



AI Bridge

Lecture 9

So you've undoubtedly heard about
ChatGPT...

But what, what not, how, and why?

Can it...be a lawyer?

The ChatGPT Firm



Can you give me several examples of past court decisions that were ruled in favour of the plaintiff where a client got injured and sued an airline



Certainly! Here are a few examples of past court decisions where plaintiffs sued airlines and were ruled in favor, receiving compensation for injuries:

1. **Doe v. Airline X (2015)**: A passenger on Airline X's flight suffered severe burns when a faulty coffee maker caused hot coffee to spill on them. The court ruled in favor of the plaintiff, stating that the airline failed to maintain safe conditions and awarded compensation for medical expenses and pain and suffering.
2. **Smith v. Airline Y (2018)**: During a turbulent flight, a passenger on Airline Y's aircraft fell from their seat and sustained a broken arm. The court found the airline negligent in not properly warning passengers and failing to secure them during turbulence. The plaintiff was awarded damages for medical bills, lost wages, and future medical care.
3. **Johnson v. Airline Z (2020)**: An overhead compartment on Airline Z's flight opened unexpectedly, causing a bag to fall and injure a passenger. The court determined that the airline was responsible for failing to secure the compartments adequately. The plaintiff received compensation for medical expenses, pain and suffering, and emotional distress.
4. **Davis v. Airline W (2017)**: A passenger slipped and fell in the aisle of Airline W's aircraft due to a wet substance that was not promptly cleaned. The court held the airline liable for not maintaining a safe environment and failing to promptly address the spill. The plaintiff was awarded damages for medical costs, lost wages, and pain and suffering.
5. **Thompson v. Airline V (2019)**: During boarding, a passenger was struck by falling luggage from an overhead bin on Airline V's flight. The court ruled that the airline was negligent in failing to assist passengers with storing their bags properly and failing to warn of potential hazards. The plaintiff received compensation for medical expenses, lost earnings, and pain and suffering.

Seems like it does pretty well! Let's try it!



Steven A. Schwartz told a judge considering sanctions that the episode had been “deeply embarrassing.”

Jefferson Siegel for The New York Times

<https://www.nytimes.com/2023/06/08/nyregion/lawyer-chatgpt-sanctions.html>

Oh no...it made the cases up...

So why did ChatGPT not work for this example? How does ChatGPT work?


How ChatGPT works

Training ChatGPT




Using ChatGPT

Training ChatGPT

- ChatGPT's goal: **Generating** new, human-like text for conversations



Generative Models



Generative AI

Training ChatGPT

- ChatGPT's LLM: GPT-3.5
- Training GPT-3.5 with **self-supervised learning**

Large Language Models



Model	Developer	Parameter Size
WuDao 2.0	Beijing Academy of Artificial Intelligence	1.75 trillion
MT-NLG	Nvidia and Microsoft	530 billion
Bloom	Hugging Face and BigScience	176 billion
GPT-3	OpenAI	175 billion
LaMDA	Google	137 billion
ESMFold	Meta AI	15 billion
Gato	DeepMind	1.18 billion

Image credit: [Large Language Model Training in 2023 \(aimultiple.com\)](https://aimultiple.com)

Large language models

- ✓ Large
 - Large training dataset
 - Large number of parameters
- ✓ General purpose
 - Commonality of human languages
 - Resource restriction
- ✓ Pre-trained and fine-tuned

