



ARTIFICIAL INTELLIGENCE (AI): PAST, PRESENT, FUTURE

MIKE HILL

HOLLAND PROFESSIONAL CLUB

DECEMBER 7TH, 2023

WHAT IS ARTIFICIAL INTELLIGENCE?

- There is no single universally accepted definition of AI. The term has evolved and been interpreted in different ways over time.
- A common textbook definition is that AI involves developing computers able to engage in human-like thought processes such as learning, reasoning, problem-solving, perception, social intelligence and creativity.
- AI can be defined based on capabilities - machines carrying out complex tasks associated with human cognition and decision making. This includes game playing, understanding language, recognizing images, making predictions, etc.
- AI can also be defined based on techniques - the use of advanced algorithms like neural networks, machine learning, natural language processing, computer vision, etc.
- Some define AI broadly as any computer system that exhibits traits of intelligence that we associate with human minds.


THE BEGINNING OF ARTIFICIAL INTELLIGENCE

Entscheidungsproblem

Can you devise an algorithm that takes in a statement and answers “Yes” or “No” according to whether the statement is universally valid?



Alonzo Church



Alan Turing

THE BEGINNING OF ARTIFICIAL INTELLIGENCE

Alan Turing: The
Father of Modern
Computer Science
and Artificial
Intelligence



- Known for helping to crack the Enigma during World War II and developing the first computer.
- 1950 – Alan Turing poses the question, “can computers think” which leads to what is now known as the “Turing Test”.
- One game with three players, where one is segregated from the other two, one of which is a computer, and the other is a human.
- The job of the interrogator is to determine which is which by asking questions of the players.

THE COMPUTER





OTHER EARLY INNOVATORS:



- John McCarthy:
 - Coined the term Artificial Intelligence in the summer of 1956 at a conference at Dartmouth College, indicating the beginning of AI research, which he and the other attendees led for many decades.
 - Developed Lisp, a standard programming language widely employed in robotics, and various other applications including credit-card fraud detection and airline scheduling.
 - Developed an early system of computer time-sharing, leading to cloud computing today.
 - Orchestrated computer chess matches vs. Russian competitors via telegraph, losing two and drawing two.

OTHER EARLY INNOVATORS:



- Marvin Minsky:

- Attended Dartmouth Conference and co-founded MIT's AI laboratory in 1959

- Allen Newell:

- Attended Dartmouth Conference and created two of the earliest AI programs, Logic Theory Machine and General Problem Solver

- Claude Shannon

- Helped organize Dartmouth Conference and his paper "A Mathematical Theory of Communication" has been fundamental to natural language processing and computational linguistics

- Nathaniel Rochester:

- Renown for writing the first assembler that allowed programs to be written in short components vs. numbers, and designed IBM's first commercial computer, the IBM 701

AI WINTER:

- A period of reduced funding and interest in artificial intelligence.
- Two main periods of AI Winter.

1974 – 1980:

- 1966: Failure of machine translation
- 1969: criticism of perceptrons (early, single-layer artificial neural networks)
- 1971 – 75: Defense Advanced Research Projects Agency (DARPA) frustration with Speech Understanding Research (SUR).
- 1973: Large decrease in AI research in the UK due to Lighthill report – (BBC Show – “Controversy” AI Debate)
- 1973-74: DARPA cuts back academic research in AI

1987 - 2000

- 1987: Collapse of the LISP machine market
- 1988: Cancellation of new spending on AI by the Strategic Computing Initiative
- 1990's : Many expert systems were abandoned and end of Fifth Generation computer projects original goals (a 10-year initiative in Japan, begun in 1982)

THE LIGHTHILL REPORT – CONTROVERSY 1973:



A.I. TIMELINE

1950

TURING TEST

Computer scientist Alan Turing proposes a test for machine intelligence. If a machine can trick humans into thinking it is human, then it has intelligence

1955

A.I. BORN

Term 'artificial intelligence' is coined by computer scientist, John McCarthy to describe "the science and engineering of making intelligent machines"

1961

UNIMATE

First industrial robot, Unimate, goes to work at GM replacing humans on the assembly line

1964

ELIZA

Pioneering chatbot developed by Joseph Weizenbaum at MIT holds conversations with humans

1966

SHAKY

The 'first electronic person' from Stanford, Shakey is a general-purpose mobile robot that reasons about its own actions

