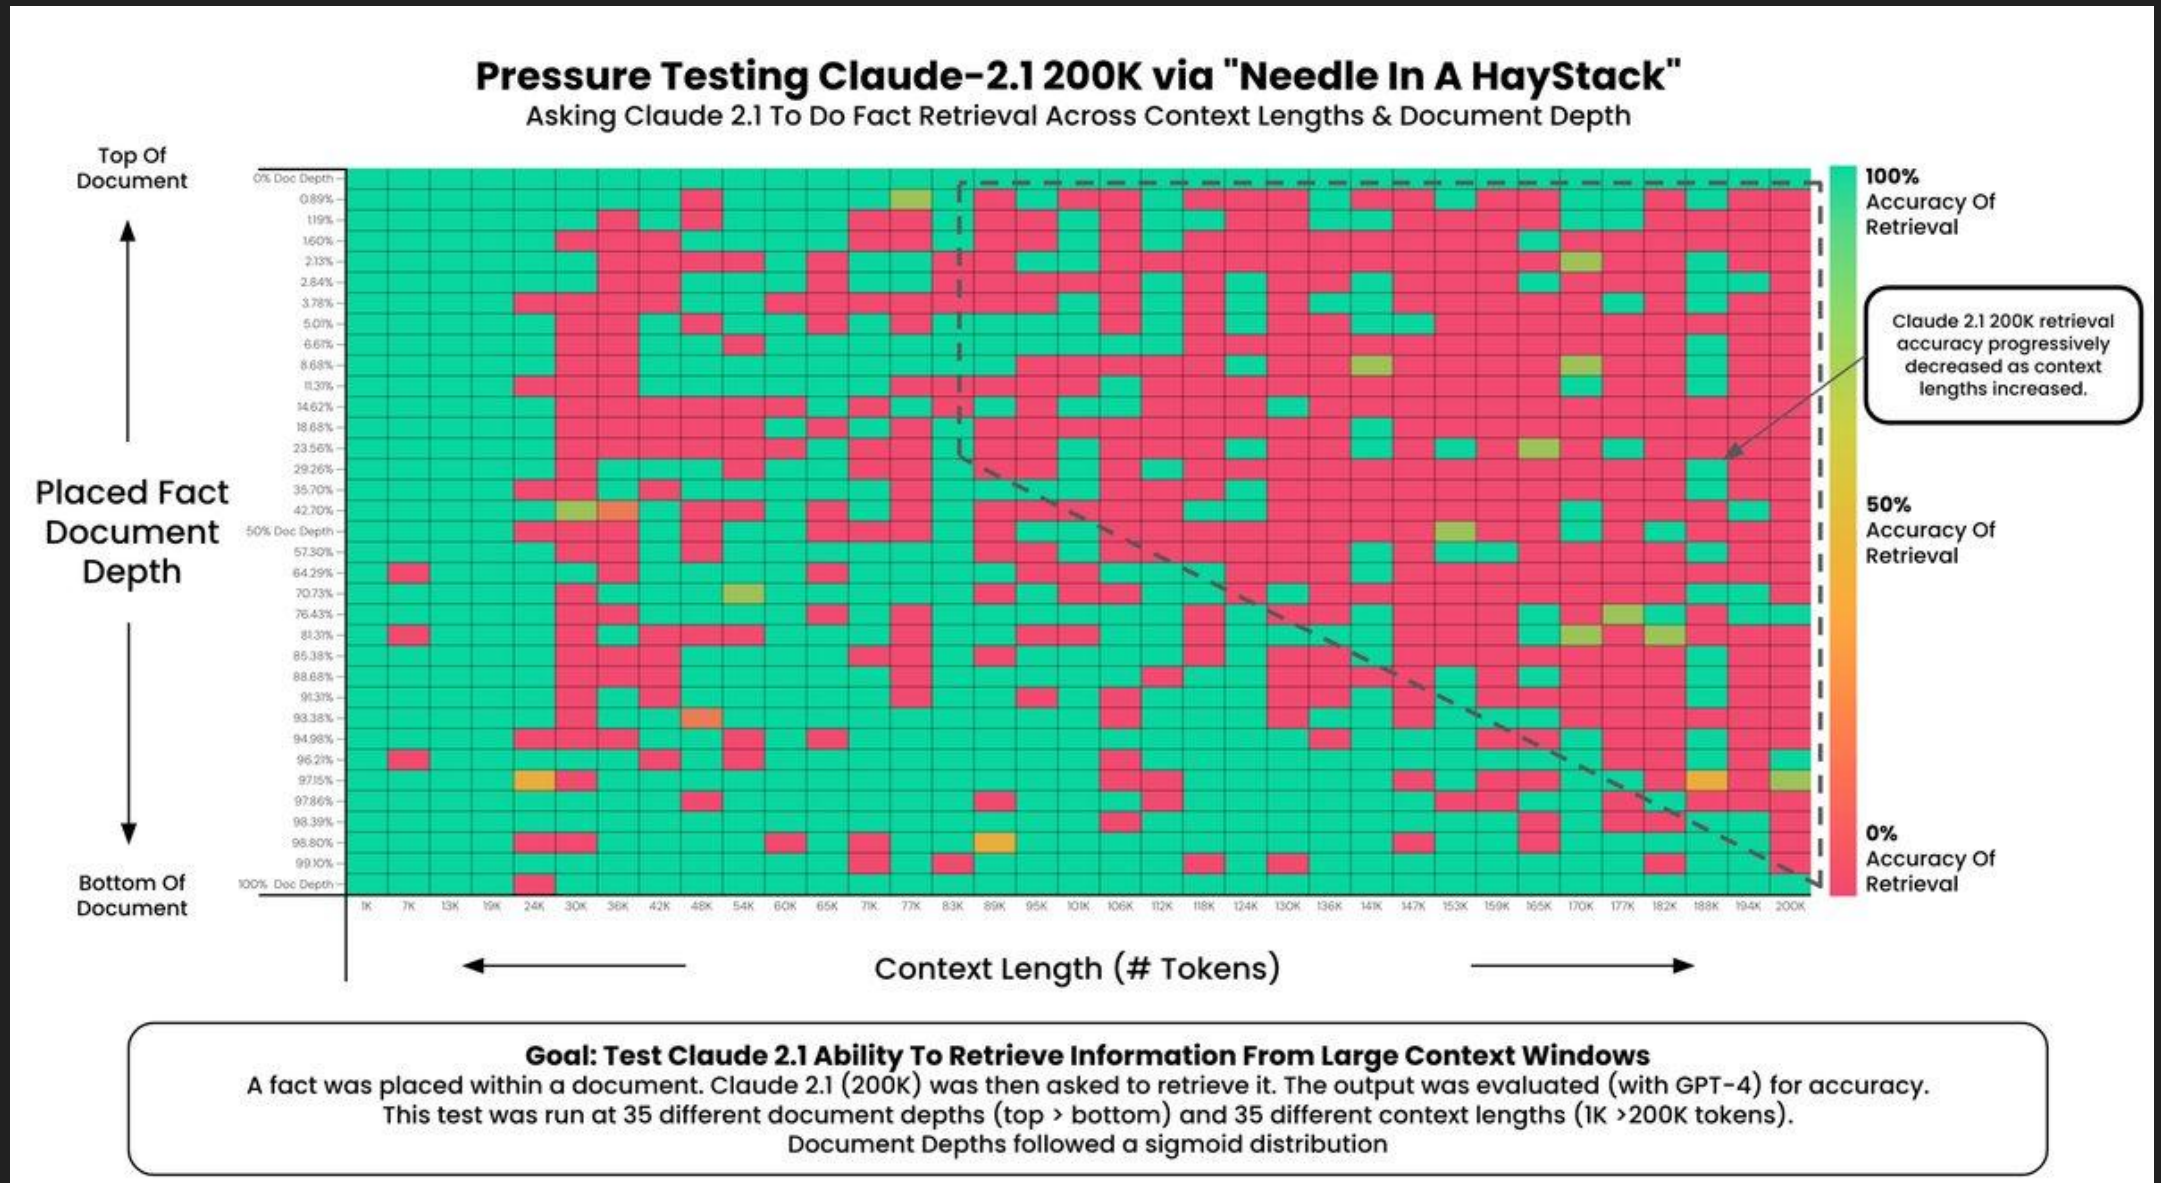
"Given these sources:
[1]: Coca cola is a drink that...

[2]: Strangely enough, the original cola recipe...
...

Answer the following question based on the data above, and cite the sources with their number

Question: What is cola made of?"

Long contexts → Less accurate retrieval of facts

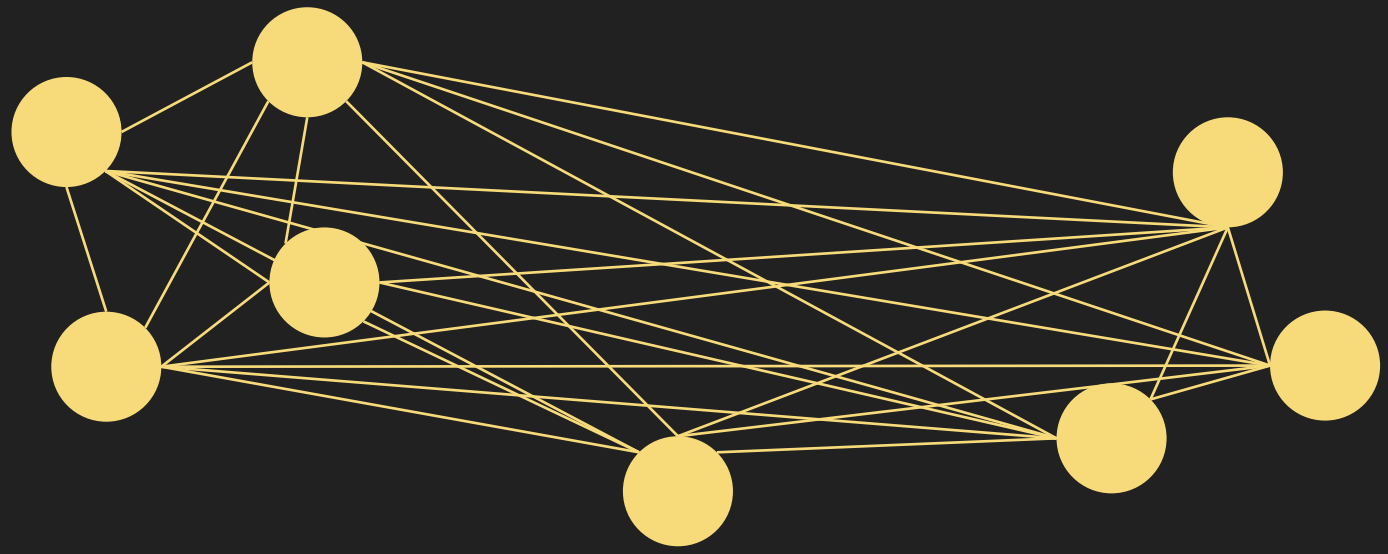




Problem 2: Unstructured textual data

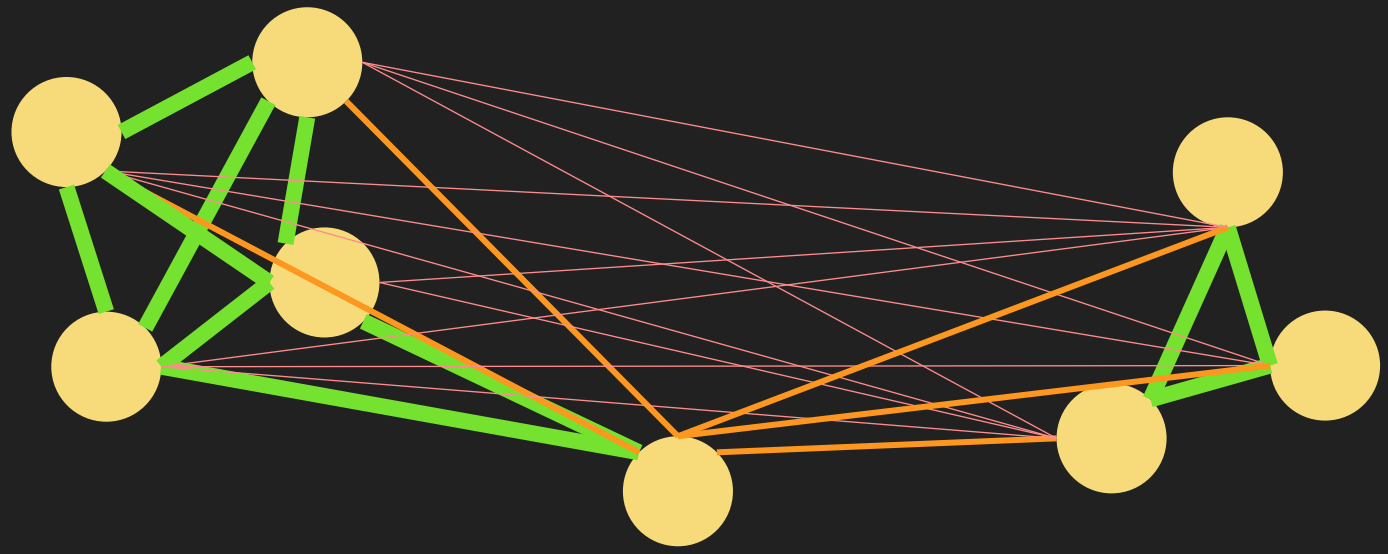
Find meaningful communities in texts

1. Take all pairs of nodes



Find meaningful communities in texts

2. Calculate similarity for all pairs

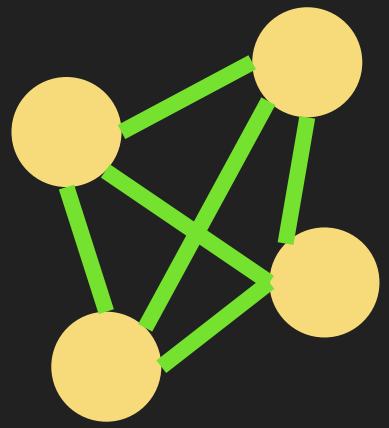


Find meaningful communities in texts

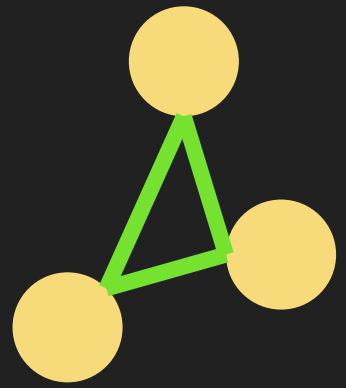
3. Filter so only above threshold similarity & find cliques

(=fully connected subgraphs)

Community 1



Community 2





Example: Rik Torfs tweets communities

Community 128: Kort cynisme

- Spaar de cynicus, want cynisme is onuitgesproken verdriet.

- Er zijn veel cynici omdat achter cynisme verdriet schuilgaat.

- Cynisme is weemoedige humor van de verliezer.

- Cynisme is een vergissing. Maar cynisch willen zijn, en er het talent voor missen... Dat is pure tragiek.

- Word nooit cynisch, zelfs al kan u soms niet anders dan cynisme als wapen te hanteren.

- Cynisme vloeit voort uit het onvermogen tot verdriet.

- Een beetje cynisme hoeft geen ramp te zijn. Zolang je er niet trots op bent.

- Geloof is de beste remedie tegen cynisme. Daarom doen sommigen er cynisch over.