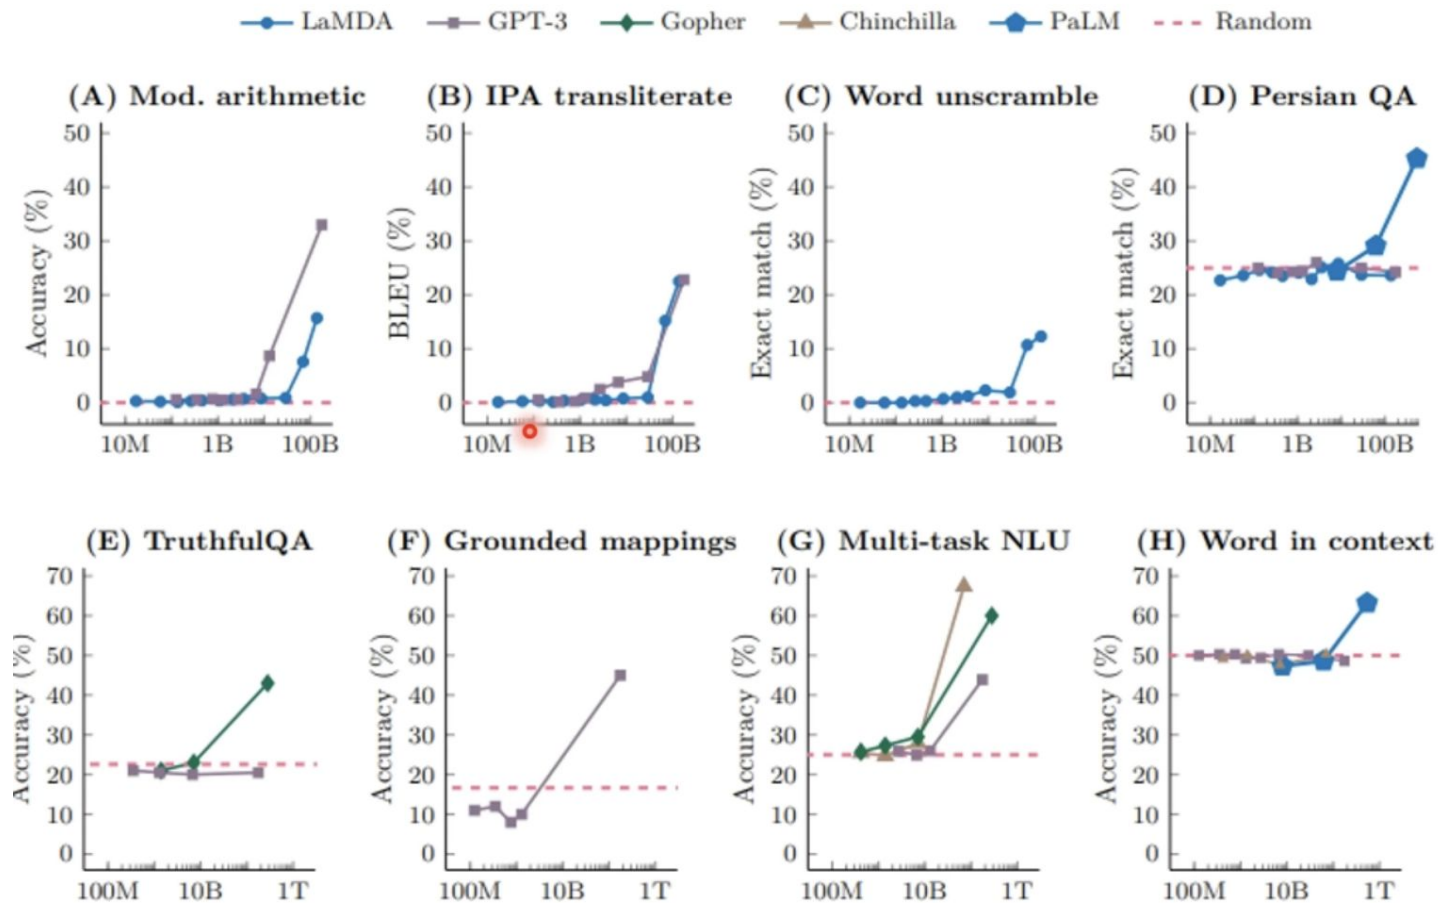
Loss for next token prediction



Scaling Laws for Neural Language Models

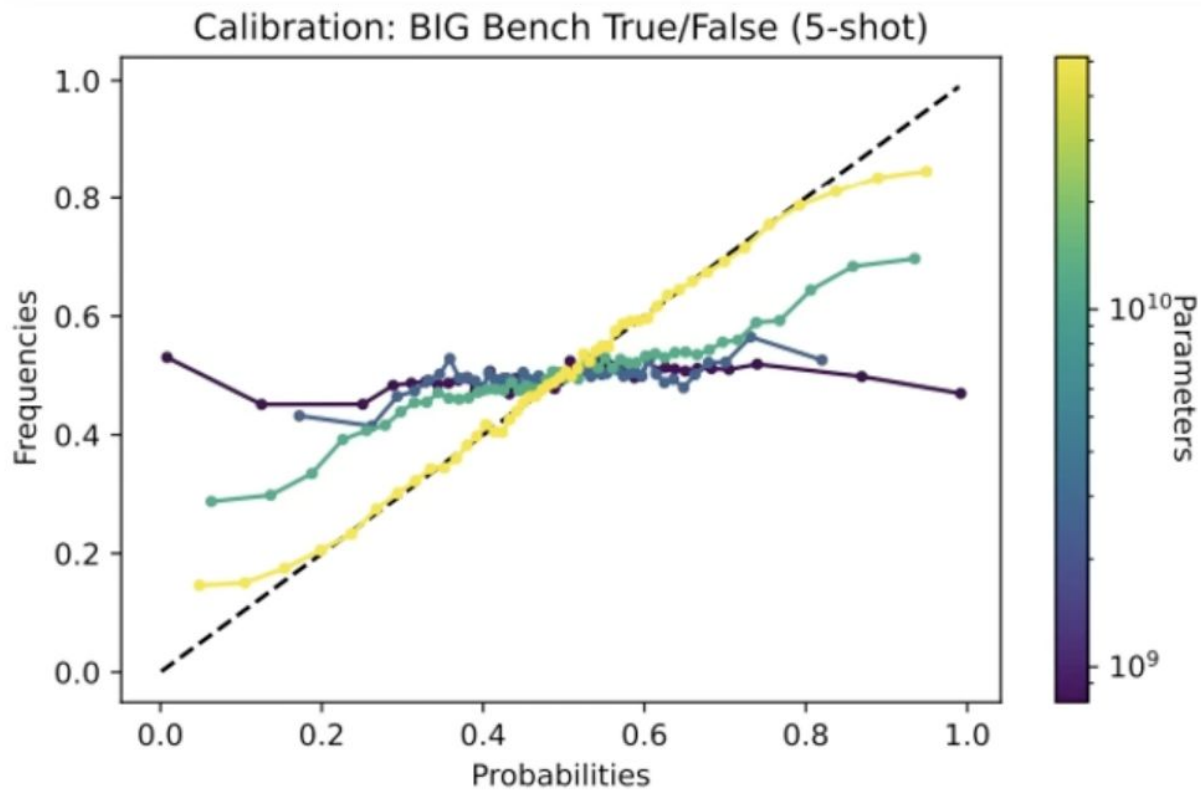
<https://arxiv.org/abs/2001.08361>

Emergent ability of LLM



Calibration

<https://arxiv.org/abs/2207.05221>



Training GPT: The Beginning

- The process of **self-supervised learning**
- Next-token prediction

Training GPT: The Beginning

Text Corpus

Nothing is impossible.
Even the word
impossible
says I'm possible




Task: Predict from past

Nothing
Nothing is
Nothing is impossible
