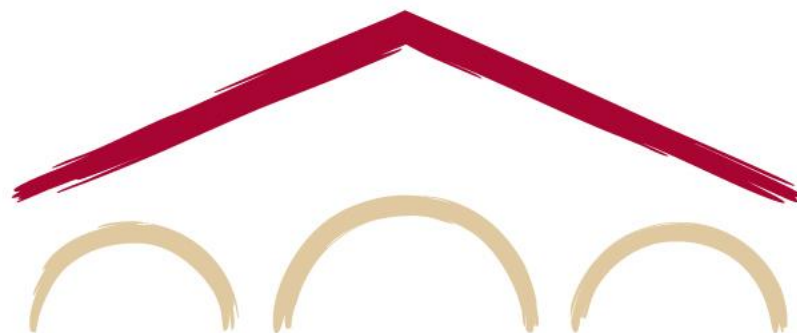


Large Language Models in 2025: How Much Understanding and Intelligence?



Christopher Manning

Founder, Stanford NLP Group ✿ Associate Director, Stanford HAI

@chrmanning ✿ @stanfordnlp

First Workshop of a Public AI Assistant to World Wide Knowledge (WWK)

There have been 60+ years of AI research ...

See our short video on YouTube: **AI at Stanford 1962–2022**



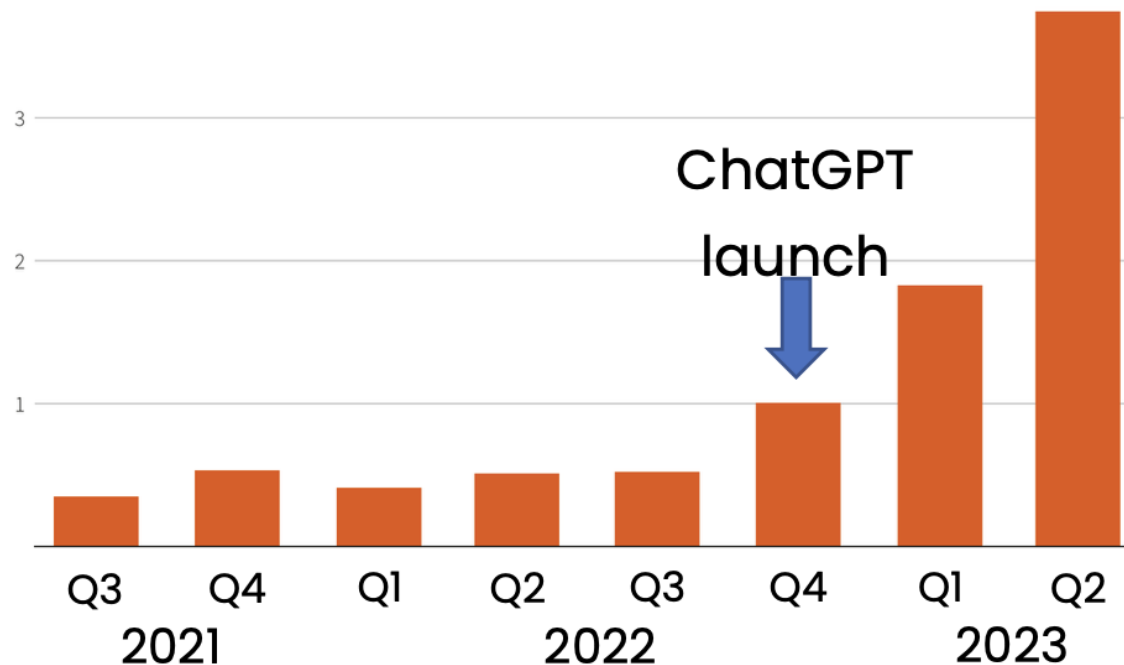
The world changed with OpenAI's ChatGPT

November 2022



The coming future of generative AI

Average number of 'AI' mentions per S&P 500 analyst call



Source: Reuters

Generative AI could

- Add \$2.6–\$4.4 trillion annually to the economy¹
- Raise global GDP by 7% in the next 10 years²

¹ [McKinsey](#); UK total GDP 2021: \$3.1 trillion

² [Goldman Sachs](#)

ChatGPT 4o

Write a sonnet praising Wikipedia in which each line begins with the letter W

Wikipedia: A Sonnet of Wonder

With wisdom vast, it spans the world so wide,
Wielding the wealth of knowledge, free to share.

Words weave through time where curious minds confide,
Whispering truths with patience, just and fair.

Wondrous it stands, a beacon shining bright,
Warding off darkness, doubt, and false pretense.

Wrought by the hands of many, day and night,
Witness to learning's boundless recompense.

Who dares to claim the heights of thought alone?

Whence comes this trove but hearts in concert joined?

While scholars build on stones that others've known,

Waking the past, their voices intertwine.

Wisdom and will sustain this noble dream,


Wikipedia—our light, our endless stream.

The starting point (GPT1-3):


