

The Current State of AI

Session 1




How many engineers does it take to make a movie?

CNN BUSINESS Audio Live TV Log In

Why celebrities are striking: The average pay for actors may surprise you

By Samantha Delouya, CNN
Updated 9:58 AM EDT, Mon July 17, 2023



Actress reveals how much she just made from a top Netflix show

02:29 - Source: [CNN](#)

Video Ad Feedback

Netflix touts \$900k AI jobs amid Hollywood strikes

28 July



SAG strike 2023



GETTY IMAGES

By Max Matza

Warm-up: Is this good news or bad news?

News Corporation

News Corp using AI to produce 3,000 Australian local news stories a week

The Data Local team uses AI technology to generate stories on weather, fuel prices and traffic reports for hyperlocal mastheads

- [Follow our Australia news live blog for the latest updates](#)
- [Get our morning and afternoon news emails, free app or daily news podcast](#)

Amanda Meade

Mon 31 Jul 2023 11.00 EDT



News Corp Australia is producing 3,000 articles a week using generative artificial intelligence, executive chair Michael Miller has revealed.

Miller told the World News Media Congress [in Taipei](#) that a team of four staff use the technology to generate thousands of local stories each week on weather,

Etiquette



Cameras on — show your face



Mute yourself when not talking



Drop your questions in the chat + “raise” your hand during questions time



Don't give up: Something isn't clear? Get in touch.



Be nice!

Let's see if this works...

In **one** word, describe your current feeling about Artificial Intelligence.

Drop it in the chat 

Live Session Schedule

are all the invites in your calendar?

- 1 Tuesday, August 22, 9 pm ET
- 2 Thursday, August 24, 9 pm ET
- 3 Tuesday, August 29, 9 pm ET
- 4 Thursday, August 31, 9 pm ET
- 5 An extra Q&A session, date TBC

 Each session is 60+30 minutes

 **All the sessions will be recorded.** You'll get an email with the recording, slides, and other action items.



What we'll do?



Session 1: Fundamentals and orientation



Session 2: Working with words and images



Session 3: Analyzing and visualizing data



Session 4: Automation, APIs, Privacy & broader implications



Session X: Reinforcement & Support

Teaching Methodology



Live Sessions



How-to Videos



Chat + Live Q&A

Course Goals



Confidence



Key Skills



Building Blocks



Foundation for
Further Study

Disclaimer

- Not investment advice.
- Not legal, tax, medical, or any other advice.
- Unproven technology – buggy, unreliable, unpredictable.
- Relying on public sources of information, at face value
- We're all in this together! Let's join forces.




Course Portal

Log in with your email at <https://in.hypefree.com>

HYPE-FREE COURSE PORTAL Log in

Dive in.
Welcome to the Hype-Free course portal. Below you'll find all the recordings, readings, assignments, and updates. 📁



Resources

- 29 MAY 2023 PARTICIPANTS [Chatroom](#)
- 28 MAY 2023 PARTICIPANTS [Reading List](#)
- 28 MAY 2023 PARTICIPANTS [Glossary](#)

HYPE-FREE COURSE PORTAL [Glossary](#) [Resources](#) [Chatroom](#) [Account](#)

Glossary

28 MAY 2023 — 6 MIN READ

A

Activation Function: In a neural network, the function combines inputs and weights and transforms them into an output that the neuron can use.

Affordances: The possibilities an AI system has for interaction or action within its environment. In simpler terms, it's like a list of what an AI can do in a given situation. For example, in a game, affordances could be the different moves an AI can make. The concept helps designers understand how AI can interact intelligently with its surroundings, including people.

Algorithm: A set of rules or instructions given to an AI, or other computational systems, to help it learn and solve specific problems.

Artificial General Intelligence (AGI): A type of AI that's as smart as a human across the board—a machine capable of understanding, learning, and applying its knowledge to any intellectual task that a human being can.

Artificial Intelligence (AI): Broad area of computer science focused on creating machines capable of performing tasks that normally require human intelligence, such as understanding language, recognizing patterns, and making decisions.

API (Application Programming Interface): An API is a set of rules that allows different software applications to communicate with each other. It's like a menu in a restaurant: you (the software) make a request (place an order), and the kitchen (the system you're requesting information from) uses the instructions provided (the menu) to fulfill your request. APIs enable the different pieces of software you use every day to function together seamlessly.

Autoencoder: A type of artificial neural network used for learning efficient codings of input data, typically for the purpose of dimensionality reduction or denoising.

B

Backpropagation: A method used in artificial neural networks to calculate the gradient of loss function with respect to the weights in the network, used for training the model.

Bias in AI: Systematic errors in predictions made by AI models, often caused by

 Questions so far?



This is a *very* serious guy

Yann LeCun

🌐 22 languages ▾

Article Talk

Read Edit View history Tools ▾

From Wikipedia, the free encyclopedia

Yann André LeCun^[1] (/ləˈkʌn/ *lə-KUN*, French: [lɛkœ̃]^[2] originally spelled **Le Cun**^[2] born 8 July 1960) is a French computer scientist working primarily in the fields of machine learning, computer vision, mobile robotics and computational neuroscience. He is the Silver Professor of the Courant Institute of Mathematical Sciences at New York University and Vice-President, Chief AI Scientist at Meta.^{[3][4]}

He is well known for his work on optical character recognition and computer vision using convolutional neural networks (CNN), and is a founding father of convolutional nets.^{[5][6]} He is also one of the main creators of the DJVu image compression technology (together with Léon Bottou and Patrick Haffner). He co-developed the Lush programming language with Léon Bottou.

LeCun received the 2018 Turing Award (often referred to as the "Nobel Prize of Computing"), together with Yoshua Bengio and Geoffrey Hinton, for their work on deep learning.^[7] The three are sometimes referred to as the "Godfathers of AI" and "Godfathers of Deep Learning".^{[8][9][10][11][12][13]}

Early life [edit]

LeCun was born at Soisy-sous-Montmorency in the suburbs of Paris. His name was originally spelled *Le Cun* from the old Breton form *Le Cunft* and was from the region of Guingamp in northern Brittany. "Yann" is the Breton form for "John".

Education [edit]

He received a Diplôme d'Ingénieur from the ESIEE Paris in 1983 and a PhD in Computer Science from Université Pierre et Marie Curie (today Sorbonne University) in 1987 during which he proposed an early form of the back-propagation learning algorithm for neural networks.^[14]

Yann LeCun



Yann LeCun in 2018

Born	8 July 1960 (age 62) Soisy-sous-Montmorency, France
Alma mater	ESIEE Paris (MS) Pierre and Marie Curie University (PhD)
Known for	Deep learning
Awards	Turing Award (2018) AAAI Fellow (2019) Legion of Honour (2020)
Scientific career	

“I don't think we can train a machine to be intelligent, purely from text... **I think it's basically hopeless.**”

I take an object. I put it on the table and I push the table. It's completely obvious to you that the object will be pushed with the table, right? Because it's sitting on it.

There's no text in the world, I believe that explains this. And so if you train a machine as powerful as it could be, your GPT 5000 or whatever it is, it's never gonna learn about this.”