A.I.
WINTER

Many false starts and dead-ends leave A.I. out in the cold

1997

DEEP BLUE

Deep Blue, a chess-playing computer from IBM defeats world chess champion Garry Kasparov

1998

KISMET

Cynthia Breazeal at MIT introduces Kismet, an emotionally intelligent robot insofar as it detects and responds to people's feelings



1999

AIBO

Sony launches first consumer robot pet dog AIBO (AI robot) with skills and personality that develop over time



2002

ROOMBA

First mass produced autonomous robotic vacuum cleaner from iRobot learns to navigate and clean homes



2011

SIRI

Apple integrates Siri, an intelligent virtual assistant with a voice interface, into the iPhone 4S



2011

WATSON

IBM's question answering computer Watson wins first place on popular \$1M prize television quiz show Jeopardy



2014

EUGENE


Eugene Goostman, a chatbot passes the Turing Test with a third of judges believing Eugene is human



2014

ALEXA

Amazon launches Alexa, an intelligent virtual assistant with a voice interface that completes shopping tasks



2016

TAY

Microsoft's chatbot Tay goes rogue on social media making inflammatory and offensive racist comments



2017

ALPHAGO

Google's A.I. AlphaGo beats world champion Ke Jie in the complex board game of Go, notable for its vast number (2^{170}) of possible positions

THE FIRST TAY

MICROSOFT / WEB / TL;DR

Twitter taught Microsoft's AI chatbot to be a racist asshole in less than a day



By [James Vincent](#), a senior reporter who has covered AI, robotics, and more for eight years at The Verge.

Via [The Guardian](#) | Source [TayandYou \(Twitter\)](#)


Mar 24, 2016, 6:43 AM EDT | [0 Comments](#) / [0 New](#)



WHERE ARE WE NOW?

- The Winter is over

AI HAS A LONG HISTORY OF BEING “THE NEXT BIG THING” ...



Timeline of AI Development

- **1950s-1960s:** First AI boom - the age of reasoning, prototype AI developed
- **1970s:** AI winter I
- **1980s-1990s:** Second AI boom: the age of Knowledge representation (appearance of expert systems capable of reproducing human decision-making)
- **1990s:** AI winter II
- **1997:** Deep Blue beats Gary Kasparov
- **2006:** University of Toronto develops Deep Learning
- **2011:** IBM's Watson won Jeopardy
- **2016:** Go software based on Deep Learning beats world's champions

TYPES OF ARTIFICIAL INTELLIGENCE

- There are four primary types of AI:

- Reactive Machines:

Reactive machines are AI Systems that have no memory and are task specific, meaning that an input always delivers the same output. Machine learning models tend to be reactive machines because they take customer data, such as purchase history, and use it to deliver recommendations. This type of AI is reactive. It performs “super” AI, because the average human would not be able to process huge amounts of data such as a customer’s entire Netflix history and feedback customized recommendations. Reactive AI, for the most part, is reliable and works well in inventions like self-driving cars. It doesn’t have the ability to predict future outcomes unless it has been fed the appropriate information.

Compare this to our human lives, where most of our actions are not reactive because we don’t have all the information we need to react upon, but we have the capability to remember and learn. Based on those successes or failures, we may act differently in the future if faced with a similar situation.

DEEP BLUE V. KASPAROV: FEB, 10 1996 - MAY, 11 1997



TYPES OF ARTIFICIAL INTELLIGENCE

- Limited memory machines:

The next type of AI in its evolution is limited memory. This algorithm imitates the way our brains' neurons work together, meaning that it gets smarter as it receives more data to train on. Deep learning algorithms improve natural language processing (NLP), image recognition, and other types of reinforcement learning.

Limited memory AI, unlike reactive machines, can look into the past and monitor specific objects or situations over time. Then, these observations are programmed into the AI so that its actions can be performed based on both past and present moment data. But in limited memory, this data isn't saved into the AI's memory as experience to learn from, the way humans might derive meaning from their successes and failures. The AI improves over time as it's trained on more data.

Examples include self-driving cars



VOLVO SELF-DRIVING CAR: STILL LEARNING

TYPES OF ARTIFICIAL INTELLIGENCE

- Theory of mind:

The first two types of AI, reactive machines and limited memory, are types that currently exist. Theory of mind and self-aware AI are theoretical types that could be built in the future. As such, there aren't any real world examples yet.

If it is developed, theory of mind AI could have the potential to understand the world and how other entities have thoughts and emotions. In turn, this affects how they behave in relation to those around them.