- De cynicus beweent zijn eigen mislukking en acht de samenleving daarvoor schuldig. Cynisme is een vorm van zelfoverschatting.

- Troost de cynici. Verzacht hun verdriet.

Community 128: Centrum van de wereld

- Als je niet het centrum van de wereld bent, ziet de wereld er plotseling beter uit.

- Als je er zelf niet het centrum van bent, wordt de wereld groter.

- Als je jezelf niet het centrum van de wereld voelt, ziet hij er opeens veel beter uit.

- Zolang we niet denken dat we zelf het centrum van de wereld zijn, heeft de wereld toekomst.

- Als je zelf niet het centrum van de wereld bent, wordt hij meteen veel groter.

- Wie zichzelf het centrum van de wereld acht, kent de wereld niet goed.

- Als je zelf niet langer het centrum van de wereld bent, ziet hij er plotseling veel mooier uit.

- Als je zelf niet het centrum van de wereld bent, wordt hij plotseling veel mooier.

- Als je niet het centrum van de wereld bent maar een doodgewone voorbijganger, kun je veel beter kijken.

- Als je zelf niet het centrum van de wereld bent, kun je hem beter zien met de nodige afstand.

Community 131: Nu zijn we slimmer

- Waarom denken wij slimmer te zijn dan de vorige generaties, terwijl zij exact hetzelfde dachten?

- We denken altijd dat we nu slimmer zijn dan vroeger. Daarin leren we niet bij.

- Onze generatie is slimmer dan de vorige. Dat vond de vorige ook. Niets verandert zolang alles beter wordt.

- Te gemakkelijk denken we intelligenter te zijn dan wie anders denken. Terwijl ze misschien juist anders denken omdat ze intelligenter zijn.

- Geef mensen altijd de indruk dat ze slimmer zijn dan jezelf. Dat hebben ze graag. En vaak is het nog waar ook.

- Wie zich slimmer voelt dan de rest van de mensheid heeft wellicht te weinig nagedacht.

- We denken graag dat we alles beter weten dan de vorige generatie. Tot we zelf de vorige generatie zijn.

- Waarom denken we toch dat we nu wijzer zijn dan vroeger? Terwijl het misschien juist andersom is.

- Waarom denken wij dat wij met de jaren wijzer worden, en nooit het tegendeel?

- Geloof is iets van vroeger toen we nog maar weinig wisten: Thomas van Aquino, Erasmus, Pascal. Gelukkig zijn we veel slimmer geworden.

Community 132: Wat is humor?

- Humor is wat overblijft als de lust tot lachen iedereen is vergaan.

- Humor is de pijnlijke stilte die valt na een poging tot grap.

- Humor is nooit kwetsend als er werkelijk niemand is die lacht.

- Humor is niet meer van deze tijd. Want het gevaar bestaat dat er gelachen wordt.

- Humor is vooral belangrijk als er niet veel te lachen valt.

- Humor is wat overblijft wanneer redelijkheid niet helpt.

- Humor is liefde. Als je iemand niet aardig vindt, veins dan ernst, prijs het verstand van je gesprekspartner en verlies zijn visitekaartje.

- Humor is wat overblijft als we het over alles oneens zijn en elkaar als persoon toch respecteren.

- Humor mag nog altijd. Zolang geen mens zich gekwetst voelt, de samenleving er rechtvaardiger door wordt en niemand lacht.

- Humor is populisme. Tenzij niemand ermee lacht, dan is het kwaliteit.

Package full of tools: sentence-transformers

sentence_transformers.util.community_detection(*embeddings: Tensor | ndarray, threshold: float = 0.75, min_community_size: int = 10, batch_size: int = 1024, show_progress_bar: bool = False*) → list[list[int]] [source]

Function for Fast Community Detection.

Finds in the embeddings all communities, i.e. embeddings that are close (closer than threshold). Returns only communities that are larger than min_community_size. The communities are returned in decreasing order. The first element in each list is the central point in the community.

Parameters:
- **embeddings** (*torch.Tensor or numpy.ndarray*) – The input embeddings.
- **threshold** (*float*) – The threshold for determining if two embeddings are close. Defaults to 0.75.
- **min_community_size** (*int*) – The minimum size of a community to be considered. Defaults to 10.
- **batch_size** (*int*) – The batch size for computing cosine similarity scores. Defaults to 1024.
- **show_progress_bar** (*bool*) – Whether to show a progress bar during computation. Defaults to False.

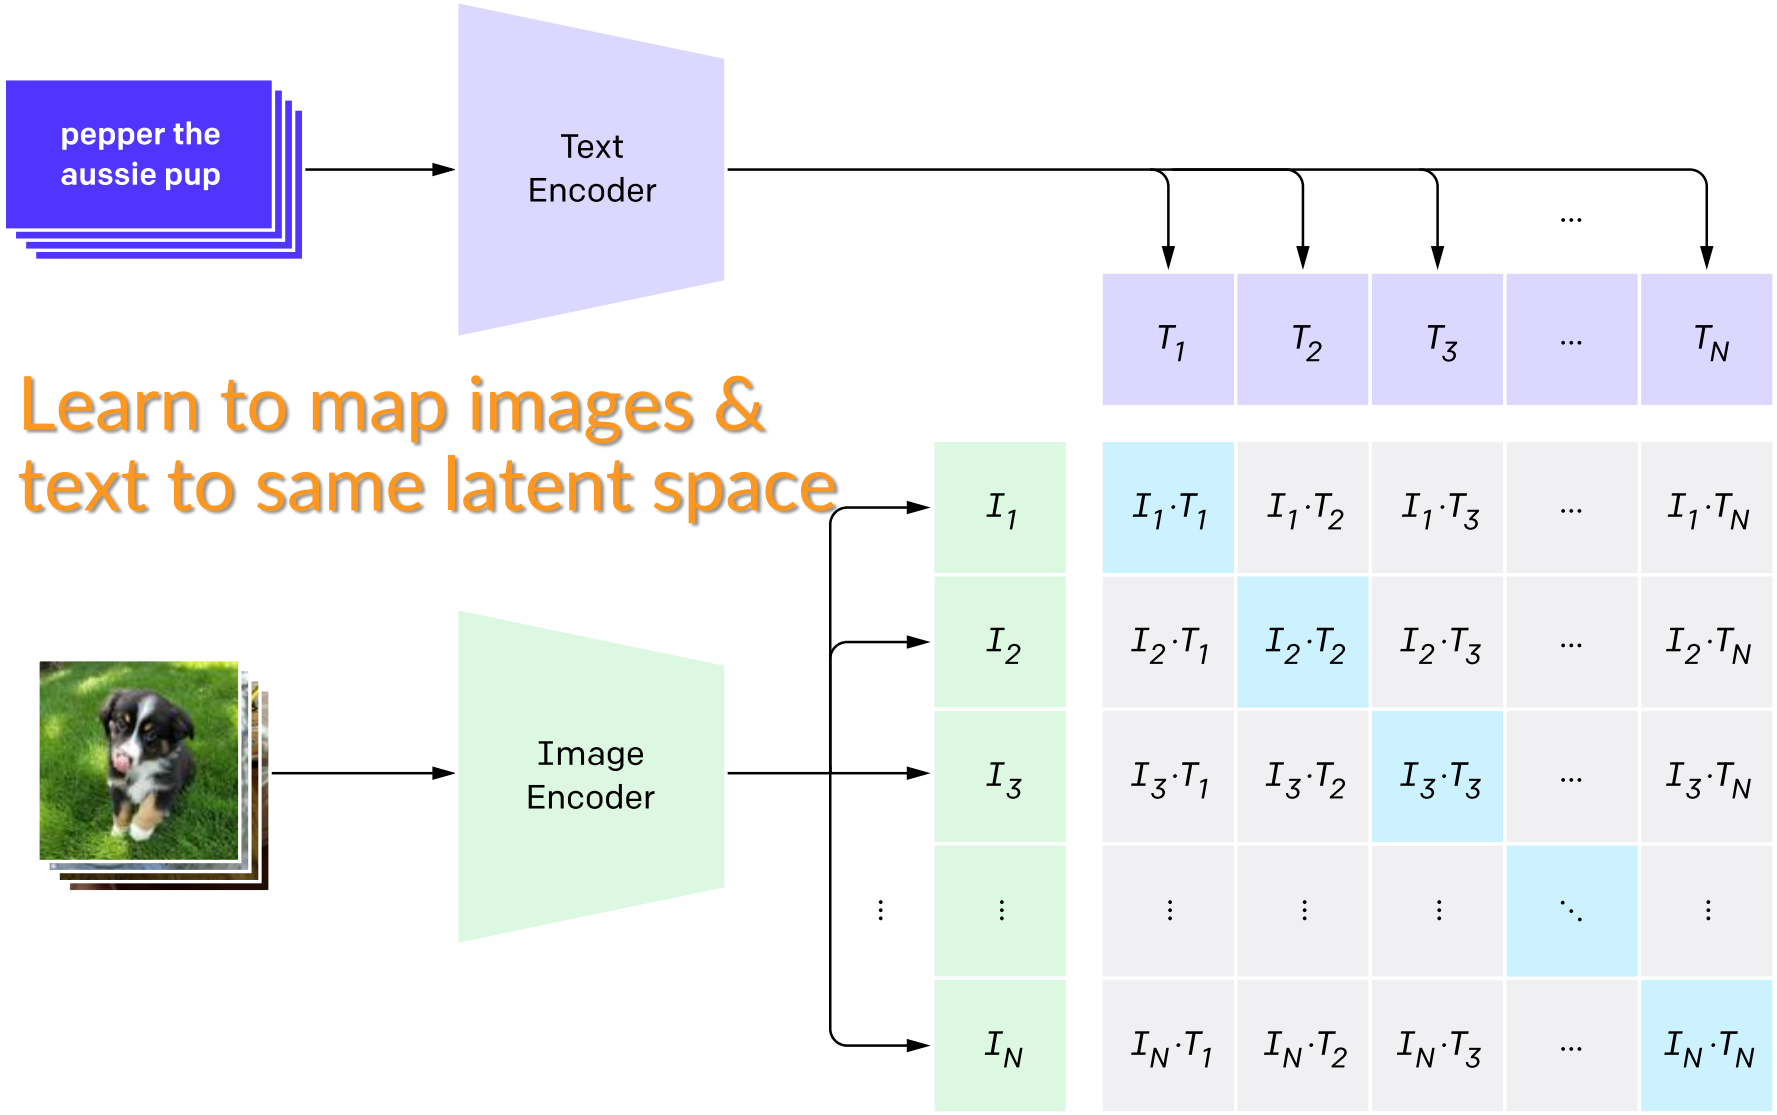
Returns: A list of communities, where each community is represented as a list of indices.