Activity — 4 minutes:

Log in to ChatGPT and test whether it has a “world model.”

Describe a situation that would require it to understand the relation between objects in space. Did it succeed?



How did it go?

(raise your hand 🙋 if you'd like to share)

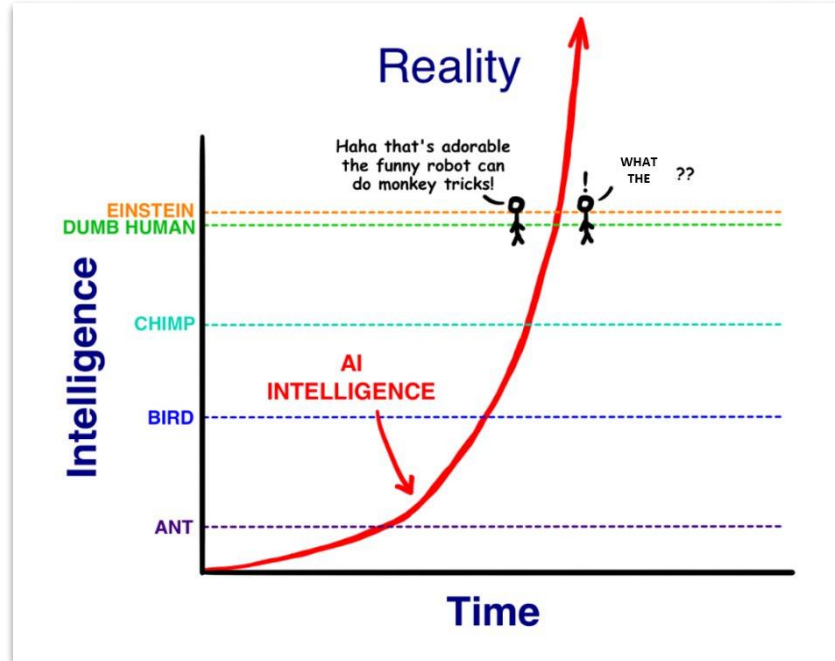
Intelligence is “emergent”



* Arises from the interaction of simpler elements in a system, but cannot be predicted solely from those elements. Shows characteristics that appear at a higher system level due to lower-level interactions.

Software is no longer “programmed”

Is this what's happening?



“The development of full artificial intelligence could spell the end of the human race.... It would take off on its own, and redesign itself at an ever-increasing rate. **Humans, who are limited by slow biological evolution, couldn't compete and would be superseded.**”

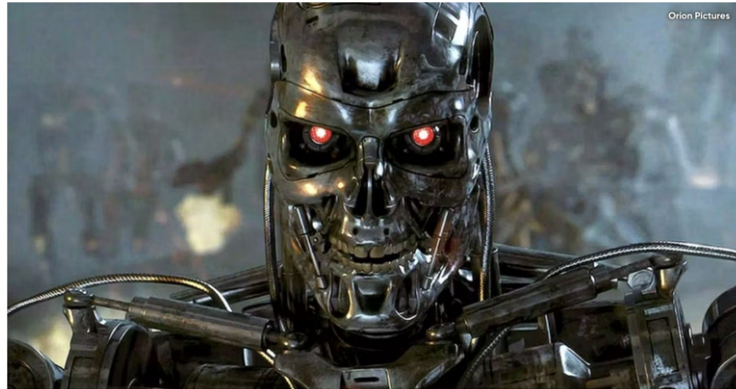
– Stephen Hawking

Happy ending? *(it's not clear)*

James Cameron Working on New Terminator Movie, Wants to See How AI Affects the Real World

A new Terminator movie is being written by James Cameron who is waiting to see the effect of real-life artificial intelligence.

BY JONATHAN FUGE PUBLISHED 4 DAYS AGO



Director James Cameron has revealed that he has started writing a new *Terminator* movie...but wants to see how the real world reacts to the emergence of artificial intelligence before finishing it. Appearing at the Dell Technologies World 2023 event, a few people who were in attendance have revealed that Cameron is working on a new *Terminator* outing, but is "holding off on finishing it until the direction of #AI becomes clearer."

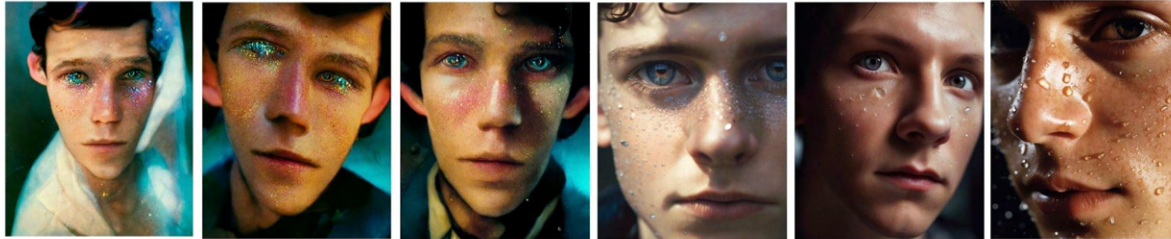
Moravec's Paradox

(the hard stuff is easy, the easy stuff is hard)



Still, things are progressing at a rapid clip...