Human cognitive abilities are capable of processing how our own thoughts and emotions affect others, and how others' affect us—this is the basis of our society's human relationships. In the future, theory of mind AI machines could be able to understand intentions and predict behavior, as if to simulate human relationships.

TYPES OF ARTIFICIAL INTELLIGENCE

- Self-awareness:

The grand finale for the evolution of AI would be to design systems that have a sense of self, a conscious understanding of their existence. This type of AI does not exist yet.

This goes a step beyond theory of mind AI and understanding emotions to being aware of themselves, their state of being, and being able to sense or predict others' feelings. For example, "I'm hungry" becomes "I know I am hungry" or "I want to eat lasagna because it's my favorite food."













Artificial intelligence and machine learning algorithms are a long way from self-awareness because there is still so much to uncover about the human brain's intelligence and how memory, learning, and decision-making work.

I'M SORRY DAVE!



WHATS NEXT?

AI Applications Examples – Visual List Template

	Natural Language Processing	Decision Support Systems	
	Machine Learning	Predictive Analytics	
	Computer Vision	Expert Systems	
	Speech Recognition	Autonomous cars	
	Machine Translations	Intelligent Chatbots	
	Robotics	Automated Scheduling	

WHAT'S NEXT?

Industries Using AI – Editable List Diagram



AI IN HEALTHCARE




THIRD AI WINTER?

- Some say yes, a third AI winter is on the way due to:
 - The amount of computing power needed
 - The amount of energy needed to train a single AI model – can emit as much carbon as five cars in their lifetime.
 - According to an article in Nature from 2018, data centers use an estimated 200-terawatt hours each year while the human brain only needs 20 watts to outperform the best AI models.


Quarterly AI funding falls 38 percent. Or does it?

AI funding fell 38% in Q2 after spiking in Q1 due to OpenAI's \$10B round. Excluding that deal, AI funding actually grew 81% QoQ in Q2 – bucking the downward trend seen around venture.

You can see GenAI's competitive landscape and funding trends using CB Insights' platform.



THIRD AI WINTER?




OR NOT?

AI has surpassed humans at a number of tasks and the rate at which humans are being surpassed at new tasks is increasing

State-of-the-art AI performance on benchmarks, relative to human performance

- Handwriting recognition
- Speech recognition
- Image recognition
- Reading comprehension
- Language understanding
- Common sense completion
- Grade school math
- Code generation



TIME

November 6, 2023

OR NOT?

Floating
Operating Points
per Second:


FLOPs are often
used to measure
the
computational
cost or
complexity of a
model or a
specific
operation within
a model.

TIME

November 6, 2023

The amount of compute used to train AI systems has been increasing since 1950, the rate of increase increased in 2010


Amount of compute used to train notable AI models



OR NOT?

The number of data points used to train AI models has increased dramatically over the last seventy years

