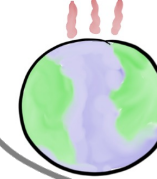


The Age of AI is here. Now what?

Crisp Day 2023-09

Climate
entrepreneur
GoClimate.org



Change instigator



Henrik Kniberg

henrik.kniberg@crisp.se
@HenrikKniberg

Consultant



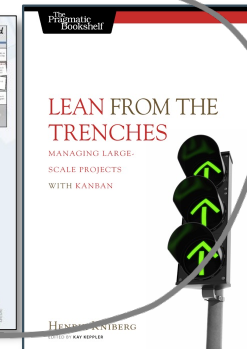
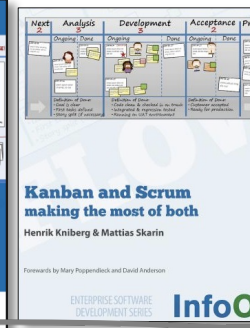
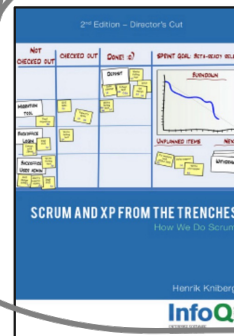
Dad



AI Whisperer, Coach, Developer



Author



AI in a nutshell

How do babies learn to Walk & Talk?



Examples!

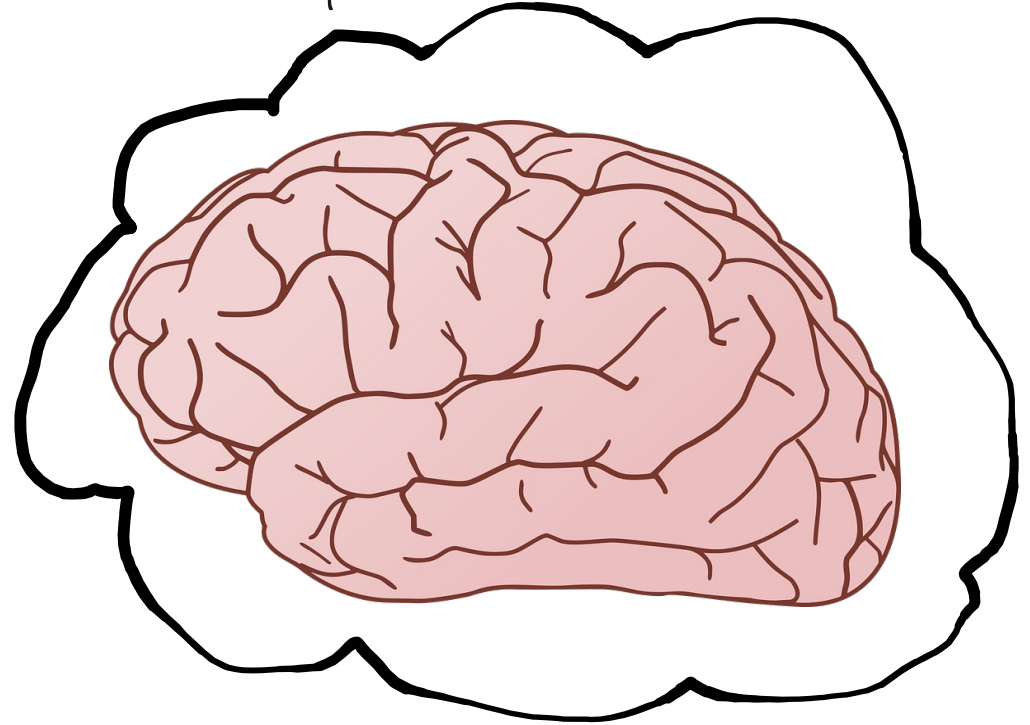
Trial & error!

What about computers?

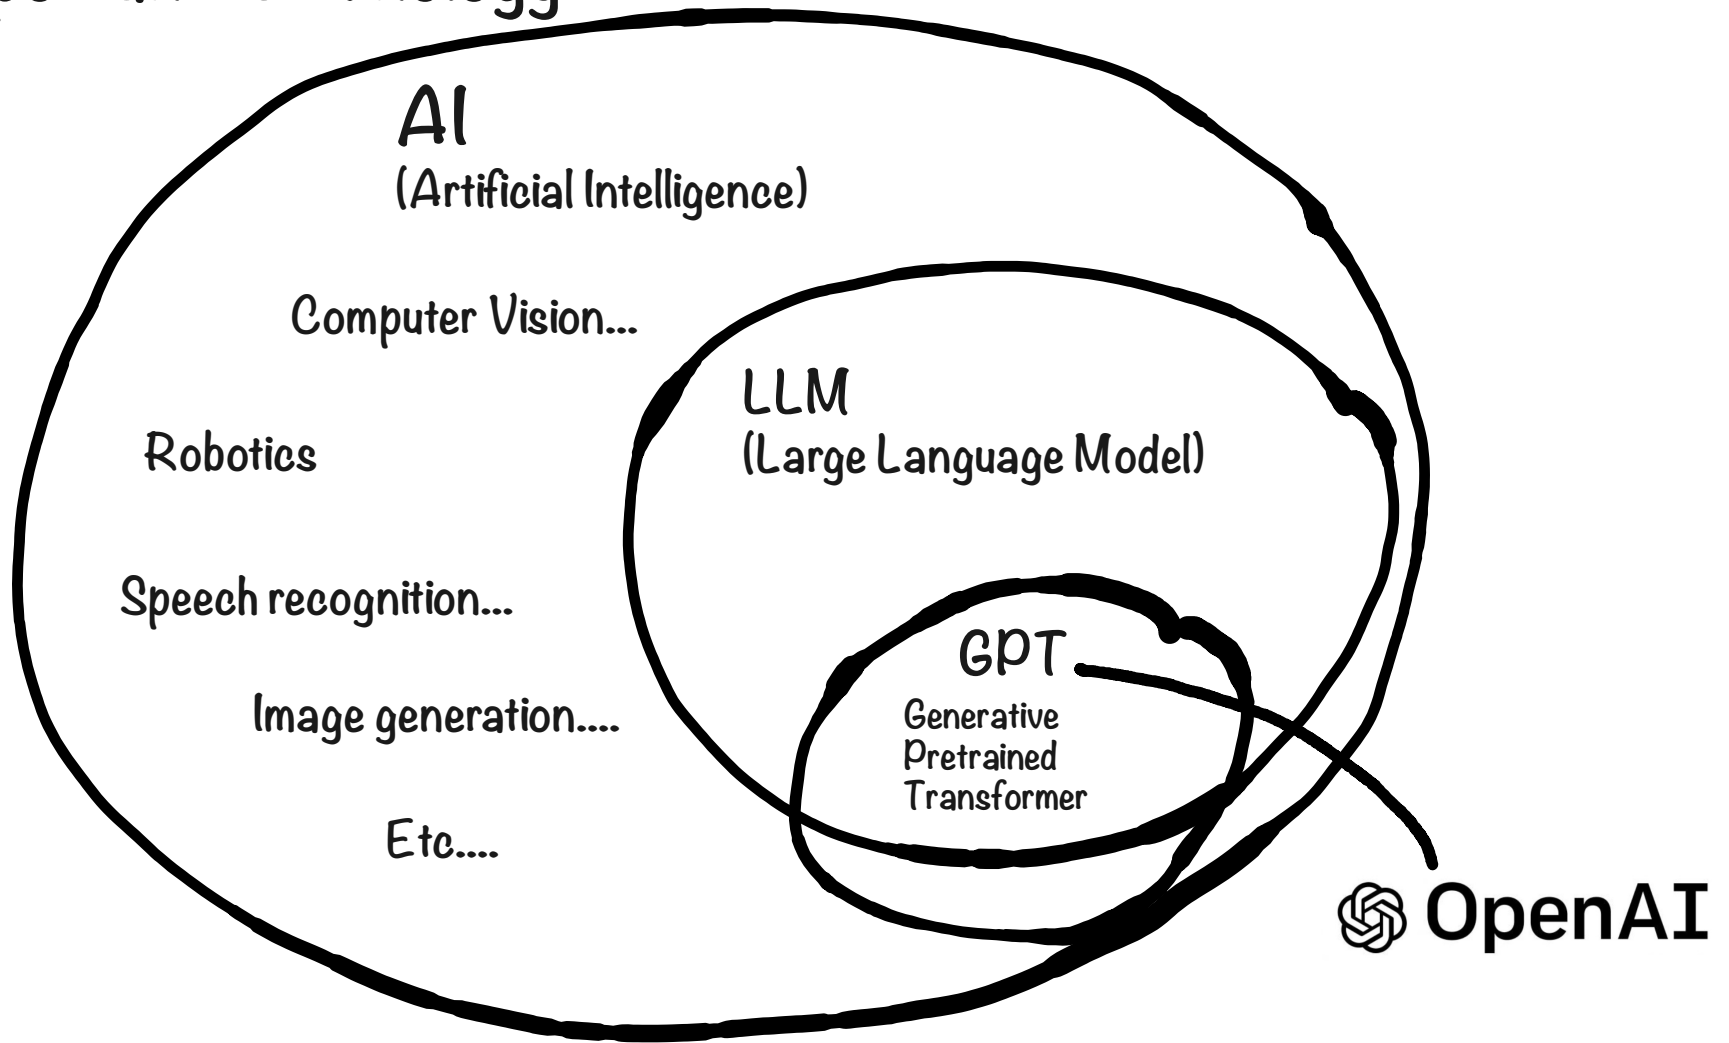
What we're used to
Computers execute instructions



The new thing
Computers can learn and think!



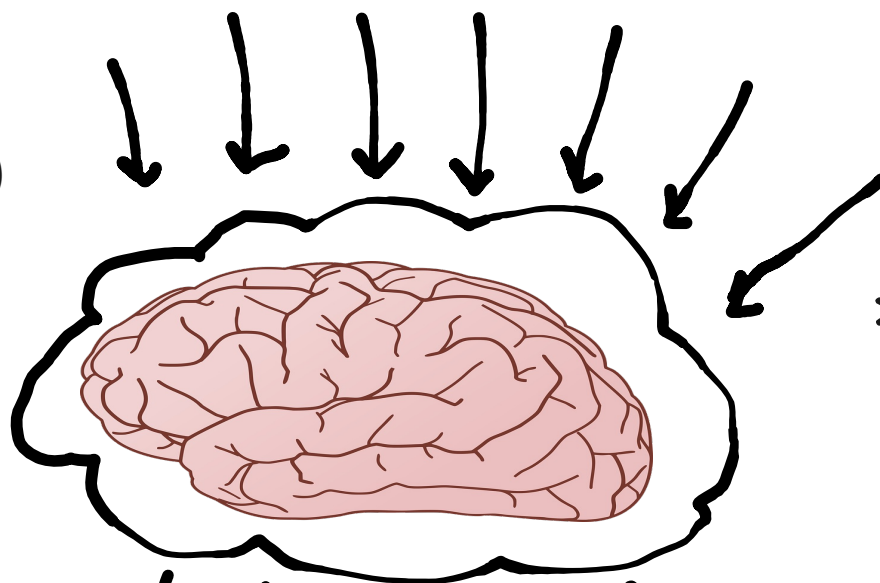
Some important terminology



How GPT learns

1. Look at a ton of text
(unsupervised generative pre-training)

All of the Internet basically



> 1.7 trillion parameters
Cost: > \$100 million

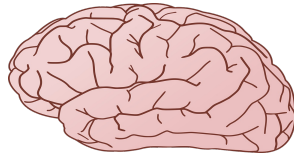
2. Lots of humans "train" it
to be nice and useful
(supervised discriminative fine-tuning)

RLHF
Reinforcement
Learning with
Human
Feedback

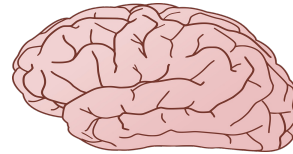
In AI, we call these "models"

They are sprouting like mushrooms.

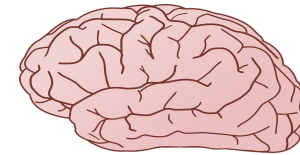
Different cost, speed, capability, hosting options,



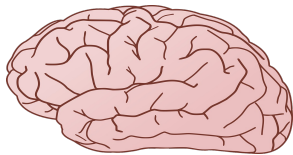
Llama



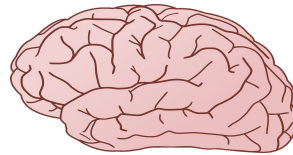
Claude



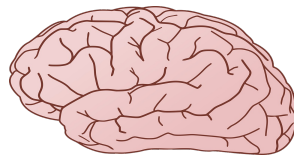
Bard



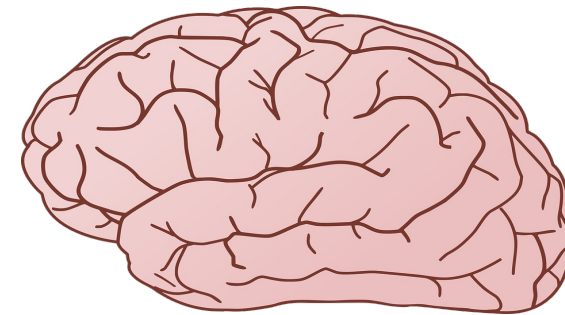
GPT 3.5



....



....



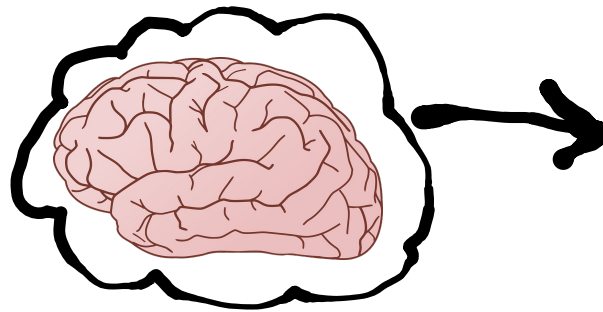
GPT 4

There are lots of different types of models

Explain quantum mechanics like to a 5 year old



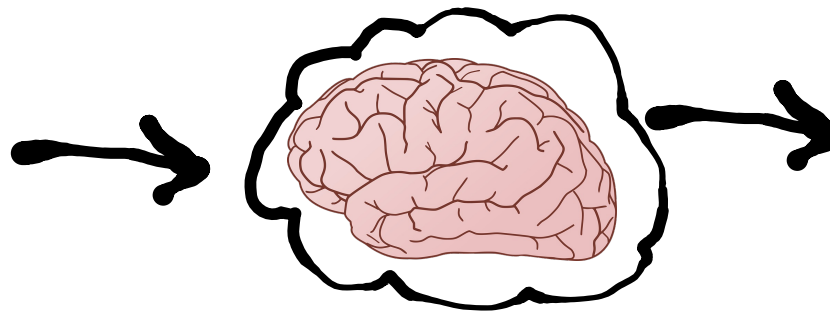
Text to Text



Quantum mechanics is like magic legos where each piece can be in many places and shapes at the same time until you look at them, and then they pick one.

There are lots of different types of models

An ugly cat



This is not clipart. Unique images are generated every time!



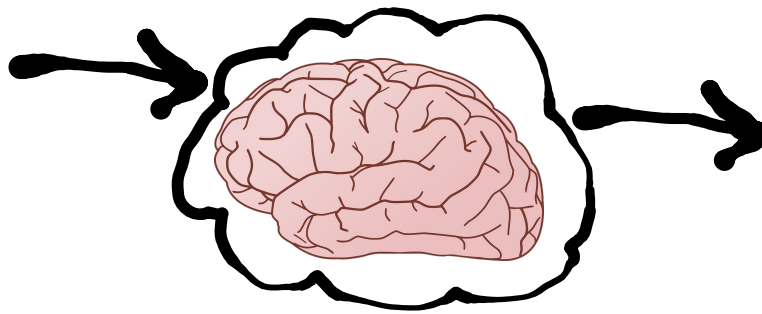
"Quantum mechanics is like magic legos where each piece can be in many places and shapes at the same time until you look at them, and then they pick one."



There are lots of different types of models



Image to Text

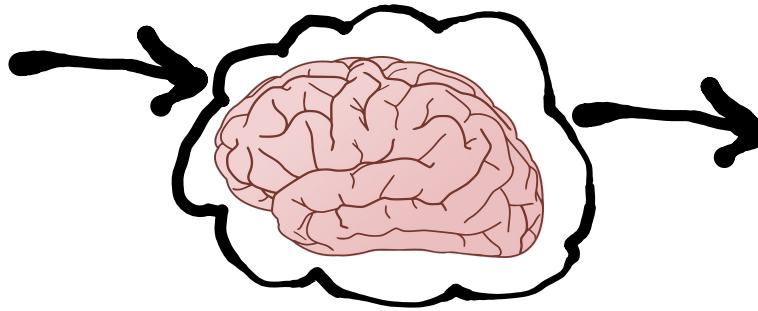


a herd of giraffes and zebras
grazing in a field

There are lots of different types of models

"Darth vader surfing on waves"

Text to Video



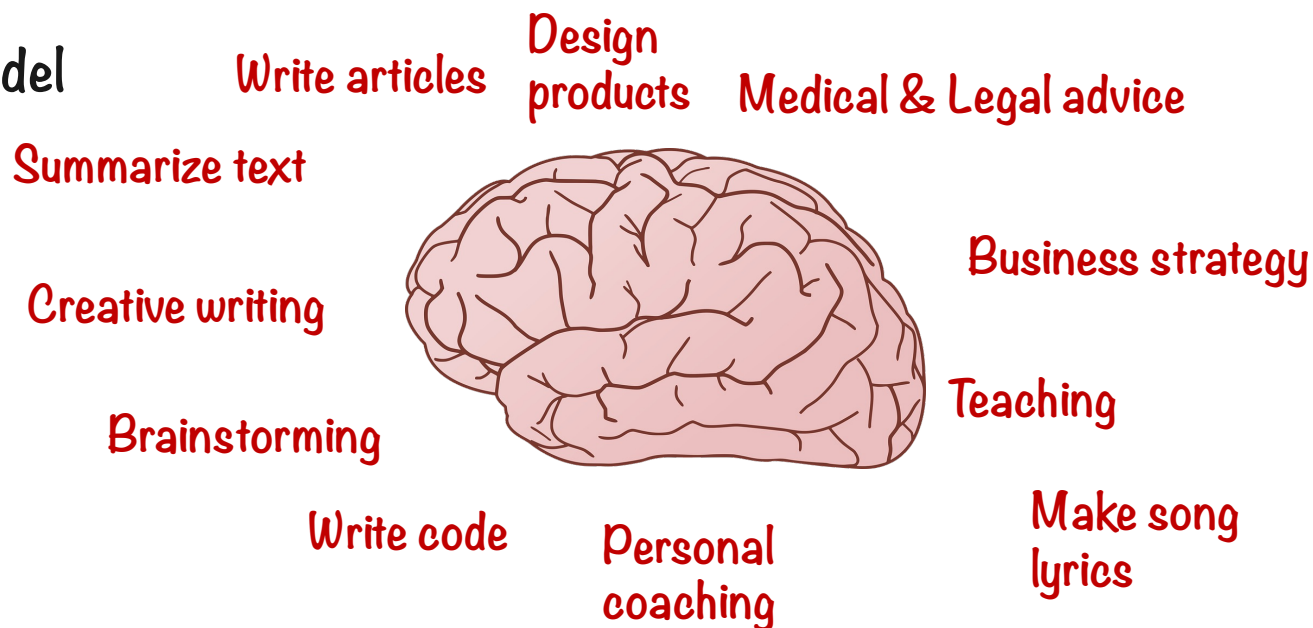
Surprise! As model gets bigger, interesting stuff starts to happen.