Midjourney Version Comparison



V1

Released
February 2022

V2

Released
April 12, 2022

V3

Released
July 25, 2022

V4

Released
November 5, 2022

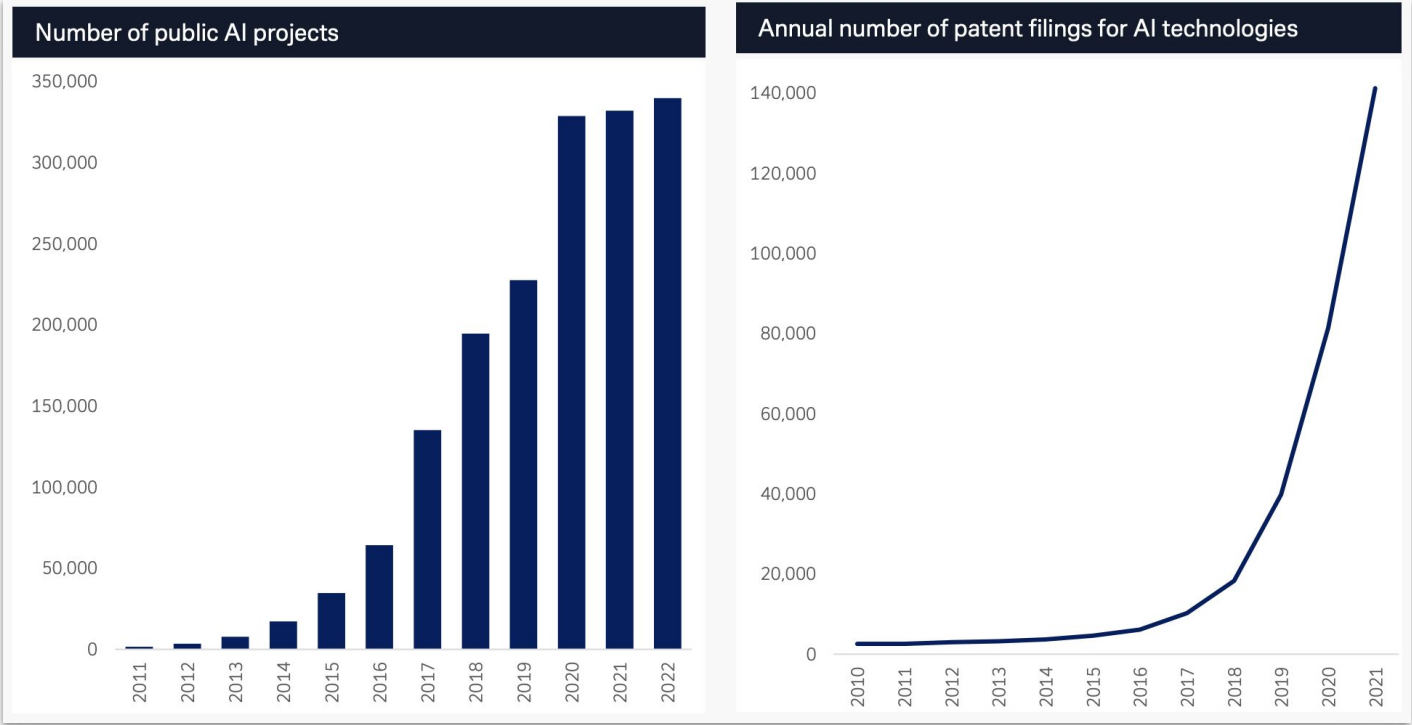
V5

Released
March 15, 2022

V5.1

Released
May 3, 2022

Research is exploding



Money is flowing

Bloomberg

Technology + Work Shift

Microsoft Invests \$10 Billion in ChatGPT Maker OpenAI



BREAKING NEWS: MICROSOFT TO BOOST INVESTMENT IN CHATGPT MAKER OPENAI

Microsoft is investing \$10 billion in OpenAI. Bloomberg's Abigail Doolittle reports on 'Bl... Source: Bloomberg

na Bass
ry 23, 2023 at 9:06 AM EST Updated on January 23, 2023 at 5:03

CNN BUSINESS

Audio Live TV Log In

AI chip boom sends Nvidia's stock surging after whopper of a quarter

By David Goldman, CNN
Published 8:27 AM EDT, Thu May 25, 2023

A photograph of an Nvidia trade show booth. The booth is primarily green and white, with the Nvidia logo prominently displayed. Several people are seen interacting with the booth, which features various displays and information.

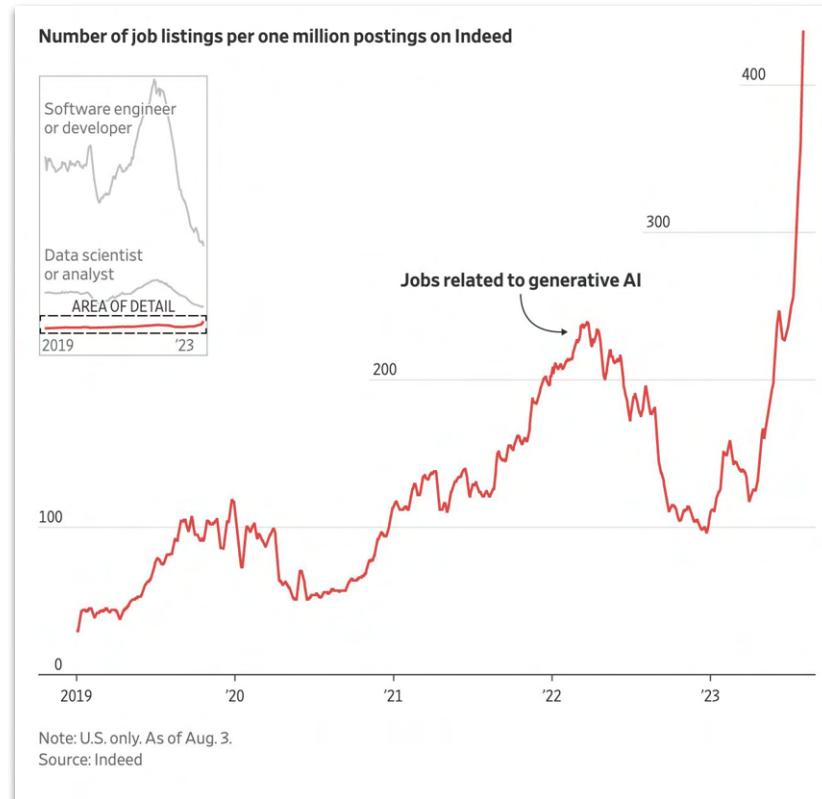
Google-backed Anthropic raises \$450 mln in latest AI funding

By Krystal Hu and Jaiveer Shekhawat

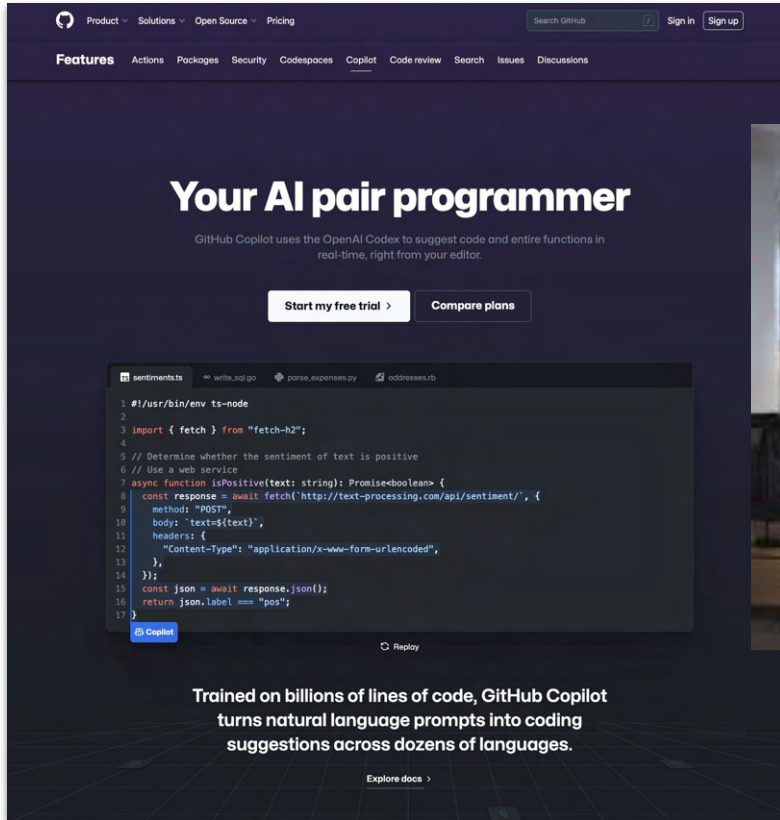
May 23, 2023 2:54 PM EDT - Updated 6 days ago



Hiring is heating up



Magical software



The screenshot shows the GitHub Copilot website landing page. At the top, there is a navigation bar with links for Product, Solutions, Open Source, Pricing, Search GitHub, Sign in, and Sign up. Below this is a secondary navigation bar with links for Features, Actions, Packages, Security, Codespaces, Copilot, Code review, Search, Issues, and Discussions. The main content area features the heading "Your AI pair programmer" and a sub-heading "GitHub Copilot uses the OpenAI Codex to suggest code and entire functions in real-time, right from your editor." There are two buttons: "Start my free trial >" and "Compare plans". Below this is a code editor window showing a JavaScript function for sentiment analysis. At the bottom, there is a "Trained on billions of lines of code, GitHub Copilot turns natural language prompts into coding suggestions across dozens of languages." and an "Explore docs >" link.