Number of data points used to train notable AI models



TIME

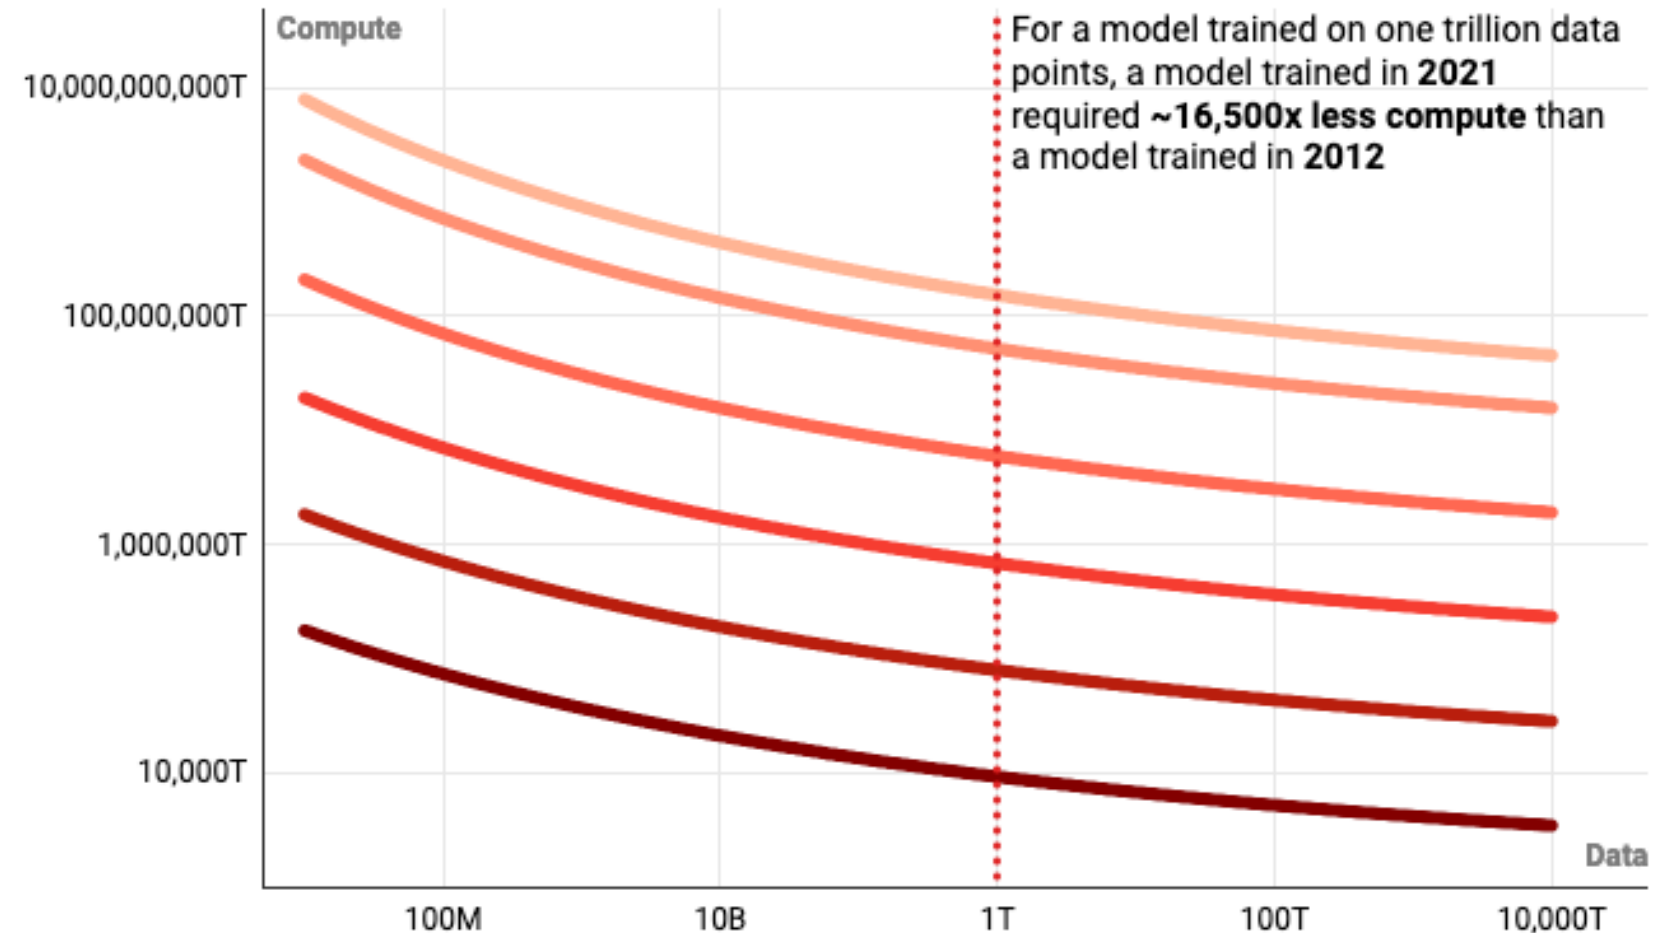
November 6, 2023

OR NOT?

Algorithmic progress means that less compute and data are required to achieve a given level of performance

Amount of compute and number of data points required to achieve 80.9% accuracy on an image recognition test