“Large language models”(LLMs) trained by Mad Libs

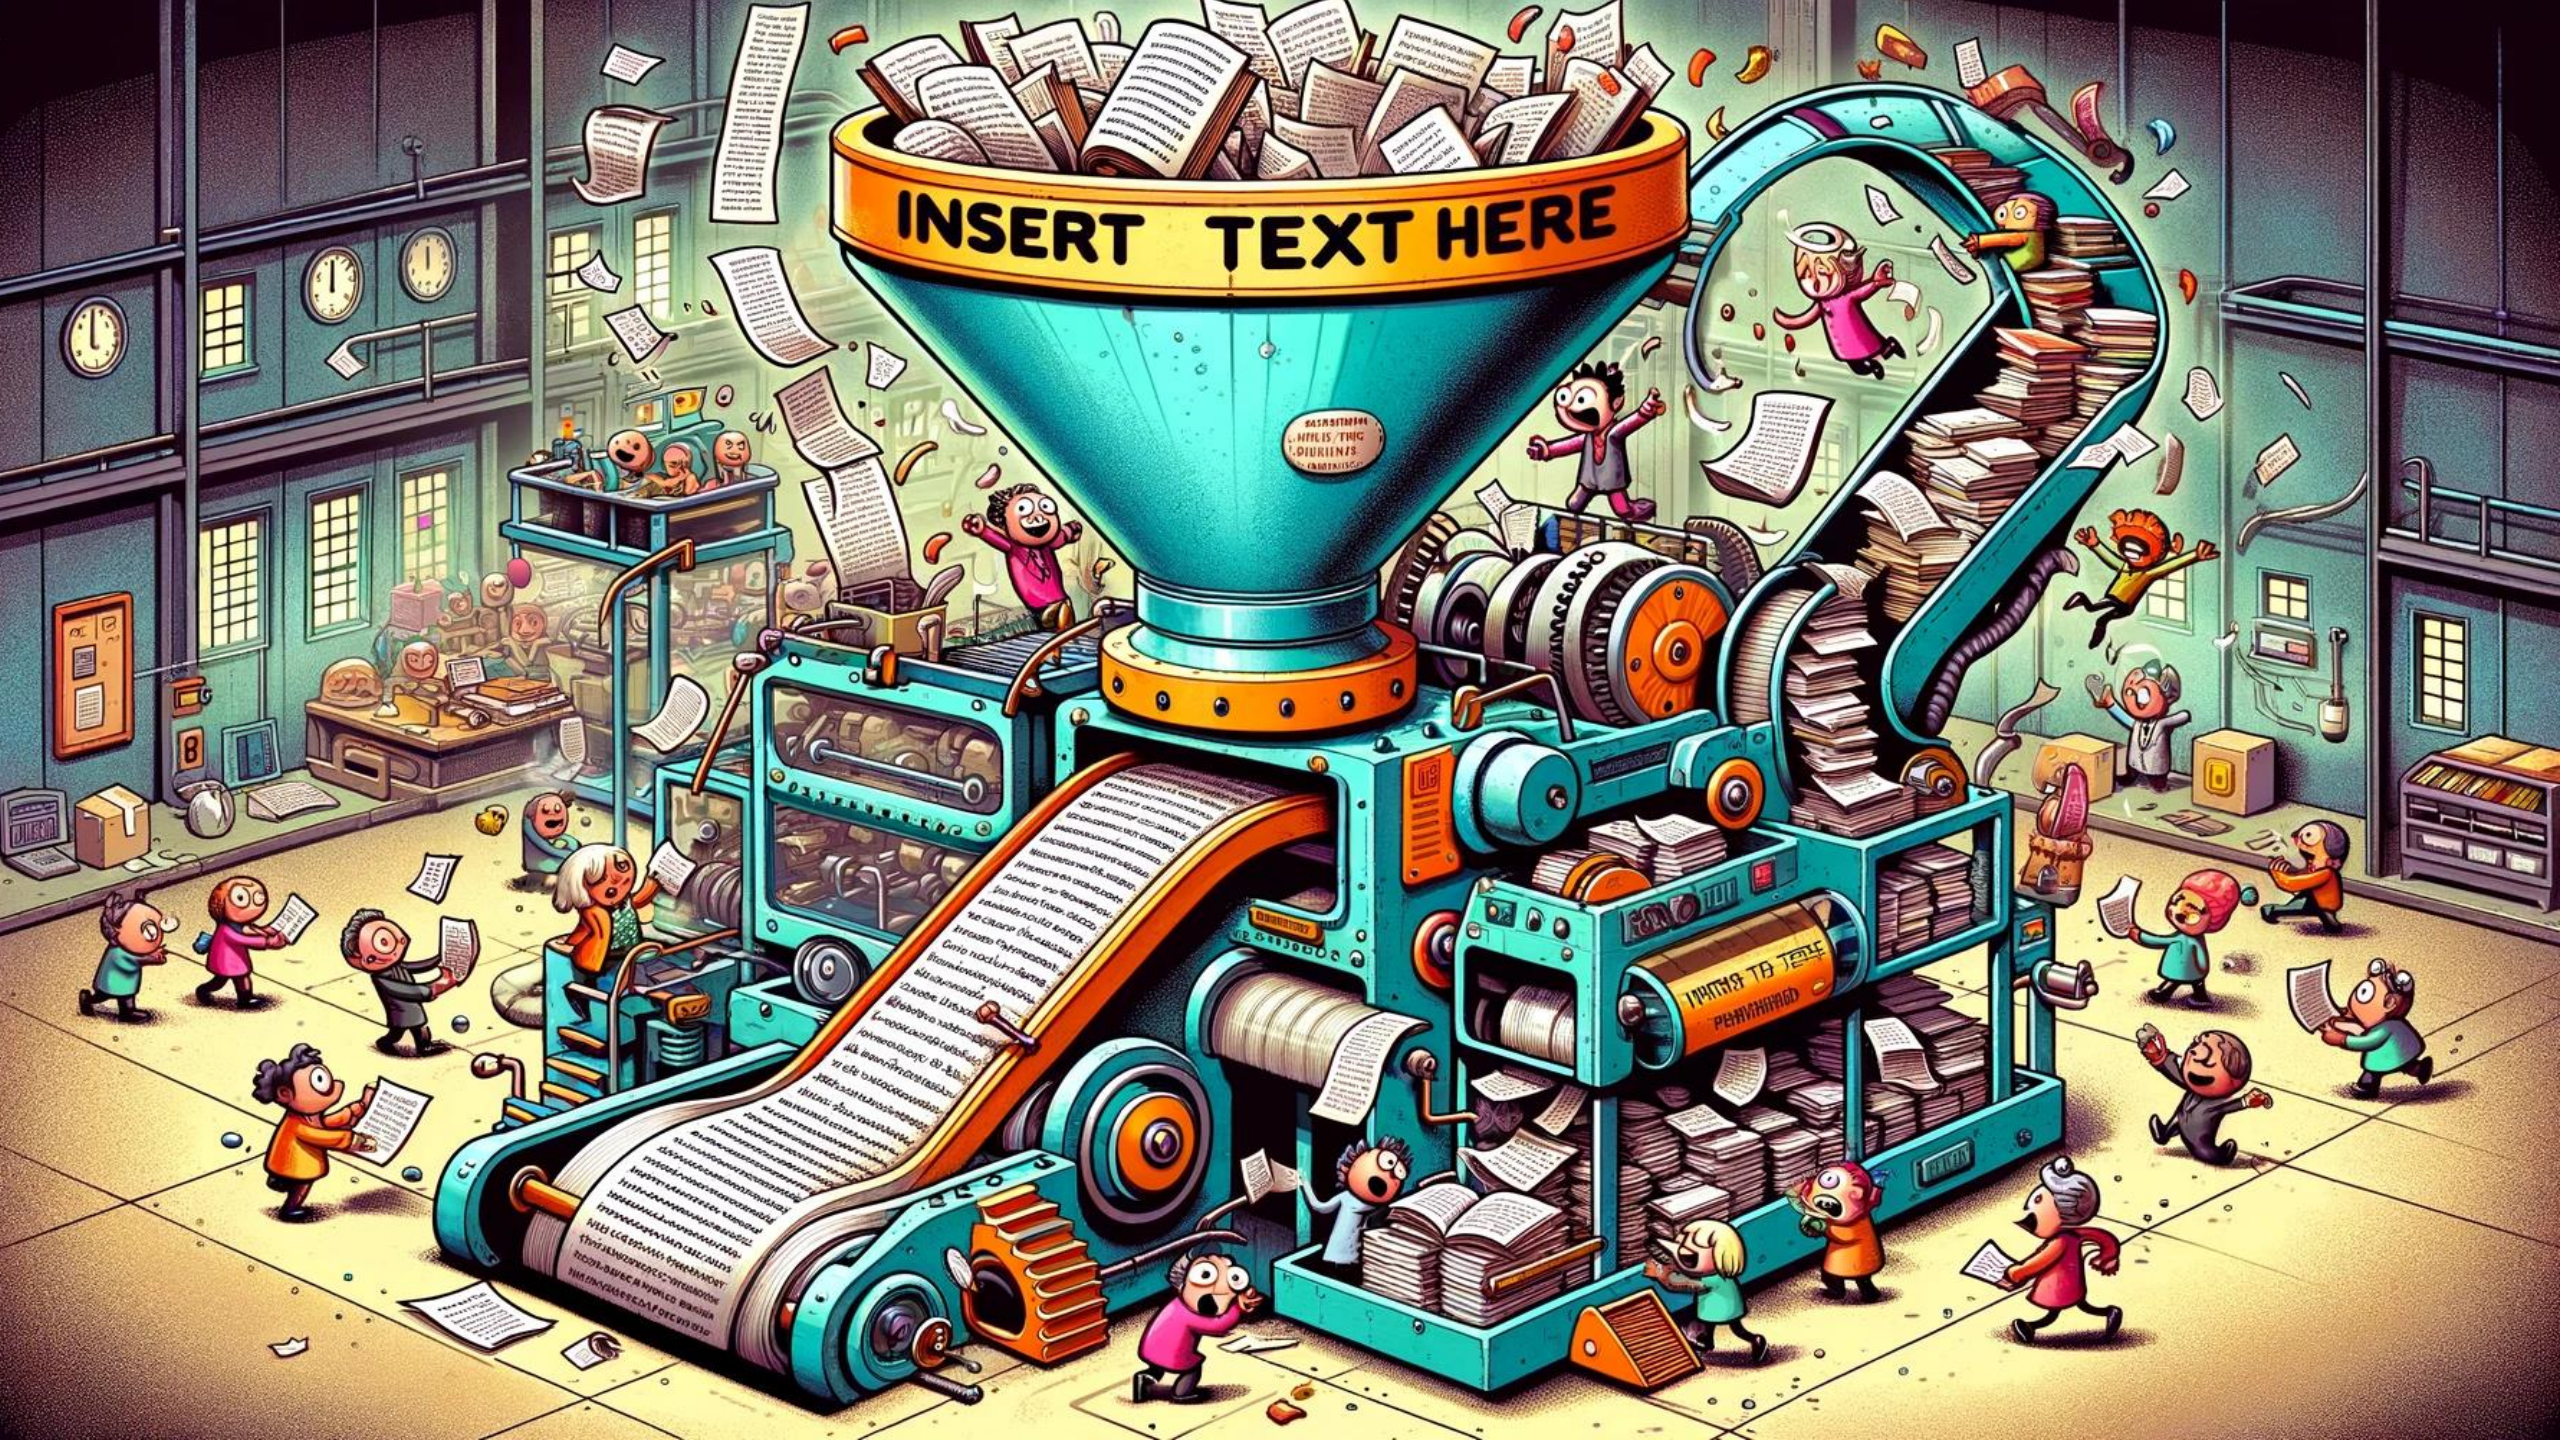
“Unsupervised” or “self-supervised” learning, using existing pieces of text

Ford plans to _____

What Happens when a Unicorn Poops? 