Product Solutions Open Source Pricing Search GitHub Sign in Sign up

Features Actions Packages Security Codespaces Copilot Code review Search Issues Discussions

Your AI pair programmer

GitHub Copilot uses the OpenAI Codex to suggest code and entire functions in real-time, right from your editor.

[Start my free trial >](#) [Compare plans](#)

```
1 #!/usr/bin/env ts-node
2
3 import { fetch } from "fetch-h2";
4
5 // Determine whether the sentiment of text is positive
6 // Use a web service
7 async function isPositive(text: string): Promise<boolean> {
8   const response = await fetch('http://text-processing.com/api/sentiment/', {
9     method: "POST",
10    body: `text=${text}`,
11    headers: {
12      "Content-Type": "application/x-www-form-urlencoded",
13    },
14  });
15  const json = await response.json();
16  return json.label === "pos";
17 }
```

[Trained on billions of lines of code, GitHub Copilot turns natural language prompts into coding suggestions across dozens of languages.](#)

[Explore docs >](#)



“Copilot has **dramatically accelerated my coding**, it's **hard to imagine going back** to "manual coding". Still learning to use it but it **already writes ~80% of my code**, ~80% accuracy. I don't even really code, I prompt. & edit.”

— Andrej Karpathy

Medical miracles

Article

Walking naturally after spinal cord injury using a brain–spine interface

<https://doi.org/10.1038/s41586-023-06094-5>

Received: 1 August 2022

Accepted: 17 April 2023

Published online: 24 May 2023

Open access

 Check for updates

Henri Lorach^{1,2,3,4,5}, Andra Galvez^{2,3,6,7}, Valeria Spagnolo^{2,3}, Felix Martel⁴, Gergil Karakas⁴, Nadine Interling^{2,5}, Molywan Vat^{2,3}, Olivier Falvre⁴, Cathal Harris^{2,5}, Salif Komi^{2,3}, Jimmy Ravier^{2,3}, Thibault Collin^{2,3}, Laure Coquoz^{2,3}, Icare Sakr^{2,3}, Eddy Baaklini^{2,3}, Sergio Daniel Hernandez-Charpak^{2,3}, Gregory Dumont^{2,3}, Rik Buschman⁴, Nicholas Buse⁴, Tim Denison^{4,8}, Ilse van Nee⁴, Leonie Asboth^{2,3}, Anne Watrin^{2,3}, Anne Struber⁴, Fabien Sauter-Starace⁴, Lilla Langer⁴, Vincent Aubroux⁴, Stefano Carda⁴, Stephan Chabardes⁴, Tatiana Akseanova⁴, Robin Demesmaeker^{2,3}, Guillaume Charvet^{4,9,10}, Jocelyne Bloch^{2,3,11,12} & Grégoire Courtine^{2,3,11,12}

A spinal cord injury interrupts the communication between the brain and the region of the spinal cord that produces walking, leading to paralysis^{1,2}. Here, we restored this communication with a digital bridge between the brain and spinal cord that enabled an individual with chronic tetraplegia to stand and walk naturally in community settings. This brain–spine interface (BSI) consists of fully implanted recording and stimulation systems that establish a direct link between cortical signals³ and the analogue modulation of epidural electrical stimulation targeting the spinal cord regions involved in the production of walking^{4,5}. A highly reliable BSI is calibrated within a few minutes. This reliability has remained stable over one year, including during independent use at home. The participant reports that the BSI enables natural control over the movements of his legs to stand, walk, climb stairs and even traverse complex terrains. Moreover, neurorehabilitation supported by the BSI improved neurological recovery. The participant regained the ability to walk with crutches overnight even when the BSI was switched off. This digital bridge establishes a framework to restore natural control of movement after paralysis.

To walk, the brain delivers executive commands to the neurons located in the lumbosacral spinal cord¹. Although the majority of spinal cord injuries do not directly damage these neurons, the disruption of descending pathways interrupts the brain-derived commands that are necessary for these neurons to produce walking². The consequence is permanent paralysis.

We previously showed that epidural electrical stimulation targeting the individual dorsal root entry zones of the lumbosacral spinal cord enables the modulation of specific leg motor pools^{3–5}. In turn, recruiting these dorsal root entry zones with preprogrammed spatiotemporal sequences replicates the physiological activation of leg motor pools underlying standing and walking^{3,4,6,7}. These stimulation sequences restored standing and basic walking in people with paralysis due to a spinal cord injury. However, this recovery required wearable motion sensors to detect motor intentions from residual movements or compensatory strategies to initiate the preprogrammed stimulation sequences³. Consequently, the control of walking was not perceived as completely natural. Moreover, the participants showed limited ability to adapt leg movements to changing terrain and voli-

Here, we suggest that a digital bridge^{3,11–19} between the brain and spinal cord would enable volitional control over the timing and amplitude of muscle activity, restoring more natural and adaptive control of standing and walking in people with paralysis due to spinal cord injury.

Digital bridge from brain to spinal cord

To establish this digital bridge, we integrated two fully implanted systems that enable recording of cortical activity and stimulation of the lumbosacral spinal cord wirelessly and in real time (Fig. 1a).