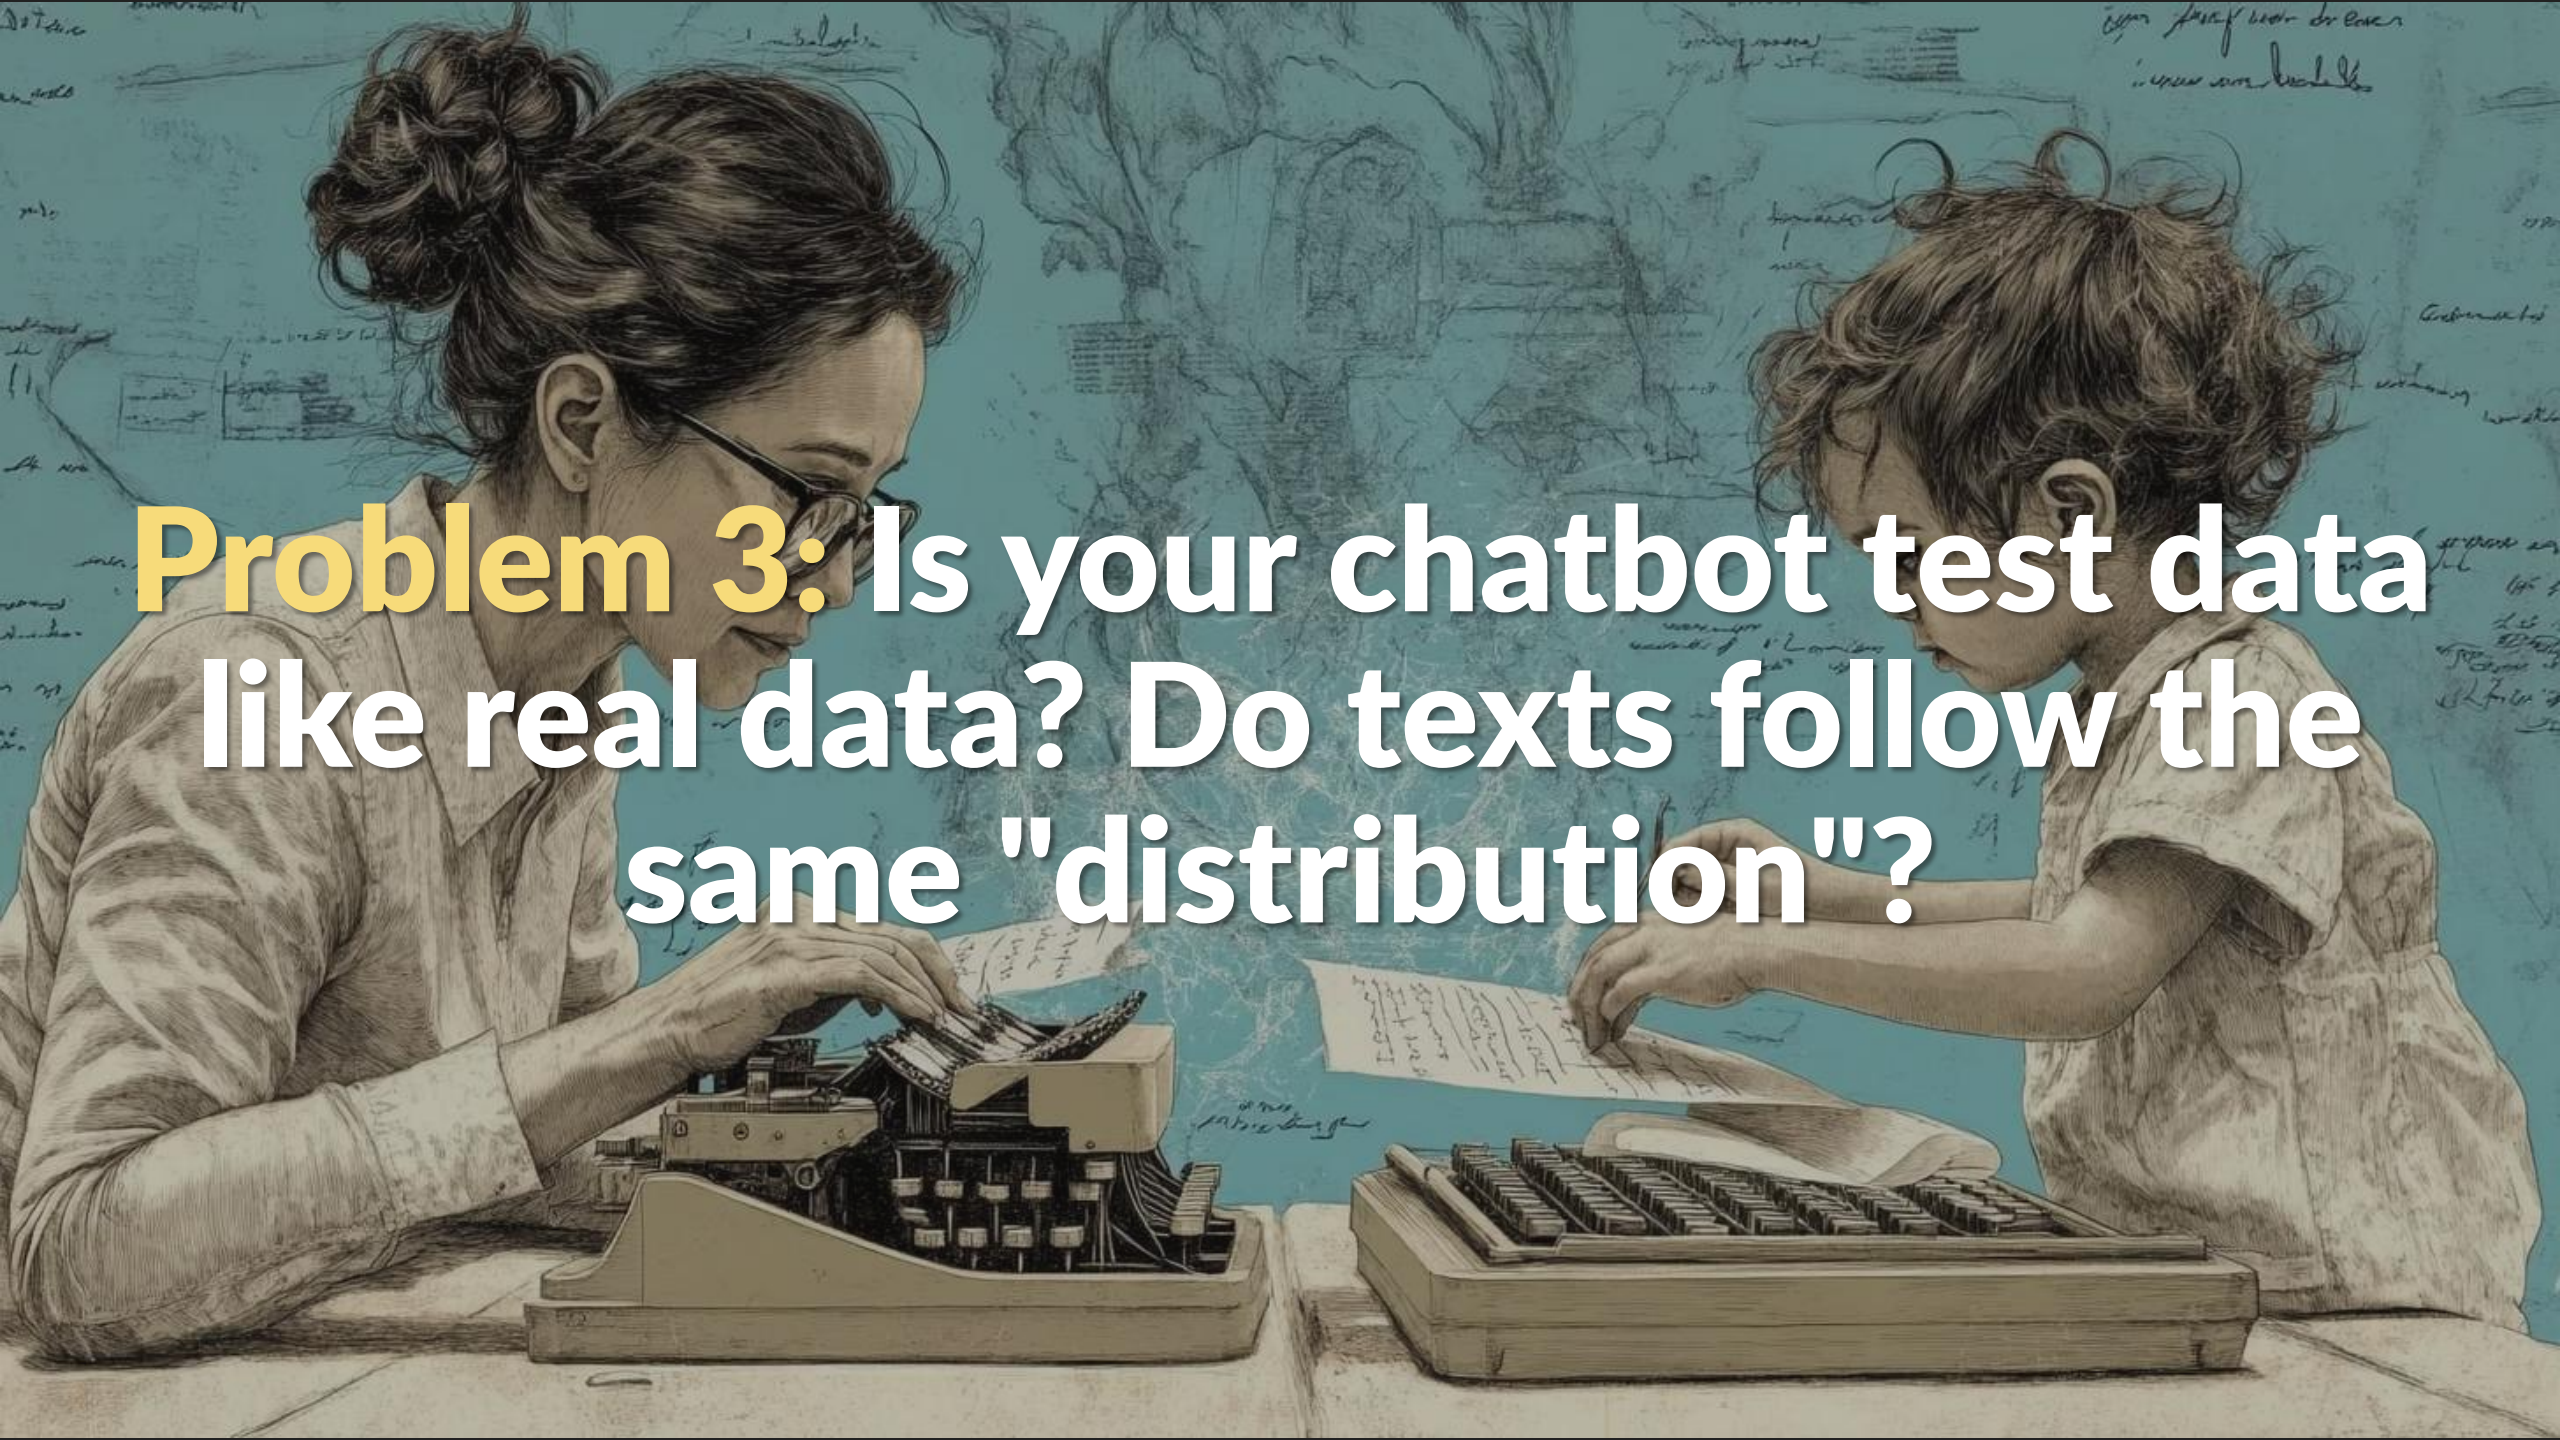
Return type: List[List[int]]

https://sbert.net/docs/package_reference/util.html#sentence_transformers.util.community_detection

CLIP: Cross-Modal Embeedings

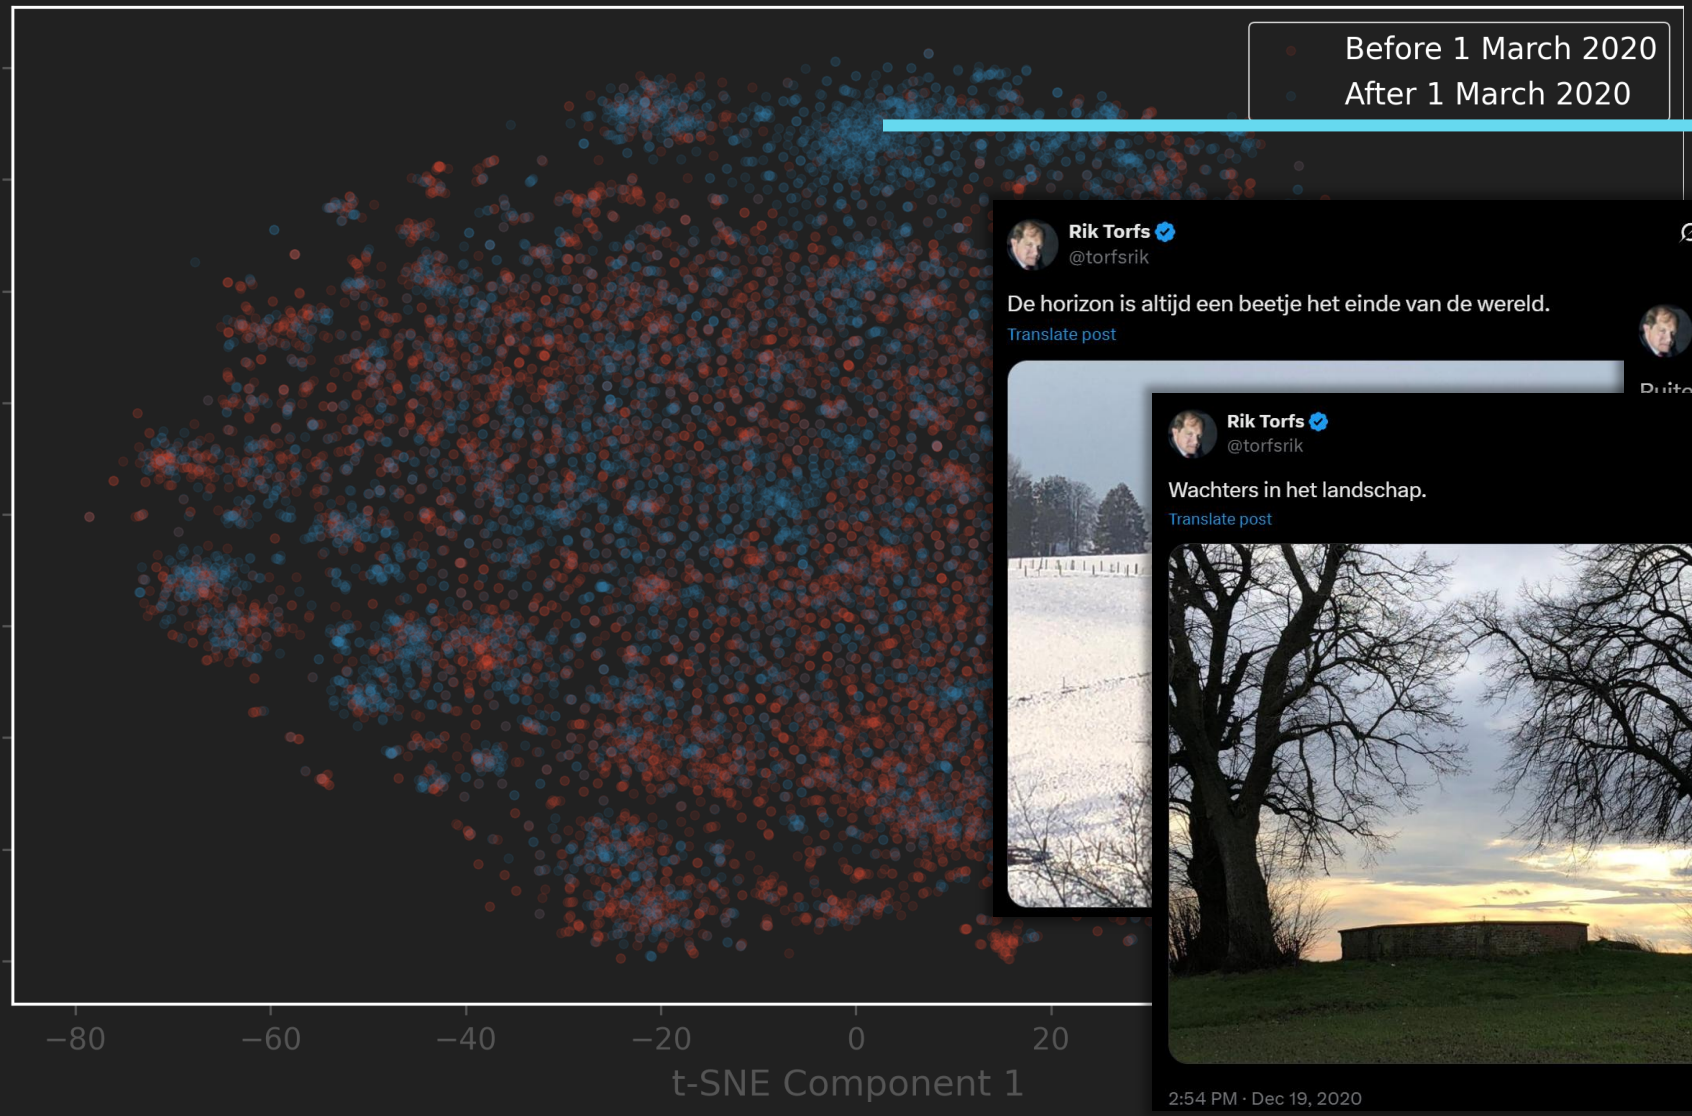
1. Contrastive pre-training




Problem 3: Is your chatbot test data like real data? Do texts follow the same "distribution"?

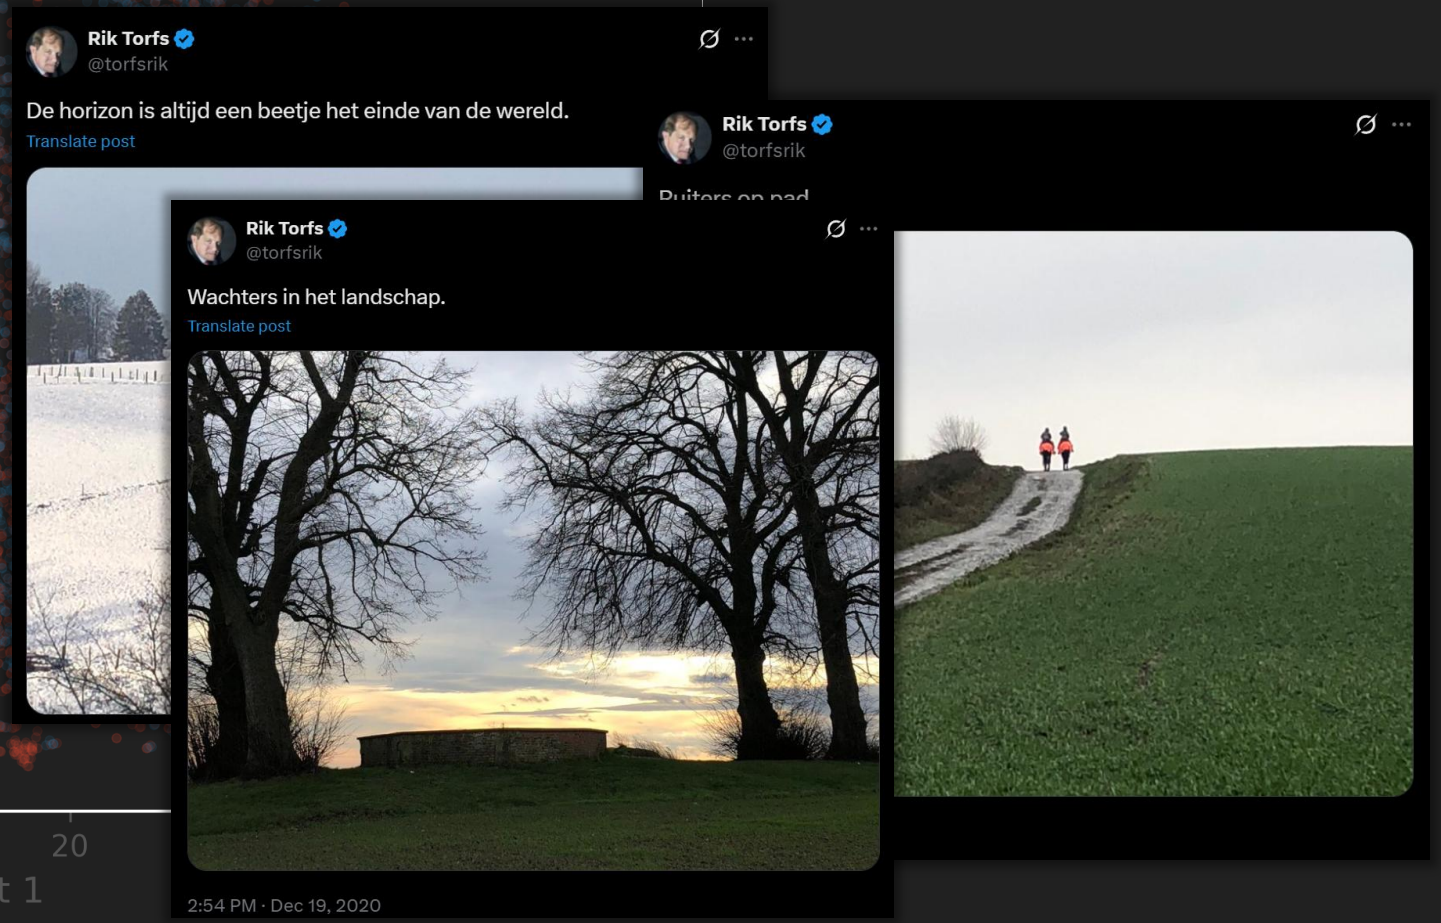
Check if similar “distribution” of texts