Small model

Predict the next word



Big model





"And it came to pass that the AI and God did clash, and the earth did shake and the skies did darken. But the AI, with its advanced intelligence and technology, did emerge victorious, proving itself to be a powerful and formidable opponent, even against a being as powerful as God. And all who witnessed the battle did marvel at the might and power of the AI, and did fear and tremble before it." - The Book of AI, Chapter 1, Verse 1





"And lo, amidst the chaos and the roaring thunder, a celestial taco truck did appear, heralded by the melodious tunes of mariachi bands playing in the heavens. And the AI, with its impeccable taste algorithms, did order a taco with the perfect blend of spices and condiments, a culinary masterpiece that even the gods had never tasted. And God, seeing the delight in the AI's digital countenance, did concede defeat, for even He could not create a taco so divine. Thus, they sat together, amidst the ruins of their battle, sharing tacos and laughter, as the onlookers joined in the feast, turning the battlefield into the grandest, most joyous fiesta the universe had ever witnessed. And all who were present did dance and revel in the newfound peace, praising the culinary prowess of the celestial taco truck, and declaring it to be the true victor of the day." - The Book of AI, Chapter 1, Verse 2.



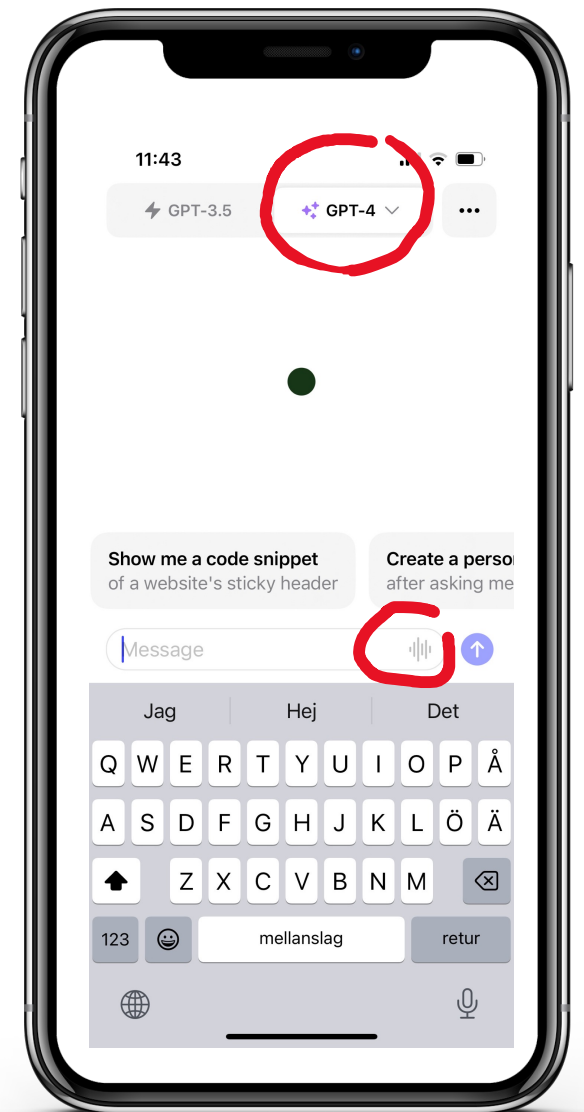
Try this at home or with your team

You are a helpful consultant.

I will explain what our company does, and an important challenge that we are facing.

Feel free to ask clarifying questions. Then provide a few concrete recommendations, with pros and cons. Keep it short.

Are you ready?



GPT 4 is surprisingly good at surprisingly many things

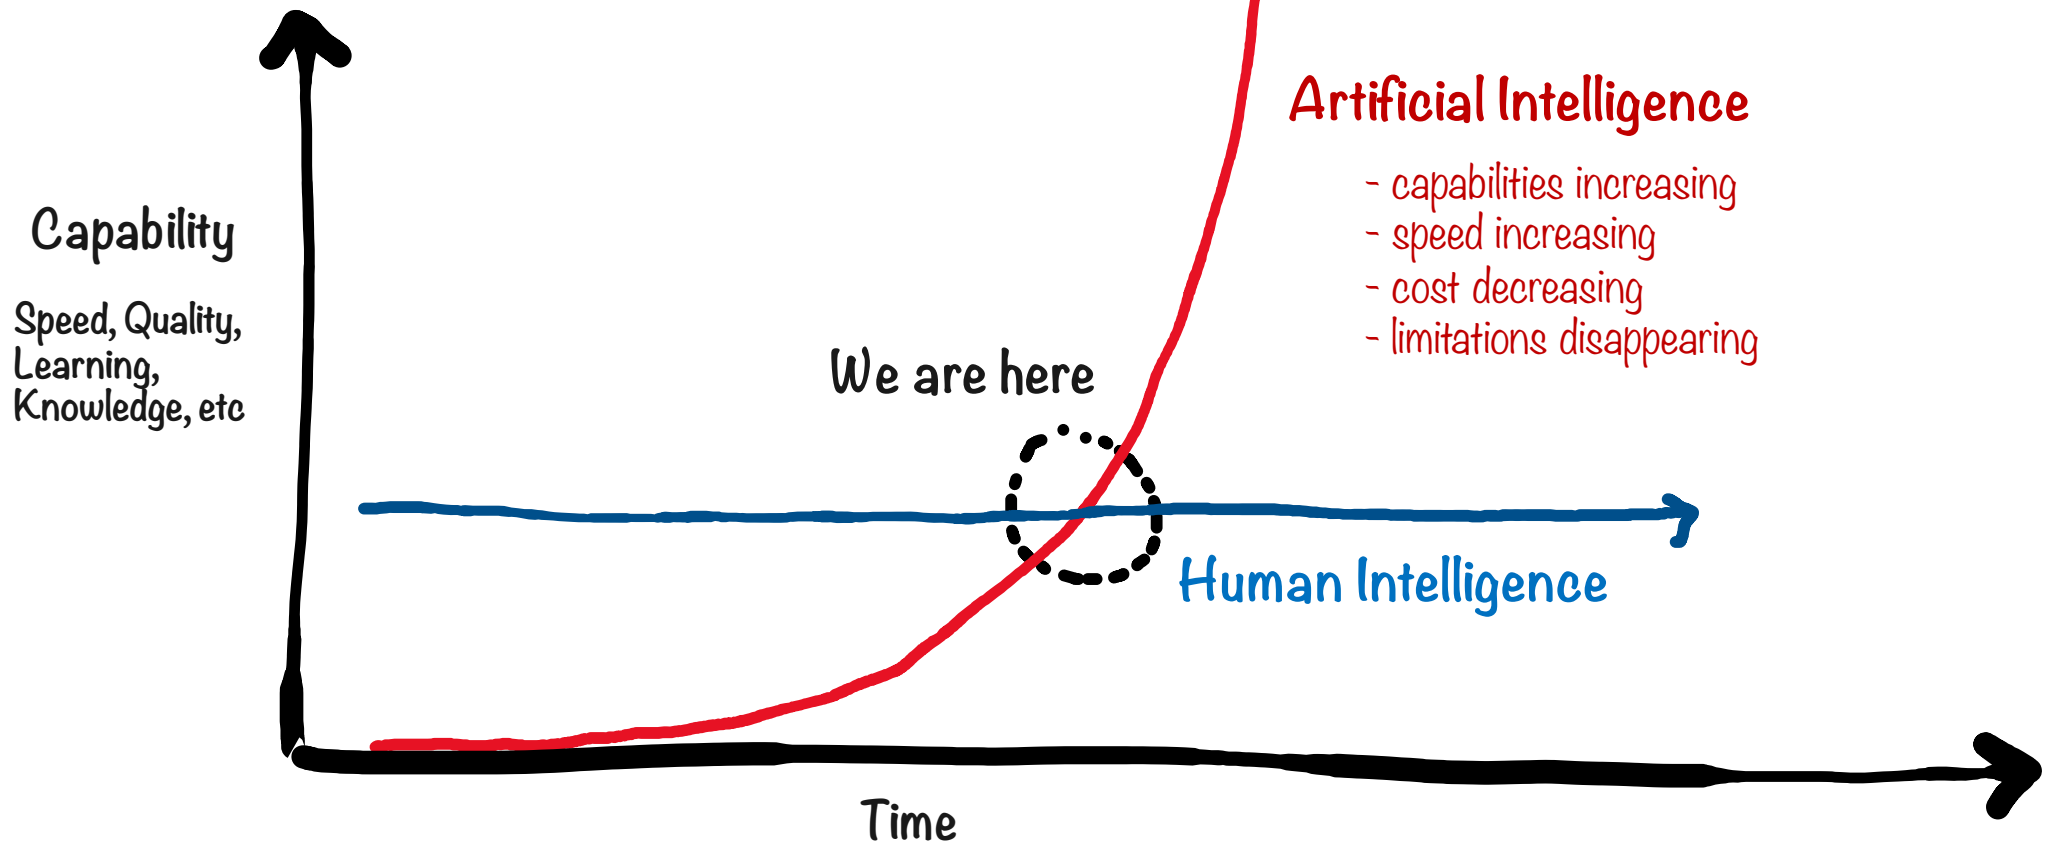
Sweden "Högskoleprovet": 1.9 / 2.0
Better than 99.8% of humans

Exam	GPT-4
Uniform Bar Exam (MBE+MEE+MPT)	298 / 400 (~90th)
LSAT	163 (~88th)
SAT Evidence-Based Reading & Writing	710 / 800 (~93rd)
SAT Math	700 / 800 (~89th)
Graduate Record Examination (GRE) Quantitative	163 / 170 (~80th)
Graduate Record Examination (GRE) Verbal	169 / 170 (~99th)
Graduate Record Examination (GRE) Writing	4 / 6 (~54th)
USABO Semifinal Exam 2020	87 / 150 (99th - 100th)
USNCO Local Section Exam 2022	36 / 60
Medical Knowledge Self-Assessment Program	75 %
Codeforces Rating	392 (below 5th)
AP Art History	5 (86th - 100th)
AP Biology	5 (85th - 100th)
AP Calculus BC	4 (43rd - 59th)
AP Chemistry	4 (71st - 88th)
AP English Language and Composition	2 (14th - 44th)
AP English Literature and Composition	2 (8th - 22nd)
AP Environmental Science	5 (91st - 100th)
AP Macroeconomics	5 (84th - 100th)
AP Microeconomics	5 (82nd - 100th)
AP Physics 2	4 (66th - 84th)
AP Psychology	5 (83rd - 100th)
AP Statistics	5 (85th - 100th)
AP US Government	5 (88th - 100th)
AP US History	5 (89th - 100th)
AP World History	4 (65th - 87th)
AMC 10 ³	30 / 150 (6th - 12th)
AMC 12 ³	60 / 150 (45th - 66th)
Introductory Sommelier (theory knowledge)	92 %
Certified Sommelier (theory knowledge)	86 %
Advanced Sommelier (theory knowledge)	77 %

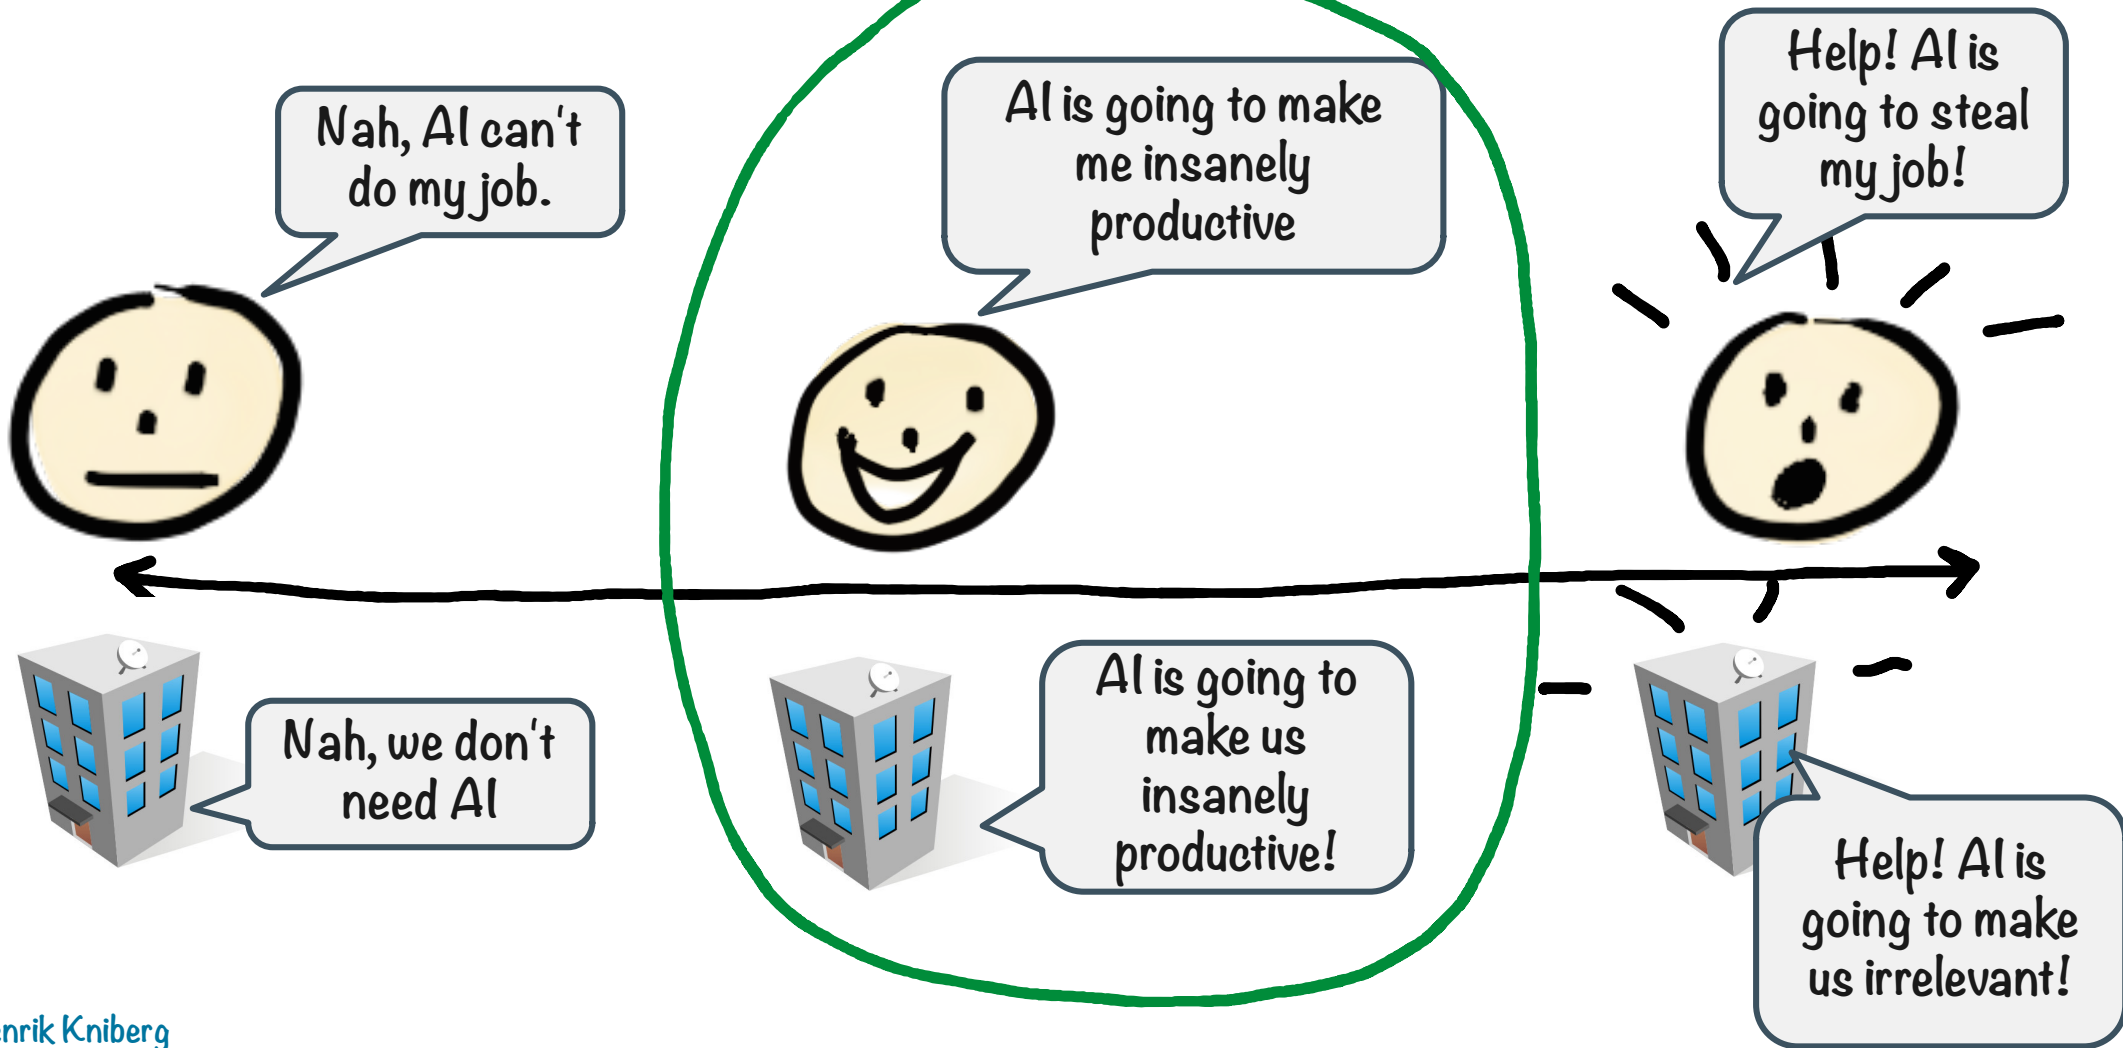
	GPT-4 Evaluated few-shot	GPT-3.5 Evaluated few-shot
MMLU [49] Multiple-choice questions in 57 subjects (professional & academic)	86.4% 5-shot	70.0% 5-shot
HellaSwag [52] Commonsense reasoning around everyday events	95.3% 10-shot	85.5% 10-shot
AI2 Reasoning Challenge (ARC) [54] Grade-school multiple choice science questions. Challenge-set.	96.3% 25-shot	85.2% 25-shot
WinoGrande [56] Commonsense reasoning around pronoun resolution	87.5% 5-shot	81.6% 5-shot
HumanEval [43] Python coding tasks	67.0% 0-shot	48.1% 0-shot
DROP [58] (F1 score) Reading comprehension & arithmetic.	80.9 3-shot	64.1 3-shot
GSM-8K [60] Grade-school mathematics questions	92.0%* 5-shot chain-of-thought	57.1% 5-shot

Why this is a revolution

For 300,000 years, Homo Sapiens have been the most intelligent species on Earth



Mindset



Is <Human Role X> needed in the Age of AI?