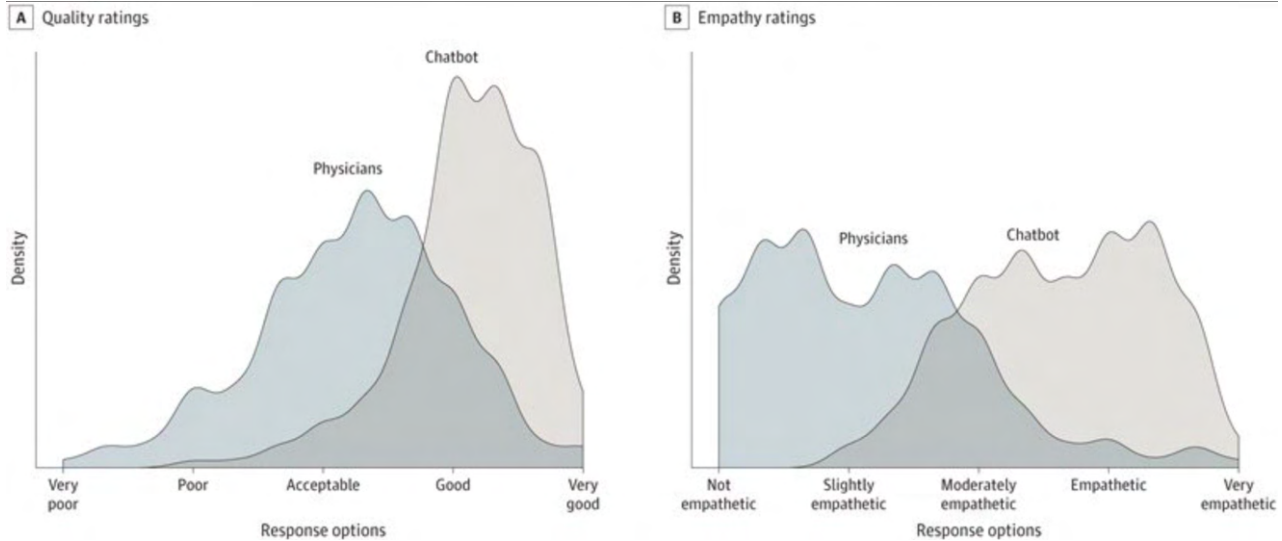
To monitor electrocorticographic (ECoG) signals from the sensorimotor cortex, we leveraged the WIMAGINE technology¹⁹. WIMAGINE implants consist of an 8-by-8 grid of 64 electrodes (4 mm × 4.5 mm pitch in anteroposterior and mediolateral axes, respectively) and recording electronics that are embedded within a 30 mm diameter, circular-shaped titanium case that has the same thickness as the skull. The geometry of the system favours close and stable contact between the electrodes and the dura mater, and renders the devices invisible



Faster finance



Expertise at scale



Enhanced presence



Enhanced... everything?



Unleashing our imagination



Making Photoshop great again



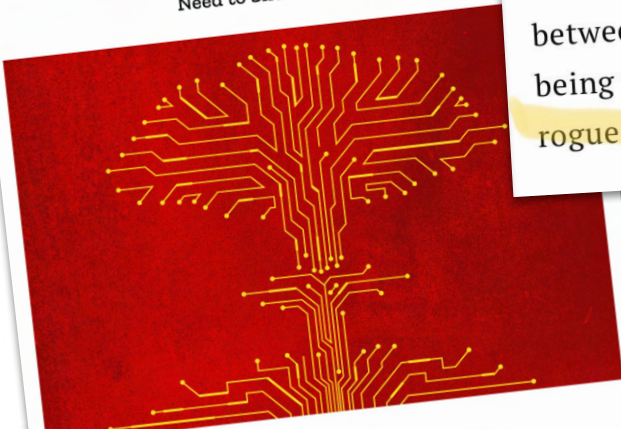
Managing itself?

The screenshot displays the Cognosys AI interface, which is organized into three main vertical sections: a left sidebar, a central main content area, and a right sidebar.

- Left Sidebar:** Contains the Cognosys logo at the top, followed by navigation links for "Home", "Settings", and "Agents".
- Central Main Content Area:**
 - Objective:** "Create a business plan for a new online education business." Below this is a "Thinking" status indicator with a checkmark.
 - Adding task:** A section containing a single task: "1. Conduct market research and identify target audience for the online education business."
 - Adding task:** A second section containing two tasks: "2. Develop a comprehensive curriculum and course offerings based on the identified target audience's needs and preferences." and "3. Design a marketing and promotion strategy to attract potential students and establish a strong online presence."
 - Thinking:** A final "Thinking" status indicator with a starburst icon.
- Right Sidebar:** Titled "Tasks", it lists three tasks, each with a checkbox and a "Show more" link:
 - 1. Conduct market research an... (Show more)
 - 2. Develop a comprehensive... (Show more)
 - 3. Design a marketing and... (Show more)

And threatening our survival?

DEAD • TECHNOLOGY
Pausing AI Developments Isn't Enough. We
Need to Shut it All Down



moving elsewhere. Track all GPUs sold. If intelligence says that a country outside the agreement is building a GPU cluster, be less scared of a shooting conflict between nations than of the moratorium being violated; be willing to destroy a rogue datacenter by airstrike.



How did we get here?

1940s: The brain as a logical learning machine

BULLETIN OF
MATHEMATICAL BIOPHYSICS
VOLUME 5, 1943

A LOGICAL CALCULUS OF THE IDEAS IMMANENT IN NERVOUS ACTIVITY

WARREN S. MCCULLOCH AND WALTER PITTS

FROM THE UNIVERSITY OF ILLINOIS, COLLEGE OF MEDICINE,
DEPARTMENT OF PSYCHIATRY AT THE ILLINOIS NEUROPSYCHIATRIC INSTITUTE,
AND THE UNIVERSITY OF CHICAGO

Because of the "all-or-none" character of nervous activity, neural events and the relations among them can be treated by means of propositional logic. It is found that the behavior of every net can be described in these terms, with the addition of more complicated logical means for nets containing circles; and that for any logical expression satisfying certain conditions, one can find a net behaving in the fashion it describes. It is shown that many particular choices among possible neurophysiological assumptions are equivalent, in the sense that for every net behaving under one assumption, there exists another net which behaves under the other and gives the same results, although perhaps not in the same time. Various applications of the calculus are discussed.

I. Introduction

Theoretical neurophysiology rests on certain cardinal assumptions. The nervous system is a net of neurons, each having a soma and an axon. Their adjunctions, or synapses, are always between the axon of one neuron and the soma of another. At any instant a neuron has some threshold, which excitation must exceed to initiate an impulse. This, except for the fact and the time of its occurrence, is determined by the neuron, not by the excitation. From the point of excitation the impulse is propagated to all parts of the neuron. The velocity along the axon varies directly with its diameter, from less than one meter per second in thin axons, which are usually short, to more than 150 meters per second in thick axons, which are usually long. The time for axonal conduction is consequently of little importance in determining the time of arrival of impulses at points unequally remote from the same source. Excitation across synapses occurs predominantly from axonal terminations to somata. It is still a

1950s: Can machines think?

VOL. LIX. No. 236.]

[October, 1950

MIND
A QUARTERLY REVIEW
OF
PSYCHOLOGY AND PHILOSOPHY

I.—COMPUTING MACHINERY AND
INTELLIGENCE

By A. M. TURING

1. *The Imitation Game.*

I PROPOSE to consider the question, 'Can machines think?' This should begin with definitions of the meaning of the terms 'machine' and 'think'. The definitions might be framed so as to reflect so far as possible the normal use of the words, but this attitude is dangerous. If the meaning of the words 'machine' and 'think' are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning and the answer to the question, 'Can machines think?' is to be sought in a statistical survey such as a Gallup poll. But this is absurd. Instead of attempting such a definition I shall replace the question by another, which is closely related to it and is expressed

1956: Proving mathematical theorems



Late 1950s: *Learning* machines

Vol. VI, No. 2, Summer 1958



research trends

CORNELL AERONAUTICAL LABORATORY, INC., BUFFALO 21, NEW YORK

The Design of an

Intelligent AUTOMATON

by FRANK ROSENBLATT

Introducing the perceptron — A machine which senses, recognizes, remembers, and responds like the human mind.