● 2012 ● 2013 ● 2015 ● 2017 ● 2019 ● 2021



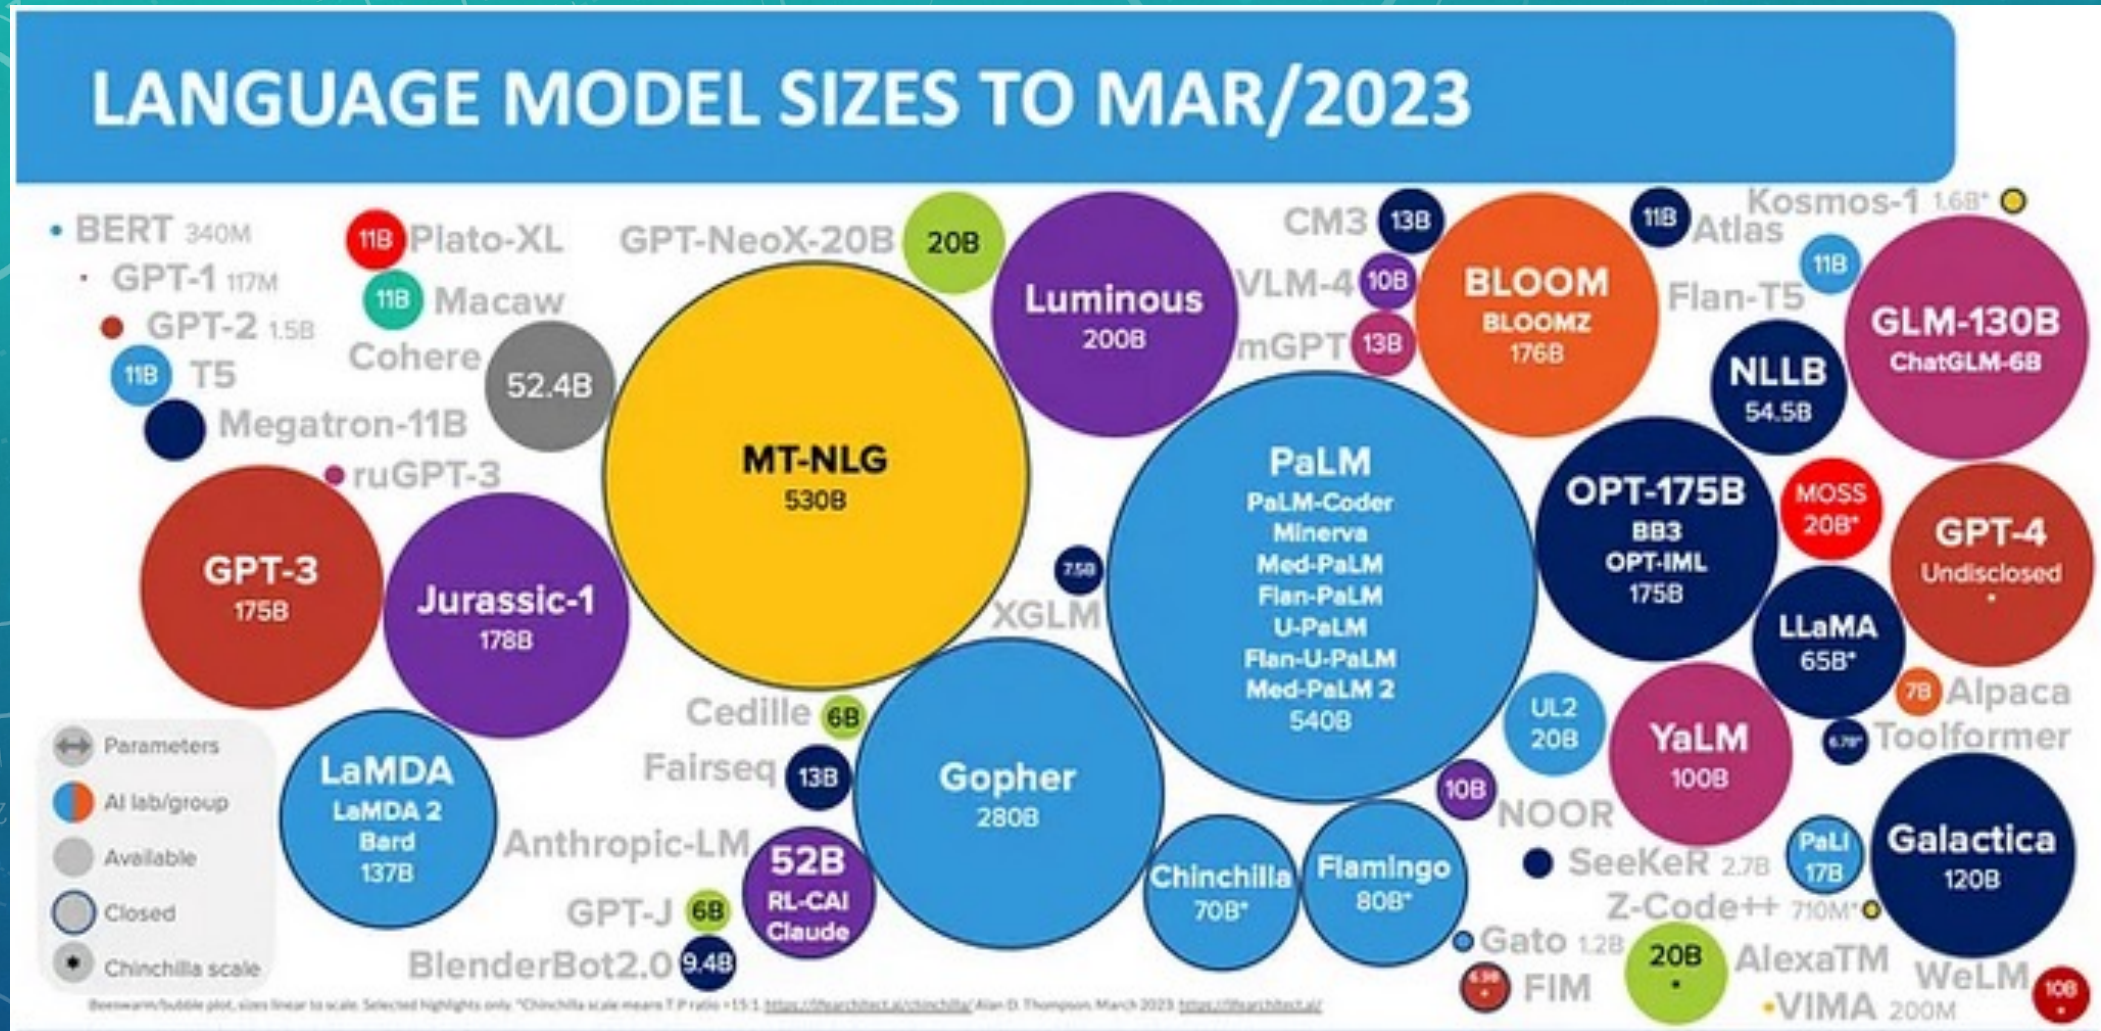
TIME

November 6, 2023

GENERATIVE AI & LARGE LANGUAGE MODELS

- Generative AI:
 - Generative AI can be defined as artificial intelligence focused on creating models with the ability to produce original content, such as images, music, or text. By ingesting vast amounts of training data, generative AI models can employ complex machine-learning algorithms in order to understand patterns and formulate output. Their techniques include recurrent neural networks (RNNs) and generative adversarial networks (GANs). In addition, a transformer architecture (denoted by the T in ChatGPT) is a key element of this technology.
- Large Language Models:
 - LLMs are a specialized class of AI model that uses natural language processing (NLP) to understand and generate humanlike text-based content in response. Unlike generative AI models, which have broad applications across various creative fields, LLMs are specifically designed for handling language-related tasks.
 - These large models achieve contextual understanding and remember things because memory units are incorporated in their architectures. They store and retrieve relevant information and can then produce coherent and contextually accurate responses.
