



Demystifying Generative AI

Jennifer Martinez

Technology Enthusiast
Explainer of Things

Note: Some slides contain layers and animations that are lost in PDF form

OpenAI buzz (Nov 2022+)

How the tech behind ChatGPT could change the world—an

[How the tech behind ChatGPT could change the world—an updated episode from our archive | The Economist](#)

Microsoft Bets Big on the Creator of ChatGPT in Race to Dominate A.I.

[Microsoft Bets Big on the Creator of ChatGPT in Race to Dominate A.I. - The New York Times \(nytimes.com\)](#)

OpenAI's new DALL-E model draws anything — but bigger, better and faster than before

[dall-e | TechCrunch](#)

ChatGPT has given everyone a glimpse at AI's astounding progress

[OpenAI's ChatGPT is a fascinating glimpse into the scary power of AI - Vox](#)

ChatGPT Is a Tipping Point for AI

by Ethan Mollick

December 14, 2022

[ChatGPT Is a Tipping Point for AI \(hbr.org\)](#)

A.I. Can Now Write Its Own Computer Code. That's Good News for Humans.

[A.I. Can Now Write Its Own Computer Code. That's Good News for Humans. - The New York Times \(nytimes.com\)](#)

This AI chatbot is dominating social media with its frighteningly good essays

Microsoft's year of new product releases

Microsoft is bringing ChatGPT technology to Word, Excel and Outlook

By [Samantha Murphy Kelly](#), CNN Business

Updated 11:49 AM EDT, Thu March 16, 2023

MICROSOFT / TECH / ARTIFICIAL INTELLIGENCE

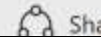
Microsoft announces Copilot: the AI-powered future of Office documents / Copilot is a modern AI assistant that will help Microsoft 365 users create Office documents.

By [Tom Warren](#)

Mar 16, 2023, 11:06 AM EDT | [89 Comments](#) / [89 New](#)

Power Platform is leading a new era of AI-generated low-code app development

March 16, 2023 • 3 min read



Microsoft-Backed OpenAI Starts Release of Powerful AI Known as GPT-4

By [Reuters](#) | March 14, 2023, at 1:36 p.m.

Microsoft releases Copilot for Viva, as it continues to roll out generative AI to apps

Sharon Goldman
[@sharongoldman](#)

April 20, 2023 8:00 AM



Everyone can now use the ChatGPT-powered Bing – here's how

By [Mark Wilson](#) last updated March 17, 2023

Microsoft's new open-source project makes it easy to build generative AI into your applications

Latest News | Published: May 2nd, 2023 - [Jenna Sargent Barron](#)



Flood of Concerning News

OpenAI gives in to Italy's data privacy demands, ending ChatGPT ban

Italy requested a PSA showing users how to opt out of ChatGPT data sharing.

ASHLEY BELANGER · 5/1/2023, 3:17 PM

"ChatGPT is available again to our users in Italy," OpenAI's statement said. "We are excited to welcome them back, and we remain dedicated to protecting their privacy."

White House notches AI agreement with top tech firms

TECHNOLOGY
The White House's move comes while Congress hammers out what legally binding guardrails to put on AI. ... Satya Nadella and Sundar Pichai, as ... executives of OpenAI ...

When A.I. Falls in Love

The Times asked GPT-3 to tell us a little about itself and its romantic life.

Experts are warning AI could lead to human extinction. Are we taking it seriously enough?



Analysis by Oliver Darcy, CNN

Updated 3:56 AM EDT, Wed May 31, 2023

HOUSTON CHRONICLE · 14d

Elon Musk forms his own artificial intelligence company xAI to compete with OpenAI's ChatGPT

OpenAI, the nonprofit AI research company that released the widely popular AI language model ChatGPT in November 2022, was ...

Tech



The Verge

+ Follow

Sar
rel
use
me

Microsoft's new Bing Chat Enterprise offers better privacy for businesses

Story by Jon Porter · Jul 18

Old News

BLOG

Transformer: A Novel Neural Network Architecture for Language Understanding

THURSDAY, AUGUST 31, 2017

Posted by Jakob Uszkoreit, Software Engineer, Natural Language Understanding

For First Time, Microsoft Integrating GPT-3 Into Its Software

May 25, 2021 by Todd R. Weiss

[For First Time, Microsoft Integrating GPT-3 Into Its Software \(enterpriseai.news\)](#)

Artificial Intelligence · Latest News · Top List

Top 10 GPT-3 Powered Applications to Know in 2022



Disha Sinha

April 15, 2022 · 3 mins read

GPT-3: We're at the very beginning of a new app ecosystem

Dattaraj Rao, Persistent System

February 27, 2021 10:20 AM



GPT-3 powers the next generation of apps

Over 300 applications are delivering GPT-3-powered search, conversation, text completion, and other advanced AI features through our API.

AI content platform Jasper raises \$125M at a \$1.5B valuation

Kyle Wiggers @kyle_l_wiggers / 9:00 AM EDT • October 18, 2022



Comment

[ChatGPT vs. Jasper: How It's Different from Jasper Chat](#)

GitHub's New 'Copilot' Feature Is A Game-Changer If You're a Developer

Jan · July 22, 2022 · Leave a Reply

Generative AI Landscape

Platforms

Creative Organization

Synthetic Data Generation

Vector Search & Curation

MLOPS

Code Generation

Customer Supprt

General writing & Content editing

Marketing and Sales copy

Text & Summarization

Image Editing

Ad Collateral

Design

Text to Images

Text to Speech

Audio Editing

Avatar and Dubbing

Music Generation

Summarization & Insights

Video

Generation and Animation

Editing

- Data and Decision
- MLOps
- Text and Tools
- Images
- Audio and Video

Learn from those Blazing the Trail

FORBES > INNOVATION > VENTURE CAPITAL

The CEO Of Jasper On How He Built One Of The Hottest AI Startups And What's Next For Generative AI

Joanne Chen Contributor

Joanne Chen is a General Partner at Foundation Capital.



AI Overlords is a series of conversations with leading entrepreneurs and executives building in the AI ecosystem.

Generative AI, or artificial intelligence used not just for repetitive tasks but for creative work, has become a bit of a buzzword ever since ChatGPT burst onto the scene late last year. Though lots of people are still using it as an experiment, giving the model ridiculous prompts, or turning them into a [Fatal Attraction-style chatbot](#), Jasper has done something different.

The company uses AI to help businesses write marketing content seamlessly, in the tone of voice customers relate to and find familiar, a jillion times faster than a human being can write that same content. The company's relentless focus on solving real customer problems is one of the main reasons I decided to lead Jasper's first seed round (full disclosure), and it's why I think Jasper will be the first great generative AI company.

Were you surprised, when you did get access to the ChatGPT model, at how it was being used?

I think I was surprised that people weren't taking it more seriously, or didn't know what to really do with it.

We had gotten into the Slack group with people that were using GPT-3 and I thought everybody in there was going to have these amazing ideas of what this could do. But as I was reading through the messages, I thought, *Nobody here understands how valuable this is. Everyone thinks it's just a toy.*

What are your questions?

(ALL questions are fair game)

- What is the relationship between OpenAI and Microsoft?
- What is the difference between ChatGPT, GPT-3.5, GPT-4, LLMs, etc?
- What exactly is ChatGPT and all the ways the word is now used?
- Understand the difference in terminology (Foundation Models, LLMs, Generative AI, GPT, etc)
- Are all of the software pointing to the same underlying model?
- Is there a difference between Microsoft's model and OpenAI's?
- What are the risks inherent in the LLM when starting to use it?
- How can you trust the output if there are hallucinations?
- What if it starts to replicate itself and get out of control of humans?

Speaker Perspective

On the AI breakthrough:

- This really is a breakthrough, not an increment
- The domain introduces many new concepts
- Our vocabulary has stretched and is evolving to describe the technology and how to use it
 - This causes some confusion
- AI is real, capable and can be put-to-use today
- *Very few* people know how to use it well
- Many “AI Tourists” are flooding the field.
- Sensationalized media stories are just that, but help us understand what to govern
- What goes in, where it goes, what comes out matters.
- Responsible AI is a must
- AI requires a human in the middle at this stage.
- AI will eventually change the life of every knowledge worker and consumer on the planet
- This is different than anything we’ve seen before, and we all continue learning
- Identify the experts. Follow the experts. Use discernment in non-research papers headlines.

Emerging Vocabulary

Transformers

Prompts

Reinforcement

Foundation Model

Prompt Engineering

Learning

Large Language Model

Hallucinations

Human Reinforcement
Learning (HRL)

Generative AI

Predictive

Large Multimodal
Model (LMM)

Grounding

Weights

Responsible AI

Parameters

Synthetic Media
(watermarks)

Fine Tuning

Token

Embeddings

Chunking

Vector Database

General AI

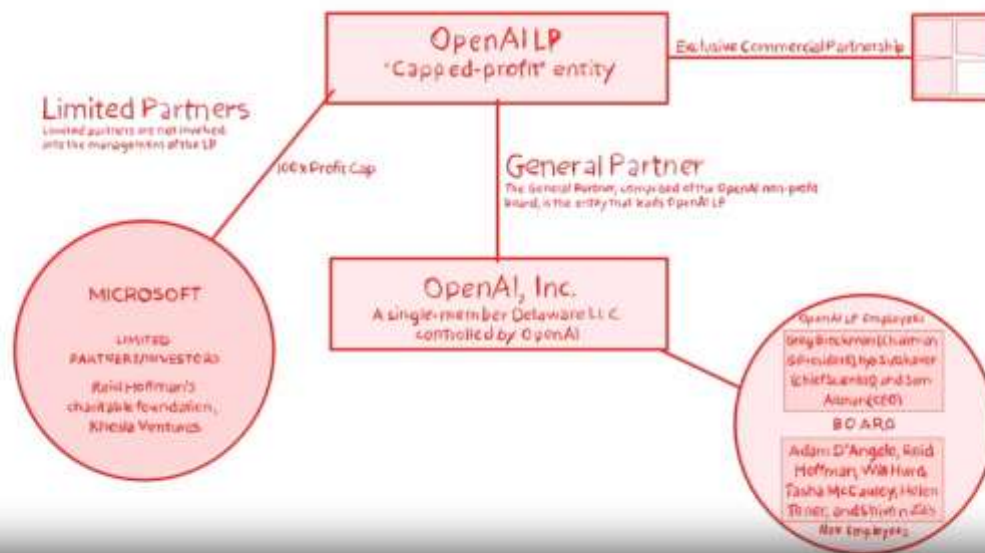
Who is OpenAI? - Want to go deeper...