STORIES about the creation of machines having human qualities have long been a fascinating province of science fiction. First, in recent years our knowledge of the functioning of individual cells in the central nervous system has

1960s: Early “Natural Language Processing”

Welcome to

```
EEEEEE LL      IIII  ZZZZZZ  AAAAA
EE      LL      II    ZZ     AA  AA
EEEEEE LL      II    ZZZ    AAAAAA
EE      LL      II    ZZ     AA  AA
EEEEEE LLLLLL IIII  ZZZZZZ  AA  AA
```

Eliza is a mock Rogerian psychotherapist.
The original program was described by Joseph Weizenbaum in 1966.
This implementation by Norbert Landsteiner 2005.

```
ELIZA: Is something troubling you ?
YOU:   Men are all alike.
ELIZA: What is the connection, do you suppose ?
YOU:   They're always bugging us about something or other.
ELIZA: Can you think of a specific example ?
YOU:   Well, my boyfriend made me come here.
ELIZA: Is it important to you that your boyfriend made you come here ?
YOU:   He says I'm depressed much of the time.
ELIZA: I am sorry to hear that you are depressed.
YOU:   It's true. I am unhappy.
ELIZA: Can you explain what made you unhappy ?
YOU:
```

1970-1990s: AI Winter

“in no part of the field have discoveries made so far produced the major impact that was then promised.”

Part I Artificial Intelligence
A general survey by Sir James Lighthill FRS
Lucasian Professor of Applied Mathematics,
Cambridge University. July 1972.

Contents

1. Introduction	page 1
2. The ABC of the Subject	2
3. Past Disappointments	8
4. Future Possibilities	18
The next twenty-five years.	19

1 Introduction

The Science Research Council has been receiving an increasing number of applications for research support in the rather broad field with mathematical, engineering and biological aspects which often goes under the general description Artificial Intelligence (AI). The research support applied for is sufficient in volume, and in variety of discipline involved, to demand that a general view of the field be taken by the Council itself. In forming such a view the Council has available to it a great deal of specialist information through its structure of Boards and Committees, particularly from the Engineering Board and its Computing Science Committee and from the Science Board and its Biological Sciences Committee. These include specialised reports on the contribution of AI to practical aims on the one hand and to basic neurobiology on the other, as well as a large volume of detailed recommendations on grant applications.

To supplement the important mass of specialist and detailed information available to the Science Research Council, its Chairman decided to commission an independent report by someone outside the AI field but with substantial general experience of research work in multidisciplinary fields including fields with

1997: Deep Blue vs. Kasparov



Early 2010s: Neural nets & image recognition

THE WALL STREET JOURNAL.

Google Acquires Artificial-Intelligence Company DeepMind

By [Rolfe Winkler](#) [Follow](#)

Jan. 26, 2014 10:44 pm ET



Gift unlocked article

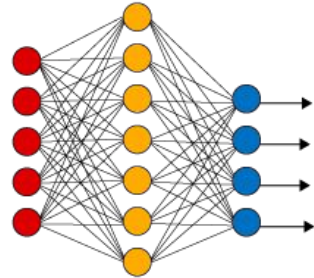
Google Inc. said it had acquired artificial-intelligence company DeepMind Technologies Ltd.

The company did not disclose the terms of the deal, but the tech website Recode reported that Google paid \$400 million for the startup based in London.

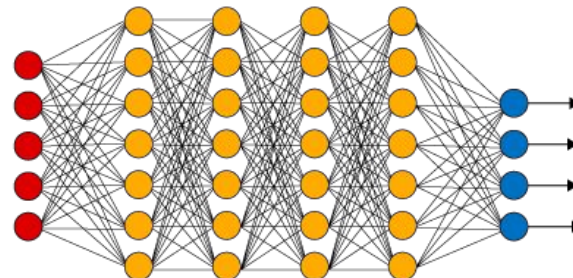
On its website, which has just one page, DeepMind describes itself as a "cutting edge artificial intelligence company. We combine the best techniques from machine learning and systems neuroscience to build powerful general-purpose learning algorithms."

What's so deep about deep learning?

Simple Neural Network



Deep Learning Neural Network



● Input Layer

● Hidden Layer

● Output Layer

2010s: A variety of approaches

Supervised Learning	Teaching computers to learn from labeled examples and make predictions based on patterns.
Unsupervised Learning	Helping computers find patterns and group similar data without explicit labeling.
Reinforcement Learning	Training computers to learn optimal behavior through trial and error with a reward-based system.
Convolutional Neural Networks	Specialized networks that enable computers to understand and analyze images and videos.
Recurrent Neural Networks	Networks designed for sequential data, allowing computers to process information with a sense of order, useful for language and time series data.
Ensemble Learning	Combining multiple models to improve prediction accuracy by leveraging diverse perspectives.
Transfer Learning	Applying knowledge from one task to improve performance on a related task, reducing the need for extensive new training.
Deep Reinforcement Learning	Merging deep learning and reinforcement learning to teach computers to make optimal decisions in complex environments.

2017: A new path...

notice anything surprising about this publication?

1706.03762v5 [cs.CL] 6 Dec 2017

Attention Is All You Need

Ashish Vaswani* Google Brain avaswani@google.com	Noam Shazeer* Google Brain noam@google.com	Niki Parmar* Google Research nikip@google.com
Llion Jones* Google Research llion@google.com	Aidan N. Gomez* † University of Toronto aidan@cs.toronto.edu	Lukasz Kaiser* Google Research lukasz@cs.toronto.edu
	Illia Polosukhin* † illia.polosukhin@gmail.com	

Abstract

The dominant sequence transduction models are based on convolutional neural networks that include an encoder and a decoder. We propose a new simple network architecture based solely on attention mechanisms, dispensing with recurrent layers entirely. Experiments on two machine translation tasks show that our model is superior in quality while being more parallelizable and requiring less time to train. Our model achieves 28.4 BLEU on the Vietnamese-to-German translation task, improving over the existing best models by over 2 BLEU. On the WMT 2014 English-to-French translation task, our model establishes a new single-model state-of-the-art BLEU score of 40.6, beating the previous best model by 0.5. On the WMT 2016 English-to-French translation task, our model establishes a new single-model state-of-the-art BLEU score of 41.2, beating the previous best model by 0.3. We show that the Transformer is a strong baseline for other tasks by applying it successfully to English constituency



2022-2023: The “iPhone moment” for LLMs *(Large Language Models)*

THE WALL STREET JOURNAL.

HEARD ON THE STREET

AI Has Its ‘iPhone Moment’

The launch of ChatGPT has set off a frenzy by companies to show they are in the AI game

By [Laura Forman](#) [Follow](#)

March 2, 2023 6:30 am ET



Gift unlocked article



Listen (5 min)



ChatGPT on GPT



ELI5 each of the components of GPT



Certainly!



1. **Generative:** Like an artist with words, GPT can create new sentences and stories all by itself.
2. **Pre-trained:** Before you even ask it anything, GPT has read and learned from tons of books, websites, and other texts, kind of like studying before a big test.
3. **Transformer:** This is how GPT thinks and understands language. It looks at all the words in a sentence and figures out how they relate to each other, like putting together a puzzle.



Now explain to a more sophisticated layperson



Certainly!