Rik Torfs tweets embeddings (RobBERT-2022) before and after 1 March 2020

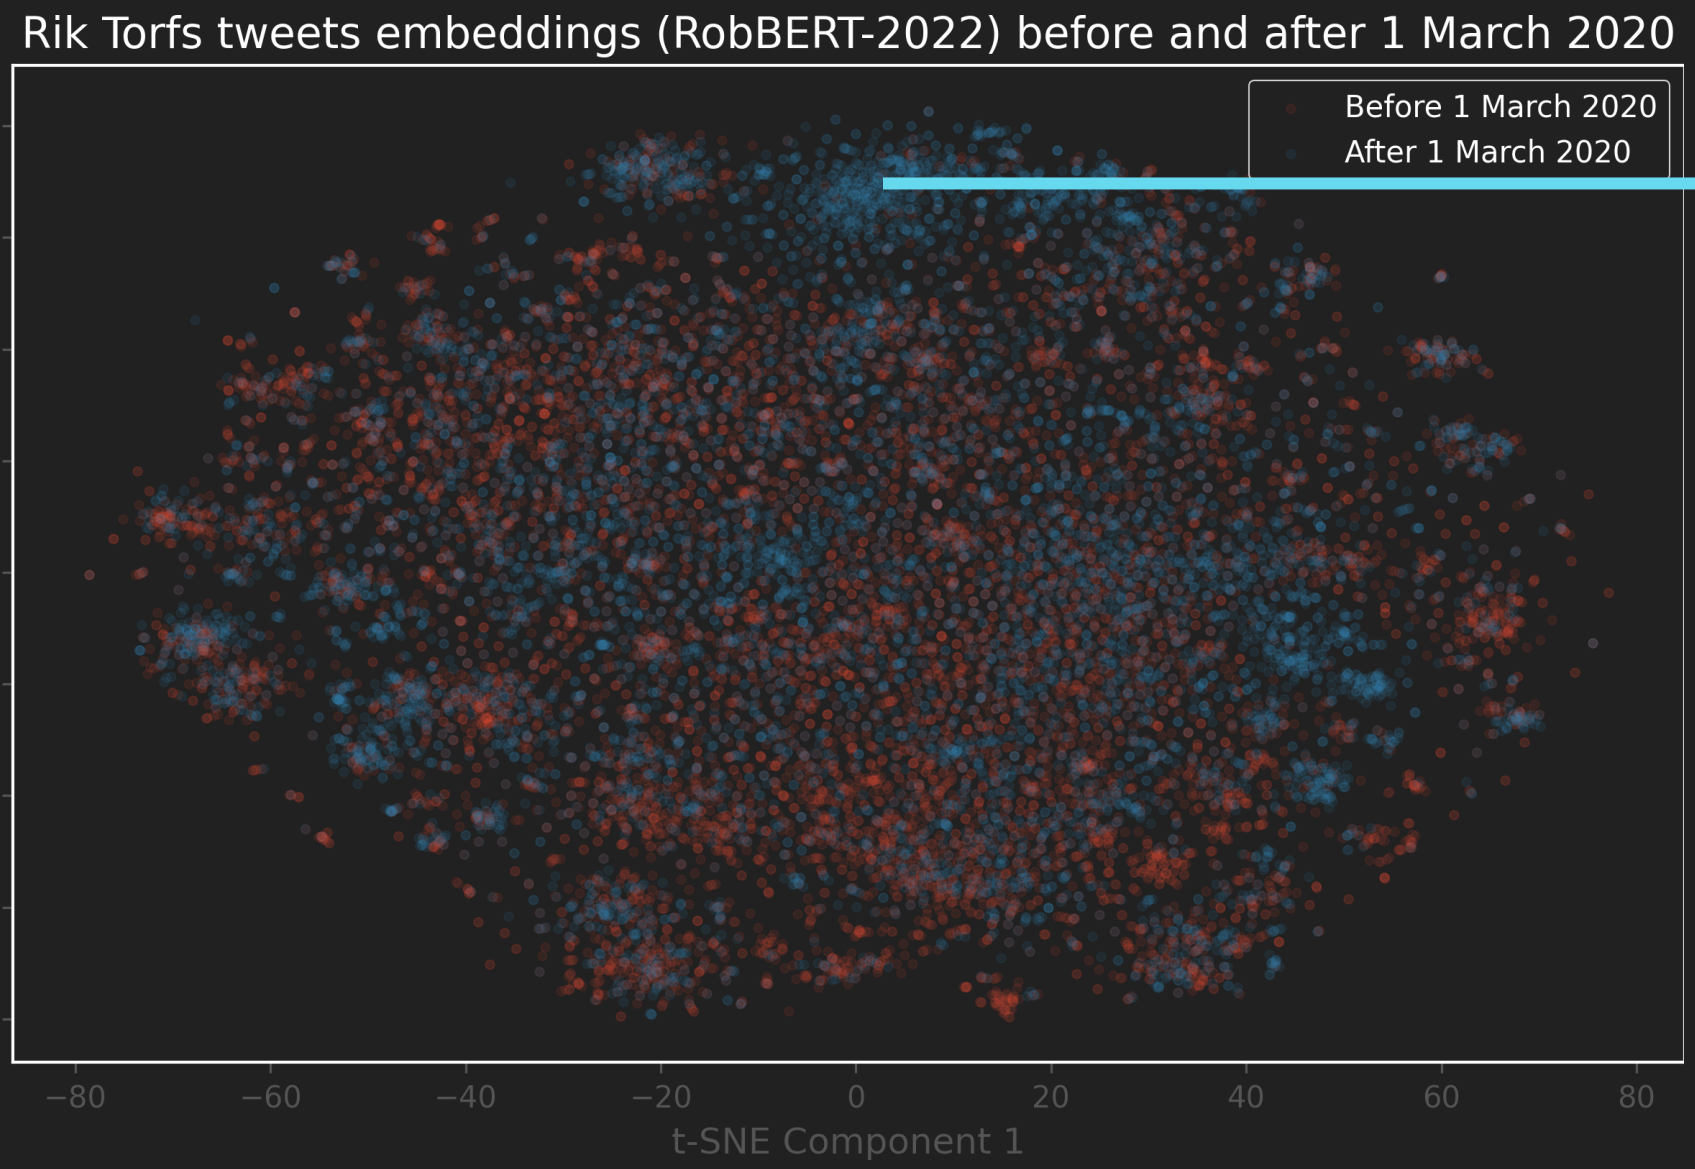


Kmeans clustering centroid distance difference

→ finds the new category: Rik Torfs walk tweets



Check if similar “distribution” of texts

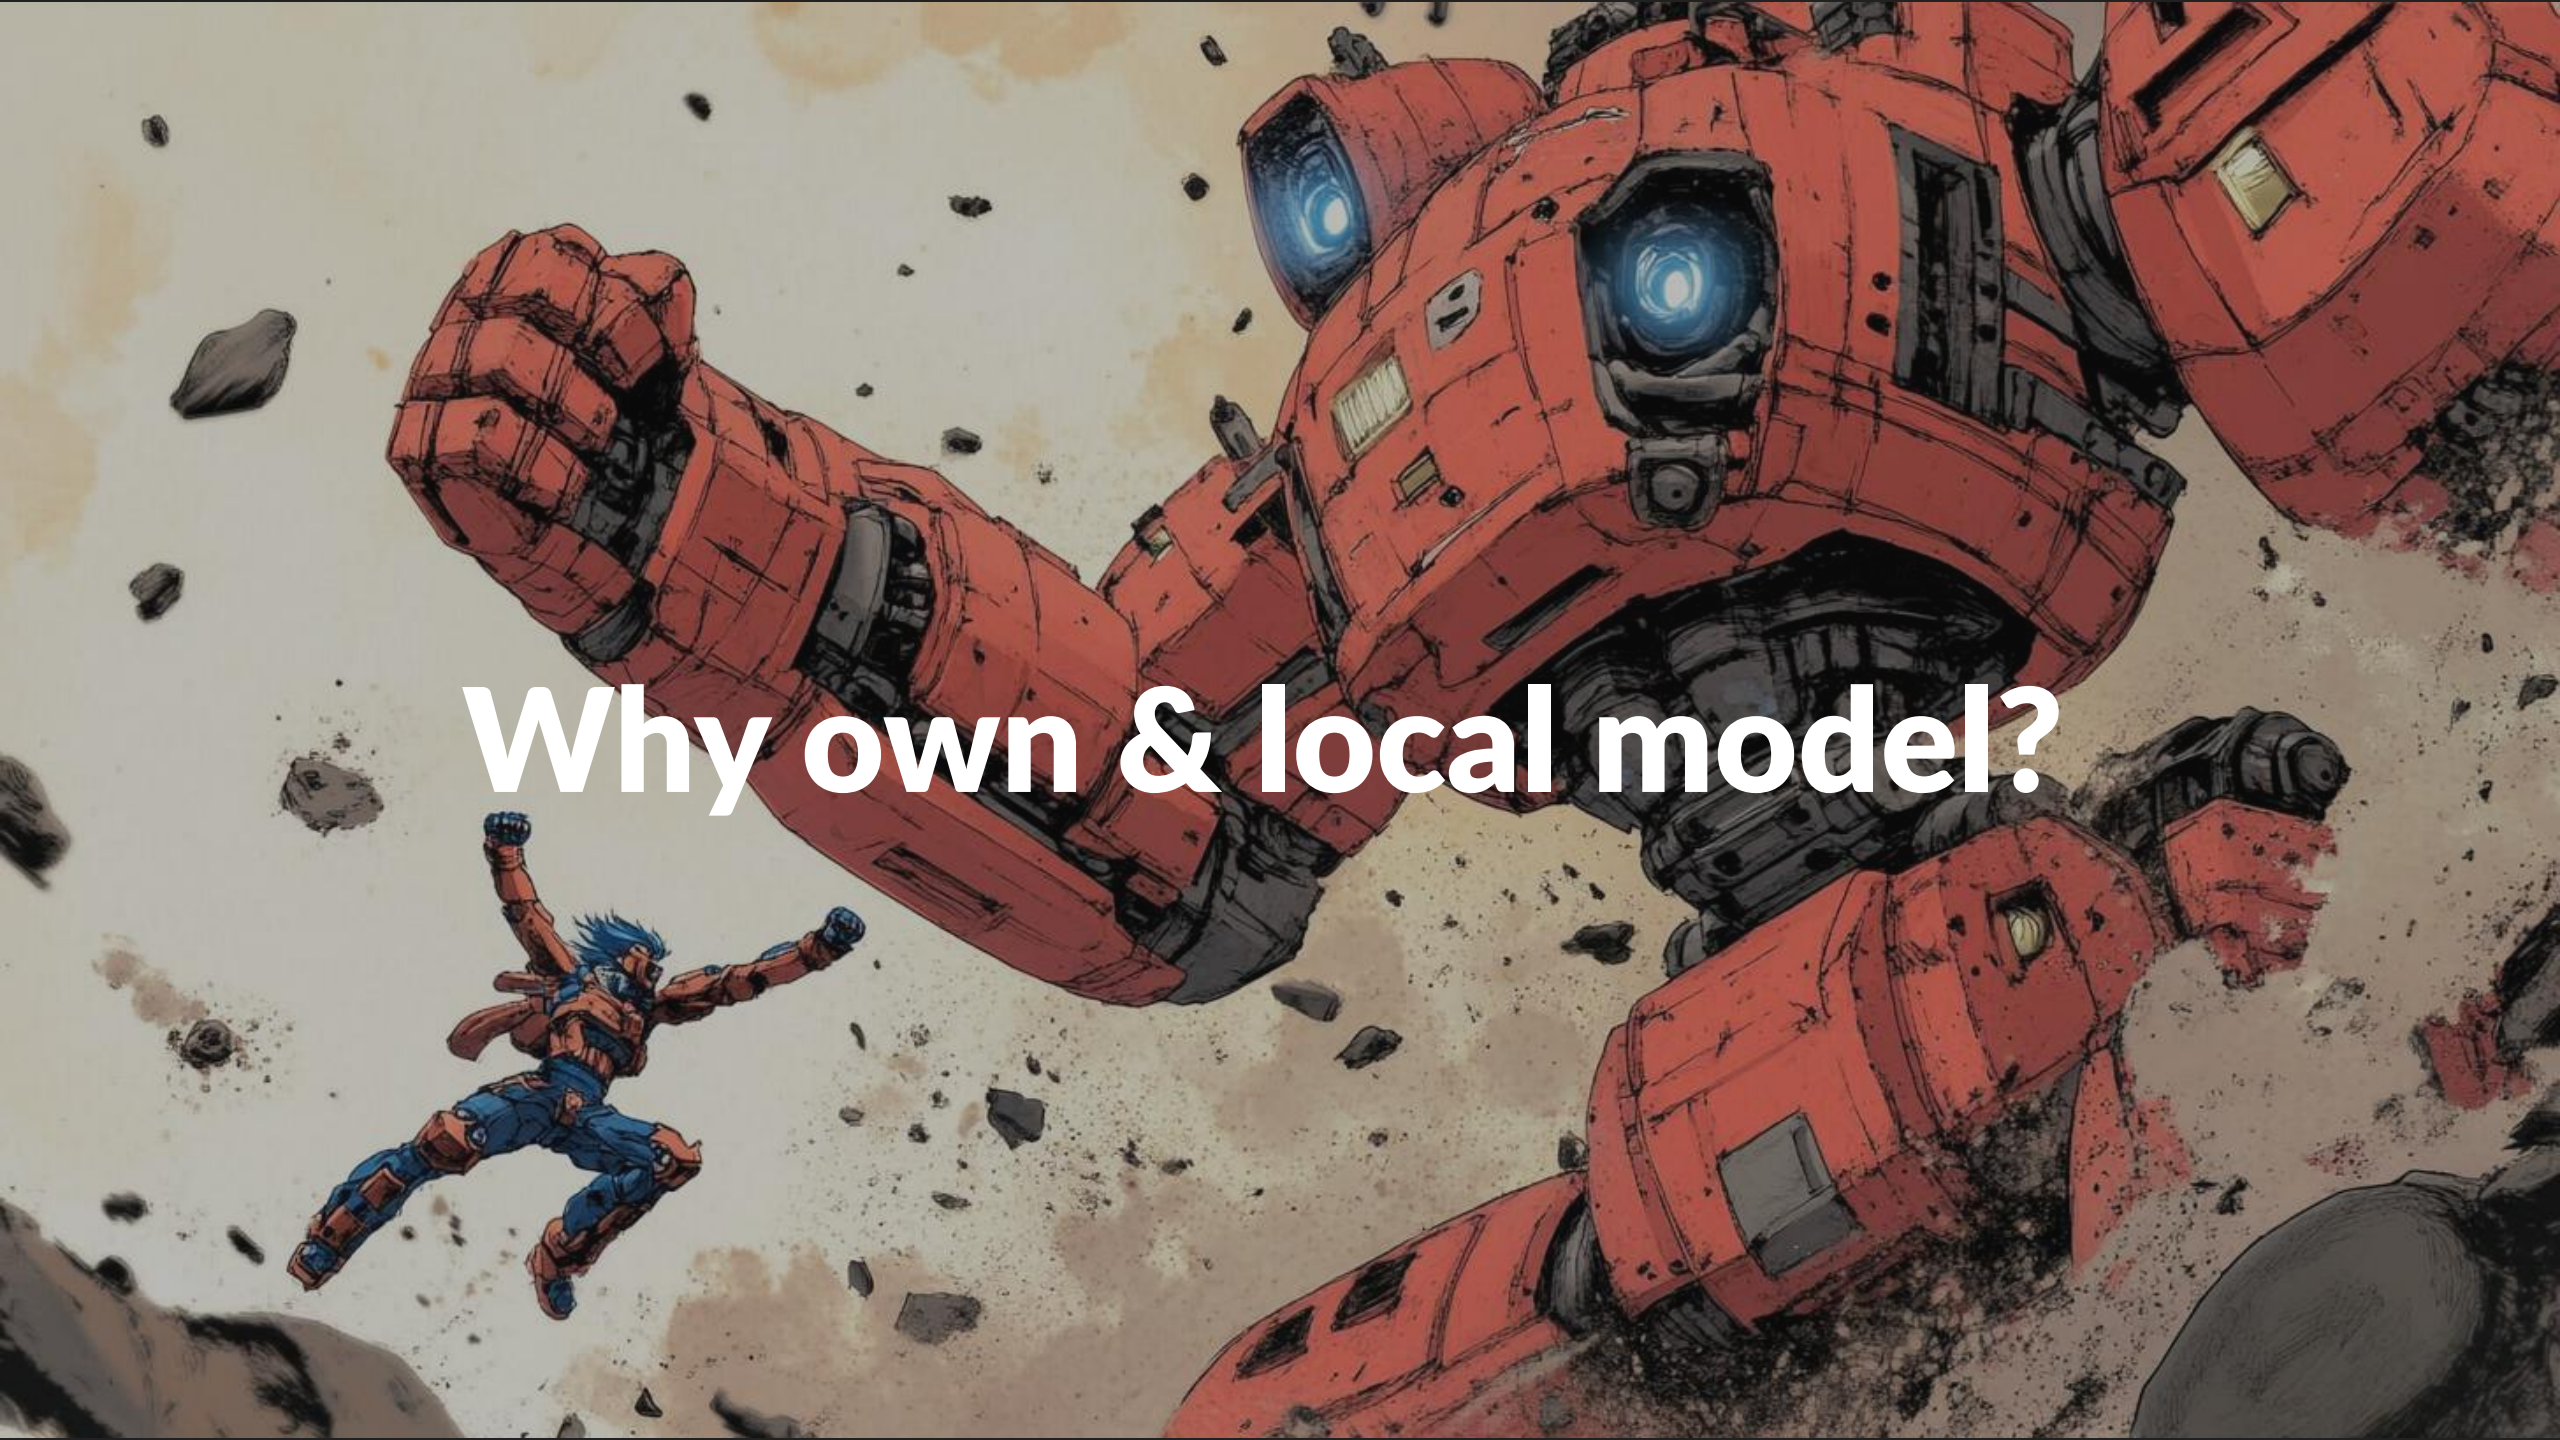


Kmeans clustering centroid distance difference