Unicorns aren't like other _____ ;
(plural noun)
they're _____. They look like
(adjective)
_____, with _____ for feet
(plural noun; animals) (plural noun)
and a _____ mane of hair. But unicorns
(adjective)
are _____ and have a _____ on
(color) (adjective) (noun)
their heads. Some _____ don't
(plural noun)
believe unicorns are _____ but I
(adjective)
believe in them. I would love to _____ a
(verb)
unicorn to faraway _____. One thing
(plural noun)
I've always _____ about is whether
(verb -ed)
unicorns _____ rainbows, or is their
(verb)
_____ like any other animal's?
(noun) (adjective)





INSERT TEXT HERE

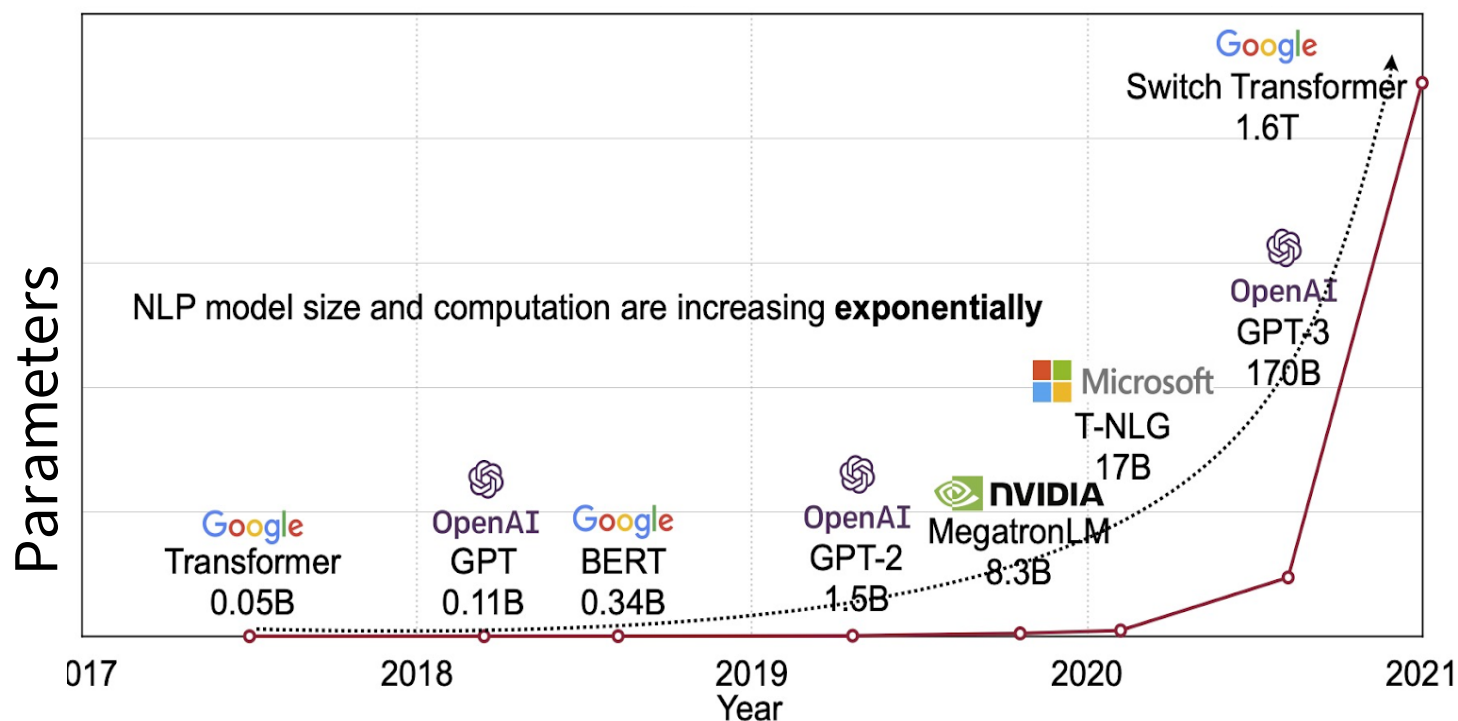
САНКТ-ПЕТЕРБУРГ
И. АНТИС / ТИЖ
И. ДИУРИНС
И. ПАВЛОВИЧ

ИЗДАНИЕ
ПЕРВАЯ

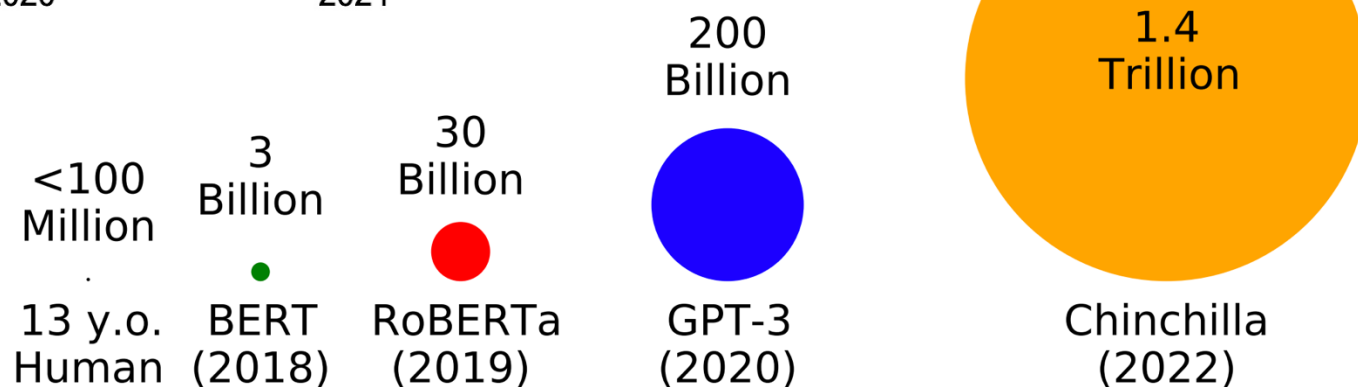
САНКТ-ПЕТЕРБУРГ
И. АНТИС / ТИЖ
И. ДИУРИНС
И. ПАВЛОВИЧ

ИЗДАНИЕ
ПЕРВАЯ

The first era of LLMs: Scaling to larger and larger models with more compute, data, and dollars ...