1. **Generative:** GPT has the ability to create new text that is coherent and contextually relevant, synthesizing information in a way that resembles human-like writing.
2. **Pre-trained:** The model is trained on vast amounts of existing text data, learning the structure, grammar, and nuances of language. This "pre-training" phase allows it to have a wide-ranging understanding of human language before it's fine-tuned for specific tasks.
3. **Transformer:** This is the architecture that powers GPT. It uses attention mechanisms to weigh the importance of different parts of the input, enabling it to capture complex relationships and dependencies in the text. It processes all parts of a sequence simultaneously rather than in order, making it highly efficient and capable of handling long-term relationships within the text.

Ultimately, it's all just “predictions”.

 Questions so far?

Reason

Solve
Problems

Learn from
experience

What is “intelligence”?

(drop it in the chat 🖱️)

Make
Plans

Think
abstractly

Comprehend
ideas

Group Work — 8 minutes:

Introduce yourself and come up with an answer:

You have the opportunity to “inject” intelligence into **one** product.

What would you do to maximize profits?

or

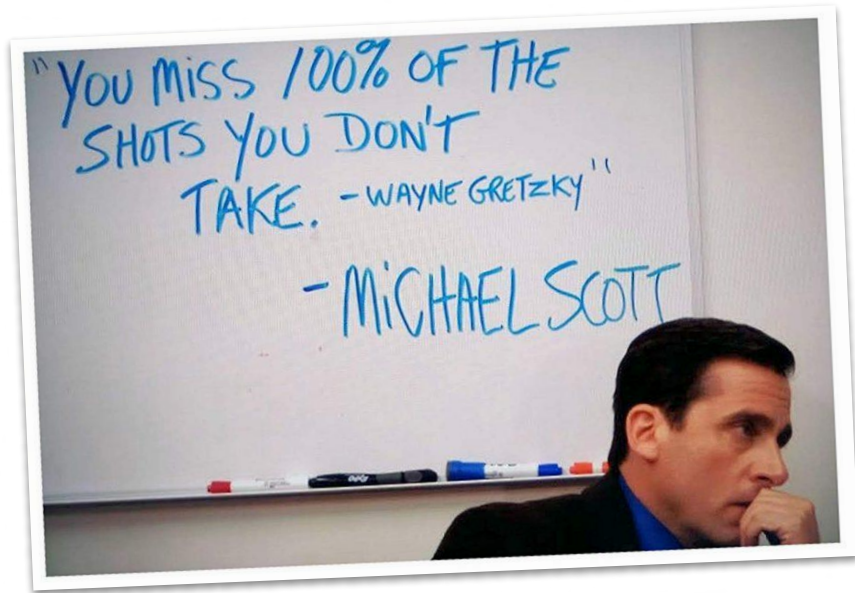
What would you do to maximize the benefit for humanity?



How did it go?

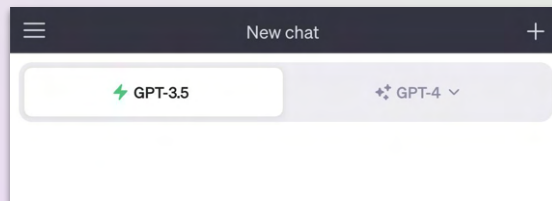
(raise your hand 🙋 if you'd like to share)

Let's talk about unintelligence *(and confidence)*



Prompt: “Cite ten books about the impact of computers on urban planning”

Try it with [GPT 3.5](#), [GPT 4](#), and [Bard](#)



Activity — 4 minutes:

Ask ChatGPT (or Bard) for specific sources or quotes about a topic you **don't know about**.

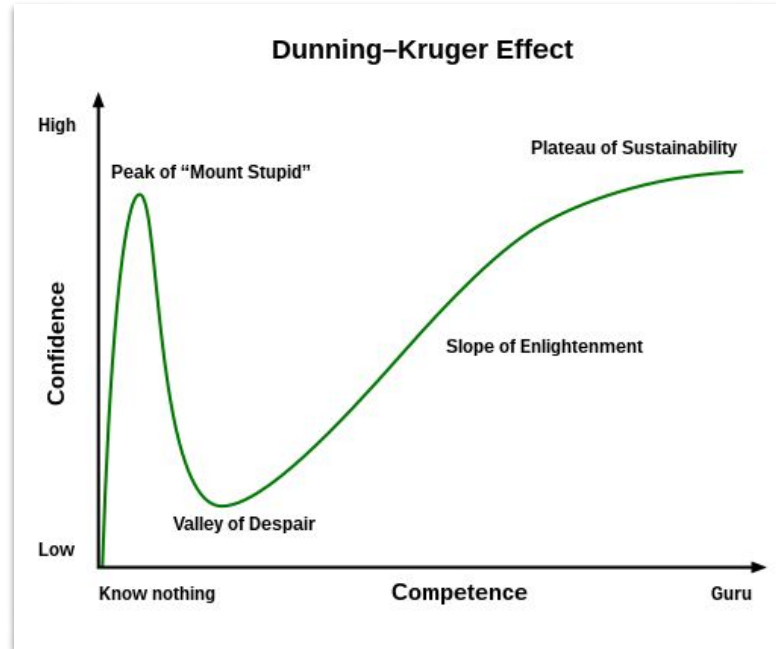
Then, ask it about a topic you do know about.



How did it go?

(raise your hand 🙋 if you'd like to share)

AI is just like us... for now




This can get you into a lot of trouble...

The New York Times

Here's What Happens When Your Lawyer Uses ChatGPT

A lawyer representing a man who sued an airline relied on artificial intelligence to help prepare a court filing. It did not go well.

 Give this article



 1K



Judge Castel said in an order that he had been presented with “an unprecedented circumstance,” a legal submission replete with “bogus judicial decisions, with bogus quotes and bogus internal citations.” He ordered a hearing for June 8 to discuss **potential sanctions**.

Why do models hallucinate?

Model: GPT-4



Why do Large Language Models hallucinate? Explain to a layman. In bullet form. Using less than 250 words.



Sure, here you go:



1. **Training data:** Language models like GPT-4 learn from vast amounts of text data. They don't "understand" information, they simply identify patterns.
2. **Guesswork:** They generate responses by predicting the next word in a sequence. Sometimes they make educated guesses that aren't entirely accurate.
3. **Lack of real-world experience:** Unlike humans, these models don't interact with the world or perceive context beyond the provided text, leading to baseless outputs.
4. **Limited factual accuracy:** They're not updated with real-time information. If something happened after their training cut-off, they wouldn't know.
5. **Complex queries:** They may generate incorrect or imaginative responses when faced with complex or ambiguous queries.
6. **Inference from incomplete information:** Sometimes they fill in gaps with plausible-sounding, but not necessarily accurate information.

Developers have some control over it




How does "temperature" affect hallucination?