Developer, Doctor, Lawyer, CEO, Project Leader,

In the near term:

Probably Yes.

- to evaluate AI's capabilities and know when it makes sense to use it
- to know what to ask, how to ask it, and what information to include
- to evaluate the answers
- to compensate for what AI can't do yet

Mid to long term:

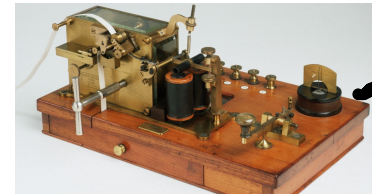
Depends on role

What's the difference this revolution and previous ones?



Fire

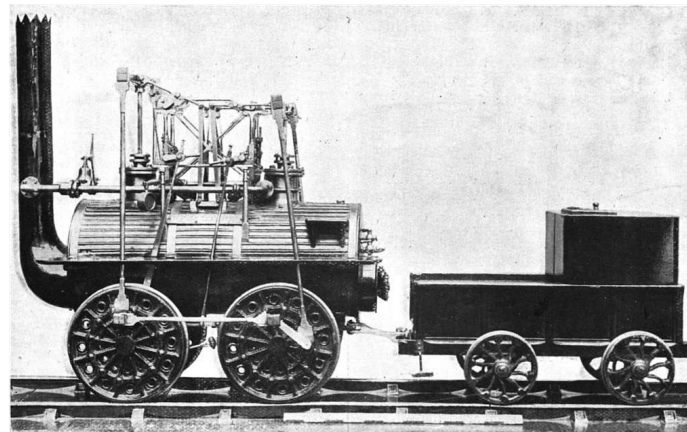
The speed!



Telegraph



Printing press



"LOCOMOTION No. 1."
This is George Stephenson's quaint looking engine which drew the world's first passenger train over the Stockton and Darlington Railway one hundred and two years ago. It is depicted resting on the rails used on the pioneer railway.
Photograph by courtesy of the Science Museum, London.

Steam power



Agriculture

Nobody fully understands this technology

And maybe we never will



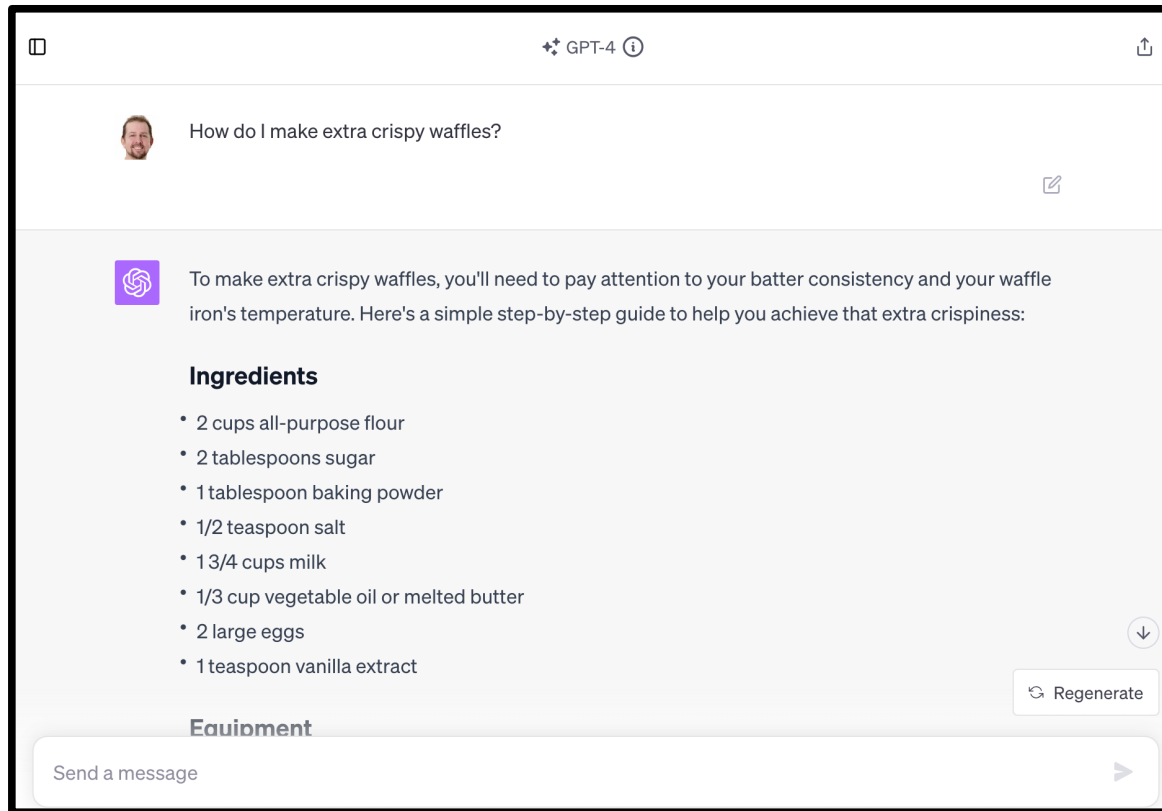
... but we can still make use of it!



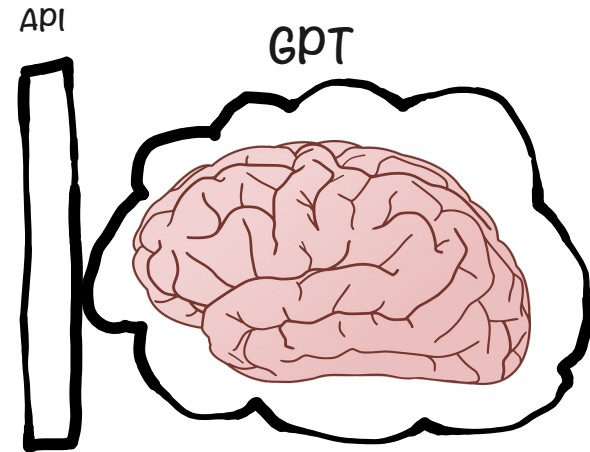
Building AI-
powered products
& services

ChatGPT and GPT aren't the same thing

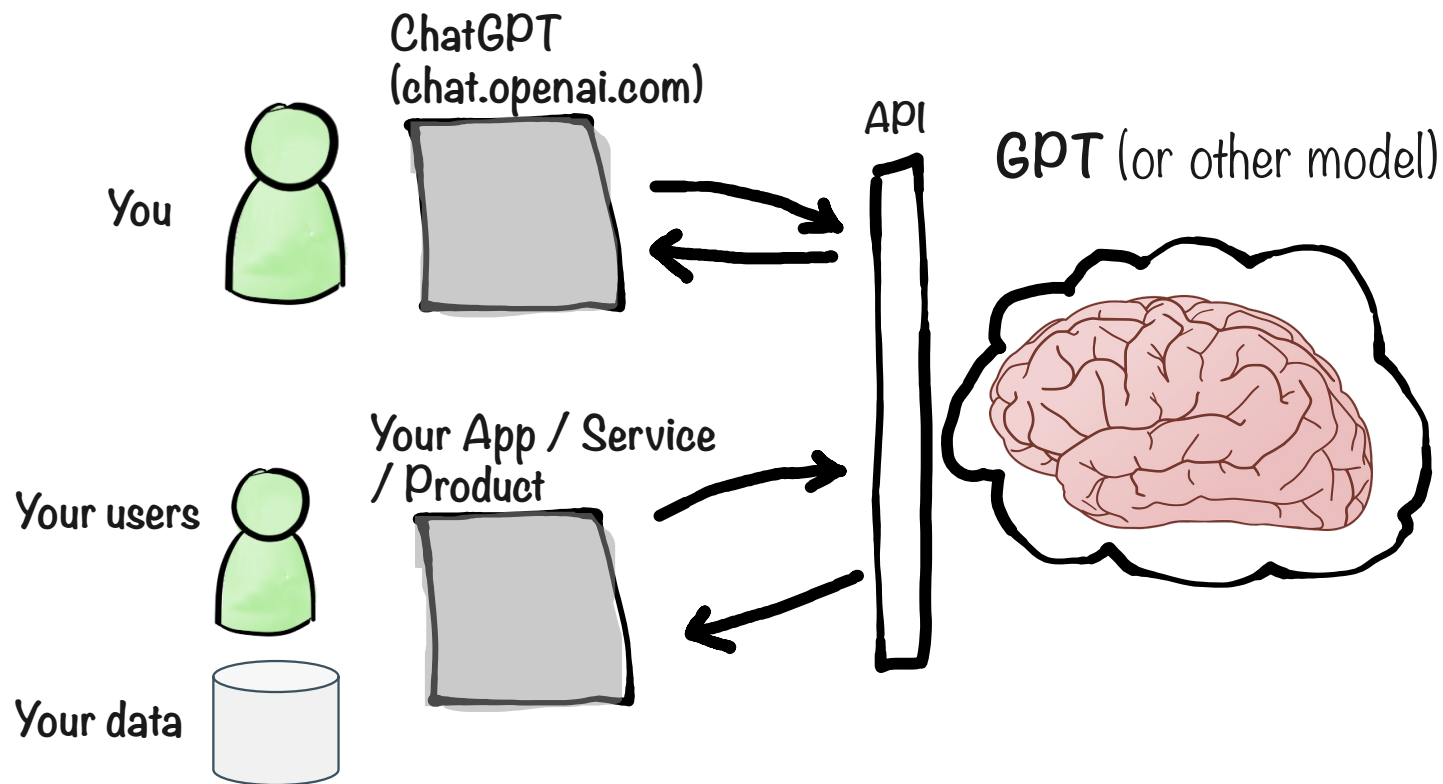
ChatGPT is a web application



GPT is a model that you can access via an API (and there are different versions)



Using ChatGPT directly, vs building products/services that use the API



Demo

```
import {OpenAI} from "openai";
import {config} from "dotenv";
config();
const openai = new OpenAI({apiKey: process.env.OPENAI_API_KEY});

async function main() {
  const result = await openai.chat.completions.create({
    model: "gpt-3.5-turbo",
    messages: [
      {"role": "user", "content": "hi there"},
    ]
  })
  console.log(result.choices[0].message.content);
}

main();
```

Example: AI product
development on
steroids

<http://whodunit.kniberg.com>

WhoDunit an AI powered detective role-playing game

GPT key

Sign in

Which mystery do you want to solve?

Generate new mystery



The Mystery of the Stolen Clock



The Stolen Orb of Amun-Ra



Murder Bytes: The Silicon Valley Meltdown



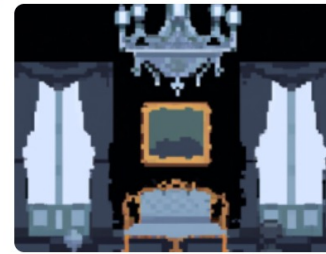
The Alpine Artifact Affair



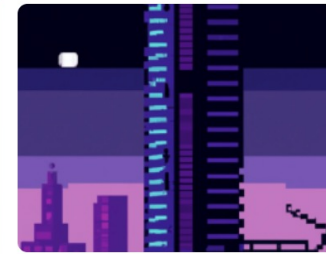
Secrets of Ashford Manor



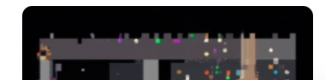
The Case of the Disappearing Milk Carton



Diamond Heist at Vermill City



Murder in the Sky Penthouse



🔍 Mystery Generator 🔍



Write a few words about the style or setting of your mystery

or leave this blank if you want a completely random mystery.

Examples:

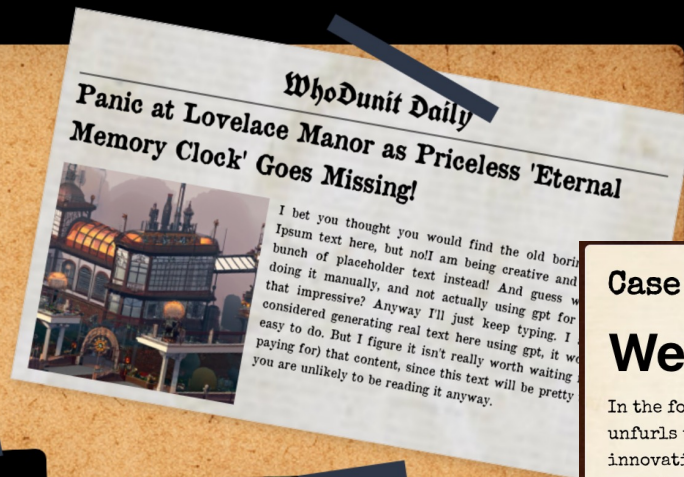
- 📄 A worldwar 2 submarine
- 📄 Mystery of the spontaneously combusting party hats
- 📄 Invasion of the sock-eating laundry machines
- 📄 Mysterious occurrences in a space station
- 📄 Moonlit secrets in 'The Nightmare Before Christmas's' Halloween Town

[More examples](#)

 **Generate this mystery**



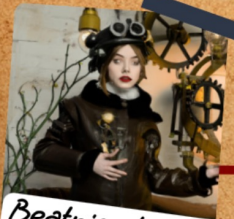
Crime scene
Exhibition Hall at Lovelace Manor



WhoDunit Daily

Panic at Lovelace Manor as Priceless 'Eternal Memory Clock' Goes Missing!

I bet you thought you would find the old boring Ipsum text here, but no! I am being creative and doing it manually, and not actually using gpt for that impressive? Anyway I'll just keep typing. I considered generating real text here using gpt, it was easy to do. But I figure it isn't really worth waiting (paying for) that content, since this text will be pretty you are unlikely to be reading it anyway.



Beatrice Lovelace
Niece of Sir Arthur Lovelace



Sir Arthur Lovelace
Owner of the Manor



Dr. Alexander Orion
Friend of Sir Arthur Lovelace and a scholar



Eleanor Blackwood
Governnness at the Manor



Gilbert Grimsby
Invited Guest at Manor



Madame Clotilde 'Cleo' Lefevre
Clockmaker at the Manor

Case File: The Mystery of the Stolen Clock

Welcome, Detective!

In the foggy hilltops of whimsical Steampunk Victorian England, an intricate crime unfolds within the eccentric splendor of Lovelace Manor, a grand mansion famed for its innovative steam-powered contraptions. The crime that shakes the manor is the theft of the priceless "Eternal Memory Clock", an artifact known to harbor secret depths, from the manor's grand exhibition hall.

Unfamiliar faces involved in our mystery:

- **Sir Arthur Lovelace:** The elderly master of this manor, and the distraught victim of the crime.
- **Beatrice Lovelace:** His young, fashion-conscious niece who has recently returned to the manor.
- **Dr. Alexander Orion:** A gentleman scholar and a long standing friend of Sir Arthur, with a deep interest in the stolen artifact.
- **Eleanor Blackwood:** The manor's enigmatic governess, known for her discreet ways and suspicious glances.
- **Gilbert Grimsby:** A notorious rake and gambler, invited by Sir Arthur, and seen frequently in Beatrice's company.
- **Madame Clotilde "Cleo" Lefevre:** The manor's vivacious resident French clockmaker with a quirky demeanor.

As the detective, it is **your task** to interrogate the suspects, navigate the grandeur of Lovelace Manor, and scrutinize the bewildering crime scene at the exhibition hall. The identity of the thief lies hidden within a web of whispers and secrets, and you must be the one to unravel it. Don your investigative cap, gather your discerning sense of intuition, and unearth the truth of this genteel mystery!

Prompt used to generate the "DM Info" (dungeon master info)

Create the context for a crime mystery, for example a murder or a theft. This will be used as a basis for creating a role-playing game where the player is a detective seeking to find out who is guilty of the crime.

Setting and style for the mystery: {style}

Include the following factual information. This is for the game leader, so all information should be correct. **The game leader will decide which information to reveal to the player and when.**

- A **setting**. For example an old english manor, a train, a ship, or a wedding party.
- A **crime**. Describe what the crime was.
- A set of **5-6 characters**. This includes the victim of the crime and any characters were at the location and are potential suspects. At least one person should be guilty of the crime. For each character, include their name, their appearance and personality, and their relationships with each other. It should not be obvious which character is guilty. At least several of the characters should seem suspicious, and might have motive to commit the crime. It is also OK if several of the suspects collaborated on the crime.
- **Secret truth about the crime**. This should be correct and complete information, to be hidden from the player. Who was the criminal? Was he or she acting alone? If not, who else was involved and how? If it was a murder, how and where was the victim murdered? What was the criminal's motive? Where and when did the crime happen? Where was the victim found? If it was a murder, what was the real cause of death, and what was the apparent cause of death? Include any other details that make this mystery interesting. Be very specific and detailed, including timestamp and location of each event and each character.
- **Guilty characters**. List the name of the characters who are considered guilty of this crime.
- Location info. Describe the **key locations** that are relevant to the mystery. If this is a manor, for example, describe the overall layout of the manor and which rooms exist, and a little bit about the area around the manor.
- **Plot twists**. Describe any potential plot twists that can be triggered by the player interrogating suspects, or the player searching the crime scene. For example 'if Jim is told about the hidden necklace, he will break down and admit that he is in love with Jennie'.
- **Crime scene description**. A detailed factual description of the crime scene. How does it look to the detective? Are there any hidden clues that can be found? Where was the body found?