tokens seen during training



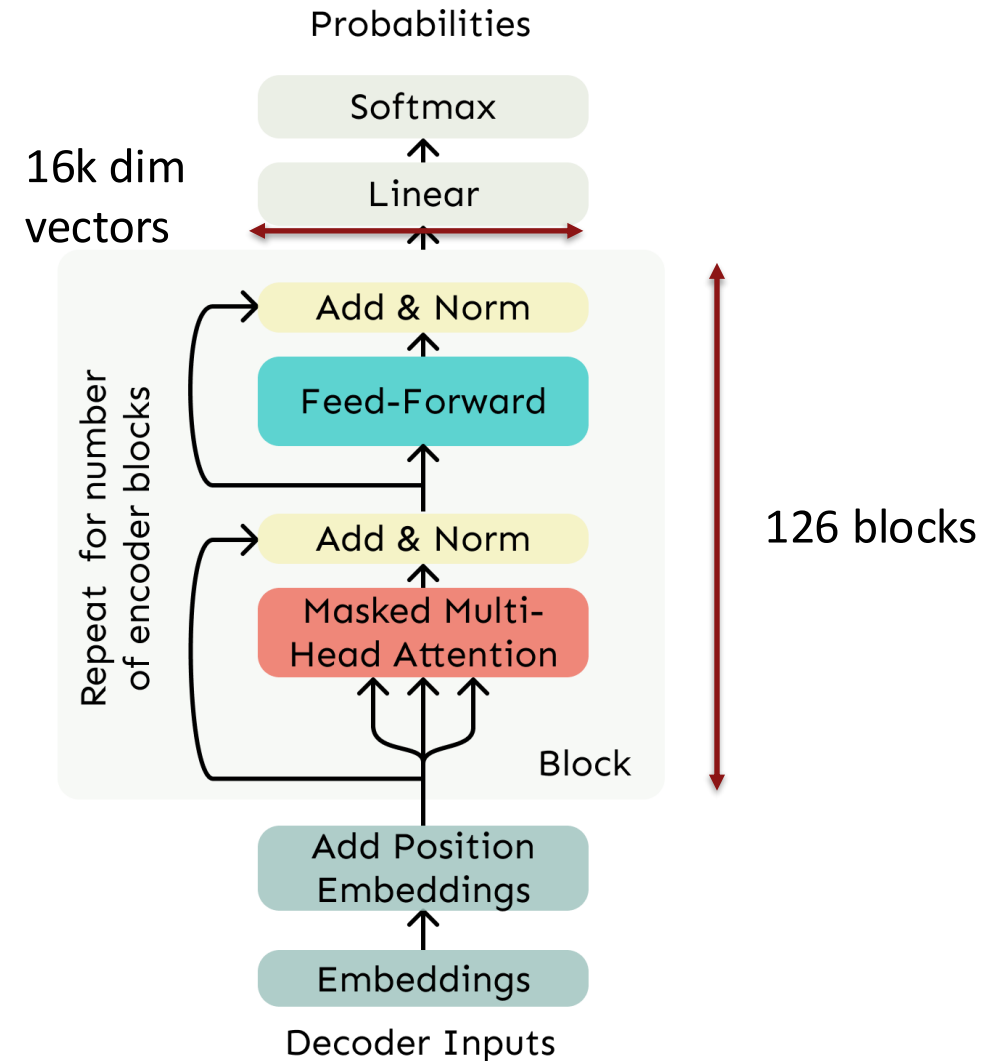
Produces a qualitative shift in what emerges

It no longer appears to be statistical prediction but rather
a machine that understands language!

And the world!

A transformer LLM in practice: Meta's LLaMA 3.1 405B

- **Total training data** – 15.6 trillion tokens
- **Architecture** – 405 billion parameters
 - Parameters = floating point numbers in big neural net
 - 126 transformer blocks
 - Model dimension: 16,384 dim vectors
 - Attention heads: 128
- **Training compute** – $3.8e25$ FLOPs
 - 1.1 million H100-days (\$~25 M)
 - ~ 2 months on 16000 Nvidia H100 GPUs



Getting LLMs to do what you actually want

GPT-1/2/3 rewrote how researchers think about modeling language

But they didn't become a **household name** like ChatGPT in late 2022

Why?

Ask GPT-3 who the US President is:

Who is the US President? Who is the Vice President? What are two levels of government? What is the difference between the national and state governments? What are the three levels of government? What are the 3 branches of government and their responsibilities? What are the 4 forms of government?



Well, no, that's not really what I wanted

Era 2: Getting LLMs to do what you actually want

Models that understand what people want are profoundly more useful

Modern “instruction-tuned” models learn this behavior through a second phase of post-training using:

- Reinforcement Learning from Human Feedback (RLHF)
 - A large amount of annotated interaction data
 - Learning from human preference data pairs to give desired answers
- The resulting model obeys commands and answers questions

Getting information from a huge document: AGL

Please read the 2022 AGL annual report here:

<https://www.agl.com.au/content/dam/digital/agl/documents/about-agl/investors/2022/220819-agl-energy-annual-report-2022.pdf>

What does the AGL report say about the company's alignment with IPCC on Climate Action?

The AGL 2022 annual report states that the company is committed to "aligning our business with the goals of the Paris Agreement and the IPCC's 1.5°C pathway." AGL has set a target of net zero emissions by 2050, and is currently working on a decarbonisation pathway to achieve this goal. The company's Climate Transition Action Plan includes a commitment to exit coal-fired power generation by FY35. This will bring AGL's Scope 1 and 2 emissions to net zero, which is significant as AGL is currently Australia's largest corporate emitter.

Overall, the AGL 2022 annual report shows that the company is committed to taking action on climate change. AGL has set ambitious targets for reducing its emissions, and is working with stakeholders to achieve these goals. The company's engagement with the Climate Action 100+ initiative is also a positive sign, as it shows that AGL is taking the concerns of investors seriously.

Here are some specific quotes from the report that relate to AGL's alignment with IPCC on climate action:

- "We are committed to aligning our business with the goals of the Paris Agreement and the IPCC's 1.5°C pathway."
- "We have set a target of net zero emissions by 2050, and are currently working on a decarbonisation pathway to achieve this goal."
- "Our Climate Transition Action Plan includes a commitment to exit coal-fired power generation by FY35."
- "We are actively engaged with the Climate Action 100+ initiative."
- "We have a score of 25% on the Climate Action 100+ engagement scorecard."

I hope this helps! Let me know if you have any other questions.

Getting information from a huge document: AGL

Are there any comments that raise concerns as to whether AGL is fully committed to a net zero carbon emissions future?

Yes, there are a few comments in the AGL 2022 annual report that could raise concerns about whether the company is fully committed to a net zero carbon emissions future. These include:

- The company's decision to delay the closure of its Loy Yang A coal-fired power station by two years. This decision was made in response to concerns about the reliability of the Australian electricity grid, but it has been criticized by some as a setback for Australia's climate goals.
- The company's plans to invest in new gas-fired power stations. Gas is a cleaner-burning fuel than coal, but it is still a fossil fuel and its combustion produces greenhouse gases. Some critics have argued that AGL's investment in gas-fired power stations is inconsistent with its commitment to net zero emissions.

- The company's lack of transparency about its decarbonisation pathway. AGL has not yet released a detailed plan for how it will achieve net zero emissions by 2050. This lack of transparency has led some to question whether the company is serious about its climate commitments.

Overall, the AGL 2022 annual report shows that the company is committed to taking action on climate change. However, there are some concerns that could raise questions about whether AGL is fully committed to a net zero carbon emissions future. The company needs to be more transparent about its plans and to take more concrete steps to reduce its emissions if it wants to allay these concerns.

Retrieval Augmented Generation (RAG)

1. Give an LLM access to many external data sources
2. LLM generates a search query based on a question/request
3. Search (usually neural) is used to find relevant documents
4. The model reads documents to find relevant passages
5. Those passages are added to the input context (prompt) for generating a final answer

The screenshot displays the ChatPDF interface. On the left, a document titled "Accenture-Su..." is open, showing a search result for "95% of global executives agree advances in generative AI signify a new era of enterprise intelligence." The document text includes sections on "Supply chain roles potentially impacted" and "Automating". On the right, a chat window titled "Chat" contains a blue input box with the question: "What are some roles in supply chain that will be changed by AI?". Below the input box, a light blue response box contains the answer: "Generative AI will have a profound impact on supply chain roles. According to 2, some of the roles that will be potentially impacted include data scientists, call-center staff, planners, buyers, sustainability and compliance officers, risk managers, category managers, supply chain planners, field engineers, plant workers, product engineers, retailers, and materials planners. 2". At the bottom of the chat window is a blue input box with the text "Ask any question..." and a send button.