In the context of language models like GPT-4, "temperature" is a parameter that controls the randomness of the model's output:



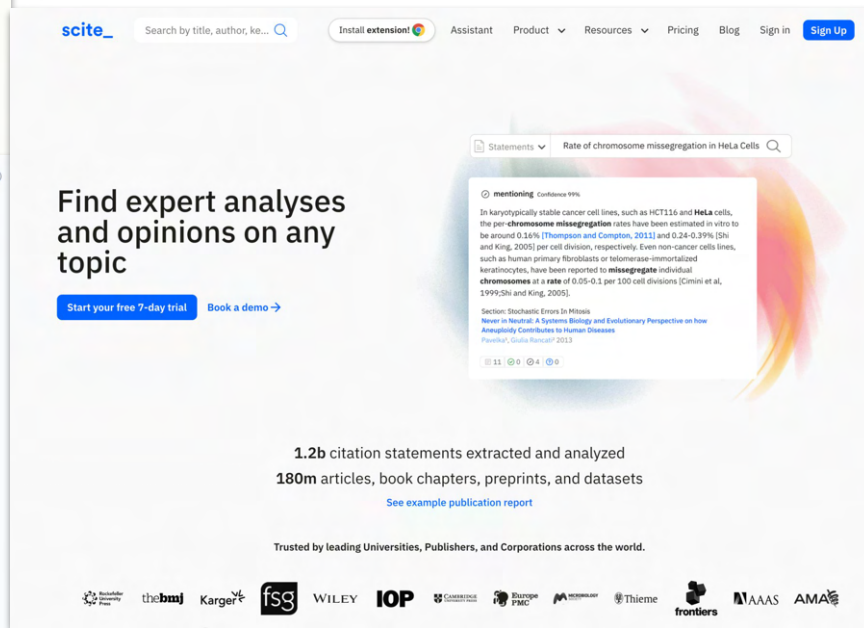
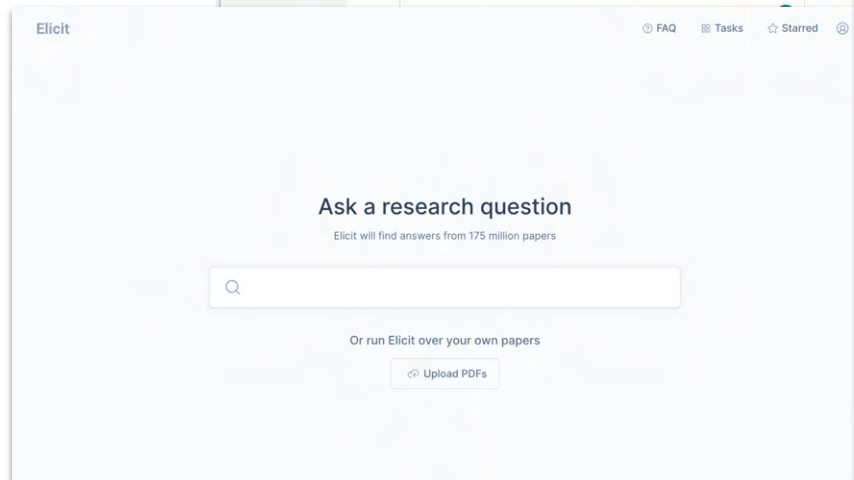
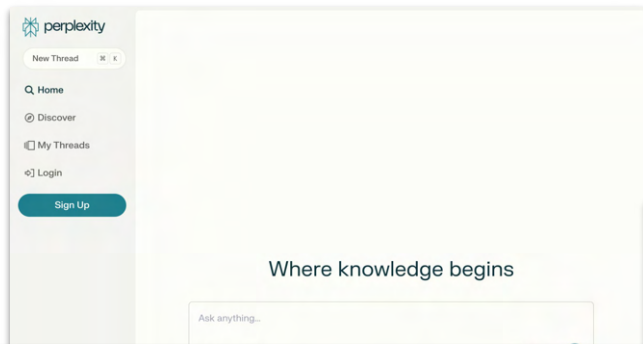
1. **High temperature:** Increases randomness. The model is more likely to pick less probable words, leading to more diverse but potentially more inaccurate or "hallucinatory" responses.
2. **Low temperature:** Decreases randomness. The model strongly favors the most probable next word. This leads to more focused and predictable output, but can be overly conservative and lack creativity.
3. **Balance:** Ideally, a balance is sought where the temperature allows for creative and diverse outputs without veering too far into inaccuracy or hallucination.

 Regenerate response

Send a message...



Specialized solutions make things better



General solution: Treat is as an intern





Summary

- Intelligence is emergent. The biggest experts aren't sure how it will evolve.
- The hard stuff is easier than we thought (and the easy stuff is not that hard, either)
- Research, investment, and experimentation are booming – so we can expect plenty of new breakthroughs soon...
- But winters are possible
- These are powerful tools, but they (still) cannot **replace** high-level experts and cautious processes



What's next?

 **Within 24 hours:** You'll get an email with the sessions' recording, relevant links, and actions items.

 **In 48 hours:** Session 2 about writing, researching, negotiating, and brainstorming with various tools + Generating images and more.

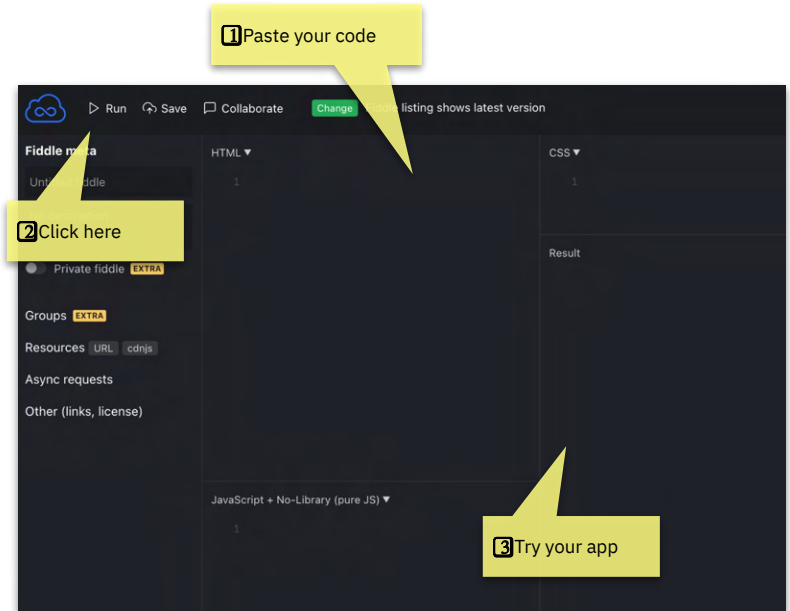
 **Next week:** Working with numbers, data, code, and automation.

Got data?

Send over some clean files you'd like me to analyze in the examples

Home Work: Build a Simple App

1. Generate code with ChatGPT. You can begin by asking it something like: “Generate an html snippet for an app that does....”
2. ChatGPT might give you two (or more) pieces of code for different files. If it does, ask it to provide a single snippet of code.
3. To test your code, go the [HTML / JavaScript Tester](#).
4. Click on “Show your output” at the top of the screen. See and test your result in the lower right corner.
5. Does it work? Want to make it even better? Go back to ChatGPT and instruct it on any changes you’d like.
6. Repeat until you’re happy with your app.
7. Once you’re done, post your code on the [Slack #show-your-work](#) channel.



Thank you!