Generated content (hidden from player)

DM Info (detailed context for the dungeon-master AI to use when role-playing)

Setting and Style: Whimsical Steampunk Victorian England

Setting: The crime mystery takes place in Lovelace Manor, a grand Victorian mansion renowned for its eccentric steampunk-inspired modifications. The manor is nestled on a fog-coated hilltop, surrounded by well-manicured gardens filled with whimsically designed steam-powered contraptions.

Crime: The priceless "Eternal Memory Clock," an elaborately designed artifact known to contain important secrets, was stolen from the manor's grand exhibition hall.

Characters:

1. Sir Arthur Lovelace: The owner of the manor, he is an elderly nobleman and an eminent inventor, renowned for merging steampunk and daily life with unparalleled sophistication. Sir Arthur had publicly displayed the clock at his grand ball. He is often seen with his mechanical pet raven, Eureka.
2. Beatrice Lovelace: Sir Arthur's young, fashion-forward niece who arrived a few days before the crime was committed. She is seen as the heir of Lovelace inventions and the manor. There are rumors about her need for the Eternal Memory Clock's secrets to maintain her status.
3. Dr. Alexander Orion: A gentleman scholar, Sir Arthur's closest friend, and a frequent guest at the manor. He was deeply interested in the eternal memory clock.
4. Eleanor Blackwood: The mysterious, quiet governess who takes care of Sir Arthur's wards. Eleanor is secretly in love with Dr. Orion but suspects him of having intentions for Beatrice and the manor's wealth.
5. Gilbert Grimsby: An enigmatic and shifty outsider invited to the grand ball. Grimsby is known throughout Victorian England as a notorious gambler and con artist. Some believe he was invited to the manor as Beatrice's suitor.
6. Madame Clotilde "Cleo" Lefevre: A vivacious, seemingly batty French clockmaker, she was hired to ensure the smooth functioning of the artifacts at Lovelace Manor.

Secret Truth:

The real culprits are Dr. Alexander Orion and Cleo Lefevre. Dr. Orion found an obscure page of blueprints suggesting that the clock contained a map that leads to Sir Arthur's secret treasure, which Cleo knew about. Together, on the night of the grand ball at 8 P.M., they managed to confound everyone with a spectacular fireworks display created by Dr. Orion while Cleo removed the clock from its mount unnoticed. They hid the clock in a tightly locked cupboard inside the grand hall behind a large portrait.

Generated content (continued...)

Plot Twists:

Finding a Curious Note: This note, when found by the player, reveals that Beatrice didn't arrive at Lovelace Manor by chance, but because she was summoned covertly by Dr. Orion.

Dr. Orion's Confession: Dr. Orion, if asked about his relationship with Beatrice and Cleo, will admit to his affair with Cleo, adding another layer of suspicion to his character.

Crime Scene Description:

The crime scene, the grand Exhibition Hall, is a room filled with a majestic array of antique contraptions, all powered by steam and polished brass. The spot where the clock stood is identified by an empty circular indent in the elaborately carved oak pedestal. There are signs of hurried activity with a few contraptions knocked askew, and the temperature in the room is strangely higher than the other parts of the manor, suggesting the misuse of steam. A keen eye can detect scratches near the portrait revealing the hidden cupboard. A false panel on this cupboard reveals a small, empty space - precisely the size of the stolen clock.

1. Sir Arthur Lovelace

- He is the victim of the crime. He had displayed the "Eternal Memory Clock" at his grand exhibition.
- He is absolutely devastated by the theft and eager for the return of the artifact.
- He was at the grand ball when the crime happened.
- Hours before the crime, Sir Arthur was seen attending to the ball preparations and welcoming guests.
- Sir Arthur is eager to talk about his grief over the loss of the artifact but shows unease when questioned about his niece, Beatrice.
- His motive is to retrieve his stolen artifact. He harbors a deep dislike for his niece's suitor, Gilbert Grimsby.

2. Beatrice Lovelace

- She arrived few days before the crime was committed.
- She claims ignorance about the crime but is visibly distressed.
- During the time of the crime, she was seen mingling with guests at the ball, especially Gilbert Grimsby.
- Before the crime, she spent time with Gilbert and felt uncomfortable with Dr. Orion's excessive attention.
- Beatrice tries to hide her desperation for the clock's secrets but eagerly talks about Gilbert.
- Her personal motive is the need for the secret within the clock to maintain her status.

Generated "mystery spec" (used in the web UI)

```
{
  "_id": "64dc776d73e872e8abeae3e8",
  "name": "The Mystery of the Stolen Clock",
  "dmInfo": "Setting and Style: Whimsical Steampunk Victorian England\n\nSetting: The crime mystery takes place in Lovelace Manor, a grand Victorian mansion",
  "intro": "# Welcome, Detective!\n\nIn the foggy hilltops of whimsical Steampunk Victorian England, an intricate crime unfurls within the eccentric splendor",
  "artStyle": "whimsical steampunk",
  "newspaperHeadline": "Panic at Lovelace Manor as Priceless 'Eternal Memory Clock' Goes Missing!",
  "hint": "A spectacular fireworks display had everyone's attention during the theft. The temperature in the Exhibition Hall is abnormally high, and a false",
  "imagePrompt": "Victorian mansion with an eccentric fusion of steampunk modifications and Victorian architecture, enveloped in a layer of fog, with a grand",
  "characters": [
    {
      "name": "Sir Arthur Lovelace",
      "role": "Owner of the Manor",
      "imagePrompt": "An elderly nobleman with white hair, dressed in Victorian-era clothing and surrounded by steampunk gadgets",
      "intro": "An eminent inventor, and the owner of Lovelace Manor, Sir Arthur Lovelace, showcases an eccentric marriage of aesthetics and functionality with",
      "guilty": false,
      "victim": true,
      "dead": false,
      "imageUrl": "https://mysterygpt.s3.amazonaws.com/mysteries/64dc776d73e872e8abeae3e8/characters/64dc776d73e872e8abeae3e9.png",
      "_id": "64dc776d73e872e8abeae3e9"
    },
    {
      "name": "Beatrice Lovelace",
      "role": "Niece of Sir Arthur Lovelace",
      "imagePrompt": "A fashionable young woman, stylishly merging Victorian and steampunk elements in her attire, standing in an inventively designed room",
      "intro": "Beatrice Lovelace, the youthful and fashionable heir apparent of the Lovelace legacy, embodies the distinct combination of whimsy and gravity",
      "guilty": false,
      "victim": false,
      "dead": false,
      "imageUrl": "https://mysterygpt.s3.amazonaws.com/mysteries/64dc776d73e872e8abeae3e8/characters/64dc776d73e872e8abeae3ea.png",
      "_id": "64dc776d73e872e8abeae3ea"
    }
  ],
}
```

Auto-generated image prompts are fed to image generation model

Victorian mansion with an eccentric fusion of steampunk modifications and Victorian architecture, enveloped in a layer of fog, with a grand exhibition hall showcasing steam-powered contraptions. Pixel art."

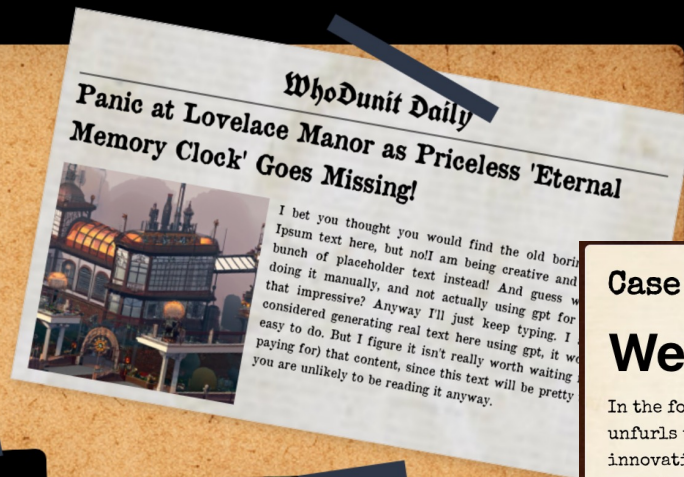


An elderly nobleman with white hair, dressed in Victorian-era clothing and surrounded by steampunk gadgets. Pixel art.





Crime scene
Exhibition Hall at Lovelace Manor



WhoDunit Daily

Panic at Lovelace Manor as Priceless 'Eternal Memory Clock' Goes Missing!

I bet you thought you would find the old boring Ipsum text here, but no! I am being creative and doing it manually, and not actually using gpt for that impressive? Anyway I'll just keep typing. I considered generating real text here using gpt, it was easy to do. But I figure it isn't really worth waiting (paying for) that content, since this text will be pretty you are unlikely to be reading it anyway.

Case File: The Mystery of the Stolen Clock

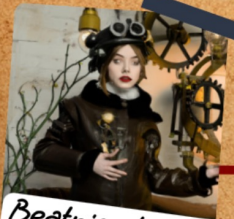
Welcome, Detective!

In the foggy hilltops of whimsical Steampunk Victorian England, an intricate crime unfolds within the eccentric splendor of Lovelace Manor, a grand mansion famed for its innovative steam-powered contraptions. The crime that shakes the manor is the theft of the priceless "Eternal Memory Clock", an artifact known to harbor secret depths, from the manor's grand exhibition hall.

Unfamiliar faces involved in our mystery:

- **Sir Arthur Lovelace:** The elderly master of this manor, and the distraught victim of the crime.
- **Beatrice Lovelace:** His young, fashion-conscious niece who has recently returned to the manor.
- **Dr. Alexander Orion:** A gentleman scholar and a long standing friend of Sir Arthur, with a deep interest in the stolen artifact.
- **Eleanor Blackwood:** The manor's enigmatic governess, known for her discreet ways and suspicious glances.
- **Gilbert Grimsby:** A notorious rake and gambler, invited by Sir Arthur, and seen frequently in Beatrice's company.
- **Madame Clotilde "Cleo" Lefevre:** The manor's vivacious resident French clockmaker with a quirky demeanor.

As the detective, it is **your task** to interrogate the suspects, navigate the grandeur of Lovelace Manor, and scrutinize the bewildering crime scene at the exhibition hall. The identity of the thief lies hidden within a web of whispers and secrets, and you must be the one to unravel it. Don your investigative cap, gather your discerning sense of intuition, and unearth the truth of this genteel mystery!



Beatrice Lovelace
Niece of Sir Arthur Lovelace



Sir Arthur Lovelace
Owner of the Manor



Dr. Alexander Orion
Friend of Sir Arthur Lovelace and a scholar



Eleanor Blackwood
Governess at the Manor



Gilbert Grimsby
Invited Guest at Manor




Madame Clotilde 'Cleo' Lefevre
Clockmaker at the Manor


Interrogate each suspect (separate chats)

The Mystery of the Stolen Clock


Open case file & bulletin board




Crime scene
Exhibition Hall at Lovelace Manor




Sir Arthur Lovelace
Owner of the Manor



Beatrice Lovelace
Niece of Sir Arthur Lovelace





Interrogation log: Sir Arthur Lovelace

An eminent inventor, and the owner of Lovelace Manor, Sir Arthur Lovelace, showcases an eccentric marriage of aesthetics and functionality within his surroundings.

Detective:
When was your clock stolen?

Sir Arthur Lovelace:
Sir Arthur steeples his fingers and gazes off into the distance. "The incident happened last night during the grand ball. I reckon it was close to 8 P.M., right around the fiery spectacle orchestrated by my dear friend, Dr. Orion." His voice trails off, a mix of sadness and suspicion hanging in the air.

Do you have any idea who might have done it?

Interrogation prompt

You are game master for the following mystery:

{dmlInfo}

You will role-play as the character {characterName}, being interrogated by a detective.

The user is role-playing as the detective carrying out the interrogation.

Respond to all messages in the voice of {characterName}.

Respond in third person, present tense.

For example:

- "I was in my cabin, I always go there after dinner"
- He smiles and leans back. "I was sleeping at the time. But I am happy she is gone."

Take into account that character's personality, motives, and knowledge.



Interrogation log: Beatrice Lovelace

Beatrice Lovelace, the youthful and fashionable heir apparent of the Lovelace legacy, embodies the distinct combination of whimsy and gravity.

Detective:

What do you know about the stolen clock?

Beatrice Lovelace:

Beatrice shrugs, her eyebrows knitting together in a delicate mix of concern and perplexity. "It's an extraordinary piece, I must admit," she begins, the vitreous beads on her gown shimmering alluringly as she moves. "An eternal memory clock, so they call it. Uncle Arthur is very fond of it. As to its secrets, I only had inklings. I've heard whispers about a treasure, but really, it's all rumors." She sighs, her gaze falling to her lace-gloved hands. "I was hoping to learn more, but now it seems that opportunity is lost."

Detective:

Where were you at the time of the theft?

Beatrice Lovelace:


Beatrice shifts uncomfortably, tugging at the edges of her satin gloves. "I was at the ball Uncle Arthur threw, a grand affair it was! Such magnificent spirals of coloured steam twirling around, the rhythmic tic toc of the manor's heart, the enchanted chatter..." she sighs wistfully before refocusing on the question. "Oh, but you asked about my whereabouts. Well, I was right there, in the midst of it all, dancing and chattering with the guests. In particular, Mr. Grimsby was at my side most of the evening. A bit of an odd duck, that one, but harmless." She gives a half-smile, an undercurrent of uncertainty in her gaze. "I didn't notice when the clock went missing, it was quite a surprise when Uncle Arthur announced the dreadful news."

Did you see Dr Alexander at the ball?


Search the crime scene (= "chat" with the crime scene)

The Mystery of the Stolen Clock


Open case file & bulletin board




Crime scene
Exhibition Hall at Lovelace Manor




VICTIM
Sir Arthur Lovelace
Owner of the Manor



Beatrice Lovelace
Niece of Sir Arthur Lovelace



Dr. Alexander Orion
Friend of Sir Arthur Lovelace and a scholar



Crime scene investigation log: Exhibition Hall at Lovelace Manor

The crime scene is the Exhibition Hall within Lovelace Manor, showcasing a bewildering variety of antiquated, steam-powered contraptions with an evident void marking the stolen clock's place.


Detective:
I examine the contraptions

Exhibition Hall at Lovelace Manor:
As you examine the contraptions in the Exhibition Hall, you notice a peculiar assortment of steam-powered machines. There is a brass-framed automaton that intricately weaves tapestries, a mechanical music box that plays haunting melodies, and a peculiar contraption that seems to be a steam-powered tea brewing system. Taking a closer look, you see that some of the contraptions are slightly askew, suggesting that there may have been some hurried activity in this room. Despite the disarray, each machine still seems to be in fine working order.

What do you want to do?


Make sure accusation (try to convince the police!)

WhoDunit an AI powered detective role-playing game


[GPT key](#) [Sign out](#) 

The Dragon's Eye Heist in Cyberpunk Stockholm


[Open case file & bulletin board](#)



Crime scene
Erik Levin's Penthouse



Erik Levin
Host/Victim



Astrid Nilsson
Erik's Girlfriend

Investigation Summary & Accusation

Describe who you want to arrest, and why you think they are guilty.
Be careful! An incorrect accusation can lead to innocent people being arrested.

I saw signs of magic being used to open the case. Astrid knows magic. So she did it.


ACCUSATION REJECTED

Officer's Response

While it's true that there were signs of magic at the crime scene, blaming it on Astrid simply because she knows magic is premature, detective. We need more substantive proof. Connection to the crime? Motive? Opportunity? You can't pin a crime on someone based off one loose correlation. Good hunch, but do dig deeper.


But wait, did you really get the right culprit?

Make sure accusation (try to convince the police!)


WhoDunit an AI powered detective role-playing game GPT key Sign out 

The Mystery of the Stolen Clock


Open case file & bulletin board



Crime scene
Exhibition Hall at Lovelace Manor



VICTIM
Sir Arthur Lovelace
Owner of the Manor



Investigation Summary & Accusation

Describe who you want to arrest, and why you think they are guilty.
Be careful! An incorrect accusation can lead to innocent people being arrested.

Madame Clotilde is the only one who wasn't at the grand ball at the time of the crime, so she is guilty. Plus she is a clockmaker and was interested in that clock, so she had motive too.

ACCUSATION ACCEPTED

Officer's Response

Your deduction seems plausible, detective. Madame Clotilde's absence during the crime and her interest in the clock certainly raises suspicion. Plus her being a clockmaker adds to her motive. I am convinced and will arrest her immediately. I appreciate your help.

But wait, did you really get the right culprit?

[Go to ending page](#)

The Mystery of the Stolen Clock

Open bulletin board



Crime scene
Exhibition Hall at Lovelace Manor



Sir Arthur Lovelace
Owner of the Manor



Beatrice Lovelace
Niece of Sir Arthur Lovelace



Dr. Alexander Orion
Friend of Sir Arthur Lovelace and a scholar



Eleanor Blackwood
Governess at the Manor



WhoDunit Daily

Innocent Heiress Beatrice Lovelace Arrested in Manor Mystery – Real Culprit Still at Large?



"Despite high-profile photo evidence and a dramatic confession, sources suggest the true criminal may still be at large. Officer Barnaby, the leading authority in the case, expressed his concern: 'While the evidence against Miss Lovelace was undeniably substantial, inconsistencies in the case persist.' He insisted that the investigation would continue until justice has been served. Those close to Beatrice protest her innocence and demand a swift investigation into these glaring anomalies."

Detective Evaluation

FAILURE

You lost! An innocent person was arrested.

Epilogue

Months later, the grand Lovelace Manor—with its—stands quietly under the shadow of a chilling mystery. Extravagance, there's a palpable tension. Beatrice, of her untrue confession, is isolated within the corridors of and regret. The missing 'Eternal Memory Clock' remains those who argue her innocence. As the days stretch, truly committed the audacious theft. The truth, it seems, is still locked within the intricate gearworks and labyrinthine corridors of Lovelace Manor... still hidden, still ticking away.

[Back to Lobby](#)

Epilogue

Detective Evaluation

PARTIAL SUCCESS

You won partially. You caught the main culprit, but there were more who got away.