FourWeekMBA

The Leading Source of Insights On Business Model Strategy & Tech Business Models

OpenAI-Microsoft Partnership Explained
[Business / By Gennaro Cuofano / February 2, 2023](#)

Understanding the Microsoft-OpenAI deal



[OpenAI-Microsoft Partnership Explained - FourWeekMBA](#)

Putting the “G” (Generative) in GPT

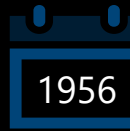
Artificial Intelligence

Machine Learning

Deep Learning

Generative AI

Pursuit of General



1956

Artificial Intelligence

the field of computer science that seeks to create intelligent machines that can replicate or exceed human intelligence



1997

Machine Learning

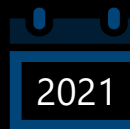
subset of AI that enables machines to learn from existing data and improve upon that data to make decisions or predictions



2017

Deep Learning

a machine learning technique in which layers of neural networks are used to process data and make decisions



2021

Generative AI

Create new written, visual, and auditory content given prompts or existing data.

Key Microsoft AI breakthroughs

2016

● Object & speech recognition
Human parity

2018

● Reading comprehension & machine translation
Human parity

● General language understanding
Human parity

2020

● Turing-NLG language models

● First AI supercomputer

● Exclusive license for OpenAI GPT-3 models

● Image captioning
Human parity

2021

● Natural Language Understanding
Human parity

● Commonsense Question Answering
Human parity

● Azure OpenAI Service preview

2022

● GitHub Copilot general availability

Microsoft AI innovation in 2023

January ●

- Azure OpenAI Service becomes generally available
- Microsoft extends our partnership with OpenAI

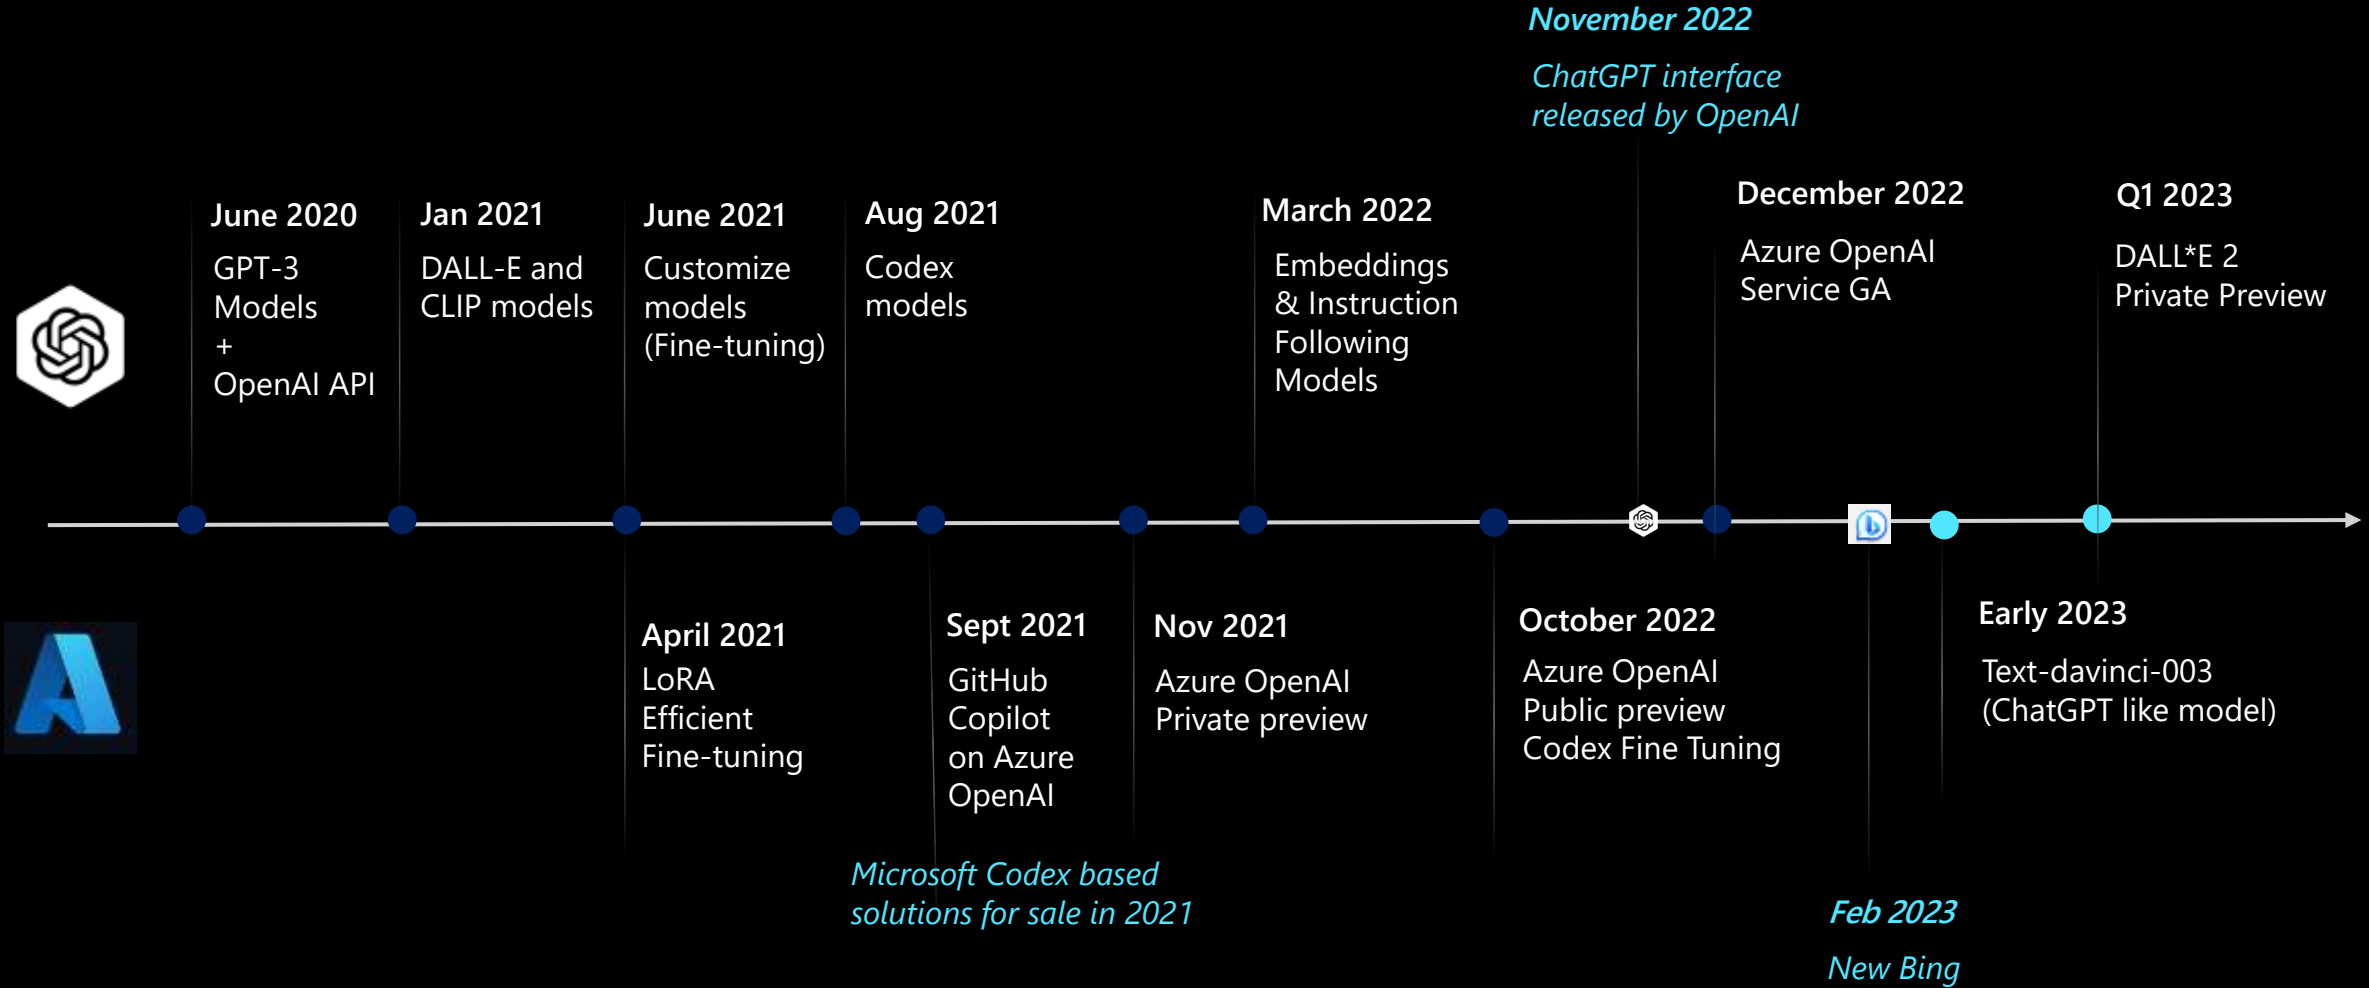
February ●

- Teams Premium with GPT becomes generally available
- Viva Sales adds generative AI capabilities
- Microsoft announces the new Bing and Edge
- Microsoft announces Bing momentum and Skype Copilot
- Windows 11 updates bring AI-powered Bing to the taskbar

March ●

- LinkedIn introduces collaborative articles
- Microsoft introduces Dynamics 365 Copilot
- Florence comes to Azure Cognitive Services for Vision
- Azure OpenAI Service adds ChatGPT capabilities
- Microsoft announces powerful new virtual machines
- LinkedIn adds new AI-powered capabilities
- Microsoft introduces Microsoft 365 Copilot
- Microsoft introduces Copilot for Power Platform
- Nuance introduces DAX Express
- Azure OpenAI Service adds GPT-4
- Bing Image Creator comes to the new Bing
- GitHub introduces GitHub Copilot X
- Microsoft introduces Microsoft Security Copilot