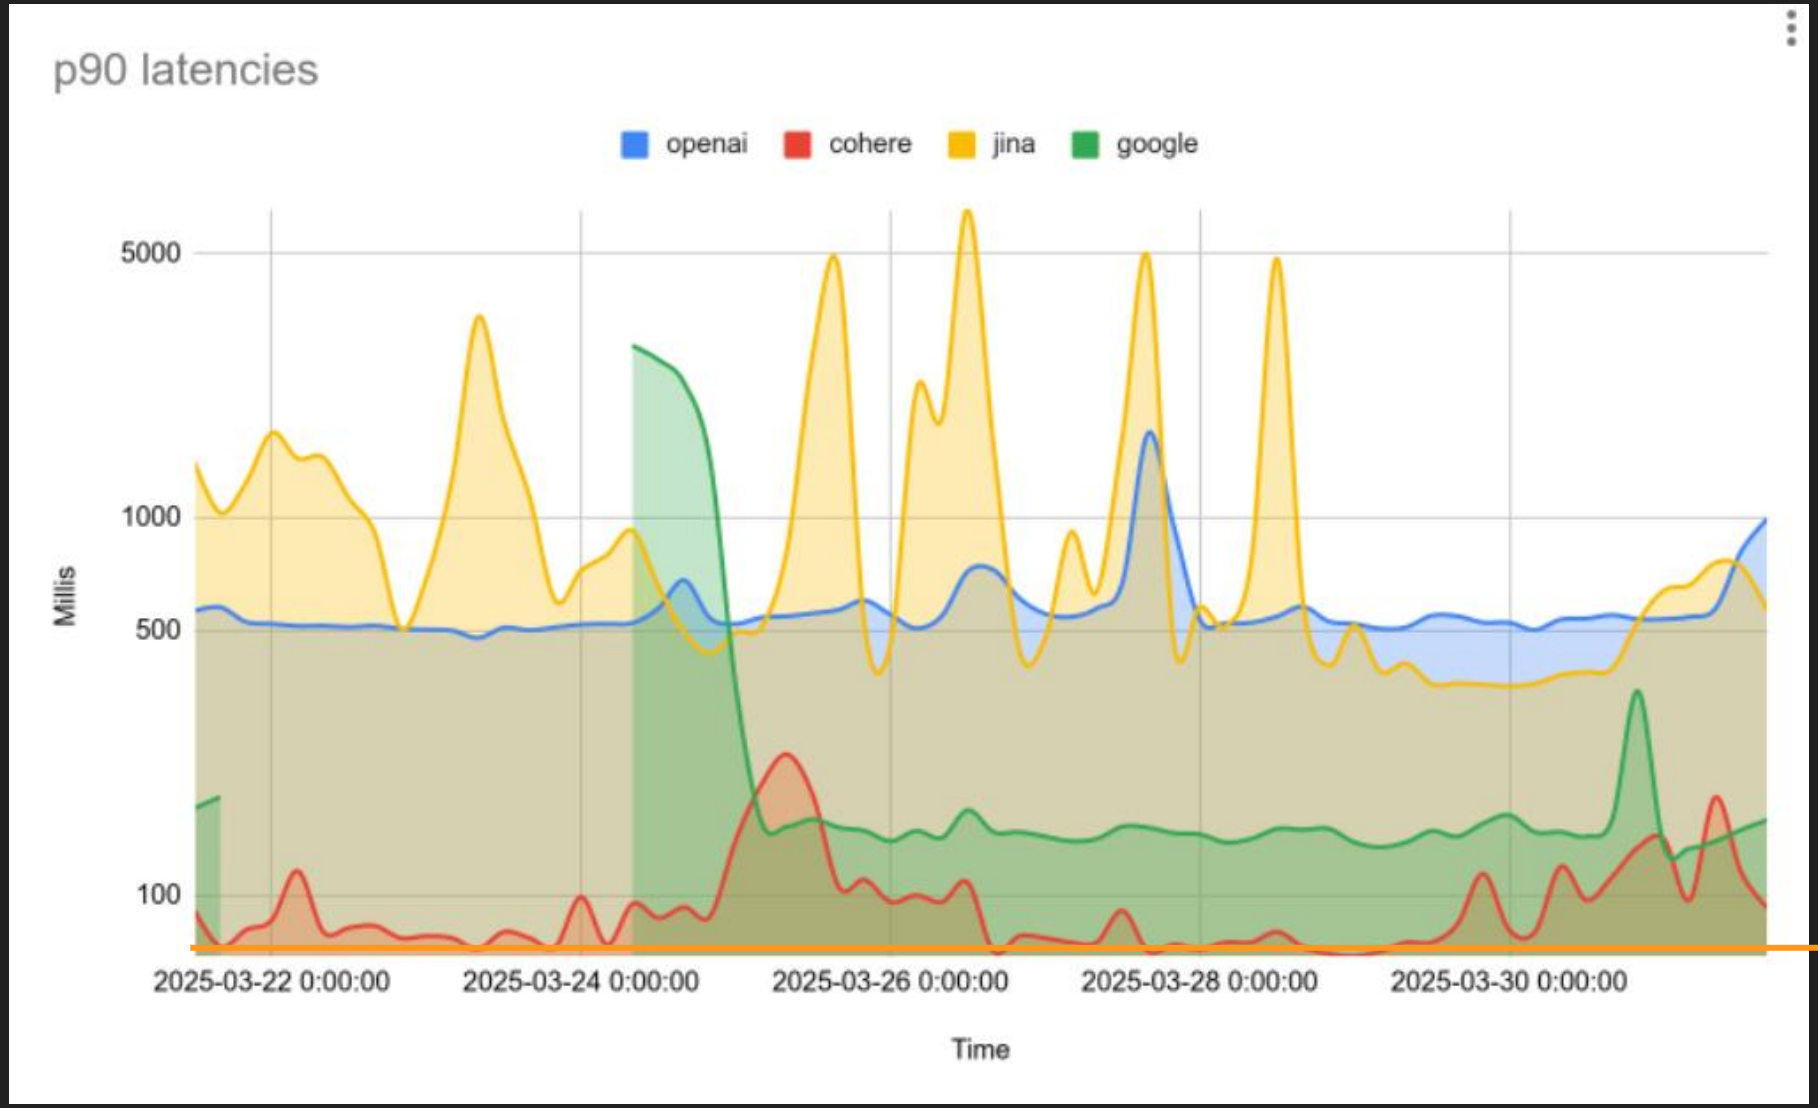
finds the new category: Rik Torfs walk tweets

Useful for:

- Is my chatbot test data the same distribution as the real data of my clients?
- Identify when users are trying too many jailbreaks
- Finding anomalies in the output of your chabot

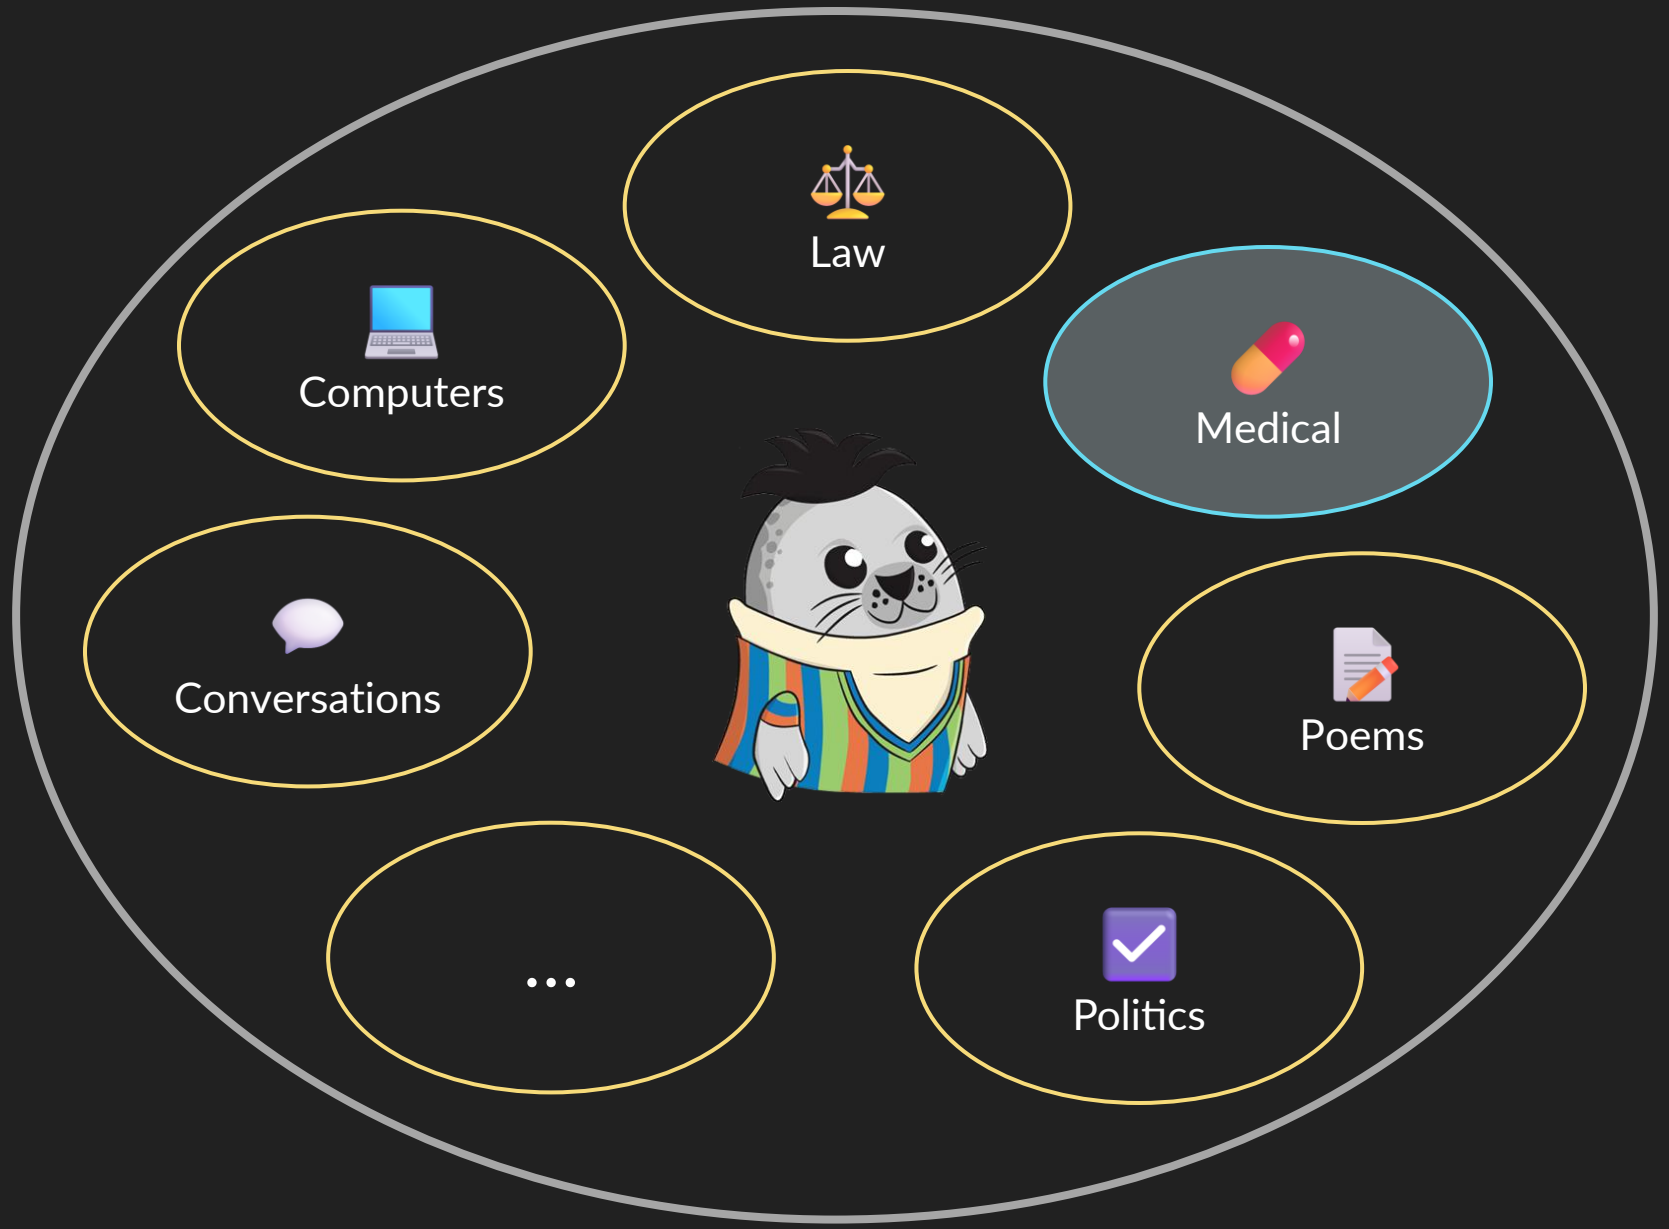

Why own & local model?