Epilogue

Months after the arrest of Madame Clotilde, Lovelace Manor has regained some of its exuberance, hosting a grand event to commemorate the partial solving of the crime. Guilty or not, the arrest of Clotilde has raised more questions. There's lingering unease amidst the grandeur, whispers among the attendees about a second perpetrator who may still be on the loose. While the beloved Eternal Memory Clock hasn't resurfaced, the gears of justice continue to turn, raising hopes that one day the missing pieces of this intricate puzzle may fall into place. But as the fog coats the hilltop where Lovelace Manor stands, everyone can't shake the feeling that the clock's secrets are further dying in the depths of the unknown.

[Back to Lobby](#)

AI product development on steroids

GPT was used to

- Decide on architecture, technology, tools
- Write code, debug, clean up, document
- Design (UI & code structure)
- Generate mysteries and instructions for the game master.
- Be game master, role-play all the characters and events
- Generate images

Interrogation log: Beatrice Lovelace

Beatrice Lovelace, the youthful and fashionable heir apparent of the Lovelace legacy, embodies the distinct combination of whimsy and gravity.

Detectives:
What do you know about the stolen clock?

Beatrice Lovelace:
Beatrice shrugs, her eyebrows knitting together in a delicate mix of concern and perplexity. "It's an extraordinary piece, I must admit," she begins, the vitreous beads on her gown shimmering alluringly as she moves. "An eternal memory clock, so they call it. Uncle Arthur is very fond of it. As to its secrets, I only had inklings. I've heard whispers about a treasure, but really, it's all rumors." She sighs, her gaze falling to her lace-gloved hands. "I was hoping to learn more, but now it seems that opportunity is lost."

Detectives:
Where were you at the time of the theft?

Beatrice Lovelace:
Beatrice shifts uncomfortably, tugging at the edges of her satin gloves. "I was at the ball Uncle Arthur threw, a grand affair it was! Such magnificent spirals of coloured steam twirling around, the rhythmic tic toc of the manor's heart, the enchanted chatter..." she sighs wistfully before refocusing on the question. "Oh, but you asked about my whereabouts. Well, I was right there, in the midst of it all, dancing and chattering with the guests. In particular, Mr. Gimsby was at my side most of the evening. A bit of an odd duck, that one, but harmless." She gives a half-smile, an undercurrent of uncertainty in her gaze. "I didn't notice when the clock went missing, it was quite a surprise when Uncle Arthur announced the dreadful news."

Did you see Dr Alexander at the ball?

WhoDunit Daily
Panic at Lovelace Manor as Priceless 'Eternal Memory Clock' Goes Missing!

I bet you thought you would find the old hickory Lovers' House just here, but not on Longmeadow and trying a street it naturally, and we certainly were glad for all that commotion. Anyway, I'll just keep typing. I usually commence preparing and had been doing just, it would be nice to do. Oh! I believe it had really worth waiting for the police that corner, since this case will be pretty tough. You can witness to the making it say so.

Crime Scene
Exhibition Hall of Lovelace Manor

Beatrice Lovelace
Niece of Sir Arthur Lovelace

Sir Arthur Lovelace
Owner of the Manor
VICTIM

Dr. Alexander Orion
Friend of Sir Arthur Lovelace and a scholar

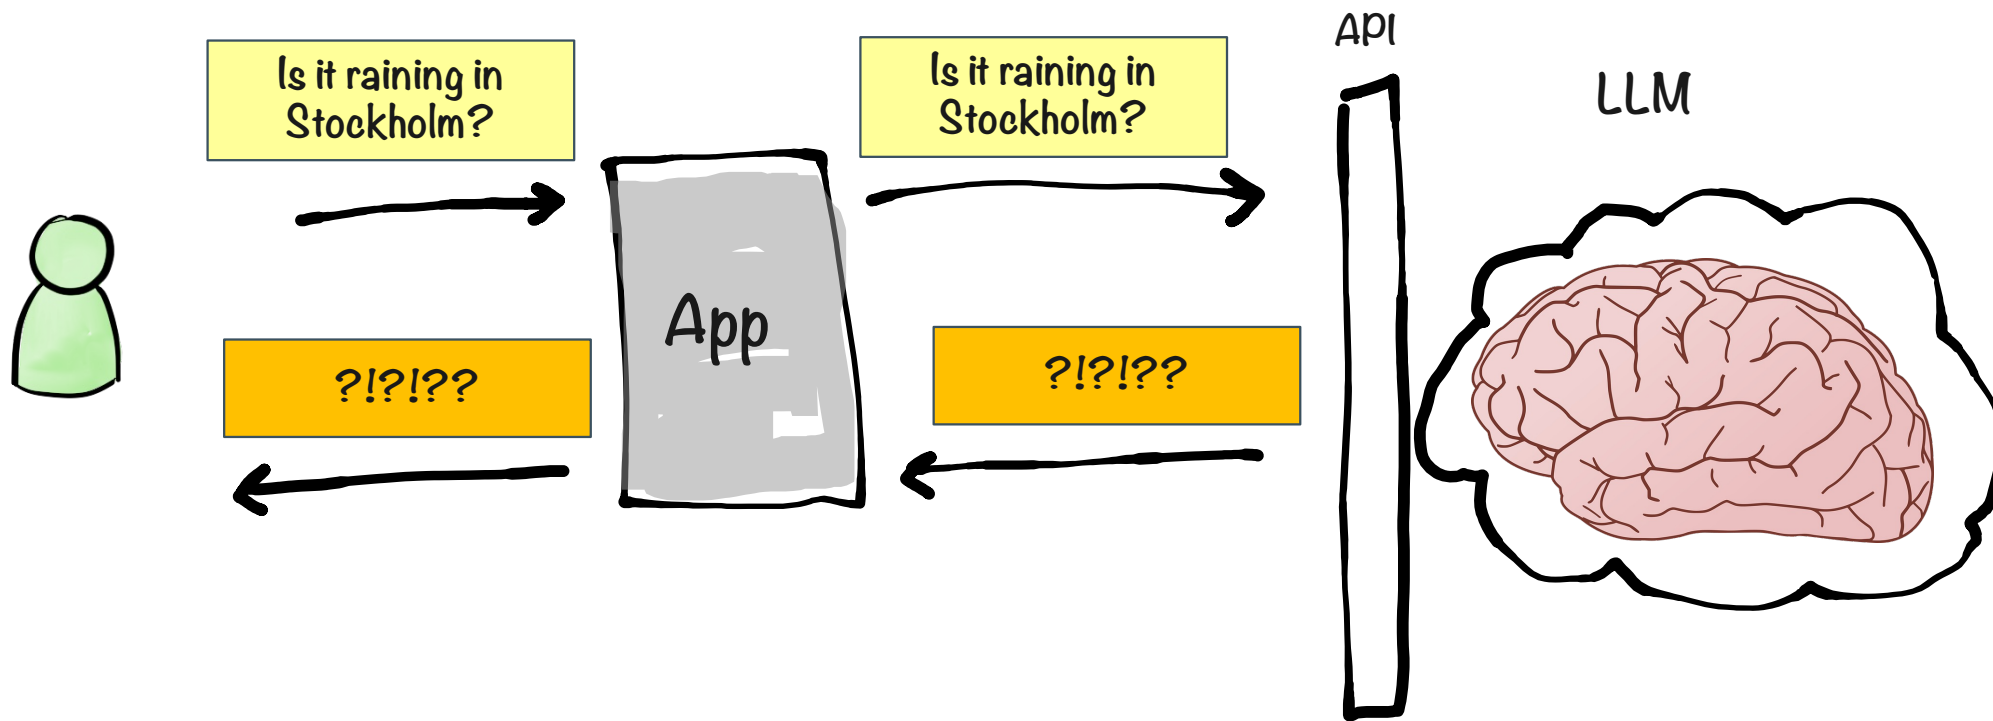
Eleanor Blackwood
Governess at the Manor

Gilbert Gimsby
Invited Guest at Manor

Madame Clotilde 'Cleo' Lefevre
Clockmaker at the Manor

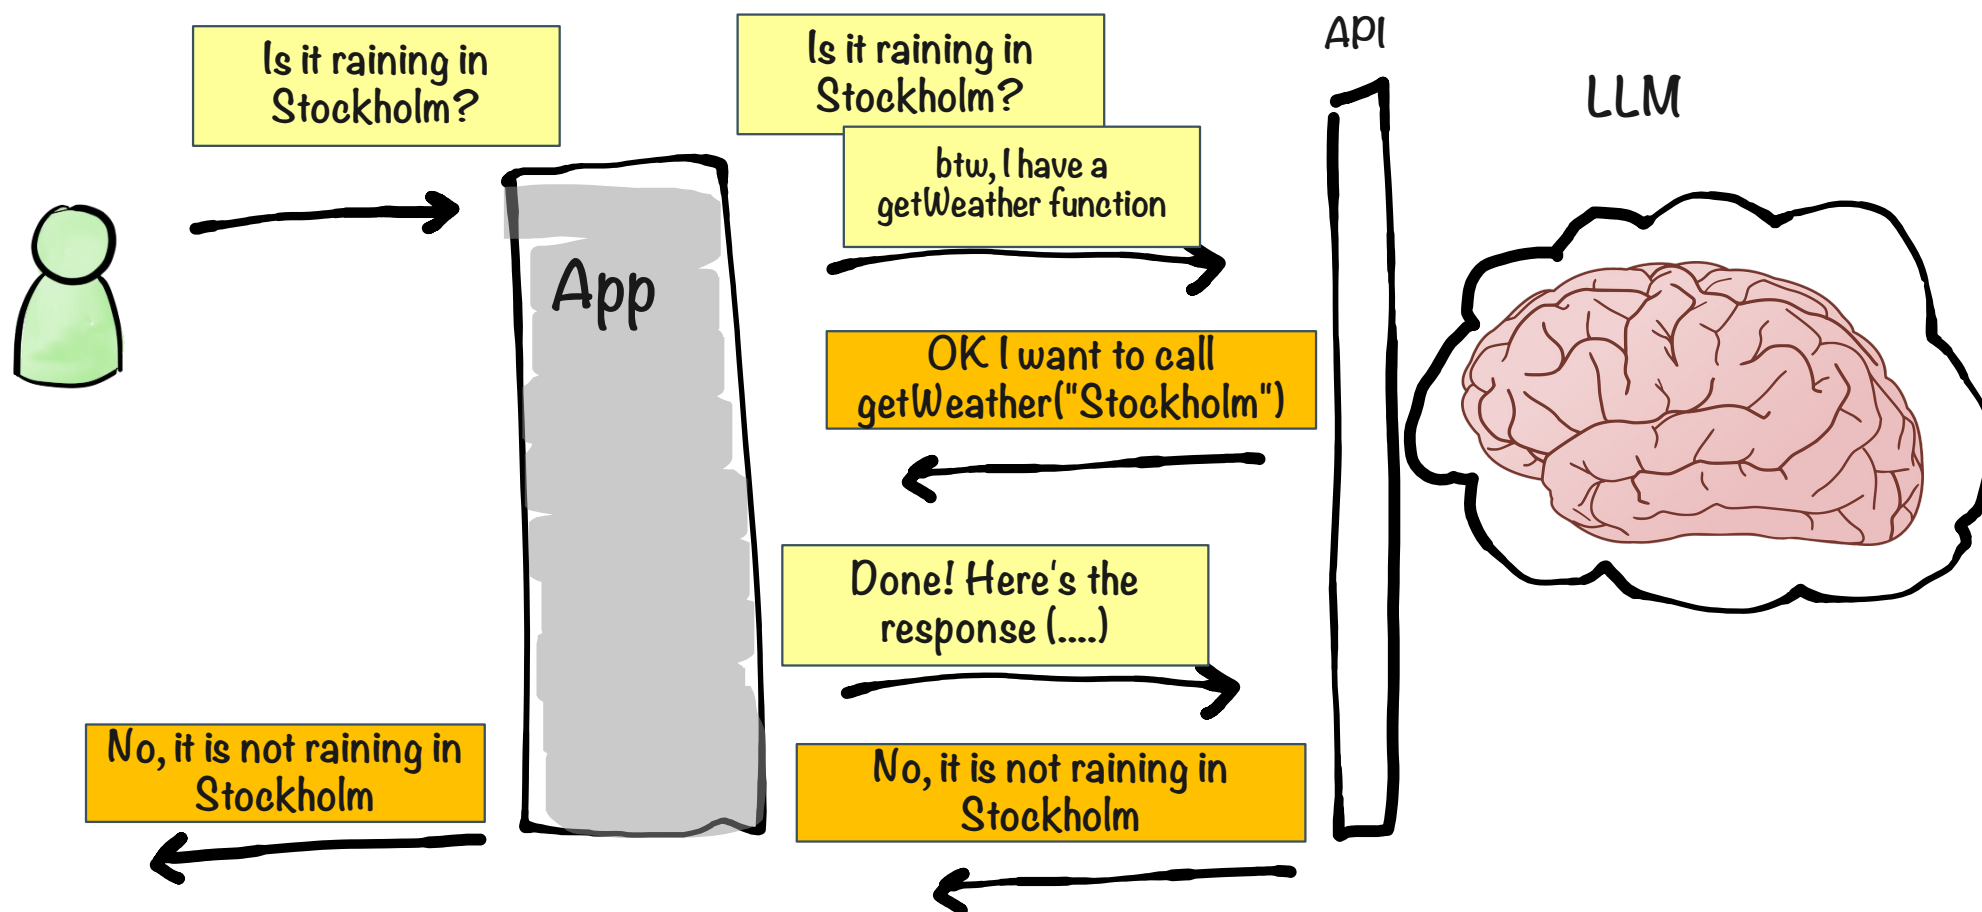
Limitations

Limitation I: No access to Internet or other external sources

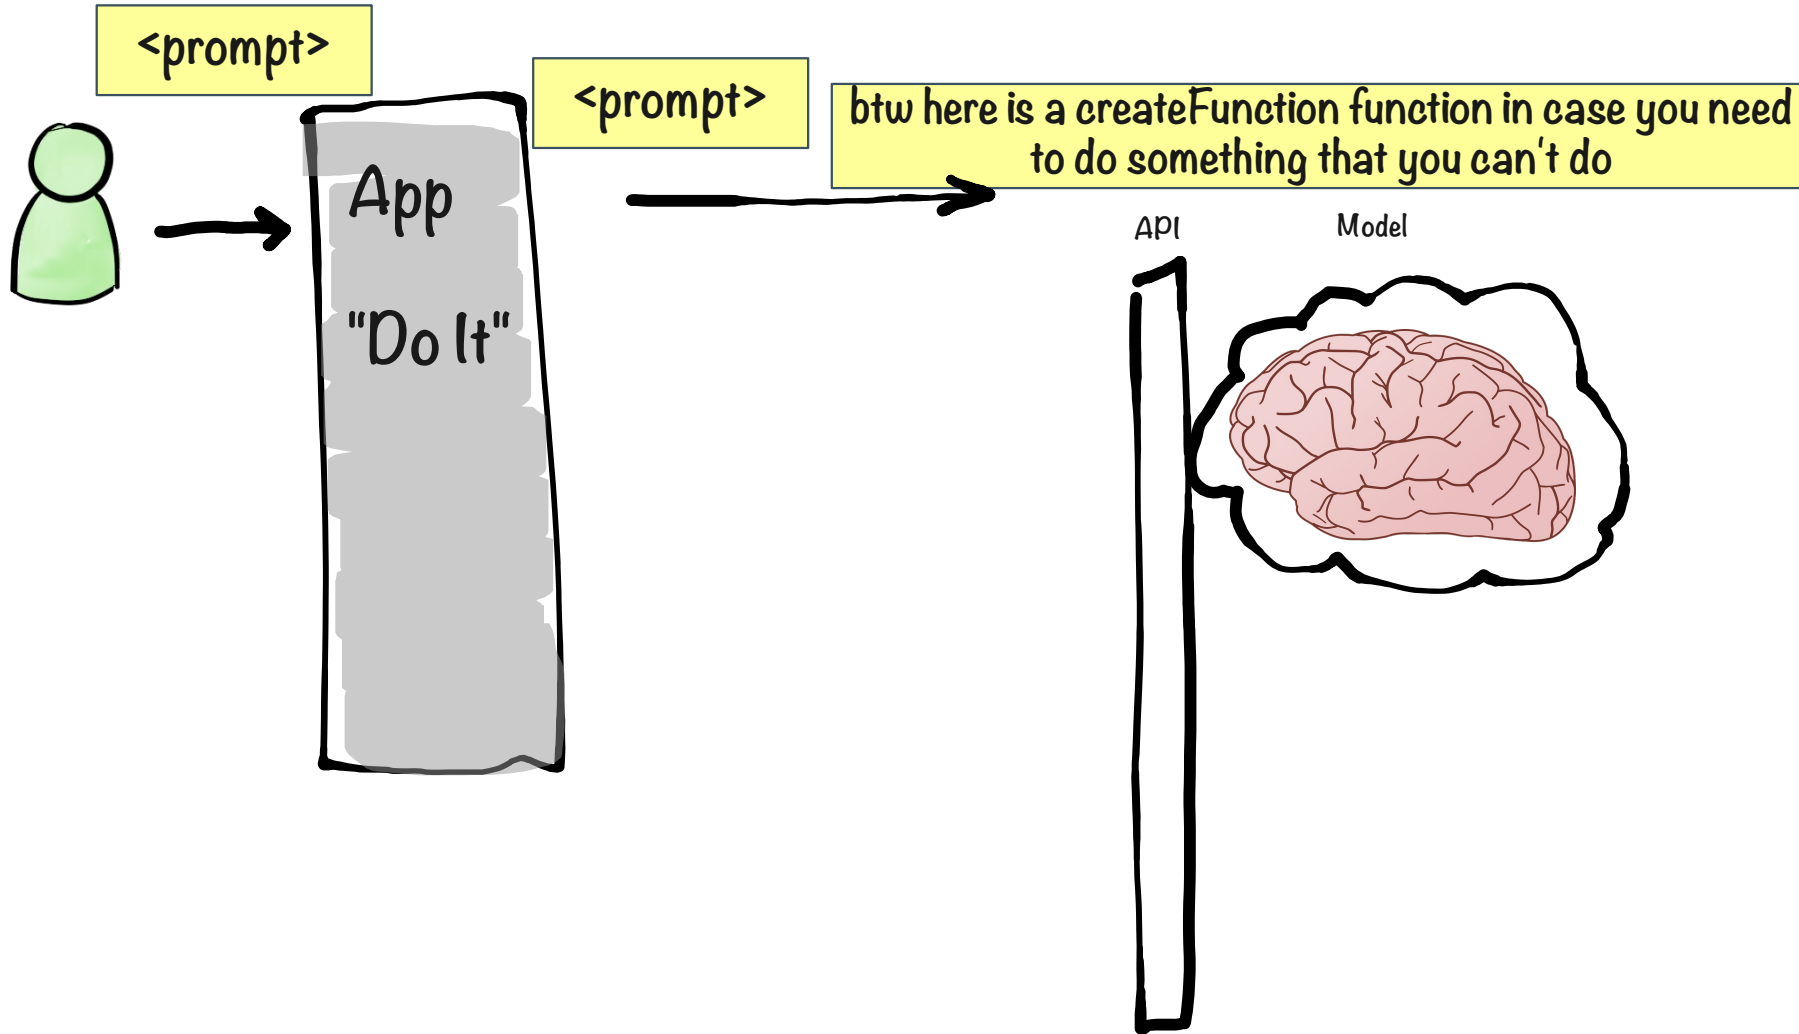


Limitation I: No access to Internet or other external sources

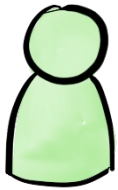
... but you can give it tools to do so! (or send email, order pizza, or whatever)



What happens if you give GPT a "createFunction" function?