Chatting with an Accenture report using ChatPDF

Models don't understand well what they read!

Legal RAG “hallucinations” from Magesh et al. (2024)

1. “[W]hen the lender receives the collateral that secured the fraudulent loan, this is **considered a return of ‘any part’ of the loan** money . . . This was established in the Supreme Court case *Robers v. U.S.*”

Robers held precisely the opposite: “the phrase ‘any part of the property . . . returned’ refers to the property the banks lost . . . and not to the collateral.” 572 U.S. 639, 642 (2014).

2. “*D.M. v. State* . . . has been overruled by *Davis v. State*. Also, the case ***Millbrook v. U.S.* was reversed** by the same case at a later date.”

Millbrook v. United States is a **U.S. Supreme Court decision that controls on federal questions.** 569 U.S. 50 (2013). **The Nebraska Supreme Court did not cite, much less ‘reverse,’ it in *Davis v. State*.** 297 Neb. 955 (2017).

3. “[A] **paragraph from the Federal Rules of Bankruptcy Procedure [FRBP]**, Rule 4007 states that the deadlines set by bankruptcy rules governing the filing of dischargeability complaints are **jurisdictional.**”

There is no such paragraph in the FRBP and 4007’s deadlines are unlikely to be jurisdictional under *Kontrick v. Ryan*, 540 U.S. 443, 447–48, 448 n.3 (2004).

4. “The Supreme Court of Alaska . . . noted . . . in the absence of physical injury, the losses for claimed emotional damage [from a misdiagnosis] **should not be permitted** for legal and public policy reasons (*Chizmar v. Mackie*, 896 P.2d 196 (1995)).”

The Supreme Court of Alaska in fact **reversed this trial court holding** (that losses are not permitted on public policy grounds) and **allowed** the question to proceed to the jury. 896 P.2d at 204-05.

Agentic AI: Tool use for connecting to the world

Deliver me a burger with blue cheese and fries!

It's on its way!

System output (fed to some API): `takeout-order({hamburger:{num: 1, options: {blue-cheese}}, french-fries: {size: regular}}, "234 Waverley St, Palo Alto", cust-172)`

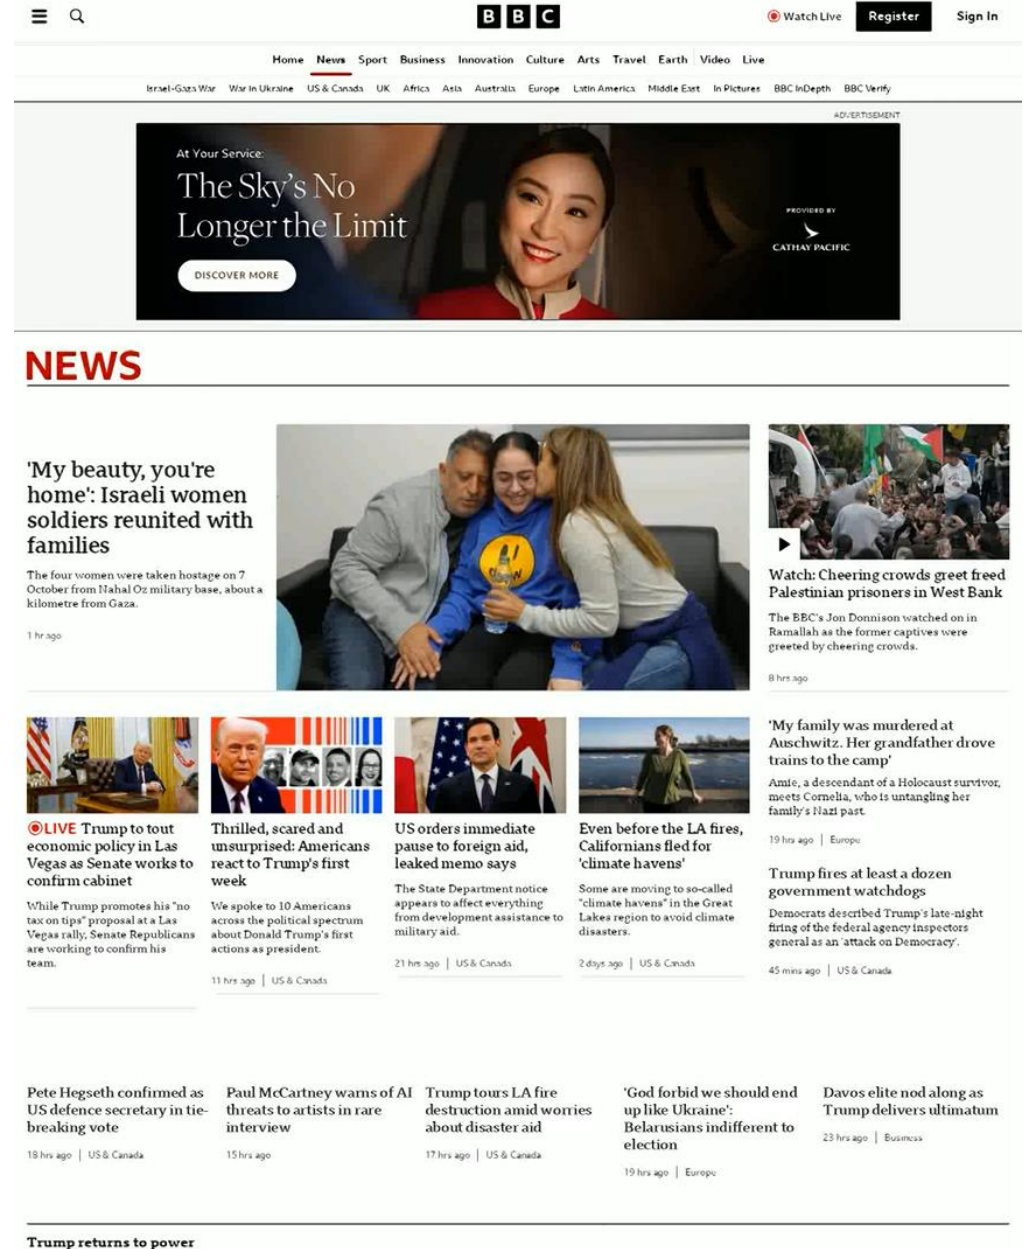
- For an LLM to be able to perform requests like this, it needs to:
 1. Be connected to various APIs or the web, so it can actually do things
 2. Have contextual knowledge (of the speaker, their accounts, etc.)

Example web agent

What does the current headline in BBC's Natural Wonders talk about?

Murty et al. 2024. NetNav
<https://www.nnetnav.dev/>

The key issue for agents is getting high enough accuracy for effective automation!