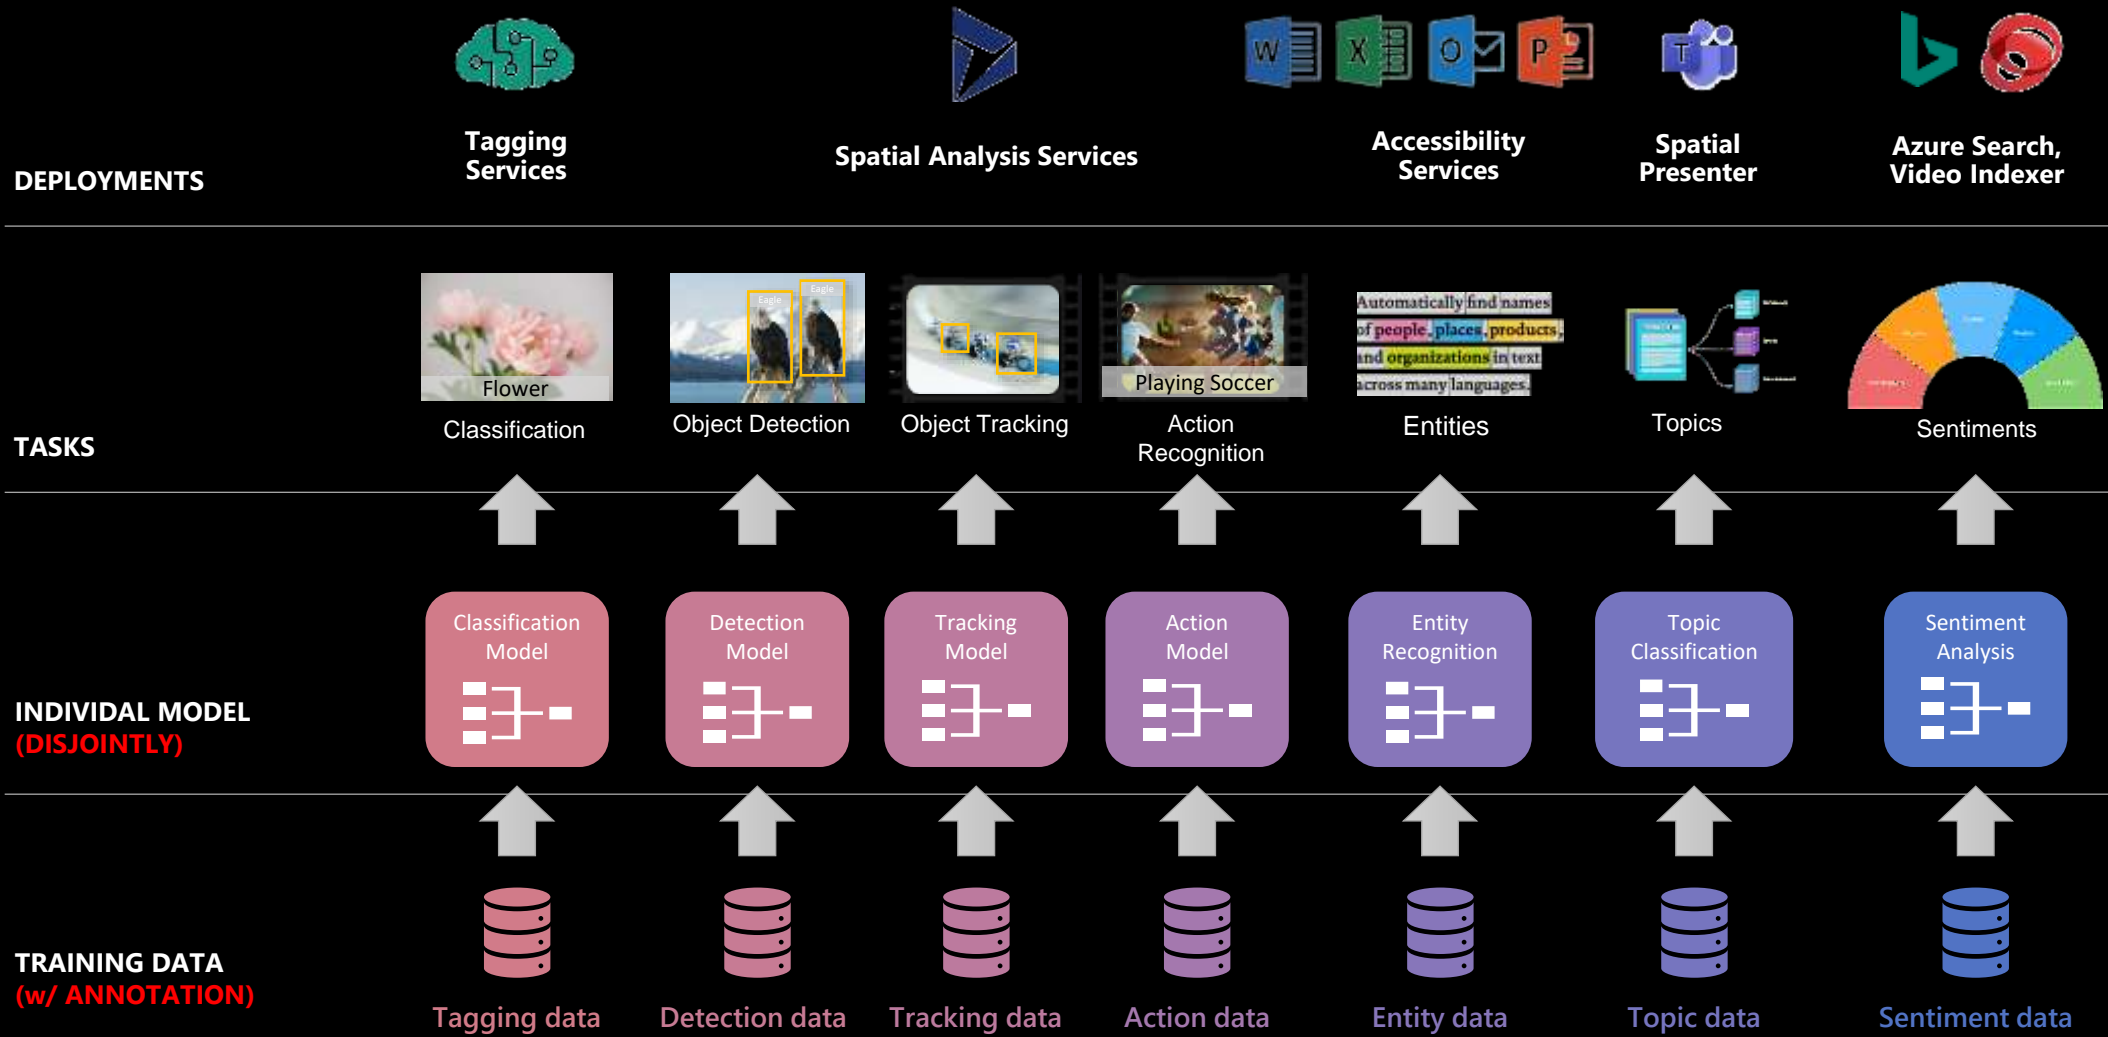
Models and platform roadmap



Putting the “P” (Pretrained) in GPT

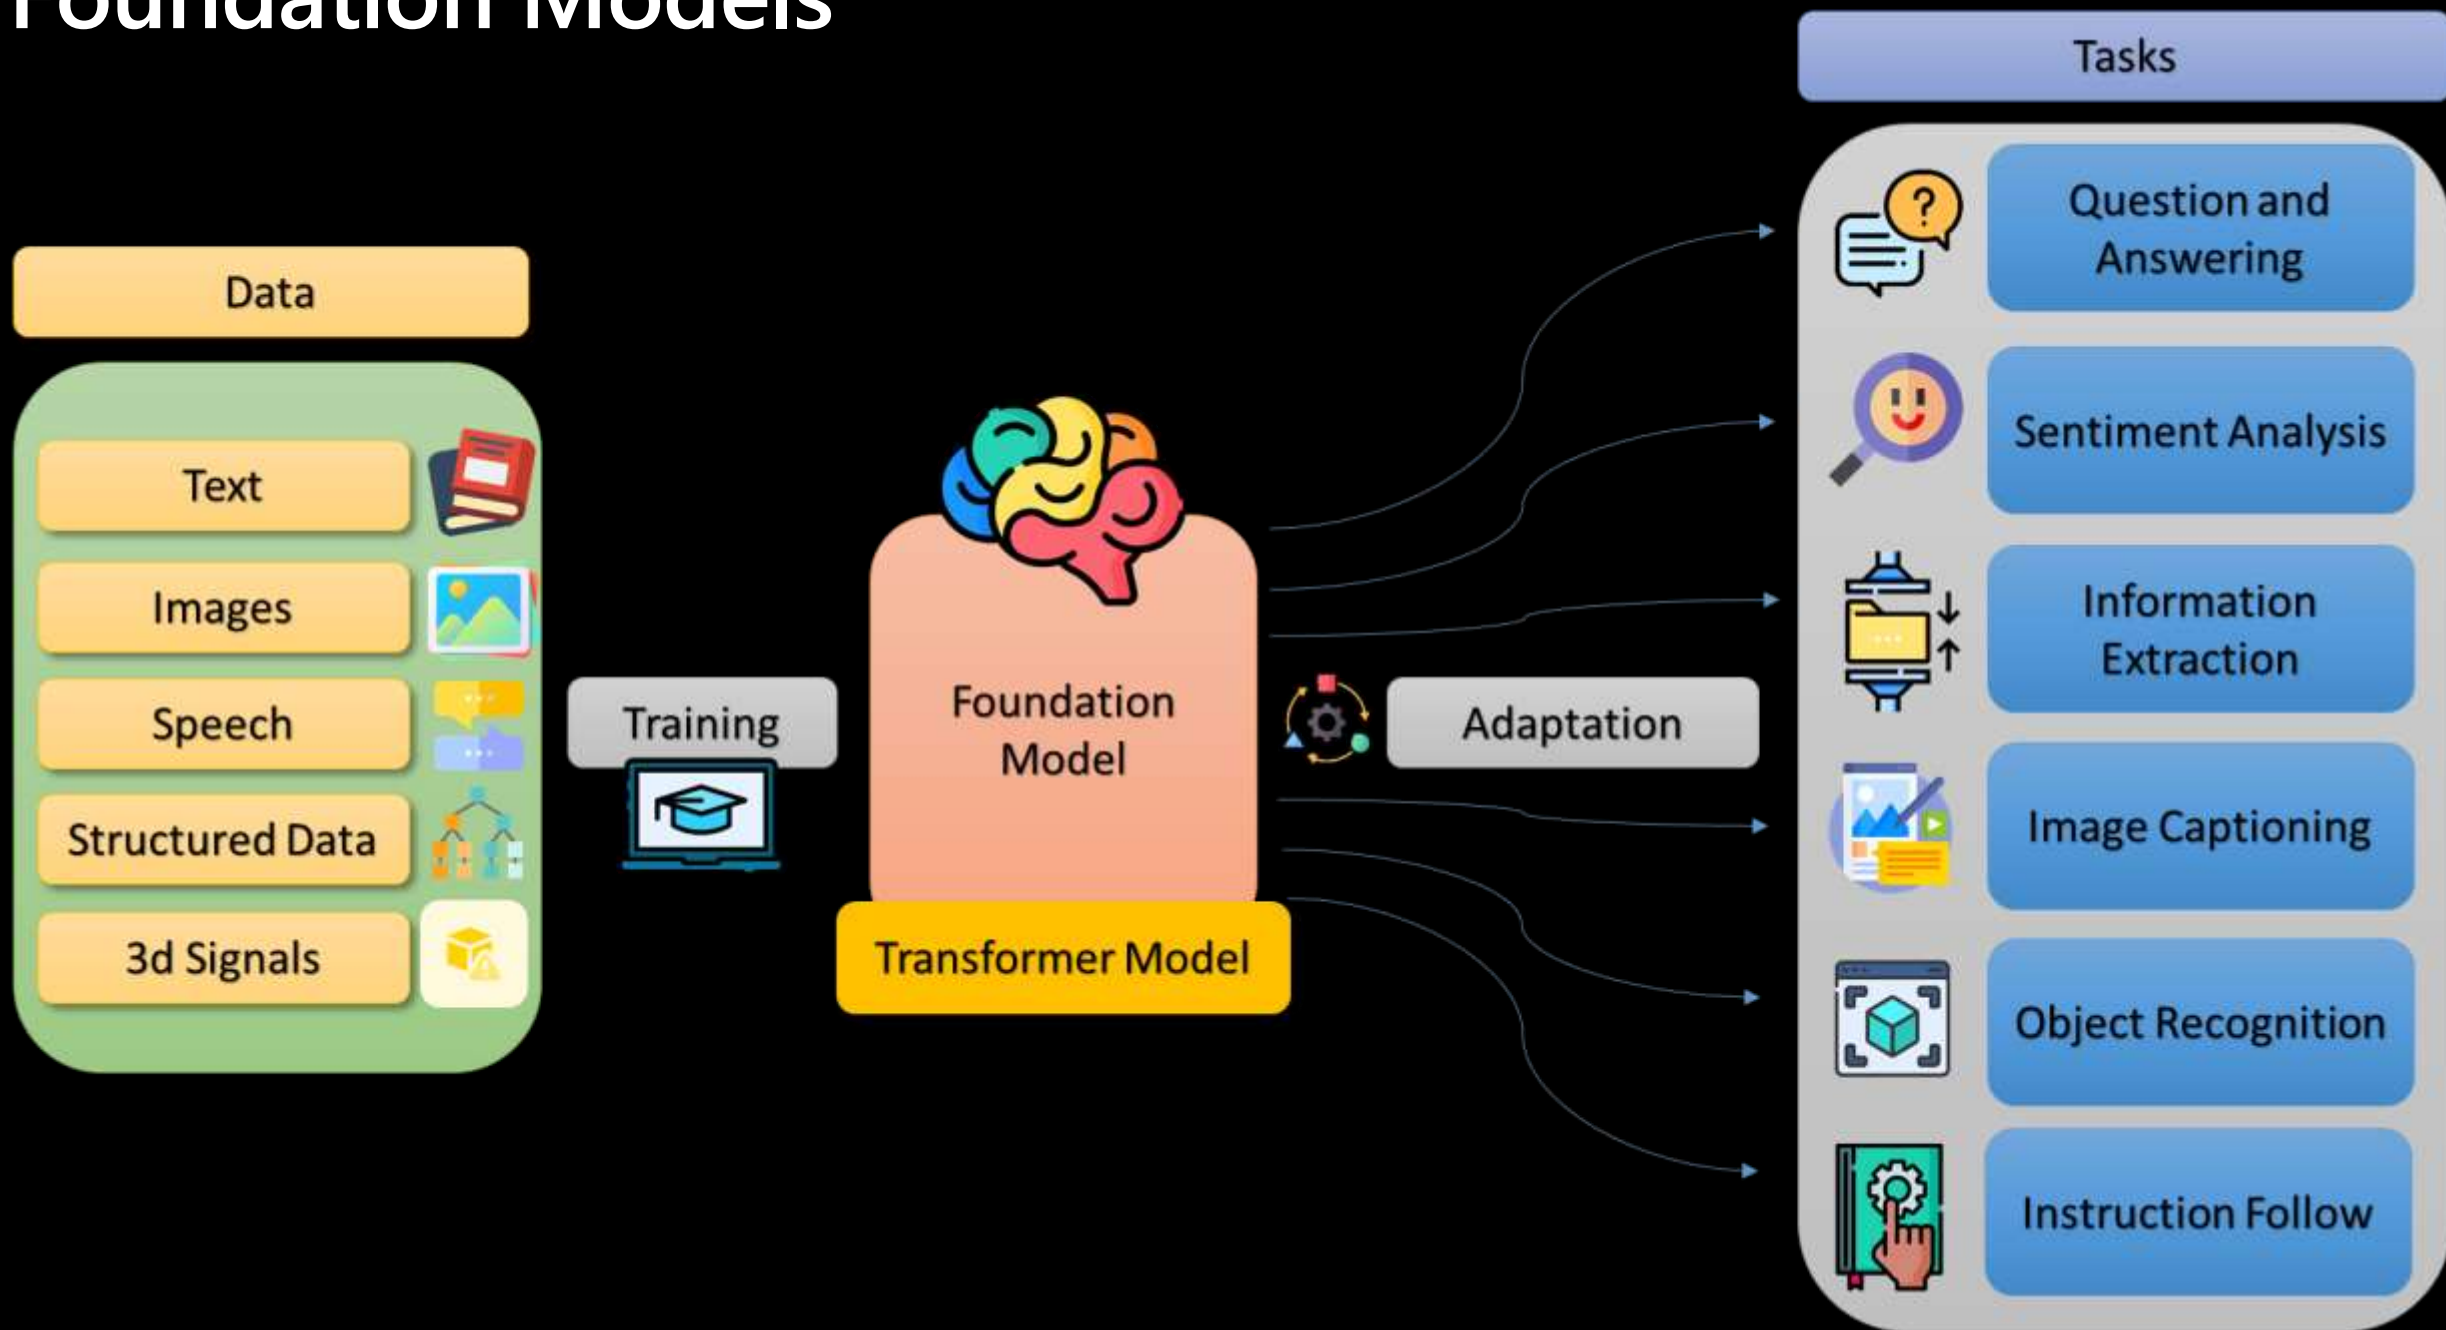
Traditional model development

High cost & slow deployment - Each service is trained disjointly



Putting the “T” (Transformers) in GPT

Foundation Models



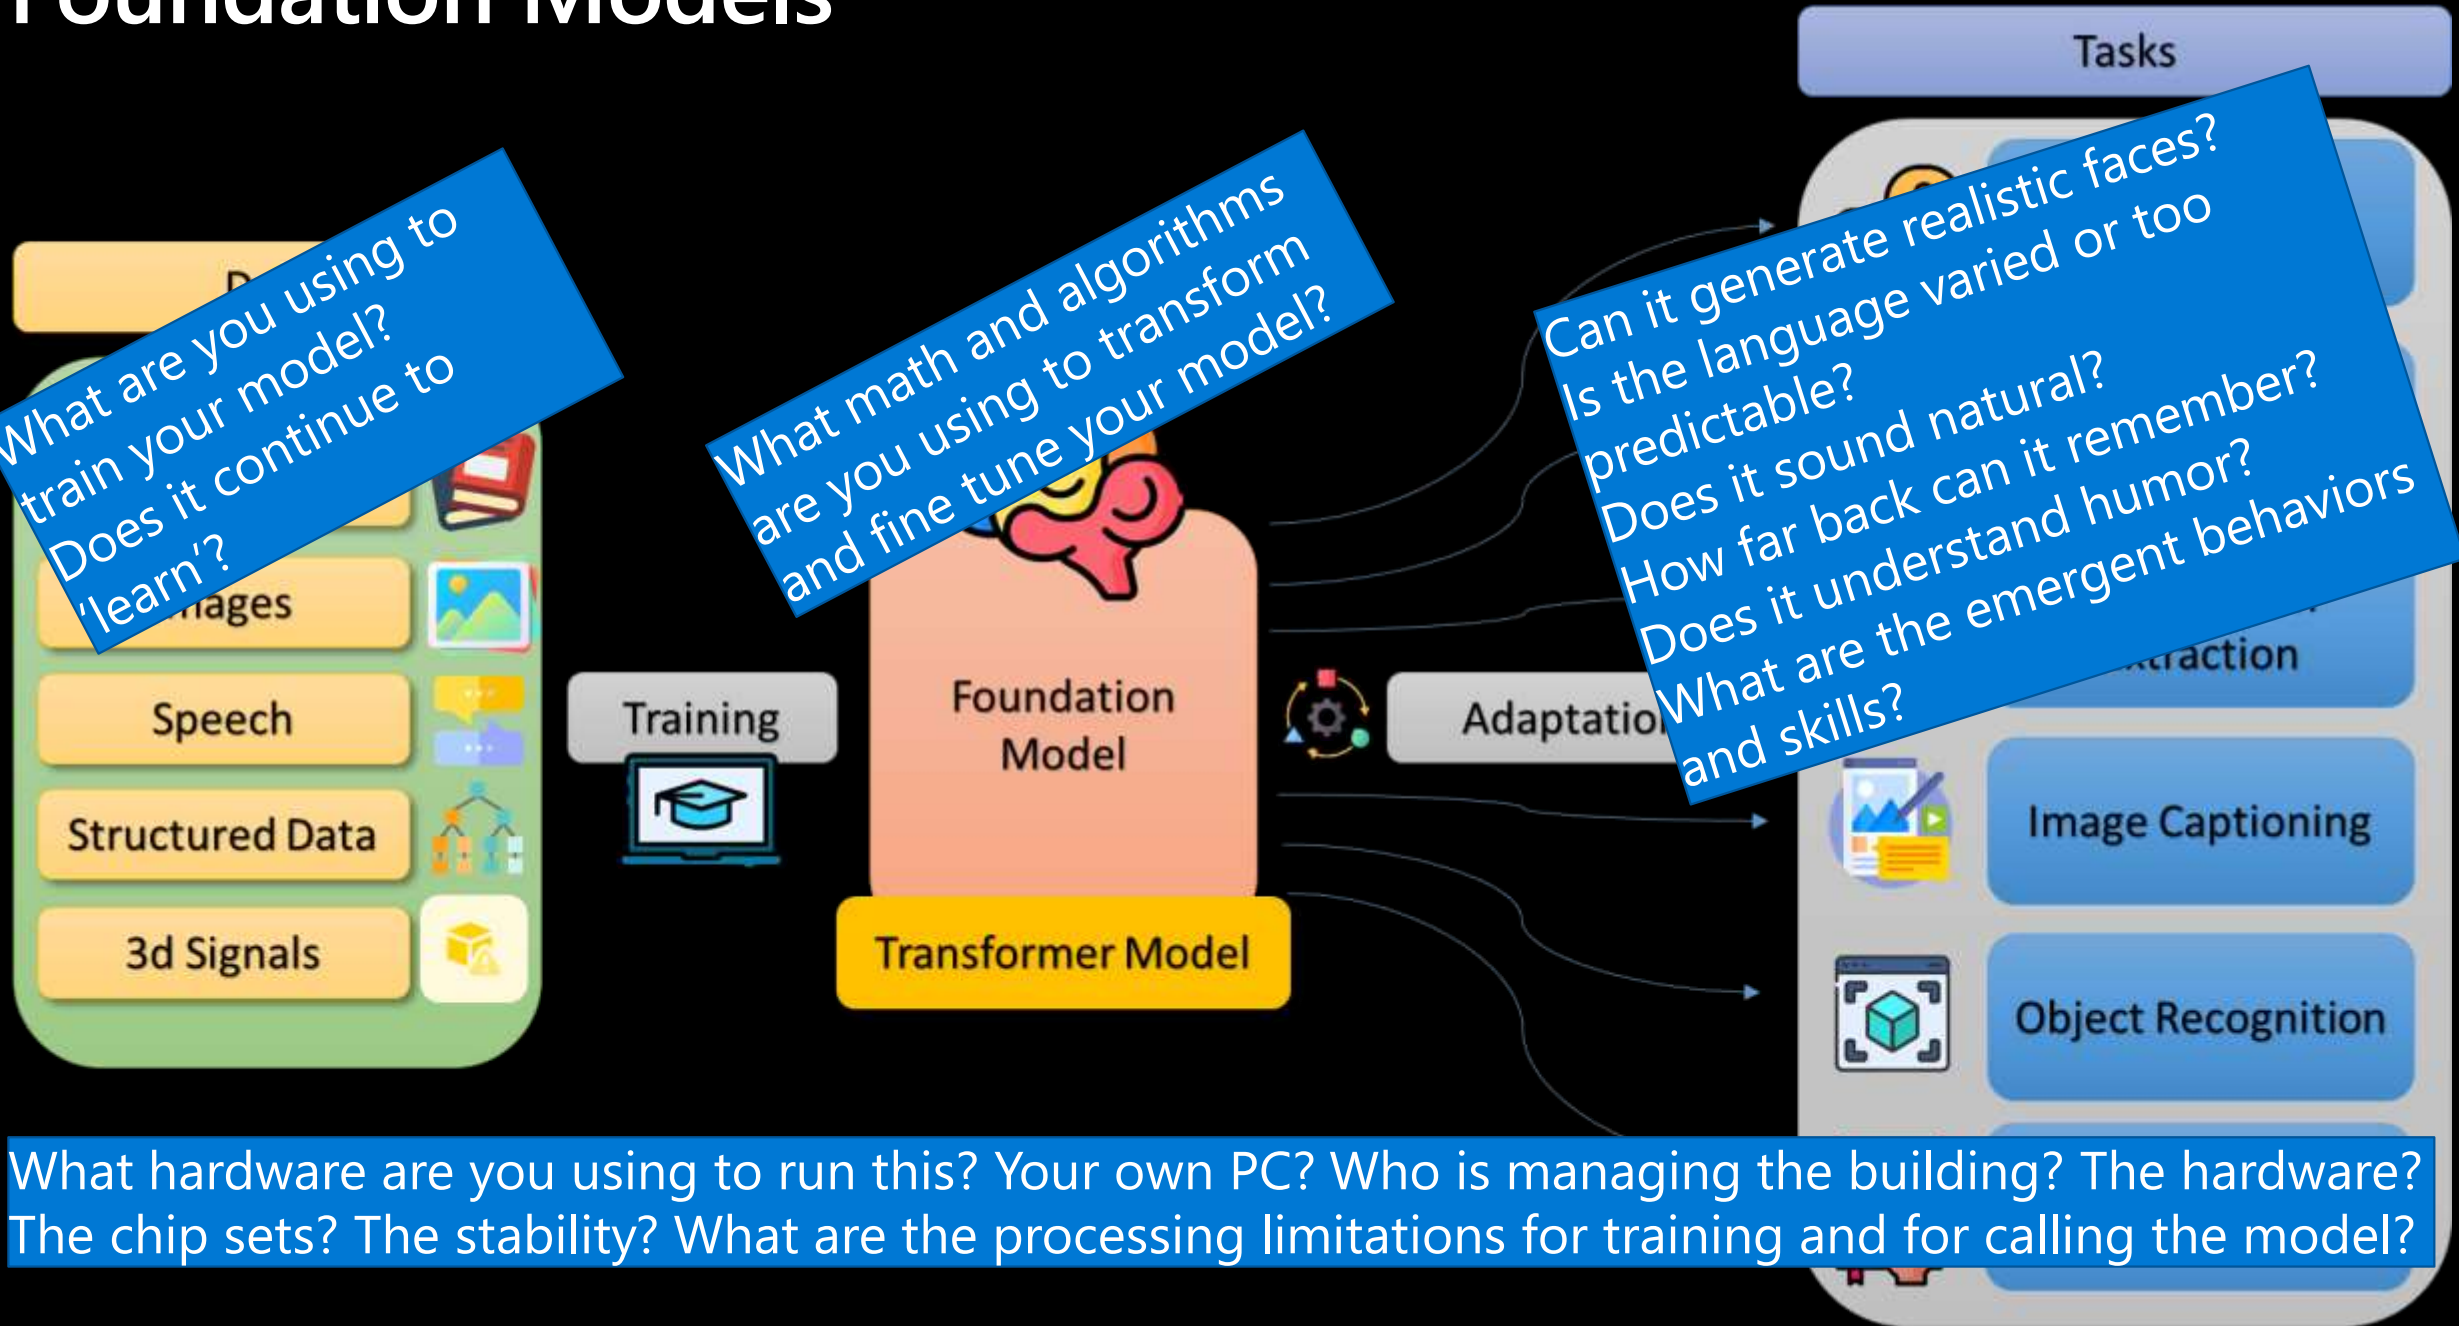
Foundation Models

What are you using to train your model?
Does it continue to 'learn'?

What math and algorithms are you using to transform and fine tune your model?

Can it generate realistic faces?
Is the language varied or too predictable?
Does it sound natural?
How far back can it remember?
Does it understand humor?
What are the emergent behaviors and skills?

What hardware are you using to run this? Your own PC? Who is managing the building? The hardware? The chip sets? The stability? What are the processing limitations for training and for calling the model?



Large Language Models and Embeddings



Trained on a massive corpus of language.

Language is converted into vectors.