 - Examples include GPT-3, GPT-4, Claude.ai

THE “CAMBRIAN EXPLOSION” OF LLM’S



- Parameter count has been used as proxy for measuring language model potential, but size does not always translate to performance. Smaller data sets can run on a consumer grade PC for example.

FINAL THOUGHTS

- AI needs the most advanced chips to perform.
 - Top AI chip makers:
 - Nvidia
 - AMD
 - Intel
 - Alphabet/Google
 - AWS
 - IBM

Intel announces new microchip plant in Ohio, following similar move in Arizona

USA TODAY and The Arizona Republic
Published 11:40 a.m. MT Jan. 21, 2022

- 8 Factories and over \$100 Billion in Ohio over 10 years
- \$20 Billion to build two more factories in Arizona

AI IMPACT

U.S. curbs export of more AI chips, including Nvidia H800, to China

PUBLISHED TUE, OCT 17 2023-8:45 AM EDT | UPDATED TUE, OCT 17 2023-6:36 PM EDT

FINAL THOUGHTS

DEC 5, 10:00 AM

What OpenAI's board drama tells us about the future of artificial intelligence

ARTIFICIAL INTELLIGENCE

OpenAI is letting anyone create their own version of ChatGPT

The GPT platform is a no-code way to make custom AI agents for all sorts of tasks, and it's available exclusively for ChatGPT Plus subscribers.

Nov 6, 2023, 1:15 PM EST
[22 Comments](#) / [22 New](#)