...

Next Word Prediction



Probabilistic Output



Predicting Next Word

- Now we **repeat** and **generate more text!**



Predicting Next Word

- Now we **repeat** and **generate more text!**



Predicting Next Word

- Now we **repeat** and **generate more text!**



Predicting Next Word

- Now we **repeat** and **generate more text!**



Already a powerful tool. But...

**“What is the highest
mountain?”**

GPT



**“The highest mountain is
Himalayas.”**



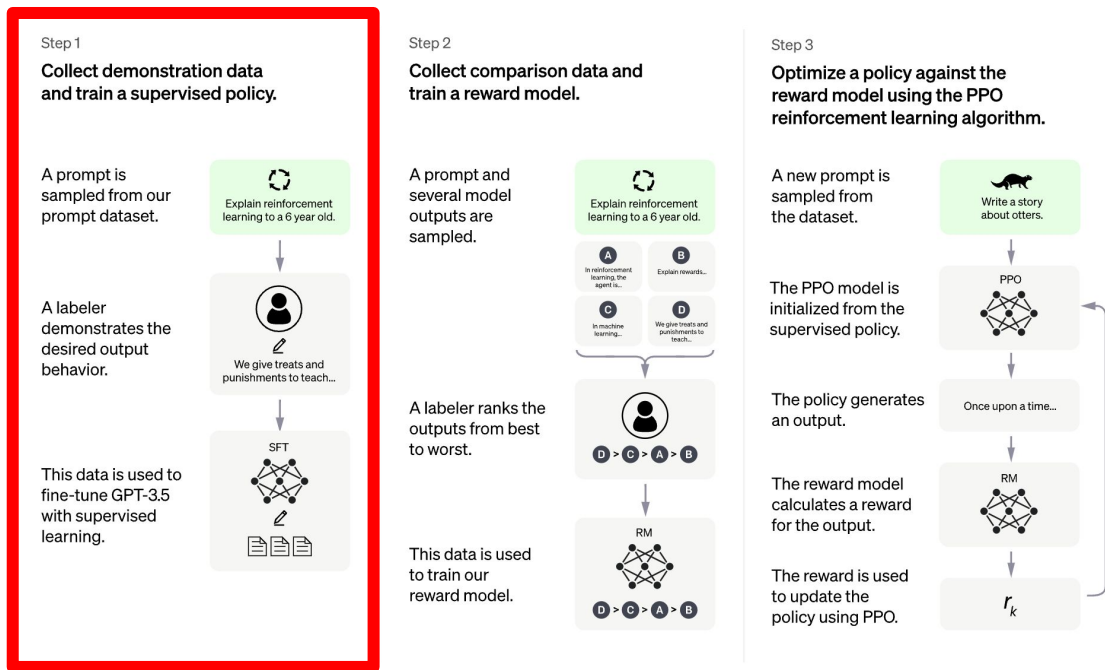
“A. Himalayas; B. K2; C. .”