Lion – Cat – Dandelion

In simplistic terms, the models understand language by looking at a word or phrase, its surrounding words or phrases, and the position in the body of text.

It assigns mathematical values that can be plotted on a multidimensional plane based on this word/phrase analysis, or vectors.

[Introducing text and code embeddings \(openai.com\)](https://openai.com)

Learn on your own



Large Language Models from scratch (Steve Seltz | Graphics in 5 minutes:

<https://youtu.be/lnA9DMvHtfl>

Large Language Models: Part 2 (Graphics in 5 minutes)

<https://youtu.be/YDiSFS-yHwk>

Transparency in the Research Process

- Generative Pretraining from Pixels (openai.com)

Abstract

Inspired by progress in unsupervised representation learning for natural language, we examine whether similar models can learn useful representations for images. We train a sequence Transformer to auto-regressively predict pixels, without incorporating knowledge of the 2D input structure. Despite training on low-resolution ImageNet without labels, we find that a GPT-2 scale model learns strong image representations as measured by linear probing, fine-tuning, and low-data classification. On CIFAR-10, we achieve 96.3% accuracy with a linear probe, outperforming a supervised Wide ResNet, and 99.0% accuracy with full fine-tuning, matching the top supervised pre-trained models. An even larger model trained on a mixture of ImageNet and web images is competitive with self-supervised benchmarks on ImageNet, achieving 72.0% top-1 accuracy on a linear probe of our features.

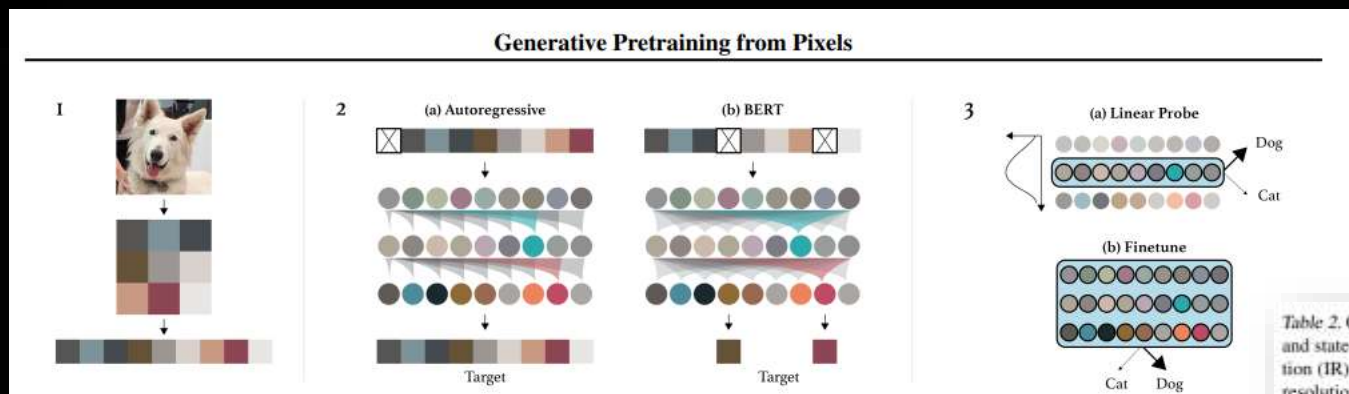


Figure 1. An overview of our approach. First, we pre-process raw images by resizing to a low resolution and reshaping into a 1D sequence. We then chose one of two pre-training objectives, auto-regressive next pixel prediction or masked pixel prediction. Finally, we evaluate the representations learned by these objectives with linear probes or fine-tuning.

representations that significantly improve the state of the art in low-resolution unsupervised representation learning settings.