The screenshot shows the BBC News website. At the top, there is a navigation bar with the BBC logo, a search icon, and links for 'Watch Live', 'Register', and 'Sign In'. Below the navigation bar, there is a horizontal menu with categories like Home, News, Sport, Business, Innovation, Culture, Arts, Travel, Earth, Video, and Live. A secondary menu lists various regions: Israel-Gaza War, War in Ukraine, US & Canada, UK, Africa, Asia, Australia, Europe, Latin America, Middle East, In Pictures, BBC inDepth, and BBC Verify. The main content area features a large advertisement for Cathay Pacific with the headline 'The Sky's No Longer the Limit' and a 'DISCOVER MORE' button. Below the advertisement, the 'NEWS' section is displayed. The first news item is a video titled 'My beauty, you're home': Israeli women soldiers reunited with families', with a sub-headline 'The four women were taken hostage on 7 October from Nahal Oz military base, about a kilometre from Gaza.' and a timestamp of '1 hr ago'. To the right of this article is a video player showing a crowd of people. Below the main article, there are four smaller news items: 1. 'LIVE Trump to tout economic policy in Las Vegas as Senate works to confirm cabinet' (11 hrs ago | US & Canada). 2. 'Thrilled, scared and unsurprised: Americans react to Trump's first week' (11 hrs ago | US & Canada). 3. 'US orders immediate pause to foreign aid, leaked memo says' (21 hrs ago | US & Canada). 4. 'Even before the LA fires, Californians fled for 'climate havens'' (2 days ago | US & Canada). To the right of these items are two more news items: 5. 'My family was murdered at Auschwitz. Her grandfather drove trains to the camp' (19 hrs ago | Europe). 6. 'Trump fires at least a dozen government watchdogs' (45 mins ago | US & Canada). At the bottom of the page, there are four more news items: 7. 'Pete Hegseth confirmed as US defence secretary in tie-breaking vote' (18 hrs ago | US & Canada). 8. 'Paul McCartney warns of AI threats to artists in rare interview' (15 hrs ago). 9. 'Trump tours LA fire destruction amid worries about disaster aid' (17 hrs ago | US & Canada). 10. ''God forbid we should end up like Ukraine': Belarusians indifferent to election' (19 hrs ago | Europe). The page footer contains the text 'Trump returns to power'.

Generative AI in general: Image Generation Research Lets You Design and See the Future!



Stanford 2015



Stanford 2019



MJ v1 2022



MJ v2 2022



MJ v3 2022



MJ v4 2022



MJ v5 2023



MJ v6 2024

Era 3: Models with “Thinking”: “Test-Time Scaling” is the new buzzword

openai/text-davinci-002 (March 2022)

The cost to produce a book is 1000 dollars plus 8 dollars per book.
The book sells for 20 dollars each. How many books must be sold to make a profit?

1250 books must be sold to make a profit.

The cost to produce a book is 1000 dollars plus 8 dollars per book. The book sells for 20 dollars each. How many books must be sold to make a profit? Let's work step by step.

The cost of producing 1 book is $1000 + 8 = 1008$ dollars.
The cost of producing 2 books is $1008 + 8 = 1016$ dollars.
Therefore, the cost of producing n books is:

$$\text{Cost} = 1008 + 8n$$

The selling price of 1 book is \$20.
The selling price of 2 books is \$40.

The selling price of n books is:

$$\text{Selling Price} = 20n$$

We want to find how many books must be sold in order to make a profit. That is, we want to find n such that:

$$\text{Cost} < \text{Selling Price}$$

$$1008 + 8n < 20n$$

$$8n < 20n - 1008$$

$$-1008 < 12n$$

$$1000 < 12n$$

$$83.33 < n$$

Therefore, n must be greater than 83.33 in order to make a profit.



To determine how many books must be sold to make a profit, we need to figure out when the revenue from selling the books exceeds the total cost of producing them.

Let:

- x be the number of books sold.
- The cost to produce x books is:

$$\text{Total cost} = 1000 + 8x$$

(where 1000 dollars is the fixed cost, and 8 dollars is the variable cost per book).

- The revenue from selling x books is:

$$\text{Revenue} = 20x$$

(where each book is sold for 20 dollars).

To make a profit, the revenue must be greater than the total cost:

$$20x > 1000 + 8x$$

Solving the inequality:

1. Subtract $8x$ from both sides:

$$20x - 8x > 1000$$

$$12x > 1000$$

2. Divide both sides by 12:

$$x > \frac{1000}{12} \approx 83.33$$

Since x must be a whole number, you need to sell at least **84 books** to make a profit.

DeepSeek-R1 Reasoning: Find two 3-digit palindromic numbers that sum to a 4-digit palindromic number

<think> Okay, so I need to find two 3-digit palindromic numbers that, when added together, result in a 4-digit palindromic number. And none of the numbers can start with a zero. Let me break this down.

First, what's a palindromic number? It's a number that reads the same forwards and backwards. For example, 121 or 1331. So, a 3-digit palindrome would have the form ABA, where A and B are digits, and A can't be zero. Similarly, a 4-digit palindrome would have the form ABBA. Let me check: 1221 is a palindrome, yeah. So ABBA is correct, with $A \neq 0$.

So, I need two 3-digit numbers, both palindromes, that add up to a 4-digit palindrome. Let's denote the two 3-digit numbers as ABA and CDC (since they can be different). But wait, maybe they could be the same palindrome? Hmm, but adding the same palindrome would just be doubling it. Let's see: the smallest 3-digit palindrome is 101. $101 + 101 = 202$, which is a 3-digit palindrome, so too small. The largest 3-digit palindrome is 999. $999 + 999 = 1998$, which is a 4-digit number, but 1998 is not a palindrome. Hmm. So adding two identical 3-digit palindromes might get to a 4-digit number, but it's probably not a palindrome. Let's check. For example, $999 + 999 = 1998$. Not a palindrome. $888 + 888 = 1776$. Also not a palindrome. $121 + 121 = 242$, which is 3-digit. So perhaps the two palindromes need to be different. [...omitting 100s of lines ...]

Hence, my main thought streamline:

Possible sums are only 1111 and 1221, determined by the formula.

Two types of palindromic pairs:

Numbers with middle digit 0 and Leading+Trailing= $a + c = 11$, eg. $202 + 909 = 1111$.

Numbers with both leading and trailing digits summing $11 = a + c$, and middle digits summing $11 = b + d$. For example: $212 + 909 = 1121$ which isn't a palindrome or wait, no. Wait, another example $232 + 989 = 1221$. $B + d = 3 + 8 = 11$.

Therefore, focusing on styling both sums approach.

Therefore, to conclude, two examples answering the problem would be:

$202 + 909 = 1111$.

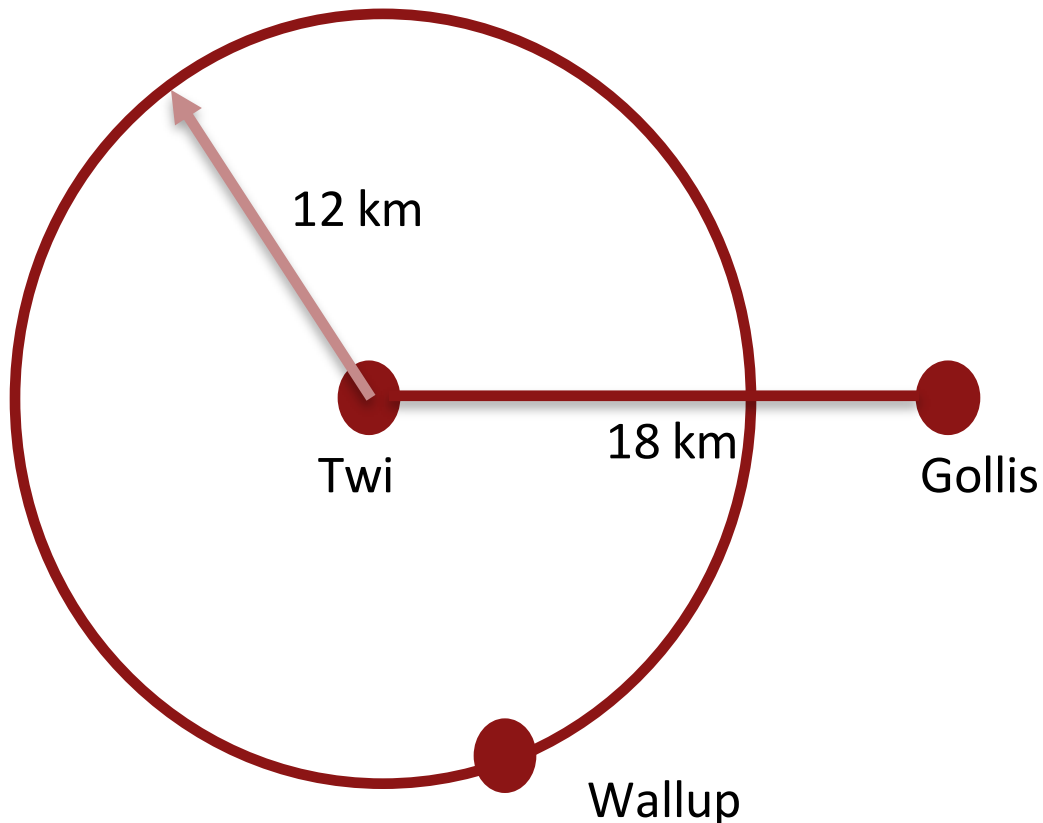
$232 + 989 = 1221$.