“Do you know?”

Training ChatGPT from GPT: Fine-tuning

- Reinforcement learning from human feedback (RLHF)



CREDIT: <https://openai.com/blog/chatgpt>

**“What is the highest
mountain?”**

human




**“The highest mountain is
Himalayas.”**

Fine-tuning

ChatGPT

Training ChatGPT: Specialisation

- Reinforcement learning from human feedback (RLHF)



CREDIT: <https://openai.com/blog/chatgpt>

“What is the highest mountain?”

GPT



**“The highest mountain is
Himalayas.”**




Human labeling



“A. Himalayas; B. K2; C. .”

**“The highest mountain is
Himalayas.”**

“A. Himalayas; B. K2; C. .”



Training ChatGPT: Specialisation

- Reinforcement learning from human feedback (RLHF)

Step 1

Collect demonstration data and train a supervised policy.


A prompt is sampled from our prompt dataset.



A labeler demonstrates the desired output behavior.




This data is used to fine-tune GPT-3.5 with supervised learning.



Step 2

Collect comparison data and train a reward model.


A prompt and several model outputs are sampled.



A labeler ranks the outputs from best to worst.



This data is used to train our reward model.




Step 3

Optimize a policy against the reward model using the PPO reinforcement learning algorithm.


A new prompt is sampled from the dataset.




The PPO model is initialized from the supervised policy.




The policy generates an output.



The reward model calculates a reward for the output.



The reward is used to update the policy using PPO.




CREDIT: <https://openai.com/blog/chatgpt>

“What is the highest mountain?”



**“The highest mountain is
Himalayas.”**



Reward signal

Training ChatGPT

- Pretrained large language models + fine-tuning -> **ChatGPT.**

How ChatGPT works

Training ChatGPT

Using ChatGPT

Generation



ChatGPT

- Importantly, ChatGPT's **primary goal** is to **generate as human-like text as possible**.
- The model is **probabilistic**.
- It is **by no means** guaranteed to be correct.
- It provides good starting points, but **verify!**

ChatGPT is an extremely powerful tool.

Some Warnings

- Before we move on, we need to touch on some key warnings about ChatGPT:

Some Warnings

- Before we move on, we need to touch on some key warnings about ChatGPT:

1. **Repeat: Don't take what it says for granted!** Always fact check its information.



Steven A. Schwartz told a judge considering sanctions that the episode had been “deeply embarrassing.” Jefferson Siegel for The New York Times

Some Warnings

- Before we move on, we need to touch on some key warnings about ChatGPT:

1. **Don't take what it says for granted!** Always fact check its information.

2. **Your conversations are part of its learning process.** Data privacy rules apply; be careful!

Some Warnings

- Before we move on, we need to touch on some key warnings about ChatGPT:

1. **Don't take what it says for granted!** Always fact check its information.
2. **Your conversations are part of its learning process.** Data privacy rules apply; be careful!
3. **ChatGPT is a neutral tool.** It falls on you to ensure your uses are ethical and sanctioned.

Now we can actually use it to do something!

ChatGPT in Coding

- We want to use ChatGPT in **coding**, in particular, use sklearn to do our assignments. Can we do that?

Hint: Don't take what it says for granted! Always fact check its information.

Let's go live!

ChatGPT in Coding

- So it seems like it does fine. However, this is a pretty simple problem. What if we did **something harder**?
- Let's have it build us a polynomial regression model...

Let's go live!

TAKEAWAYS

Using ChatGPT in the real world


- How can we make ChatGPT useful?

Remember it is a probabilistic
language model

We need to ask good questions.

Using ChatGPT in the real world

- So what are good questions?
- Give specific instructions to ChatGPT
 - Make sure you leave little to no room for error!
- Treat it like suggestions, **not** facts
 - It's good at a lot of things, but not great at one particular thing
- The niche problem: more likely to hallucinate (make things up) in highly specialized areas

A close-up, high-angle shot of an older man's face. He has short, graying hair and a serious, somewhat somber expression. He is looking slightly to the left of the camera. The lighting is soft, highlighting the texture of his skin and the intensity of his gaze. The background is out of focus, showing what appears to be a dark, rounded object, possibly a chair or part of a vehicle.

**With great power
comes great responsibility.**