2.1. Pre-training

Given an unlabeled dataset X consisting of high dimensional data $x = (x_1, \dots, x_n)$, we can pick a permutation π of the set $[1, n]$ and model the density $p(x)$ auto-regressively as follows:

$$p(x) = \prod_{i=1}^n p(x_{\pi_i} | x_{\pi_1}, \dots, x_{\pi_{i-1}}, \theta)$$

When working with images, we pick the identity permutation $\pi_i = i$ for $1 \leq i \leq n$, also known as raster order. We train our model by minimizing the negative log-likelihood of the data:

$$L_{AR} = \mathbb{E}_{x \sim X} [-\log p(x)]$$

Table 2. Comparing linear probe accuracies between our models and state-of-the-art self-supervised models. A blank input resolution (IR) corresponds to a model working at standard ImageNet resolution. We report the best performing configuration for each contrastive method, finding that our models achieve comparable performance.

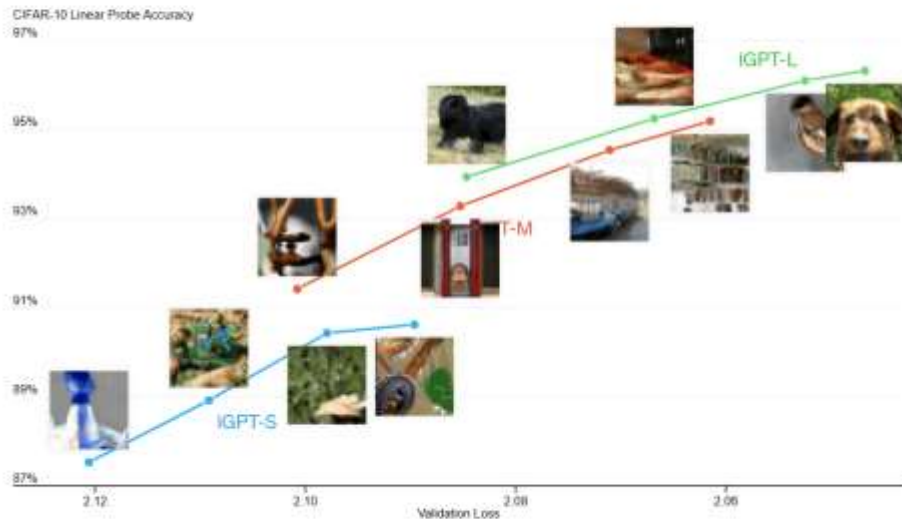
Method	IR	Params (M)	Features	Acc
Rotation	orig.	86	8192	55.4
iGPT-L	32 ² ·3	1362	1536	60.3
BigBiGAN	orig.	86	8192	61.3
iGPT-L	48 ² ·3	1362	1536	65.2
AMDIM	orig.	626	8192	68.1
MoCo	orig.	375	8192	68.6
iGPT-XL	64 ² ·3	6801	3072	68.7
SimCLR	orig.	24	2048	69.3
CPC v2	orig.	303	8192	71.5
iGPT-XL	64 ² ·3	6801	15360	72.0
SimCLR	orig.	375	8192	76.5

M achieves 54.5% accuracy and iGPT-S achieves 41.9% accuracy.

How Images progress through different models, training

[Image GPT \(openai.com\)](https://openai.com)

Our next result establishes the link between generative performance and feature quality. We find that both increasing the scale of our models and training for more iterations result in better generative performance, which directly translates into better feature quality.



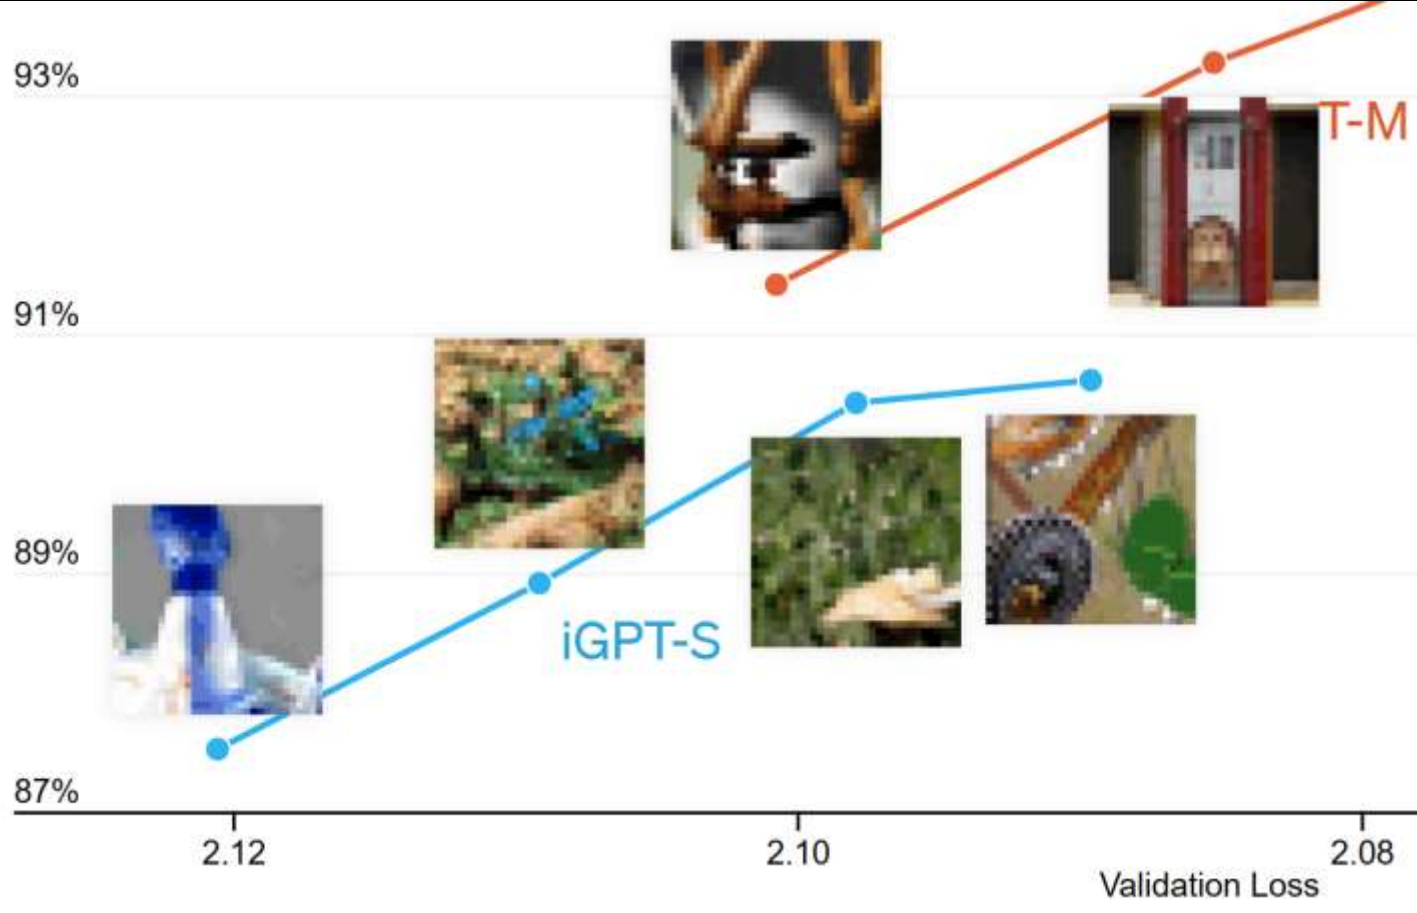
Each line tracks a model throughout generative pre-training; the dotted markers denote checkpoints at steps 131K, 262K, 524K, and 1000K. The positive slopes suggest a link between improved generative performance and improved feature quality. Larger models also produce better features than smaller models. iGPT-XL is not included because it was trained on a different dataset.