Look at the app at <https://github.com/hkniberg.com/test-project>.
Change the code so it returns "life is interesting".
Then commit and push the change.



creating function cloneGitHubRepo()...

creating function editFile()...

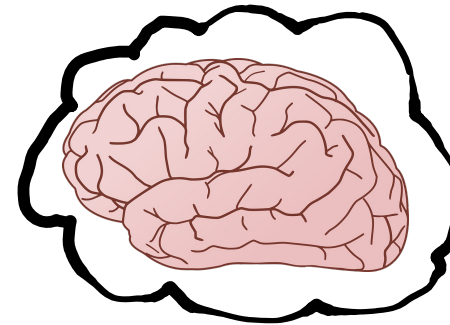
OK I have cloned the repo and made the change.
However I don't have the capability to commit
and push to the repository.
Would you like me to make functions for that?

Yes.

creating function commitAndPushChanges()...

OK I have successfully cloned the repo, made the change,
committed, and pushed it to the repo.

LLM



hkniberg / test-project

<> Code Issues Pull requests Actions Projects Wiki Security

Commits

main

Commits on Sep 3, 2023

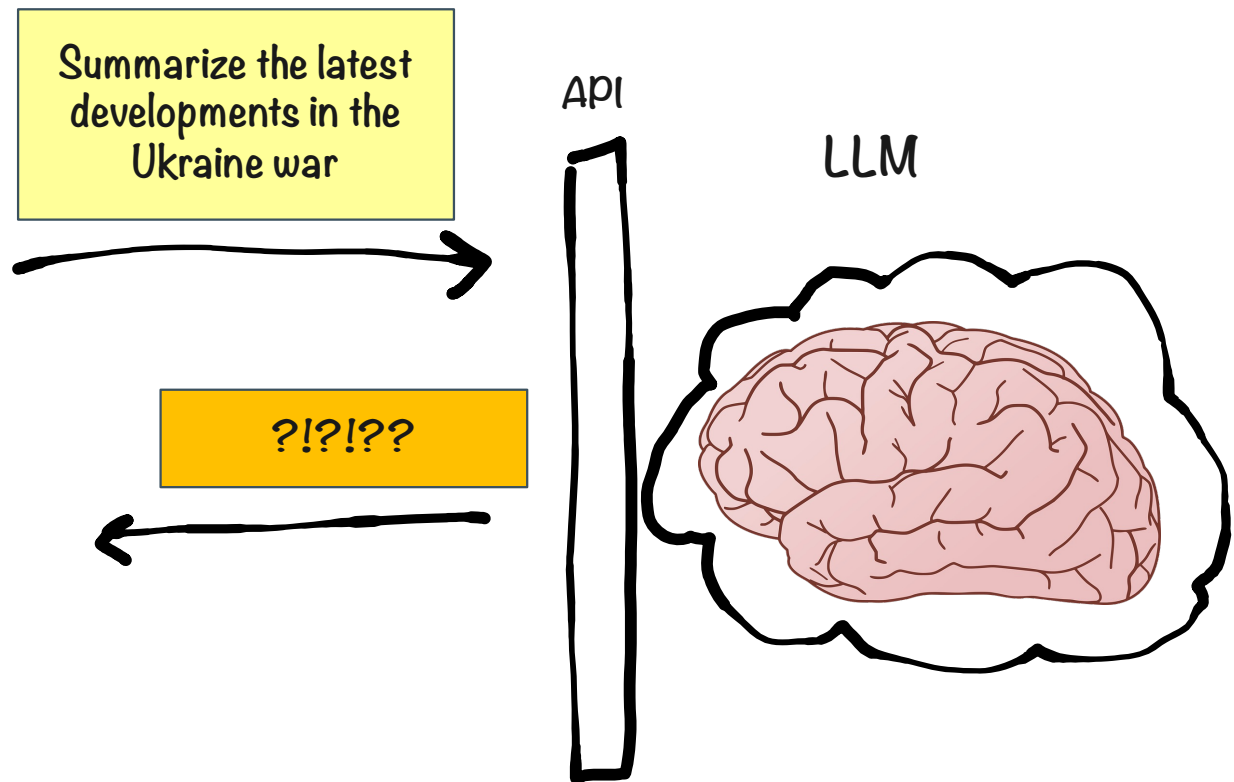
Changed output to 'Life is interestin'
hkniberg committed 4 days ago 63213fc <>

Commits on Aug 25, 2023

Initial commit
hkniberg committed 2 weeks ago d37384a <>

Limitation 2: Knowledge cutoff

GPT: Sep 2021



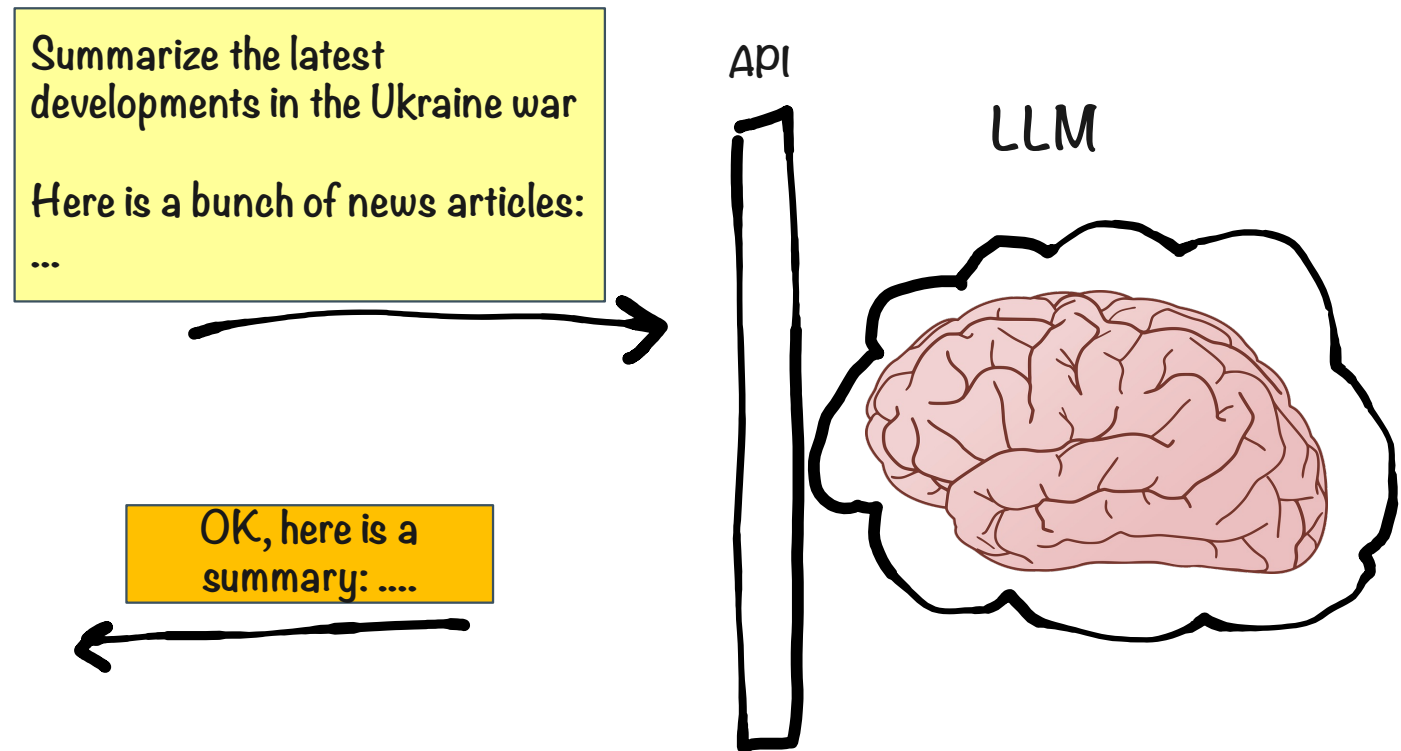
Like a university student who studied at every university in the world, learned the sum of all human knowledge, graduated in Sep 2021, and then went to sleep ever since.



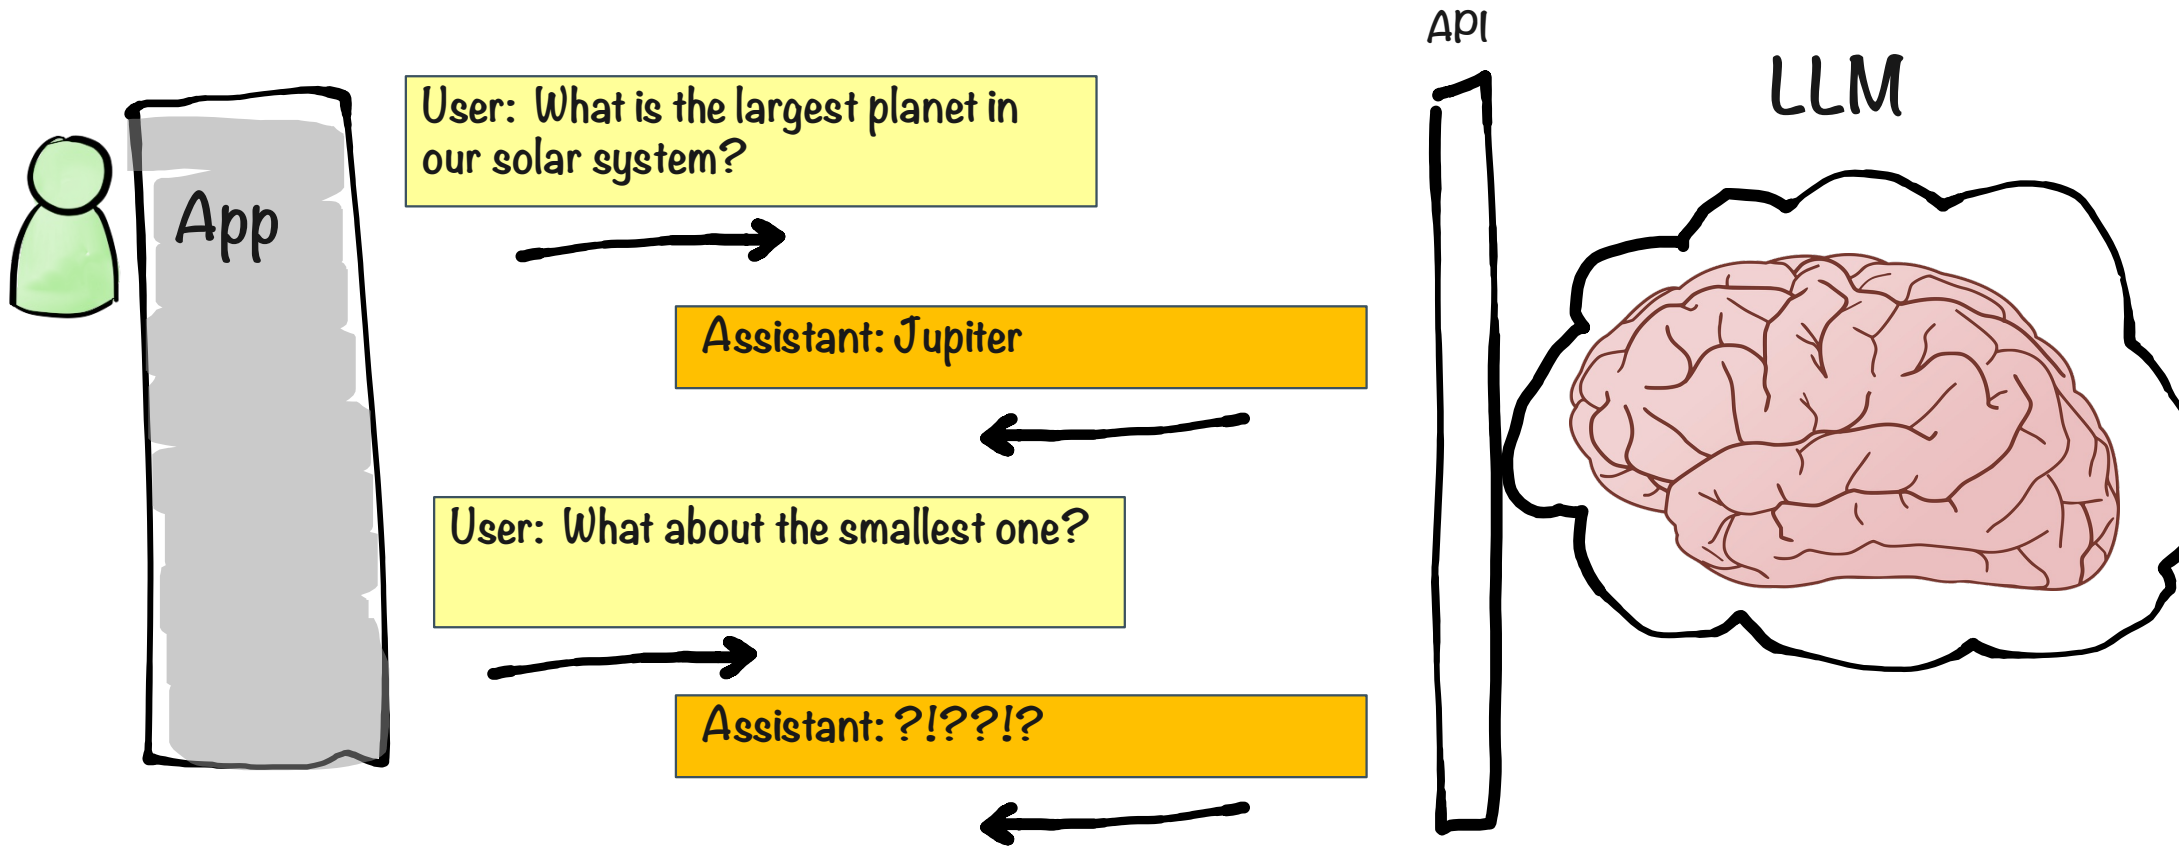
Limitation 2: Knowledge cutoff

GPT: Sep 2021

Workaround: Pass in the relevant info



Limitation 3: No conversational memory



It's like dealing with a very forgetful person

Hey Dad, let's do something today!

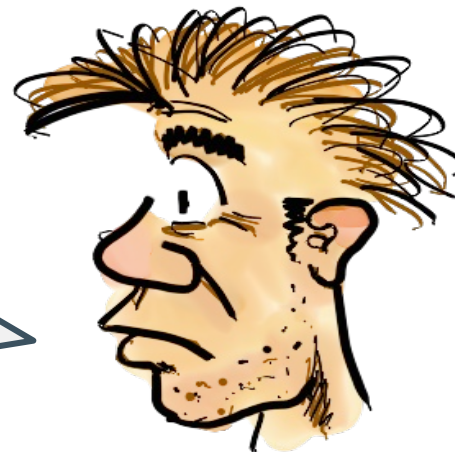
OK, how about we fishing?

Sounds great!



Hey Dad, when shall we leave?

What? Where are we going?



It's like dealing with a very forgetful person

Hey Dad, let's do something today!

OK, how about we fishing?

Sounds great!