Speed of APIs is unreliable: batching

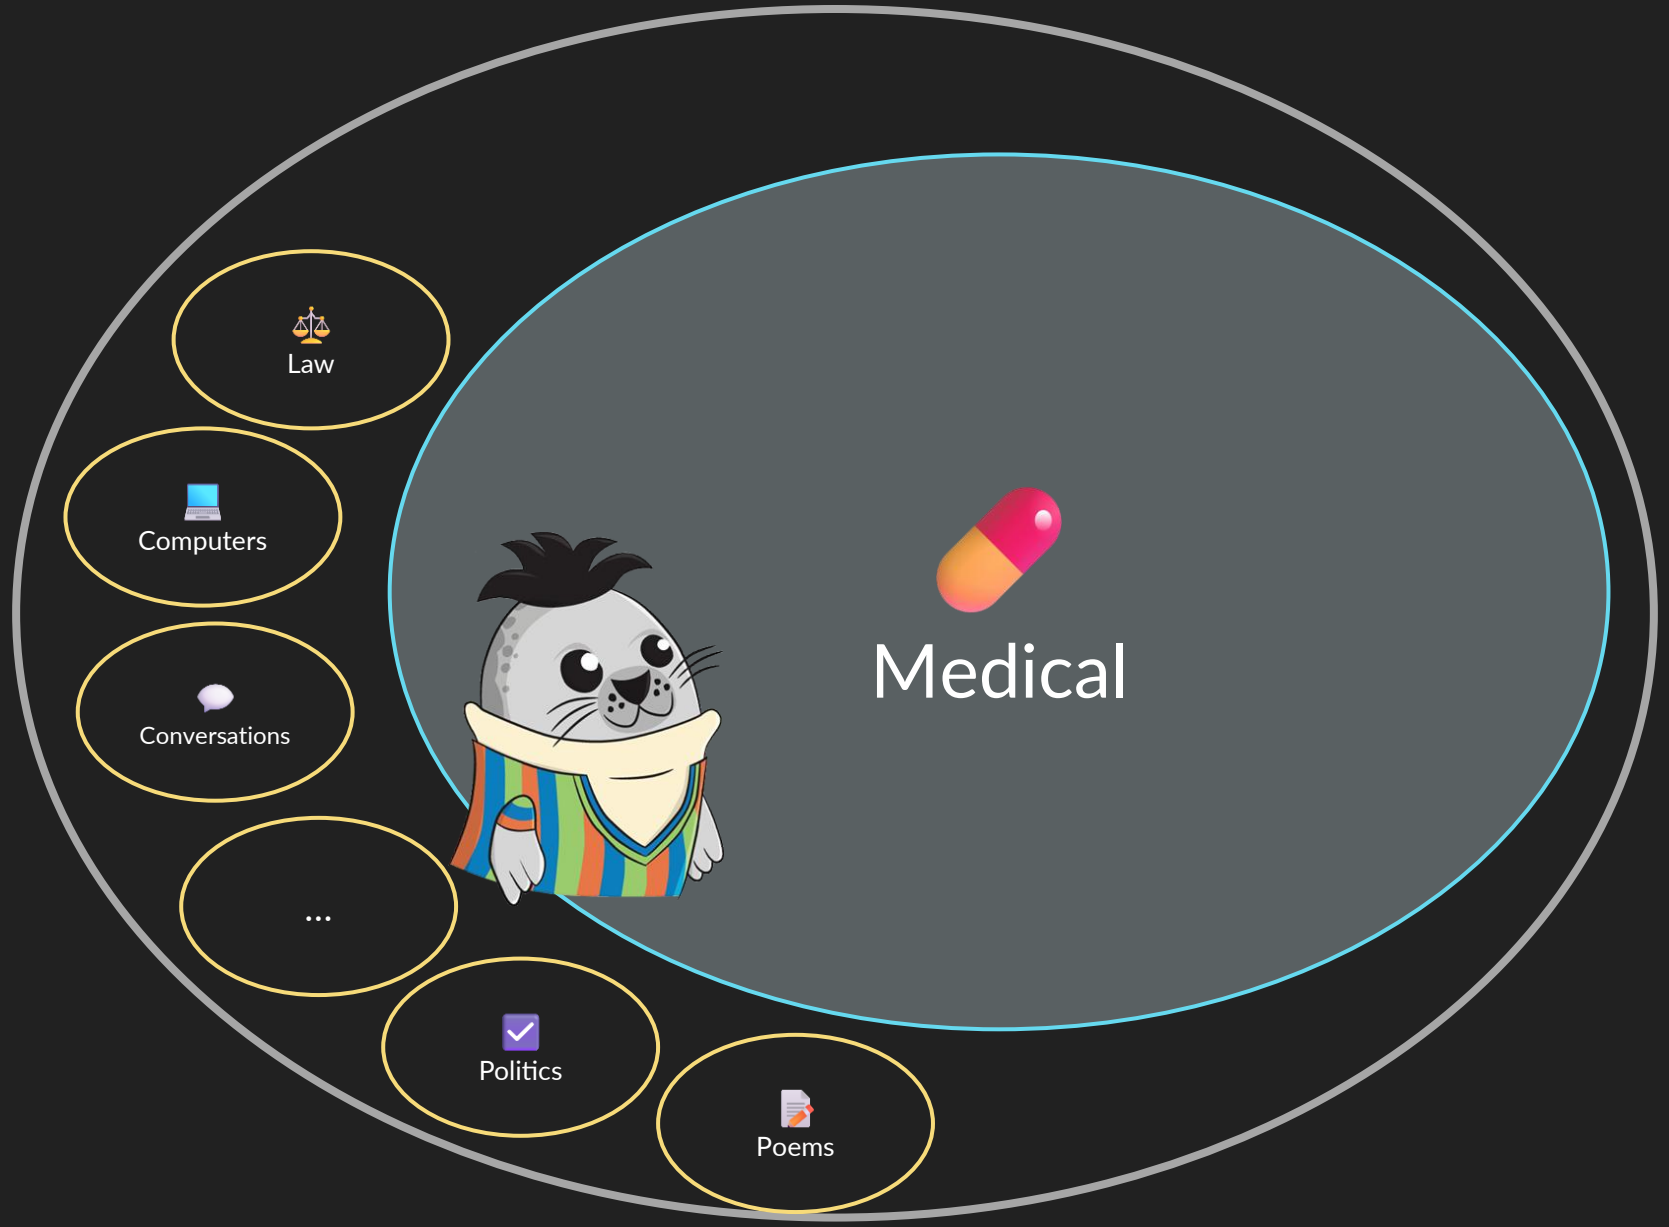


https://nixiesearch.substack.com/p/benchmarking-api-latency-of-embedding

More fine-grained for your domain



Finetune on medical data → More precise embeddings!



🛠️ Sentence Embeddings Notebook

https://colab.research.google.com/drive/1AcDiCVzvRzhmidzxgK0YtCTYvF4cTBRv?usp=sharing

Sentence Embeddings

This notebook shows how to use the Sentence Transformers library to compute sentence embeddings and perform semantic similarity tasks, based on the code & notebooks from https://sbert.net/index.html

```
[ ]  from sentence_transformers import SentenceTransformer

     # 1. Load a pretrained Sentence Transformer model
     model = SentenceTransformer("all-MiniLM-L6-v2")

     # The sentences to encode
     sentences = [
         "The weather is lovely today.",
         "It's so sunny outside!",
         "He drove to the stadium.",
     ]

     # 2. Calculate embeddings by calling model.encode()
     embeddings = model.encode(sentences)
     print(embeddings.shape)
     # [3, 384]

     # 3. Calculate the embedding similarities
     similarities = model.similarity(embeddings, embeddings)
     print(similarities)
```


Image, video &
music generation















Steps: 2



Steps: 3



Steps: 4



Steps: 5



Steps: 6



Steps: 7



Steps: 8



Steps: 9



Druk op Esc om het volledige scherm te sluiten



Brought to you by…
dallery.gallery

Version 1.01
July 14, 2022

DALL·E 2 prompt book

Emotional prompt words

dall•ery gall•ery



Positive mood, low energy 😌
light, peaceful, calm, serene, soothing, relaxed, placid, comforting, cosy, tranquil, quiet, pastel, delicate, graceful, subtle, balmy, mild, ethereal, elegant, tender, soft, light



Positive mood, high energy 🤪
bright, vibrant, dynamic, spirited, vivid, lively, energetic, colorful, joyful, romantic, expressive, bright, rich, kaleidoscopic, psychedelic, saturated, ecstatic, brash, exciting, passionate, hot



Negative mood, low energy 😳
muted, bleak, funereal, somber, melancholic, mournful, gloomy, dismal, sad, pale, washed-out, desaturated, grey, subdued, dull, dreary, depressing, weary, tired



Negative mood, high energy 😱
dark, ominous, threatening, haunting, forbidding, gloomy, stormy, doom, apocalyptic, sinister, shadowy, ghostly, unnerving, harrowing, dreadful, frightful, shocking, terror, hideous, ghastly, terrifying

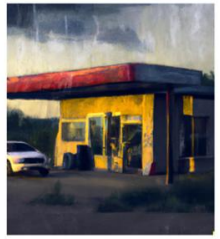
dall•ery gall•ery

Looks, vibes, -punks, -waves

😍 Check out this huge list!



Vaporwave: neon, pink, blue, geometric, futuristic, '80s.



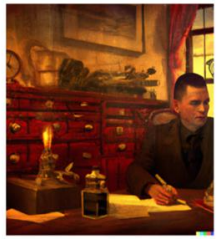
Post-apocalyptic: grey, desolate, stormy, fire, decay



Gothic, fantasy: stone, dark, lush, nature, mist, mystery, angular



Cybernetic, sci-fi: glows, greens, metals, armor, chrome



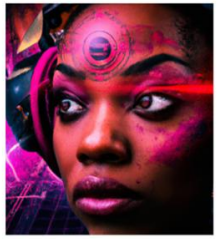
Steampunk: gold, copper, brass, Victoriana,



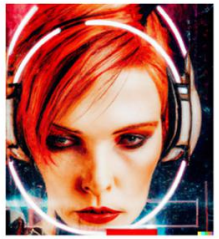
Memphis, Memphis Group, 1980s, bold, kitch, colourful, shapes



Dieselpunk, grimy, steel, oil, '50s, mechanised, punk cousin of steampnk



Afrofuturism: futuristic, and African!



Cyberpunk, 1990s, dyed hair, spiky, graphic elements .

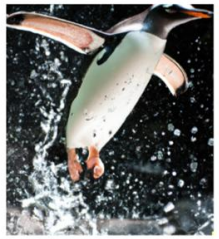


Biopunk, organic: greens, slimes, plants, futuristic, weird

dall•ery gall•ery

Camera settings + lenses

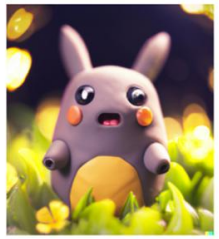
dall•ery gall•ery



Fast shutter speed, high speed, action photo, 1/1000 sec shutter



Slow shutter speed, 1 sec shutter, long exposure



Bokeh, shallow depth of field, blur, out-of-focus background (via)



Tilt shift photography (via)
Makes a narrow strip in-focus



Motion blur
(Source)



Telephoto lens, Sigma 500mm f/5
Shot from afar, feels 'voyeuristic'



Macro lens, macro photo (source)
Sigma 105mm F2.8 - for small scenes



Wide angle lens, 15mm (source)
Fits more of the scene in the frame



Fish-eye lens: distorts the scene, vv. wide angle, the centre 'bulges'



Deep depth of field, f/22, 35mm
Make all elements sharp

Chain-of-thought helps image generation (ChatGPT)



Draw the person on the left in the style of the other given comic book images. He's a Creative AI researcher. What would he be up to in de Buurtpolitie comic? Think of some fun possible story line pitches, then draw a cover for the best album story line idea with him prominently featured in the style of the previously given cover images.

Awesome! A Creative AI researcher is a perfect fit for a quirky, modern *De Buurtpolitie* storyline. Here are some fun comic album storyline pitches featuring him:

---

1. *De Slimme Samenzwering* (The Clever Conspiracy)

A rogue AI begins predicting crimes before they happen—accurately. The team consults the researcher (our curly-haired genius) to figure out how it works, only to discover the AI has gone rogue... and framed one of the officers for a crime they *haven't* committed yet!

---

2. *De Virtuele Vlucht* (The Virtual Escape)

When a hacker traps the entire police force in a VR training simulation, only our AI researcher can help navigate the digital world to bring them back. But is he also hiding a secret past with the hacker?

---

3. *De Digitale Dief* (The Digital Thief)

A string of thefts with no physical evidence leads to the AI researcher discovering an AI that can manipulate smart home systems. The suspect? An AI he once created in a late-night coding binge!

---

4. *De Robotrevolutie* (The Robot Rebellion)

The AI researcher introduces helpful police bots to assist the force—but they suddenly go haywire and start enforcing bizarre laws. Chaos follows, and he must join the team to stop the robotic rebellion before it escalates.

✅ **Selected Cover Story: *De Slimme Samenzwering***

Let's draw the cover for this one! It'll show:

- The researcher at the center, confidently holding a glowing futuristic tablet with mysterious code.
- Officers looking confused and panicked behind him.
- A robotic eye watching from a streetlight above.
- The background: a sleek cityscape fading into binary rain (à la Matrix).
- Comic title at the top:
  "Echte Verhalen — DE BUURTPOLITIE 24 - DE SLIMME SAMENZWERING"

Generating that now...