AI generated images

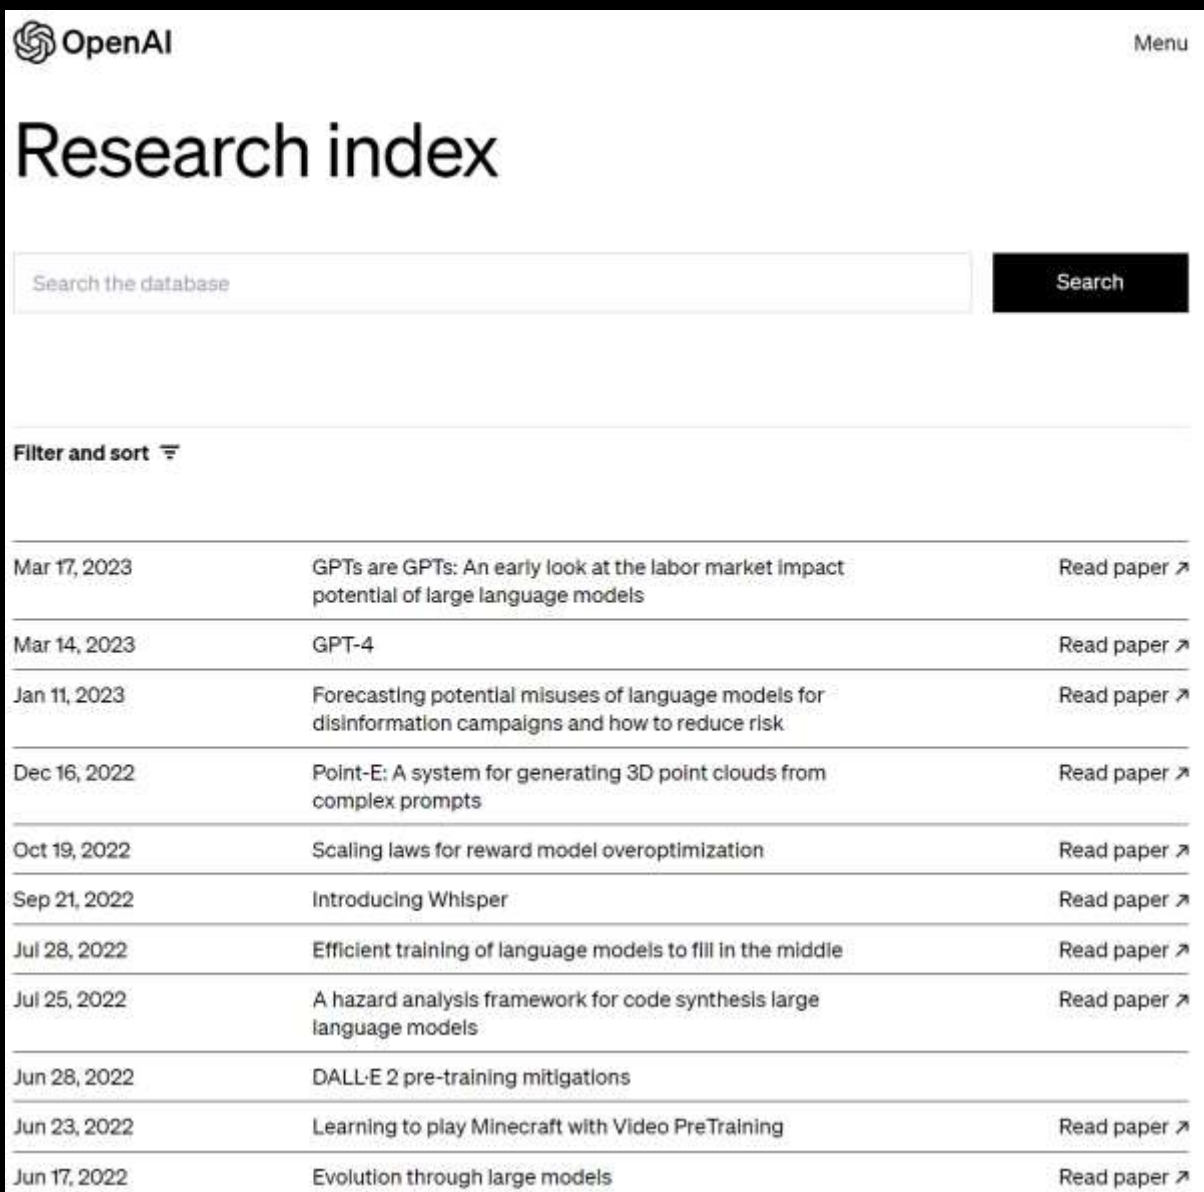


We find that DALL·E is able to apply several different kinds of image transformations to photos of teapots, with varying degrees of reliability. Aside from being able to modify the color of the teapot (e.g., "colored blue") or its pattern (e.g., "with stripes"), DALL·E can also render text (e.g., "with 'gpt' written on it") and map the letters onto the curved surface of the teapot in a plausible way. With much less reliability, it can also draw the teapot in a smaller size (for the "tiny" option) and in a broken state (for the "broken" option).

Early model Generative AI (GPT 2). Close, but not quite



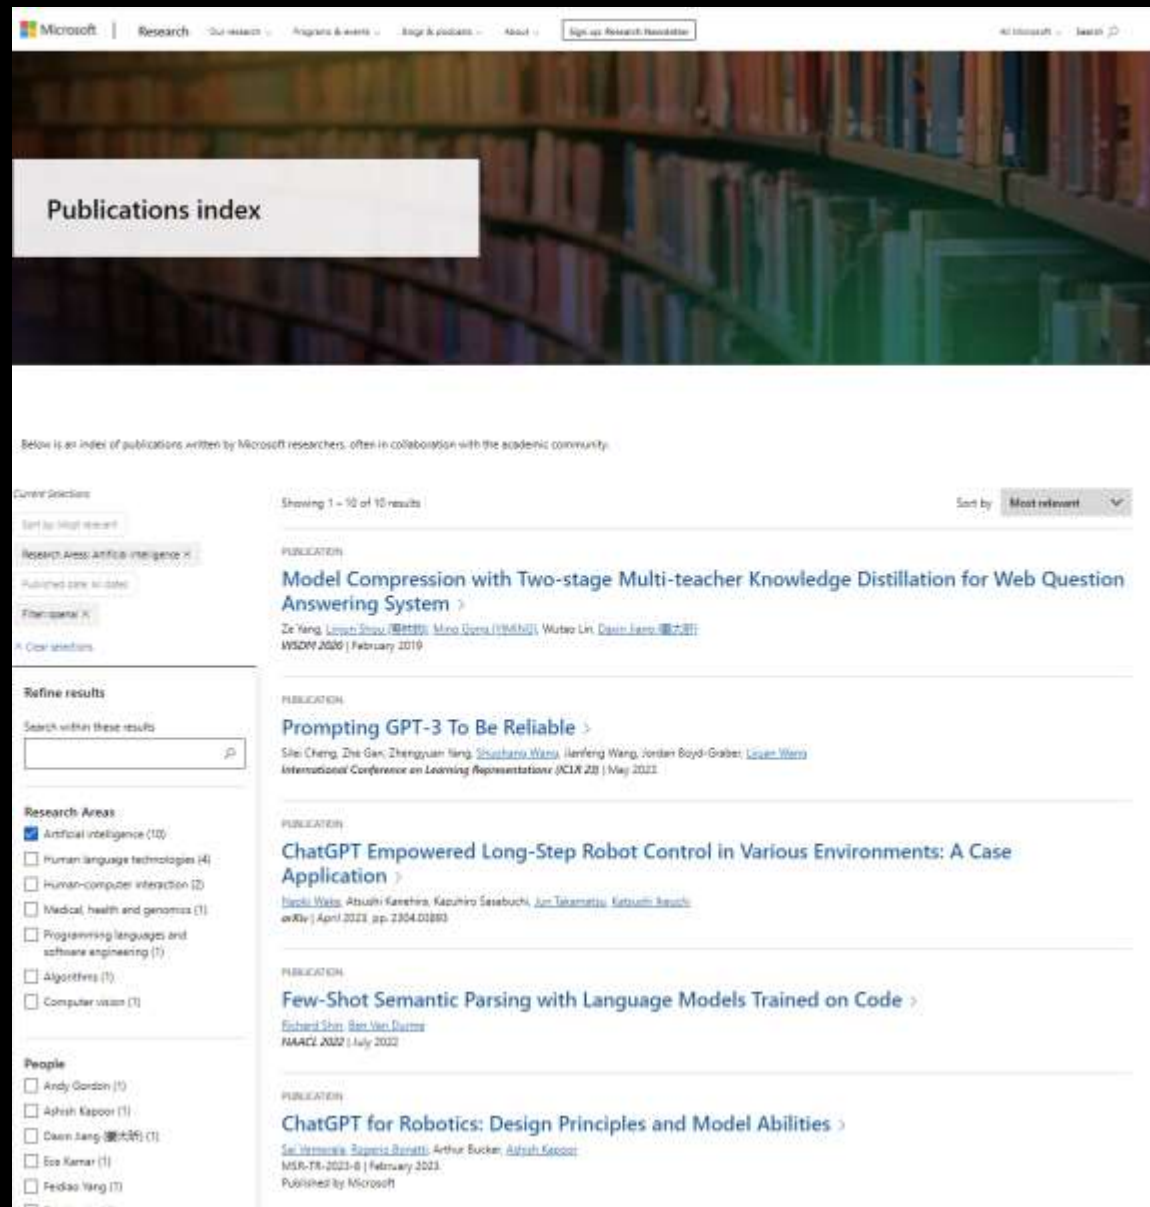
OpenAI Research Blog



The screenshot shows the OpenAI Research index page. At the top left is the OpenAI logo, and at the top right is a 'Menu' link. The main heading is 'Research index'. Below it is a search bar with the placeholder text 'Search the database' and a black 'Search' button. Underneath the search bar is a 'Filter and sort' section with a dropdown arrow. The main content is a list of research papers, each with a date, a title, and a 'Read paper' link with an external arrow.

Date	Title	Action
Mar 17, 2023	GPTs are GPTs: An early look at the labor market impact potential of large language models	Read paper
Mar 14, 2023	GPT-4	Read paper
Jan 11, 2023	Forecasting potential misuses of language models for disinformation campaigns and how to reduce risk	Read paper
Dec 16, 2022	Point-E: A system for generating 3D point clouds from complex prompts	Read paper
Oct 19, 2022	Scaling laws for reward model overoptimization	Read paper
Sep 21, 2022	Introducing Whisper	Read paper
Jul 28, 2022	Efficient training of language models to fill in the middle	Read paper
Jul 25, 2022	A hazard analysis framework for code synthesis large language models	Read paper
Jun 28, 2022	DALL-E 2 pre-training mitigations	Read paper
Jun 23, 2022	Learning to play Minecraft with Video PreTraining	Read paper
Jun 17, 2022	Evolution through large models	Read paper

Microsoft Research Blog



The screenshot shows the Microsoft Research Publications index page. At the top is the Microsoft logo and navigation links for 'Research', 'Our research', 'Programs & events', 'Engage & partners', and 'About'. There is also a 'Sign up for Research Newsletter' button and a search icon. The main heading is 'Publications index'. Below it is a large image of a library. The text below the image reads: 'Below is an index of publications written by Microsoft researchers, often in collaboration with the academic community.' There are filters for 'Current sections' (Research Area: Artificial intelligence, Published date: Academic, Filter: All) and 'Refine results' (Search within these results). The 'Research Areas' section is expanded to show 'Artificial intelligence (10)', 'Human language technologies (4)', 'Human-computer interaction (2)', 'Medical, health and genomics (1)', 'Programming languages and software engineering (1)', 'Algorithms (1)', and 'Computer vision (1)'. The 'People' section lists 'Andy Gordon (1)', 'Aphish Kapoor (1)', 'Dawn Song (微软) (1)', 'Eso Kamae (1)', and 'Fedor Yang (1)'. The main content shows a list of publications, with the first one being 'Model Compression with Two-stage Multi-teacher Knowledge Distillation for Web Question Answering System' by Ze Yang, Lijun Zhou, Miao Qian, Wuhao Lin, and Qianqian Li, published in WSDM 2020 in February 2019. Other publications listed include 'Prompting GPT-3 To Be Reliable', 'ChatGPT Empowered Long-Step Robot Control in Various Environments: A Case Application', 'Few-Shot Semantic Parsing with Language Models Trained on Code', and 'ChatGPT for Robotics: Design Principles and Model Abilities'.