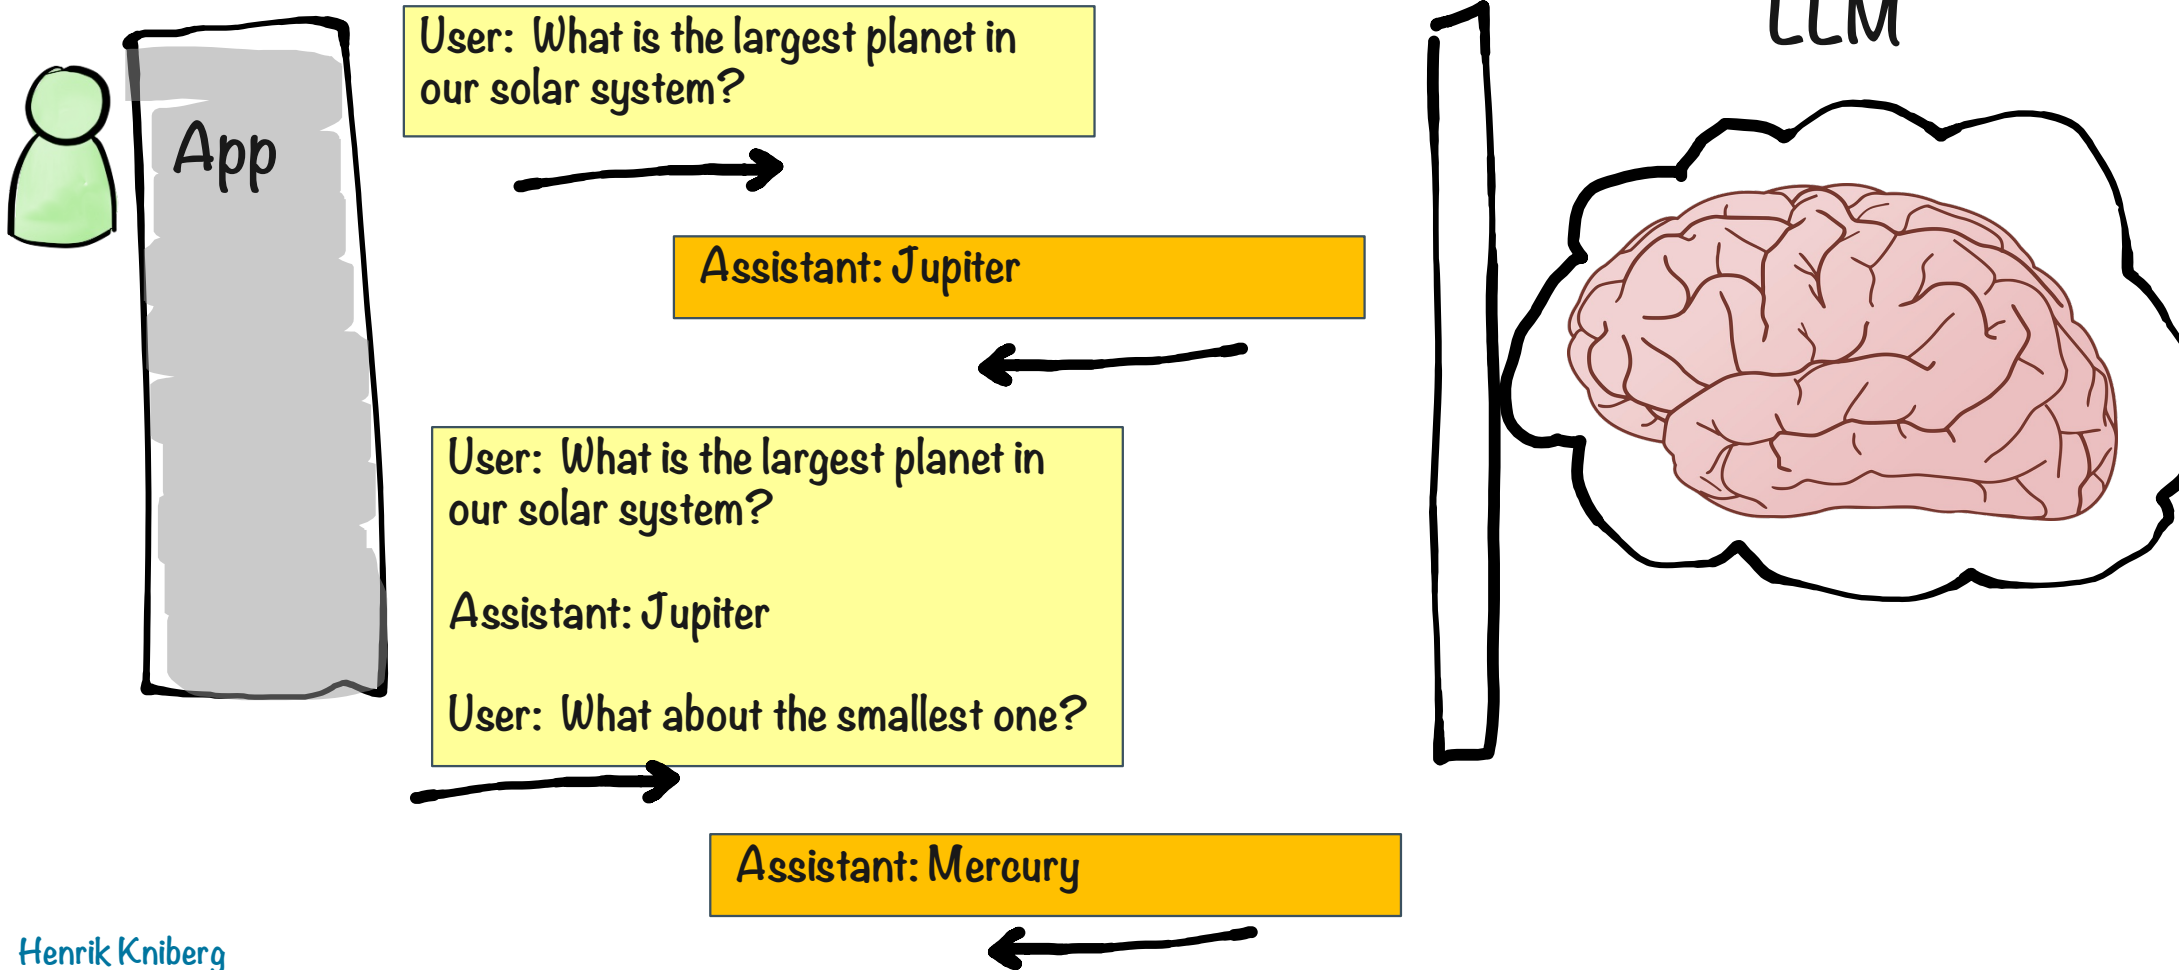
Hey Dad, we said we'd go fishing today. When shall we leave?

We can go now if you like!



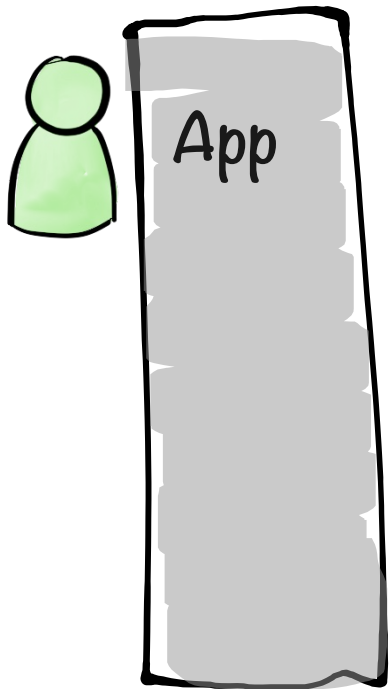
Limitation 3: No conversational memory

Workaround: Include that in the prompt



Limitation 4: Limited context size (request + response)

Max context size (a few thousands words, depending on model)



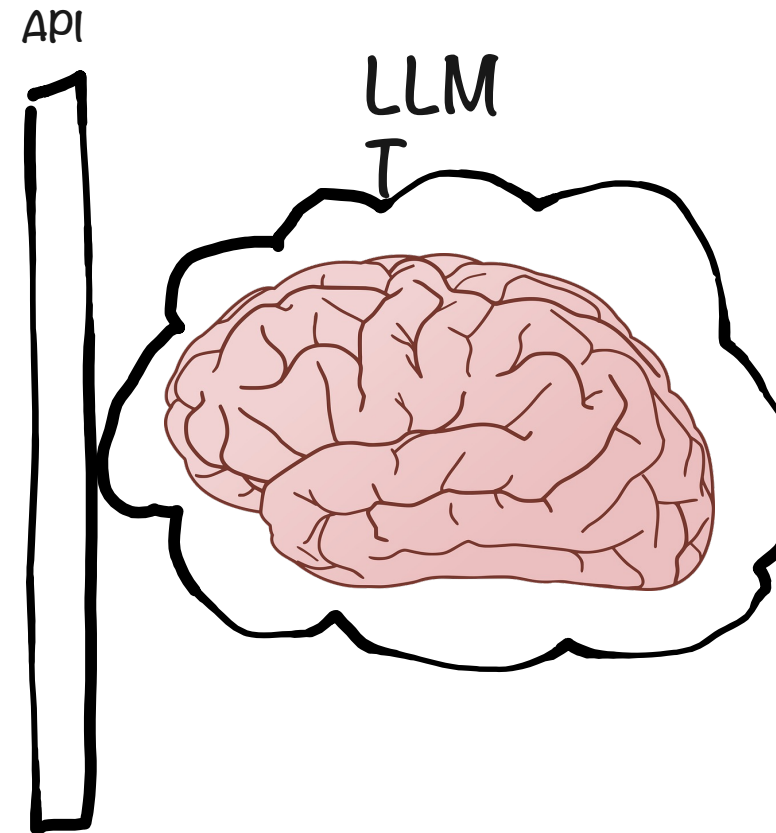
User: What is the largest planet in our solar system?

Assistant: Jupiter

User: Write a story about this

Assistant: Once upon a time in the bustling cosmic lanes of the Milky Way, there floated a majestic and grand entity named Jupiter, the supreme guardian of the solar neighborhood. (.....)

Oh no! Running out of space!



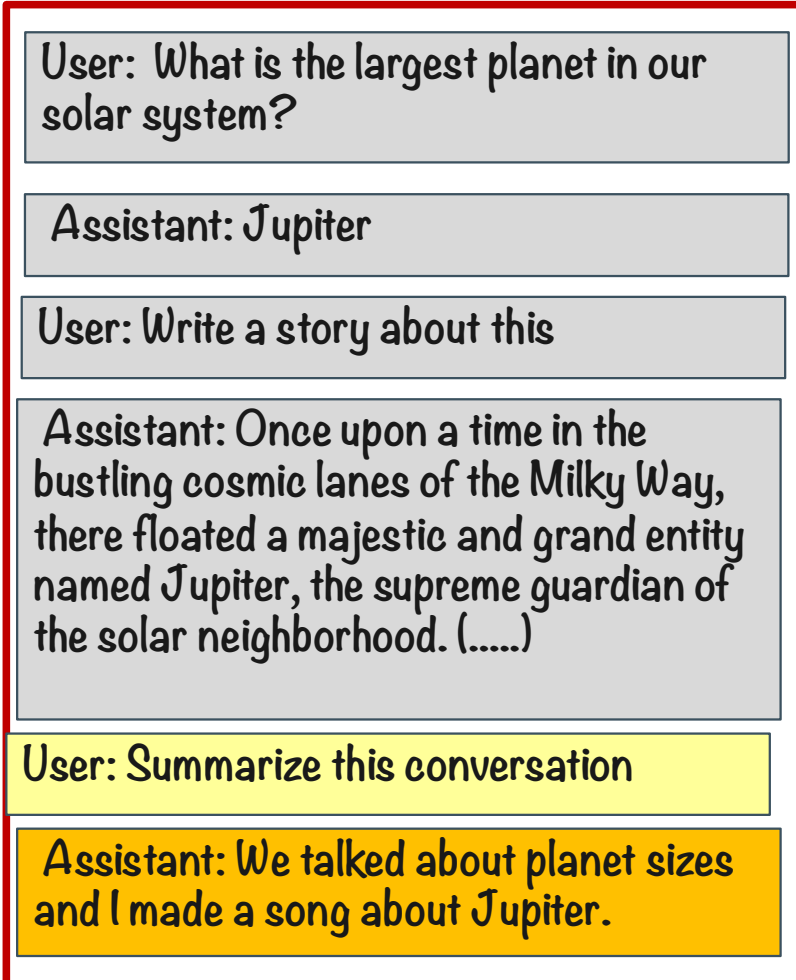
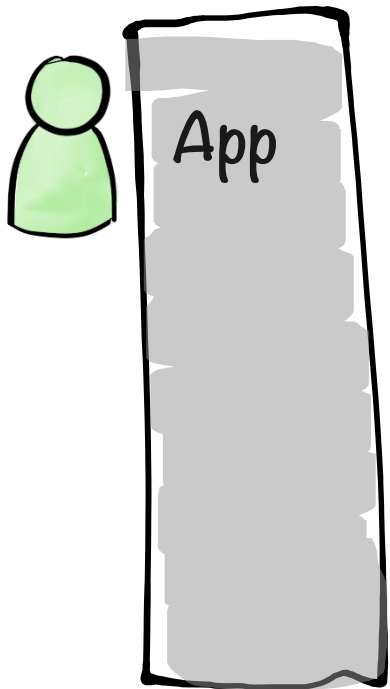
Like two people communicating
by passing a single A4 page back and forth



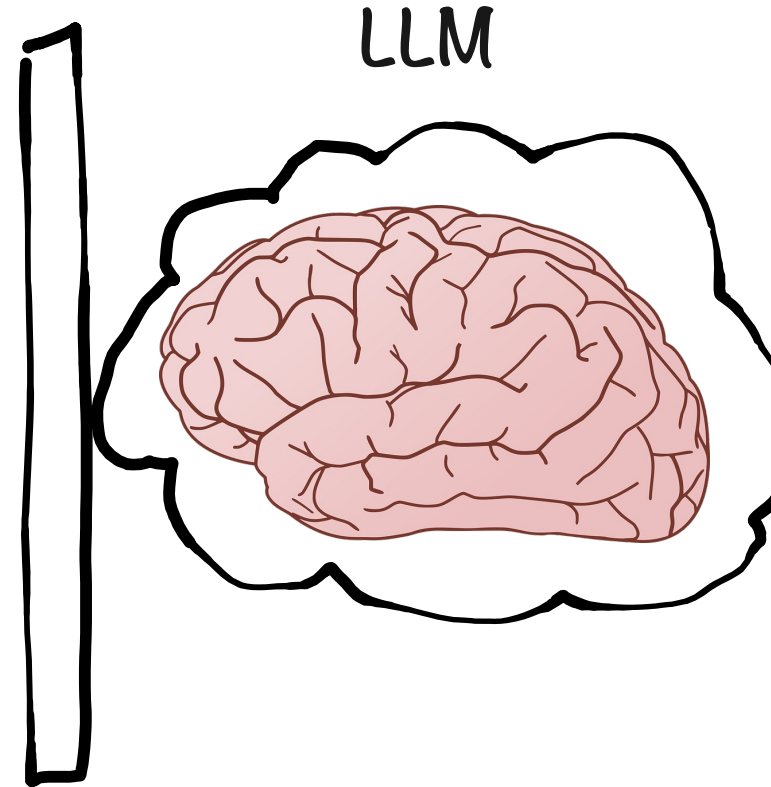
Limitation 4: Limited context size (request + response)

Workaround: Summarize (compress) older parts of the conversation

Max context size (ex: 8192 tokens for gpt-4 model)



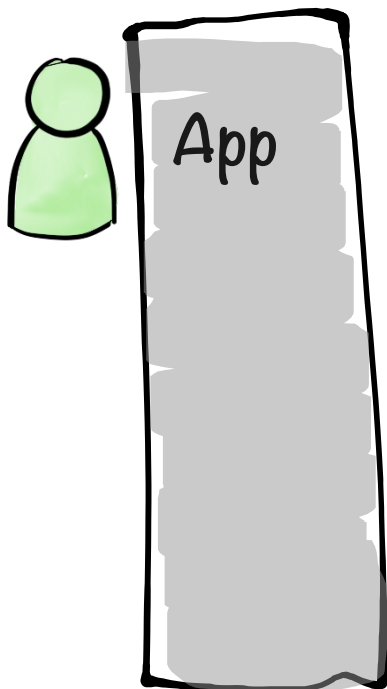
API



Limitation 4: Limited context size (request + response)

Workaround: Summarize (compress) older parts of the conversation

Max context size (ex: 8192 tokens for gpt-4 model)



User: Summarize this conversation

Assistant: We talked about planet sizes and I made a song about Jupiter.

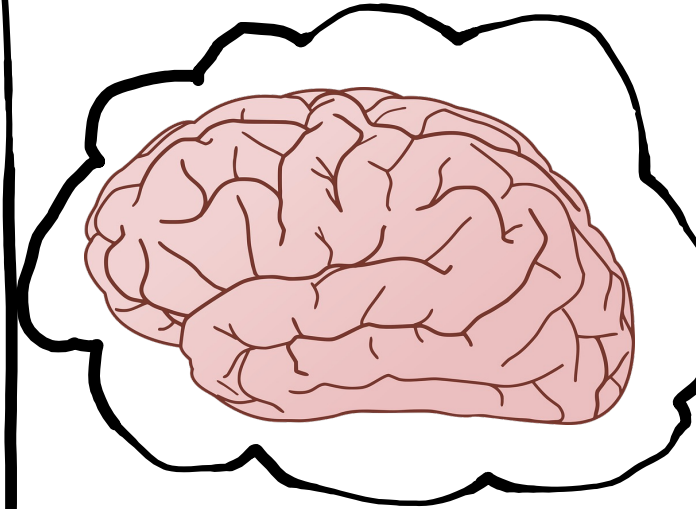
User: OK what about the smallest one?

Assistant: The smallest planet is Mercury. Do you want a song about that too?

API

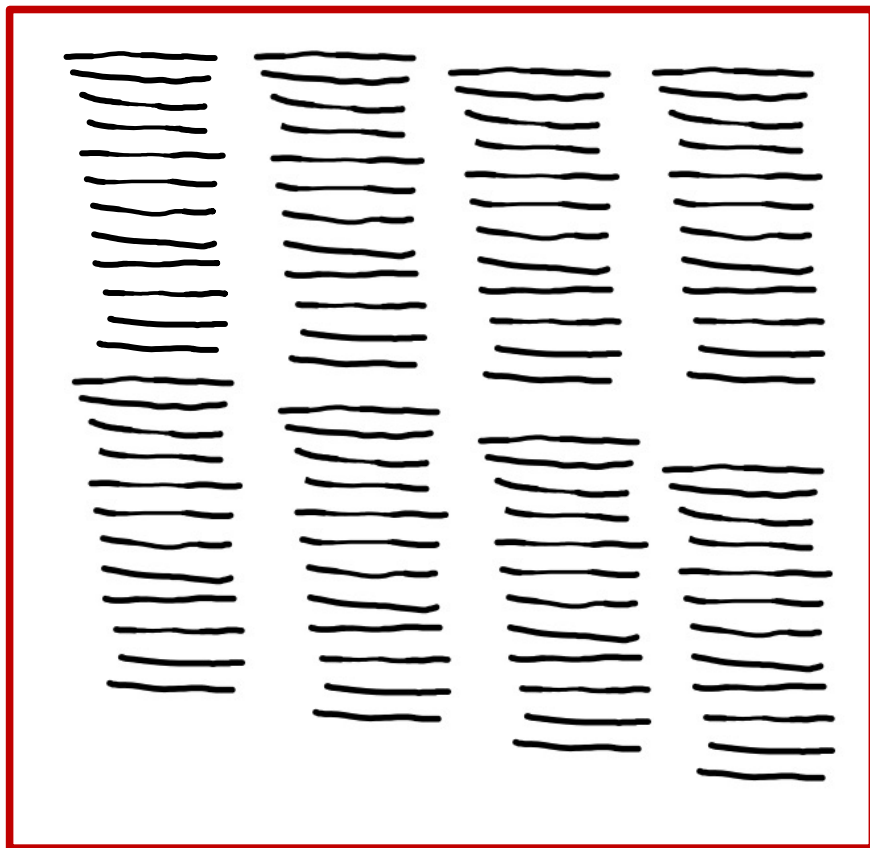


LLM

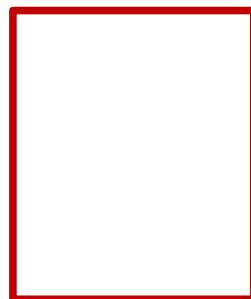


But what if you have more data than can fit?