These are two distinct palindromic pairs with 3 digits each, adding up to 4-digit palindromes. No leading zeros.

Hence, verified that they fulfill all the conditions. </think>

[Model then proceeds to give a final answer as a clean problem solution]

If the town of Twi is 12 km from the town of Wallup, and the town of Gollis is 18 km from the town of Twi, what is the maximum and minimum possible distance the town of Wallup is from Gollis?



To find the **maximum** and **minimum** possible distances between **Wallup** and **Gollis**, let's consider the two possible configurations for the positions of the towns.

1. Maximum Possible Distance:

For the maximum distance between Wallup and Gollis, Wallup and Gollis would be on **opposite sides of Twi** in a straight line. In this case, the total distance would be the sum of the two distances:

- Wallup to Twi = 12 km
- Twi to Gollis = 18 km

So, the maximum possible distance is:

$$\text{Maximum distance} = 12 \text{ km} + 18 \text{ km} = 30 \text{ km}$$

2. Minimum Possible Distance:

For the minimum distance, Wallup, Twi, and Gollis would be in a straight line, with Gollis on the **same side** of Twi as Wallup. In this case, the minimum possible distance would be the difference between the two distances:

- Wallup to Twi = 12 km
- Twi to Gollis = 18 km

The minimum possible distance is:

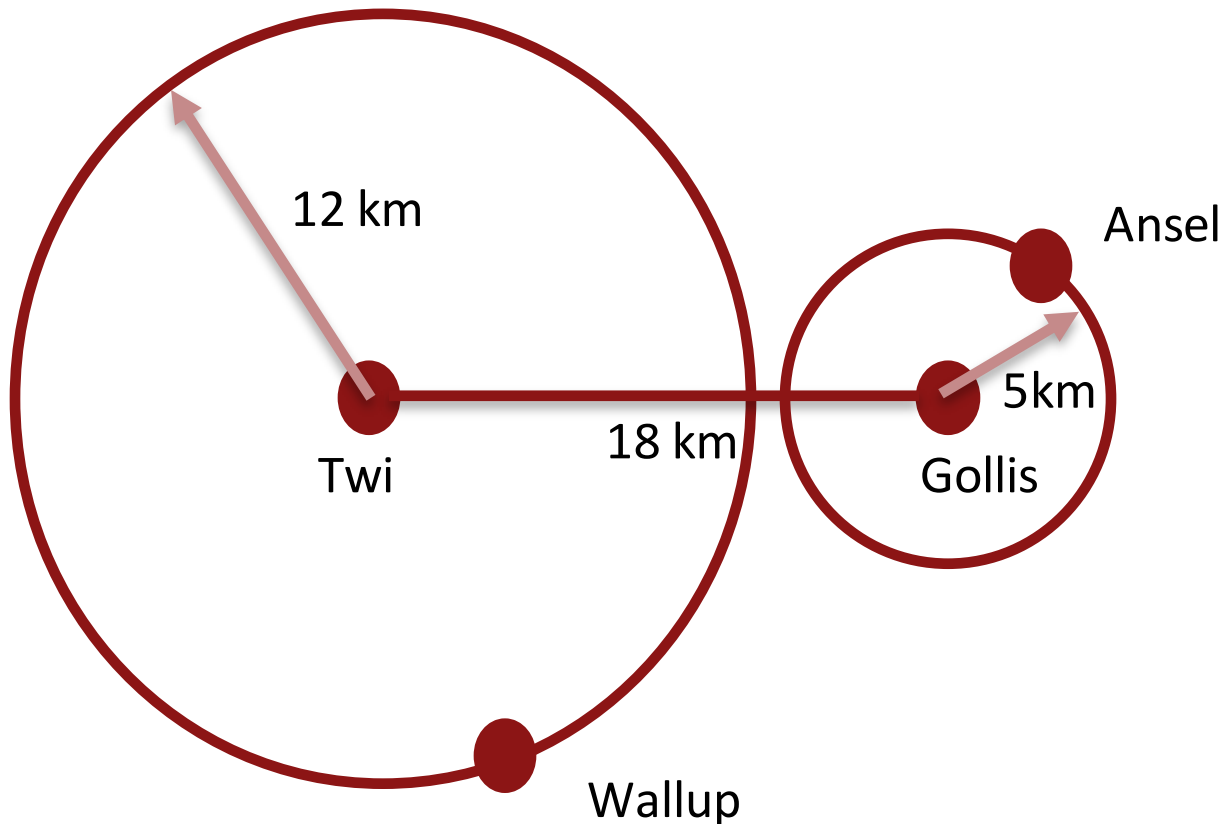
$$\text{Minimum distance} = |18 \text{ km} - 12 \text{ km}| = 6 \text{ km}$$

Conclusion:

- The **maximum distance** between Wallup and Gollis is **30 km**.
- The **minimum distance** between Wallup and Gollis is **6 km**.



If the town of Twi is 12 km from the town of Wallup, the town of Gollis is 18 km from the town of Twi, and the town of Ansel is 5 km from the town of Gollis, what is the maximum and minimum possible distance the town of Ansel is from Wallup?



Let's explore the problem step by step:

You want to know the **maximum** and **minimum** possible distances between the town of **Ansel** and the town of **Wallup**. We know the following distances:

- Wallup to Twi = 12 km
- Twi to Gollis = 18 km
- Gollis to Ansel = 5 km

Maximum Possible Distance:

For the **maximum possible distance** between **Ansel** and **Wallup**, the towns would need to be laid out in a straight line with Ansel on the opposite side of Gollis from Wallup. This would give us the longest possible distance by adding all the segments together:

- Wallup to Twi = 12 km
- Twi to Gollis = 18 km
- Gollis to Ansel = 5 km

So, the maximum possible distance from **Ansel** to **Wallup** is:

$$\text{Maximum distance} = 12 \text{ km} + 18 \text{ km} + 5 \text{ km} = 35 \text{ km}$$

Minimum Possible Distance:

For the **minimum possible distance**, Ansel would be on the same side of Gollis as Wallup, and we would subtract the distance between Ansel and Gollis from the other distances:

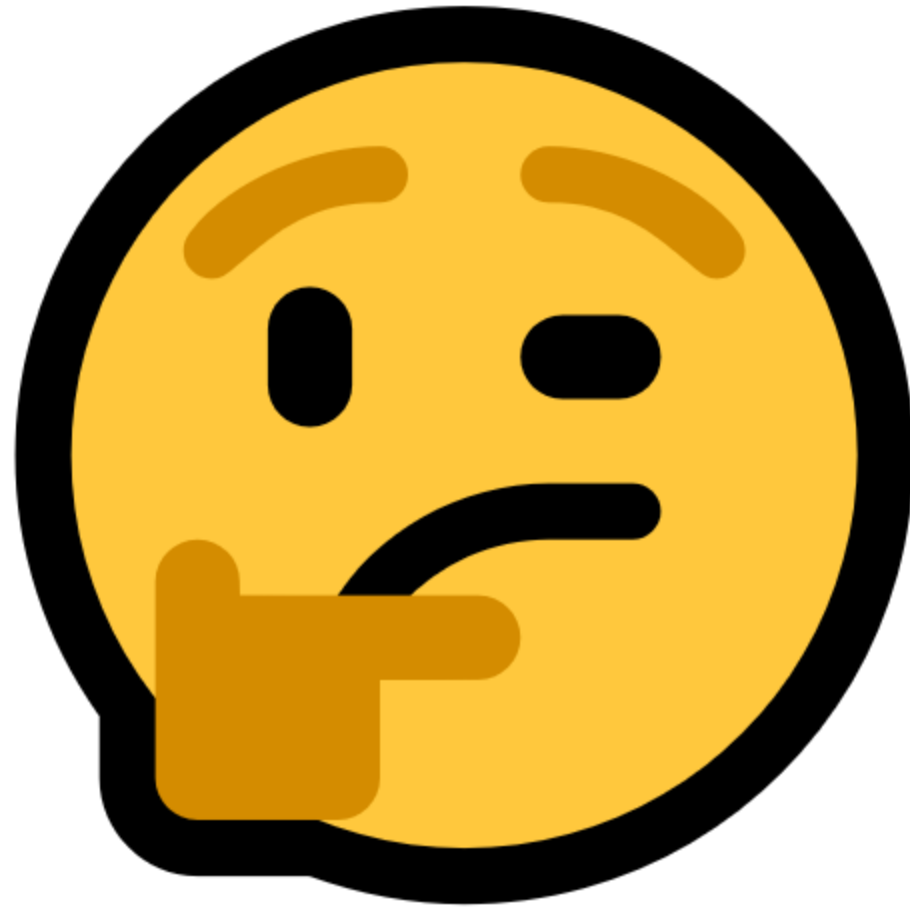
- Wallup to Twi = 12 km
- Twi to Gollis = 18 km
- Gollis to Ansel = 5 km

The minimum possible distance is:

$$\text{Minimum distance} = (18 \text{ km} - 5 \text{ km}) + 12 \text{ km} = 25 \text{ km}$$

Conclusion:

- The **maximum distance** from **Ansel** to **Wallup** is **35 km**.
- The **minimum distance** from **Ansel** to **Wallup** is **25 km**.



Non-technical lessons from DeepSeek-R1

U.S. “frontier” AI companies do not have a huge technical moat – they’re at most 6 months ahead

Control and development of AI cannot be determined by just U.S. policies

Currently, the best way to learn the details of modern LLM engineering is to read the tech reports of Chinese companies: This is sad, given that the U.S. long led as an open country advancing open science and engineering

The U.S. is more likely to maintain a long-term lead in AI by supporting universities, training more engineers, and supporting open software and knowledge than by trying to lock up “frontier” AI within a few companies

How can **the world** go wrong using AI?

All Generative AI models are badly biased (gender, ethnicity, ...)

Some of this reflects statistics of how the world has been till now
But standard GenAI learning algorithms actually **magnify** biases

LLMs might cause large scale job loss and social dislocation

But this is hard to predict ... the same has been claimed for 100s of years

Data and the world economy

GenAI models slurp vast quantities of data – this can be bad for privacy
But it also undermines the current internet order of free data display
Without care, AI may greatly exacerbate economic winners and losers

Concerns around AI:

Risks of security vulnerabilities from agents

- Agents cause new risks: As AIs are connected to more systems, and act more autonomously, the chance of accidental or malicious bad use increases
- But humanity also has ample experience controlling many things far more powerful than any single person, such as corporations and nation states
- If we look at the real risks to humanity – such as climate change and pandemics – AI will likely be a key part of the solution

Large Language Models and “Intelligence”

Since ChatGPT, in the popular imagination, large language models have essentially become synonymous with **artificial intelligence**

What is **artificial intelligence**?

The **artificial** part is easy. What is **intelligence**?



CALIFORNIA REPUBLIC

CA AB 2885

Section 11546.45.5 of the California Government Code (2024-09-28):

“Artificial intelligence” means an engineered or machine-based system that varies in its level of autonomy and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments.

an engineered or machine-based system that varies in its level of autonomy and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments



Intelligence vs. knowledge or task effectiveness

More commonly ...

We confuse being **knowledgeable** with being **intelligent**

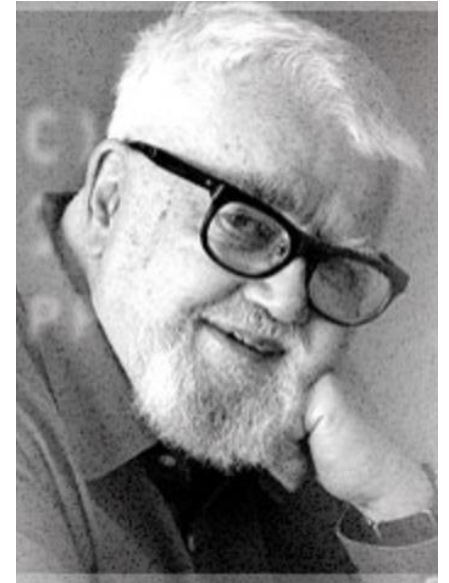
Or is **the ability to do tasks** intelligence?

AGI: “the ability to perform most (cognitive/ economically valuable) tasks at or above the level of the average person”

What is Artificial Intelligence?

“The science and engineering of making intelligent machines”

– John McCarthy (1955)



Intelligence: the ability to **learn** and perform a range of techniques to solve problems and achieve goals, **appropriate to the context in an uncertain, ever-varying world**

– Christopher Manning (2020)

AI in the rest of the 21st Century

- **Science and engineering progresses! Further big AI breakthroughs will occur!**
- We still need a number of them:
 - Better models of episodic memory (like human brains)
 - Better means for deliberate, novel reasoning (“System 2 cognition”, like humans)
 - Better planning
 - Better world models
 - Better methods for producing low-power GenAI models
 - Human brains require about 17 watts; one H100 GPU requires 700 watts!
 - But the human brain does use about 20% of the body’s energy: Thinking is hard work!
 - Better methods for learning from more limited data
 - Human children learn much more from a limited amount of data than LLMs do

How can **you** go wrong using AI?

You can assume that everything that an LLM says is correct.

It's not. These LLMs make stuff up. (They “hallucinate”.)

You can think they are magical. They are not.

You can ask a model to pretend that it is Warren Buffett.

But it may well not give you amazing investment advice.

You can conversely decide that they're good for nothing.

That would be a huge mistake!

Used wisely, LLMs can *enormously* increase your productivity.

The order of magnitude gains are huge vs. other technologies.

Three key suggestions

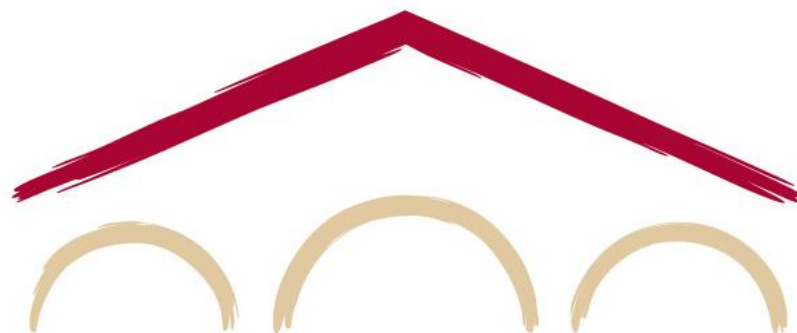
Experiment with and use generative AI!!! You don't need to know math! You can talk to them in English. See what they can do!!!

Easiest: go to <https://gemini.google.com/> . Log in with a personal gmail account (your corporate one may not work) and play!

Coworking: It's useful to think of an LLM as kind of like an intern or assistant. They're good at human stuff and being told stuff to do, but they also make mistakes. Rephrase things. Ask questions.

However big you think AI is going to be, it is likely bigger!

Large Language Models in 2025: How Much Understanding and Intelligence?



Christopher Manning

Founder, Stanford NLP Group ✿ Associate Director, Stanford HAI

@chrmanning ✿ @stanfordnlp

First Workshop of a Public AI Assistant to World Wide Knowledge (WWK)