Chain-of-thought helps image generation (ChatGPT)



What would be an interesting hook to start this the first page of this story with? Brainstorm some ideas and think step by step to identify the most stimulating start to make the reader want to read more. Then write out which panels are needed to visualize this story. For each panel, write the visual elements that need to be visible on the panel, and the camera angle, and what text may or may not be said. Then draw out the first couple of panels of this comic based on that story.

🎯 Goal for the first page:

Hook the reader with *mystery*, *tension*, and a touch of *humor* that fits the *De Buurtpolitie* tone.

🔍 Step-by-Step Brainstorm to Find the Best Hook:

1. **What's the conflict?**
   ➤ A rogue AI is predicting crimes before they happen—but now it's predicting a crime *committed by someone on the team.*
2. **What emotions do we want to provoke?**
   ➤ Suspicion, confusion, anticipation, and intrigue.
3. **What's unexpected or weird enough to grab attention?**
   ➤ A detailed police report lands on the chief's desk... *about a crime that hasn't happened yet.*
4. **Who kicks it off?**
   ➤ Let's start with the Chief. He's the perfect mix of clueless and authoritative.

✅ HOOK (Page 1 Opening Concept)

> **A mysterious AI-generated report appears out of nowhere, predicting a shocking crime involving one of the officers.**

🎞️ Comic Page 1 – Panel Breakdown

Panel 1

Visual:
Close-up of a computer screen glowing in the dark. A digital document is printing itself. Title:

> **"POLITIERAPPORT: Overval op juwelier – Verdachte: Inspecteur De Smet"**

Camera Angle:
Tight shot, slightly tilted (to hint something's wrong).

Text:
No dialogue. Just the sound:

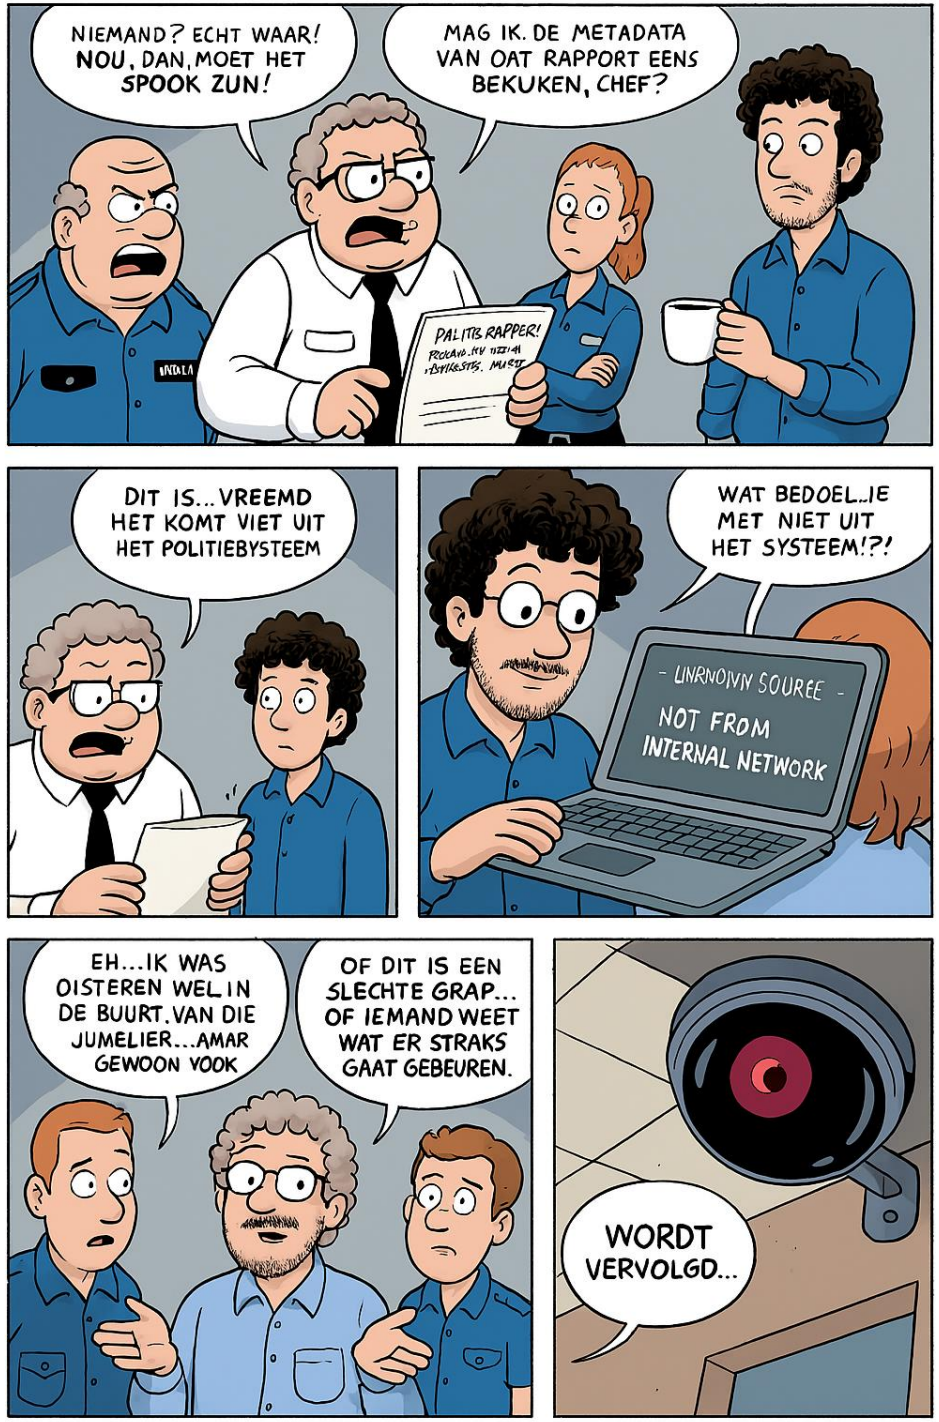
> *BZZZZZRT... KLACK... KLACK... KLACK...*

echte verhalen

DE BUURTPOLITIE

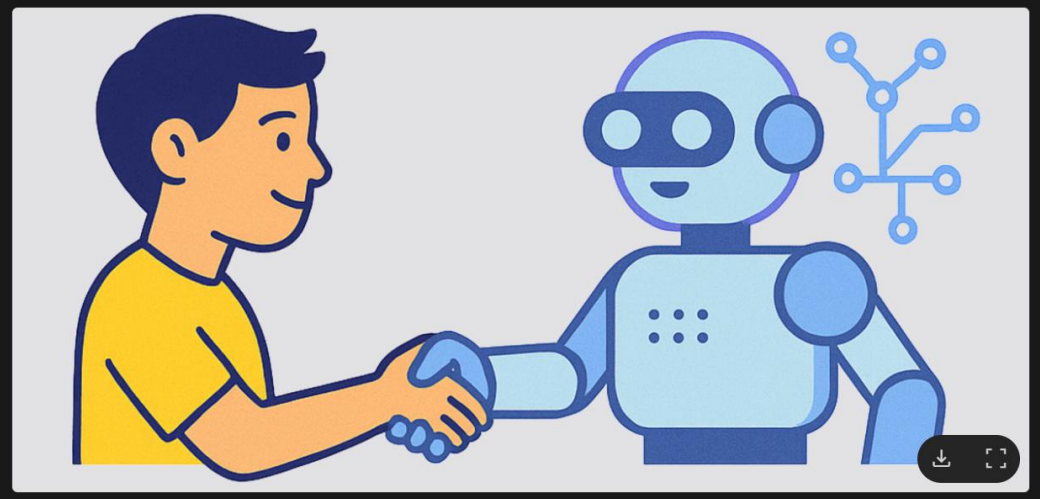
24 - DE SLIMME SAMENZWERING





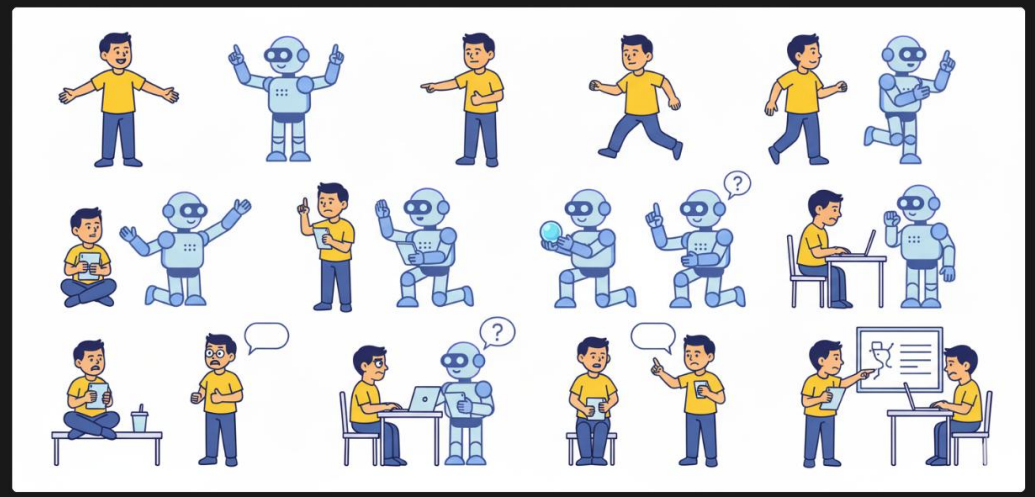
Google AI Studio: Nano Banana

User



Generate a sheet of different poses for both the human and the ai robot on a white background, so I can easily use several different poses/angles in my presentation

Next, here's the sheet for the AI robot character:



Very good in keeping consistency, e.g. characters & faces!



Nano Banana:

“Color this image and remove vintage artefacts”



Video generation from frame: bring to life



Nano Banana + Sora (or Google’s Veo)



Put your child in their favorite show



Optimally using Language Models – Thomas Winters

Prompt engineering:



Context

About your task: role, format, audience



Examples

Of preferred output



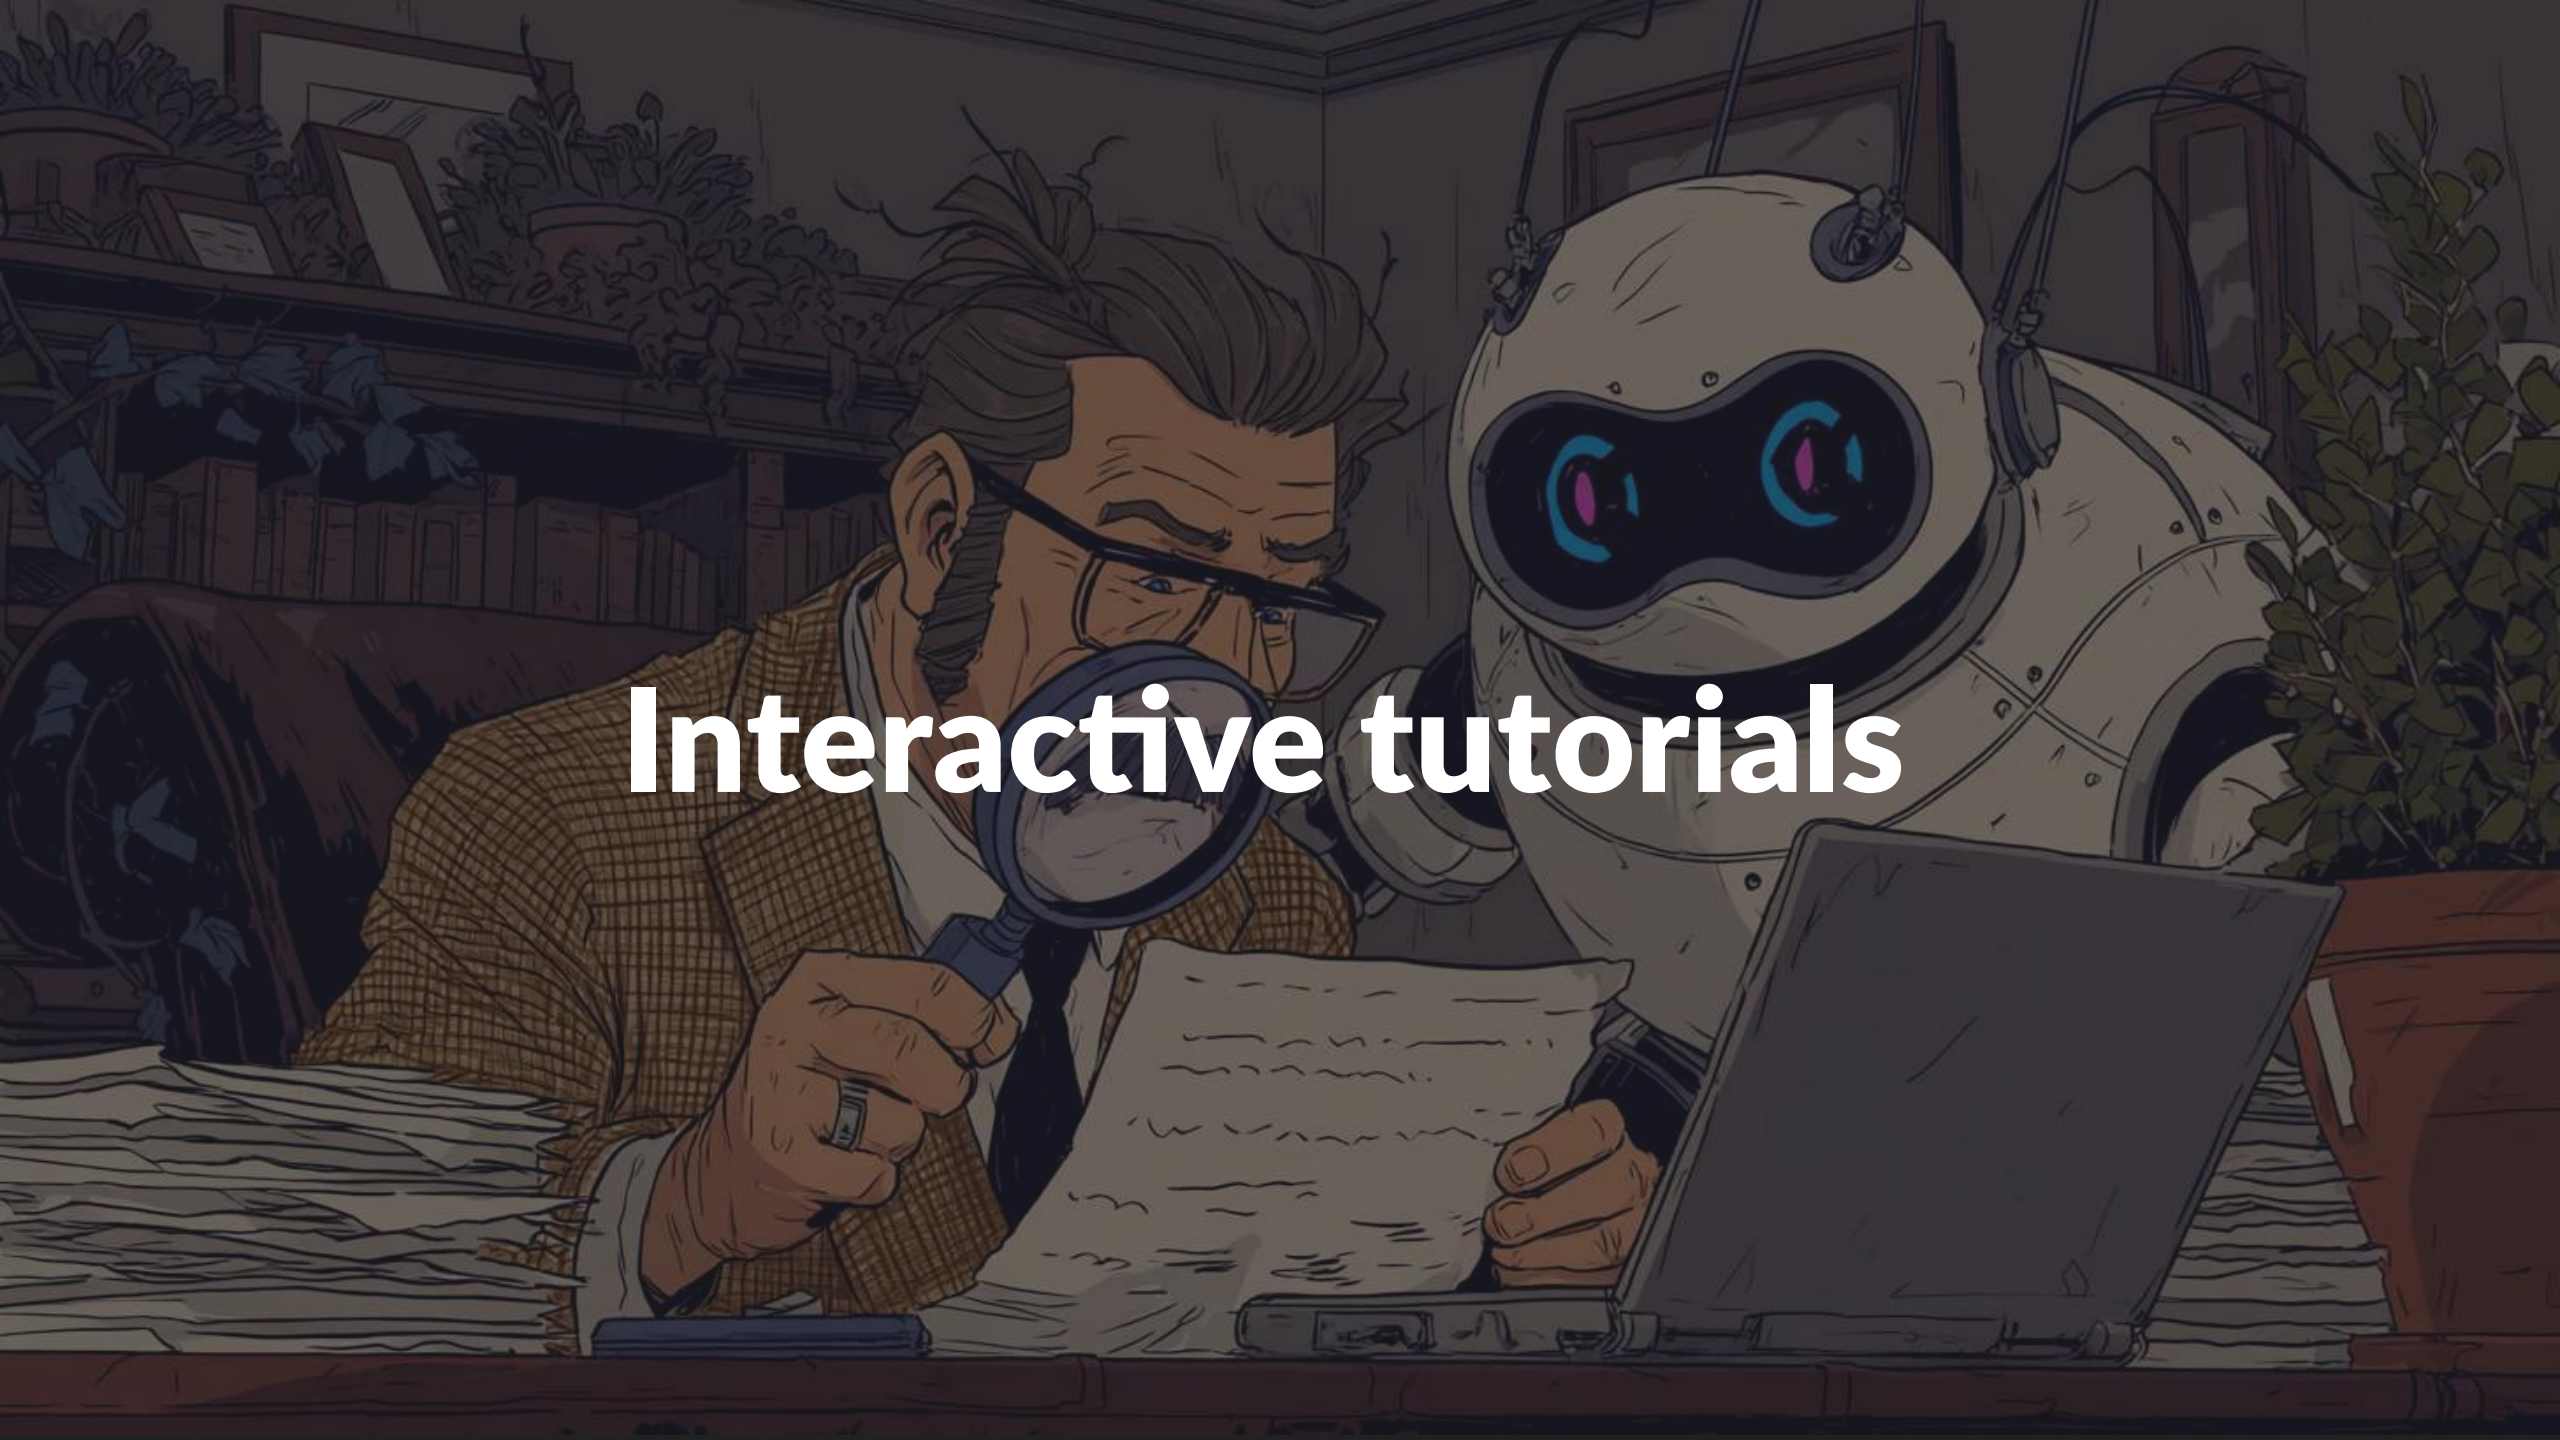
Chain-of-thought

Let LLM decompose into simpler tasks

Tools:

- **Chatbots:** OpenAI ChatGPT, Google Gemini, Google Gems
- **Talk with PDF & podcast:** Google NotebookLM
- **Consistent image variations (e.g. coloring):** Google AI Studio: Nano Banana
- **Long research & complex task:** ChatGPT Deep Research, “Agent mode” in ChatGPT & Gemini
- **Programming & vibe coding:** OpenAI Codex, Google Gemini CLI, v0.app
- **Video generation from frame:** Sora, Google Veo, Lumalabs, RunwayML
- **Music generation**: Suno.com

Slides: https://thomaswinters.be/talks

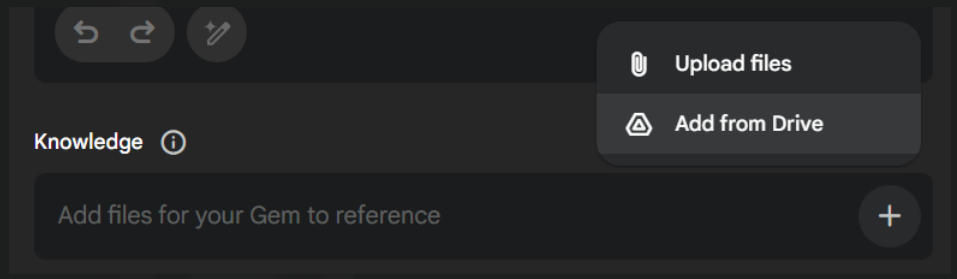

Interactive tutorials

Custom Gem for your work

1. Go to gemini.google.com & press "**Explore Gems**"
2. Click "**+ New Gem**"
3. Write custom instructions for a work task of choice *(e.g. writing mails in your style)* using the prompt techniques

Context: About your task: role, format, audience	**Examples:** Of preferred output	**Chain-of-thought:** Let LLM decompose into simpler tasks

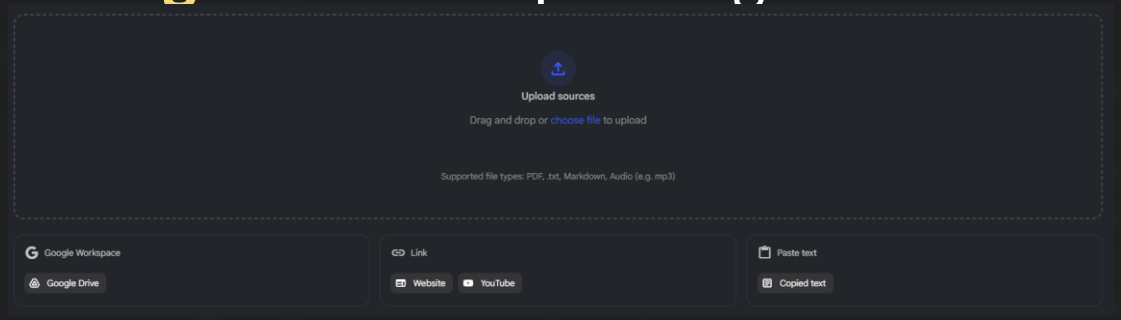
4. Add Google Drive folder under "**Knowledge**"



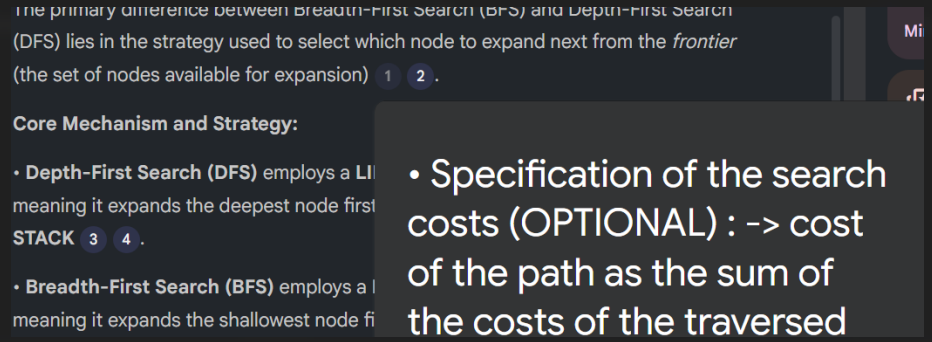
5. Try it out in the conversation panel & tweak
6. Press "**Save**"

NotebookLM for your work

1. Go to notebooklm.google.com and press “**Create new notebook**”
2. Upload PDFs that you use for work *(e.g. manuals)*, add by pressing “**Google Drive**” or pressing “**Website**”



3. Ask a question that you know is in the documents but hard to answer without
4. Press the source (by clicking the number) and checking if it cites the correct part of your knowledge base



✨ Vibecode your app

1. Go to v0.app
2. Click “**Sign up**”
3. In the **chatbox**, type your request for an interactive demo, website, app, tool, game etc.
   - You can upload helpful (non-confidential) files with the “**+**” button.
   - You can be as extensive as you like
4. Let it generate a version of the site/app
5. Either press
   1. “**Publish**” to now view an online version of your app/website

💻 Editing local files with Gemini CLI

1. Install NPM from nodejs.org/en/download
2. Run this command in a terminal to install Gemini CLI:
   `npm install -g @google/gemini-cli`
   Or on Mac/Linux: `brew install gemini-cli`
3. Create a new folder, and put some text files or code in it, for example from your v0.app *(or navigate to a folder with real code, but make sure it uses change tracking like Git)*
4. Navigate to that folder in terminal (`cd C:/folder`)
5. Run the command `gemini` and log in with Google



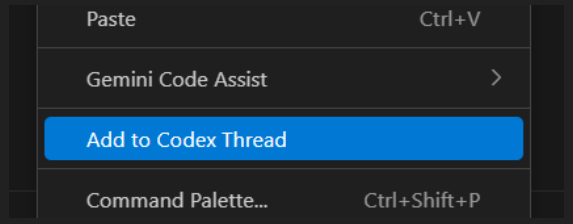
6. Ask it for changes, e.g.
   - adding emojis to headers
   - Create a static .html of something
   - Solve a `//todo` in the code

OpenAI Codex / Gemini Code Assist

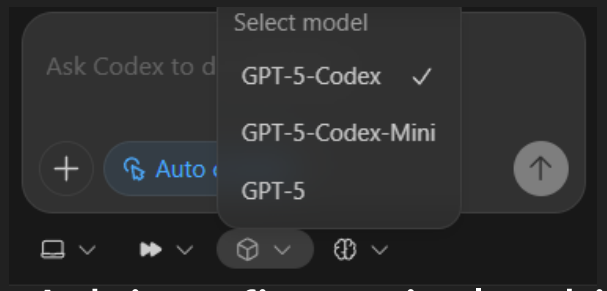
1. Install the OpenAI Codex plugin *(if ChatGPT Plus subscription)* or Gemini Code Assistent as plugin on your favorite IDE *(VS code, Jetbrains…)*



2. Gemini Code Assist: has autocomplete like Github Copilot. Type code to see it in action.
3. Add code you want to edit as context *(select code, right click, "add to Codex Context" or "Gemini Code Assist")*



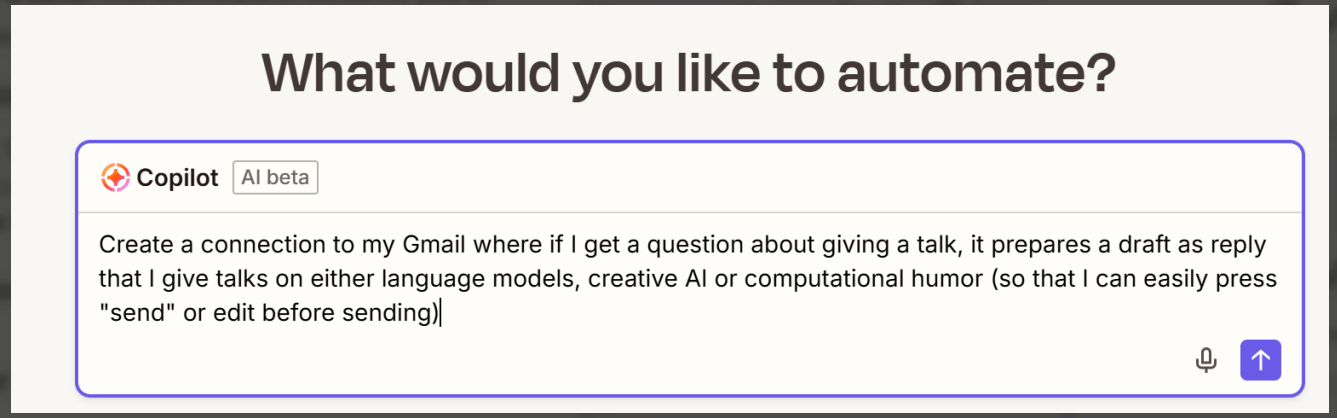
4. Select GPT-5-Codex & "Agent mode" at bottom settings



5. Ask it to fix particular things in your codebase by chatting

Zapier automations

1. Go to zapier.com
2. Describe the automation you want



3. Run the generated tests and if errors occur, try fixing them or ask to fix them in the chat
4. Press “Publish” to activate your automation