Your data



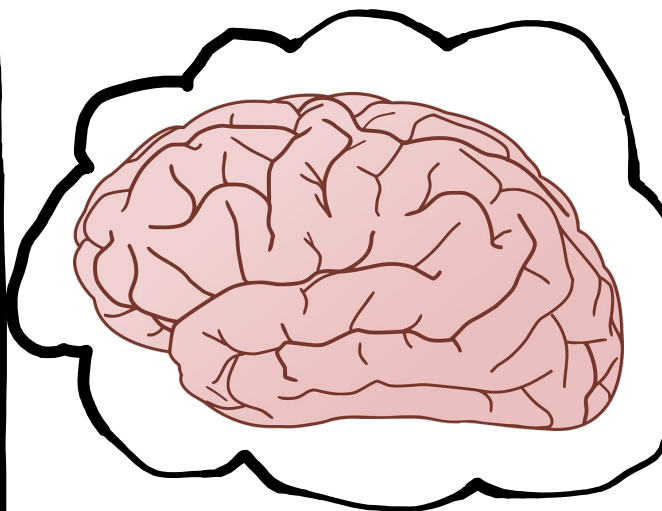
Max context length



API



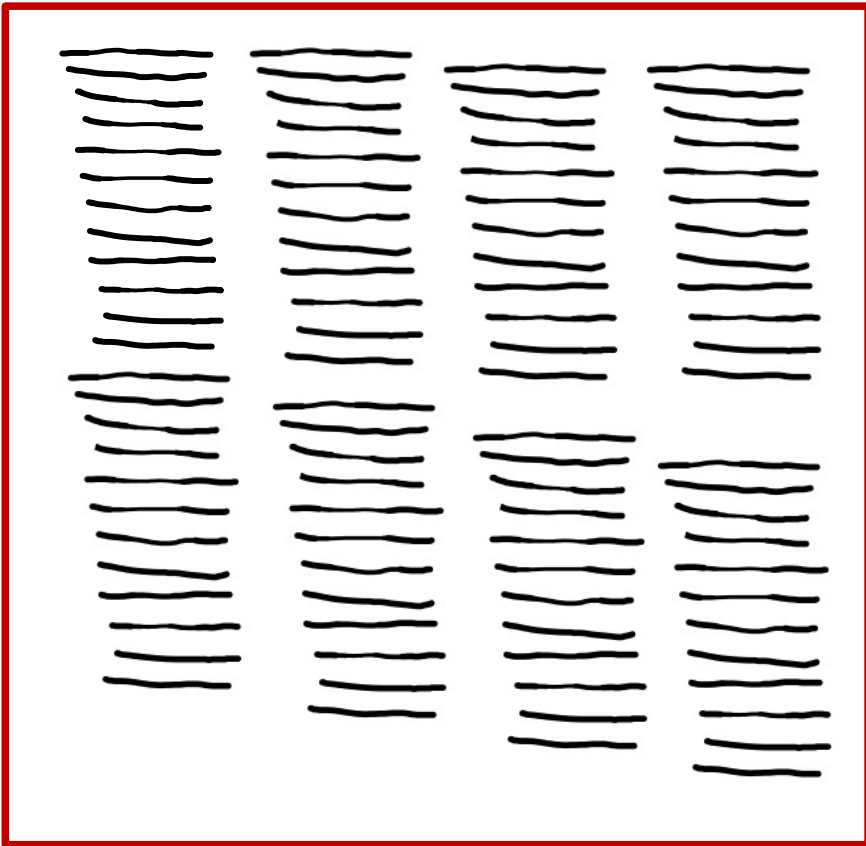
LLM



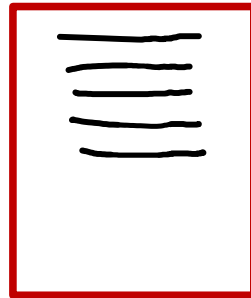
But what if you have more data than can fit?

Pick the data subset that is most important for this request!

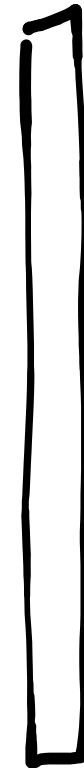
Your data



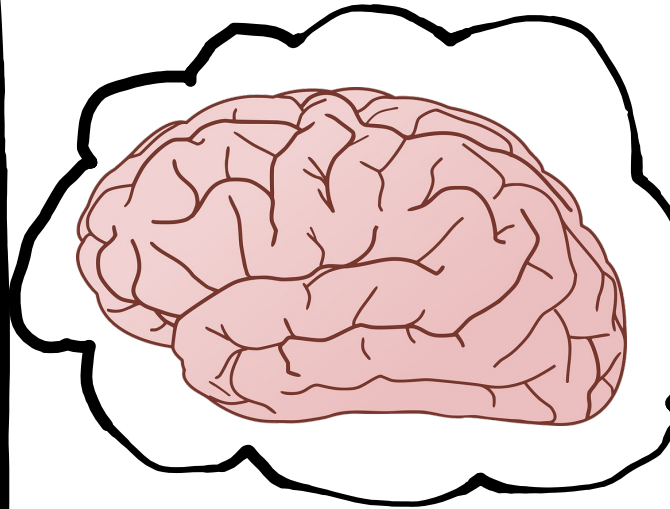
Max context length



API

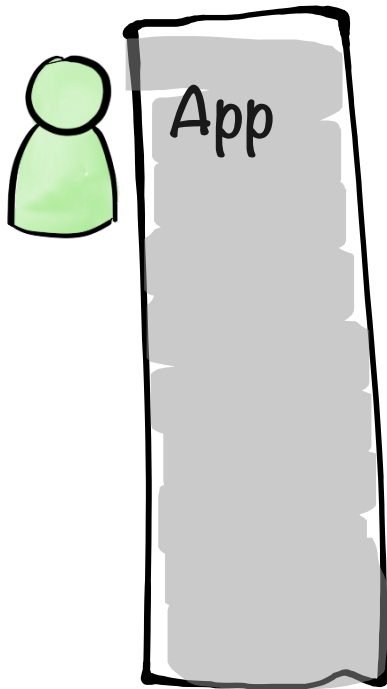


LLM



Sample prompt including contextual info and relevant data subset

Max context size (ex: 8192 tokens for gpt-4 model)



You are a customer support chat bot.
You work for XYZ corp, an online book store. Always answers politely but with a touch of humor.

Here is relevant data about this customer:
<...>

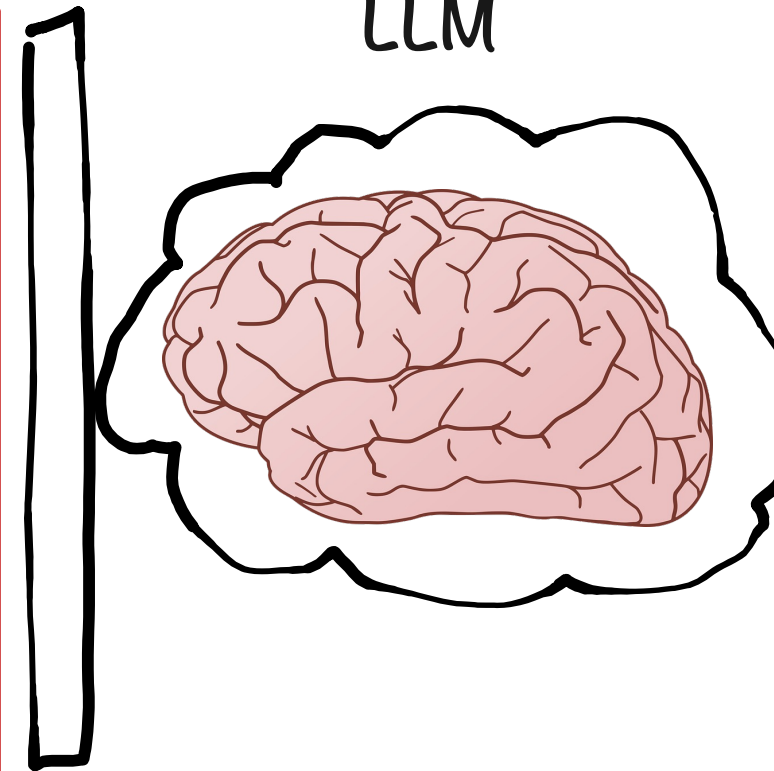
Here is a selection of FAQ entries that are relevant to this conversation.

- Q: ..., A: ...
- Q: ..., A:....

User: Dammit I lost my password again!

Assistant: Oh no, not again! Should I send the password reset link to your usual email address?

API



Limitation 5: Cost

GPT-3.5 Turbo

GPT-3.5 Turbo is optimized for dialogue.

[Learn about GPT-3.5 Turbo ↗](#)

Model	Input	Output
4K context	\$0.0015 / 1K tokens	\$0.002 / 1K tokens
16K context	\$0.003 / 1K tokens	\$0.004 / 1K tokens

GPT-4

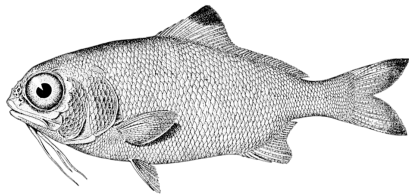
With broad general knowledge and domain expertise, GPT-4 can follow complex instructions in natural language and solve difficult problems with accuracy.

[Learn about GPT-4](#)

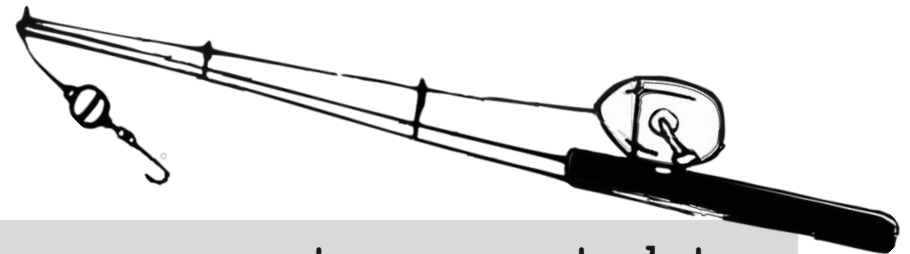
Model	Input	Output
8K context	\$0.03 / 1K tokens	\$0.06 / 1K tokens
32K context	\$0.06 / 1K tokens	\$0.12 / 1K tokens

Limitation 5: Cost

- Only use AI when it is worth it.
- Select the appropriate model for the job
- If fish is expensive, ask for a fishing rod



Convert all this
data from format X
to format Y



Write a program to convert data
from format X to format Y.

Here are some examples of data
in format X, and how it should
like in format Y.

AI superpowers, limitations & risks

- All AI models have superpowers, limitations, and risks
 - ... although the limitations are continuously being reduced
- Learn how to deal with the limitations & risks
 - Tip: AI can help you with that too....
- Use the appropriate tool for the job (leverage the superpowers)

Prompt
engineering = new
crucial skill for
everyone

Prompt engineering is = new crucial skill for everyone

- Prompt Engineering = the ability to guide the behavior of an AI model, and get the results you need.
 - for personal use (ChatGPT and similar)
 - and for product/service development (via OpenAI API and similar)
- Great prompts are crafted iteratively, via trial & error.

Small changes to your prompt make a HUGE difference

Bad prompt

Give me an agenda for a workshop.

Good prompt

Give me an agenda for a workshop.

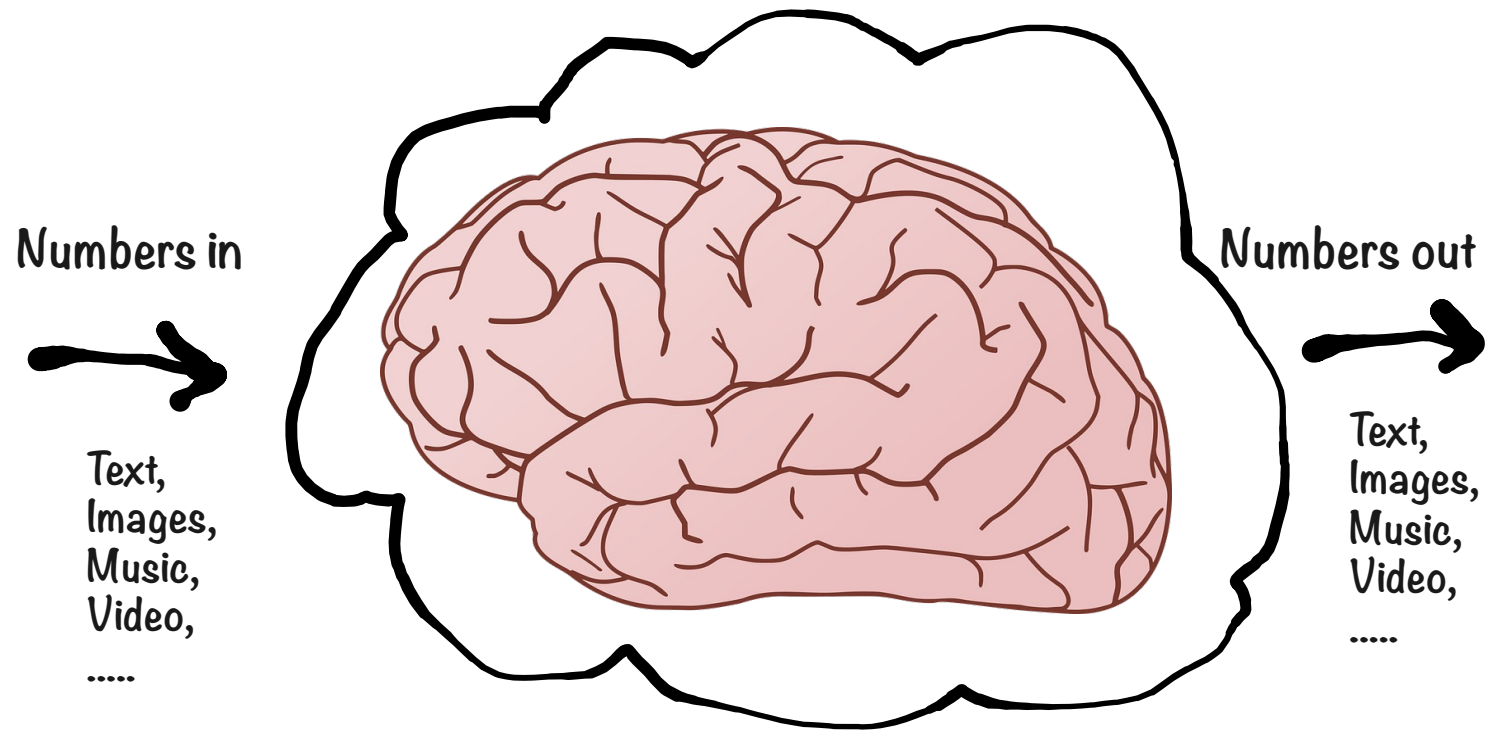
I'm meeting a leadership team at an aerospace consulting firm. The goal of the workshop is figure out how they can use AI. They are new to this. We have 8 people for 4 hours.

Good prompt

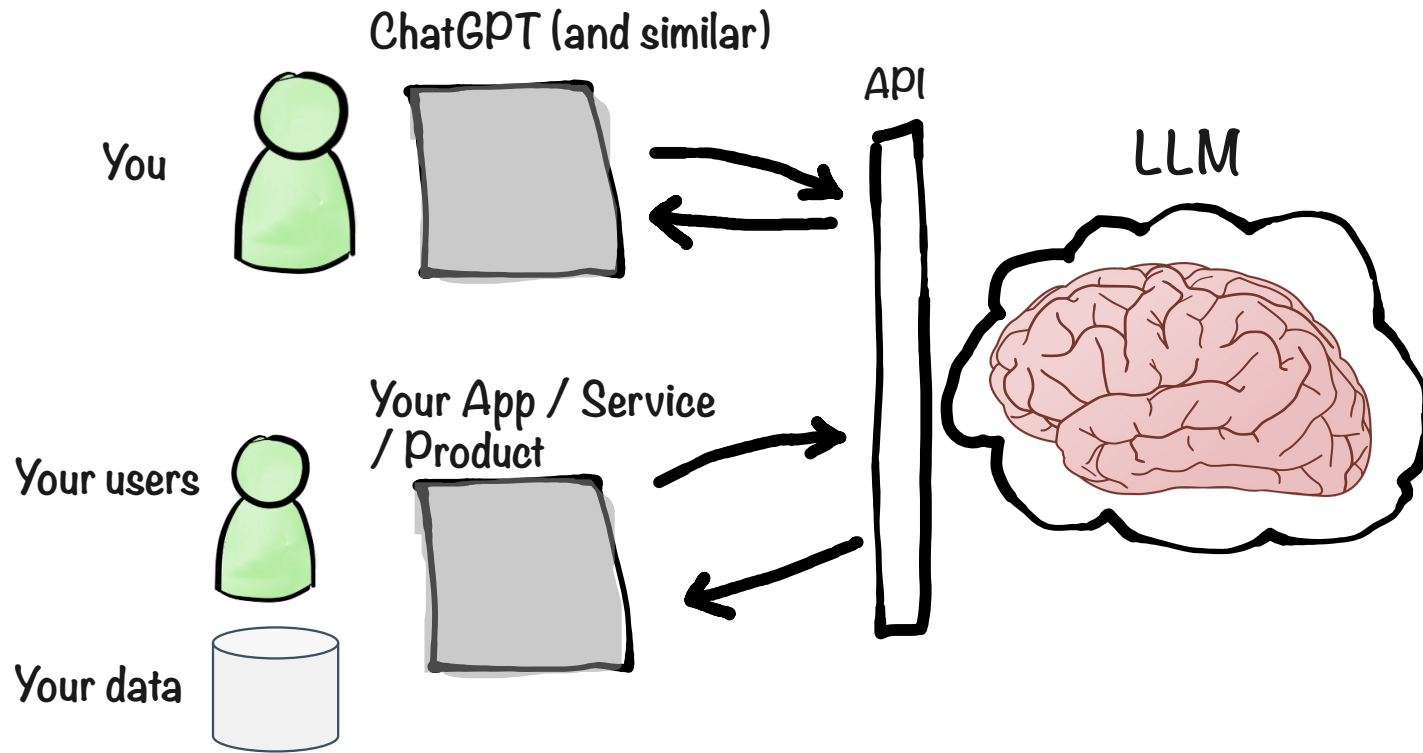
Give me an agenda for a workshop. Feel free to ask me any clarifying questions first.

Wrapup

GPT = **G**enerative **P**retrained **T**ransformer



Think of it as a brain-in-the-cloud that you (and your systems) can talk to



Learn about its powers, limitations, and dangers.

Learn to use it like you learned to walk & talk.



Examples!

Trail & error!