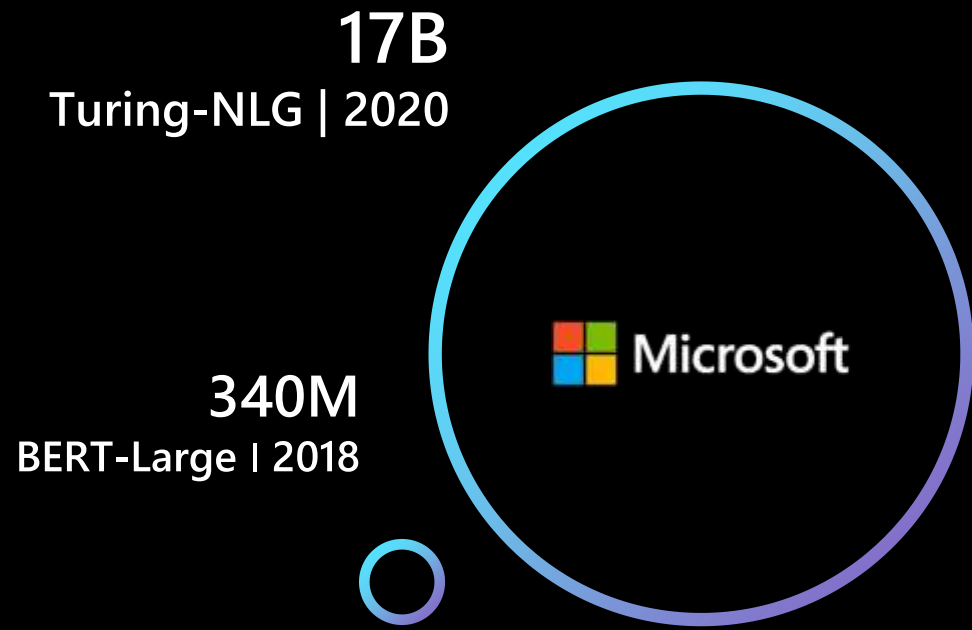
Large Language Models (LLMs)

Size, speed, training material matters

Does everyone have one?

There is no going back

Large Language Foundation models



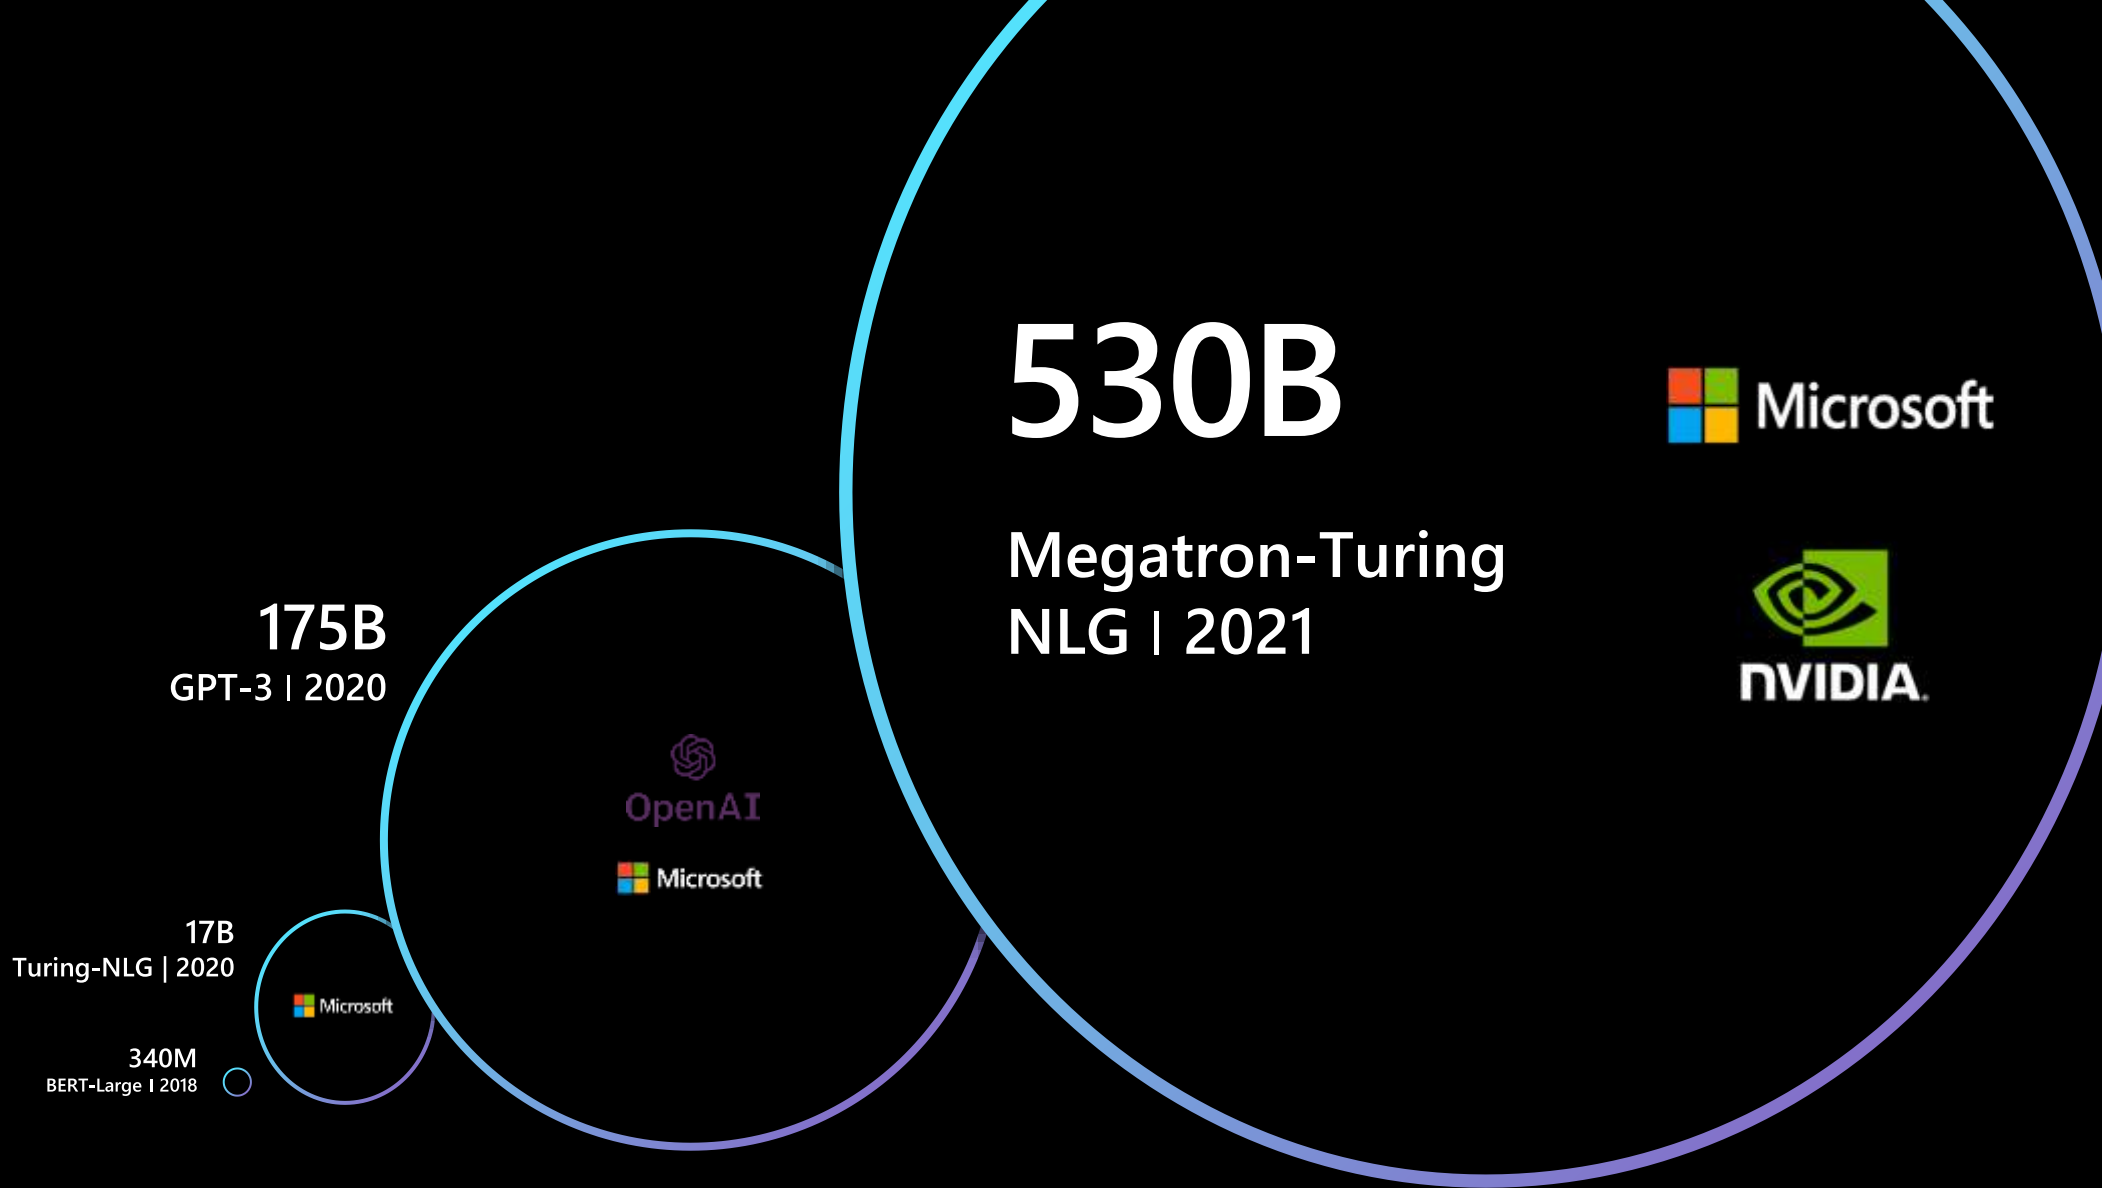
175B
GPT-3 | 2020

17B
Turing-NLG | 2020

340M
BERT-Large | 2018



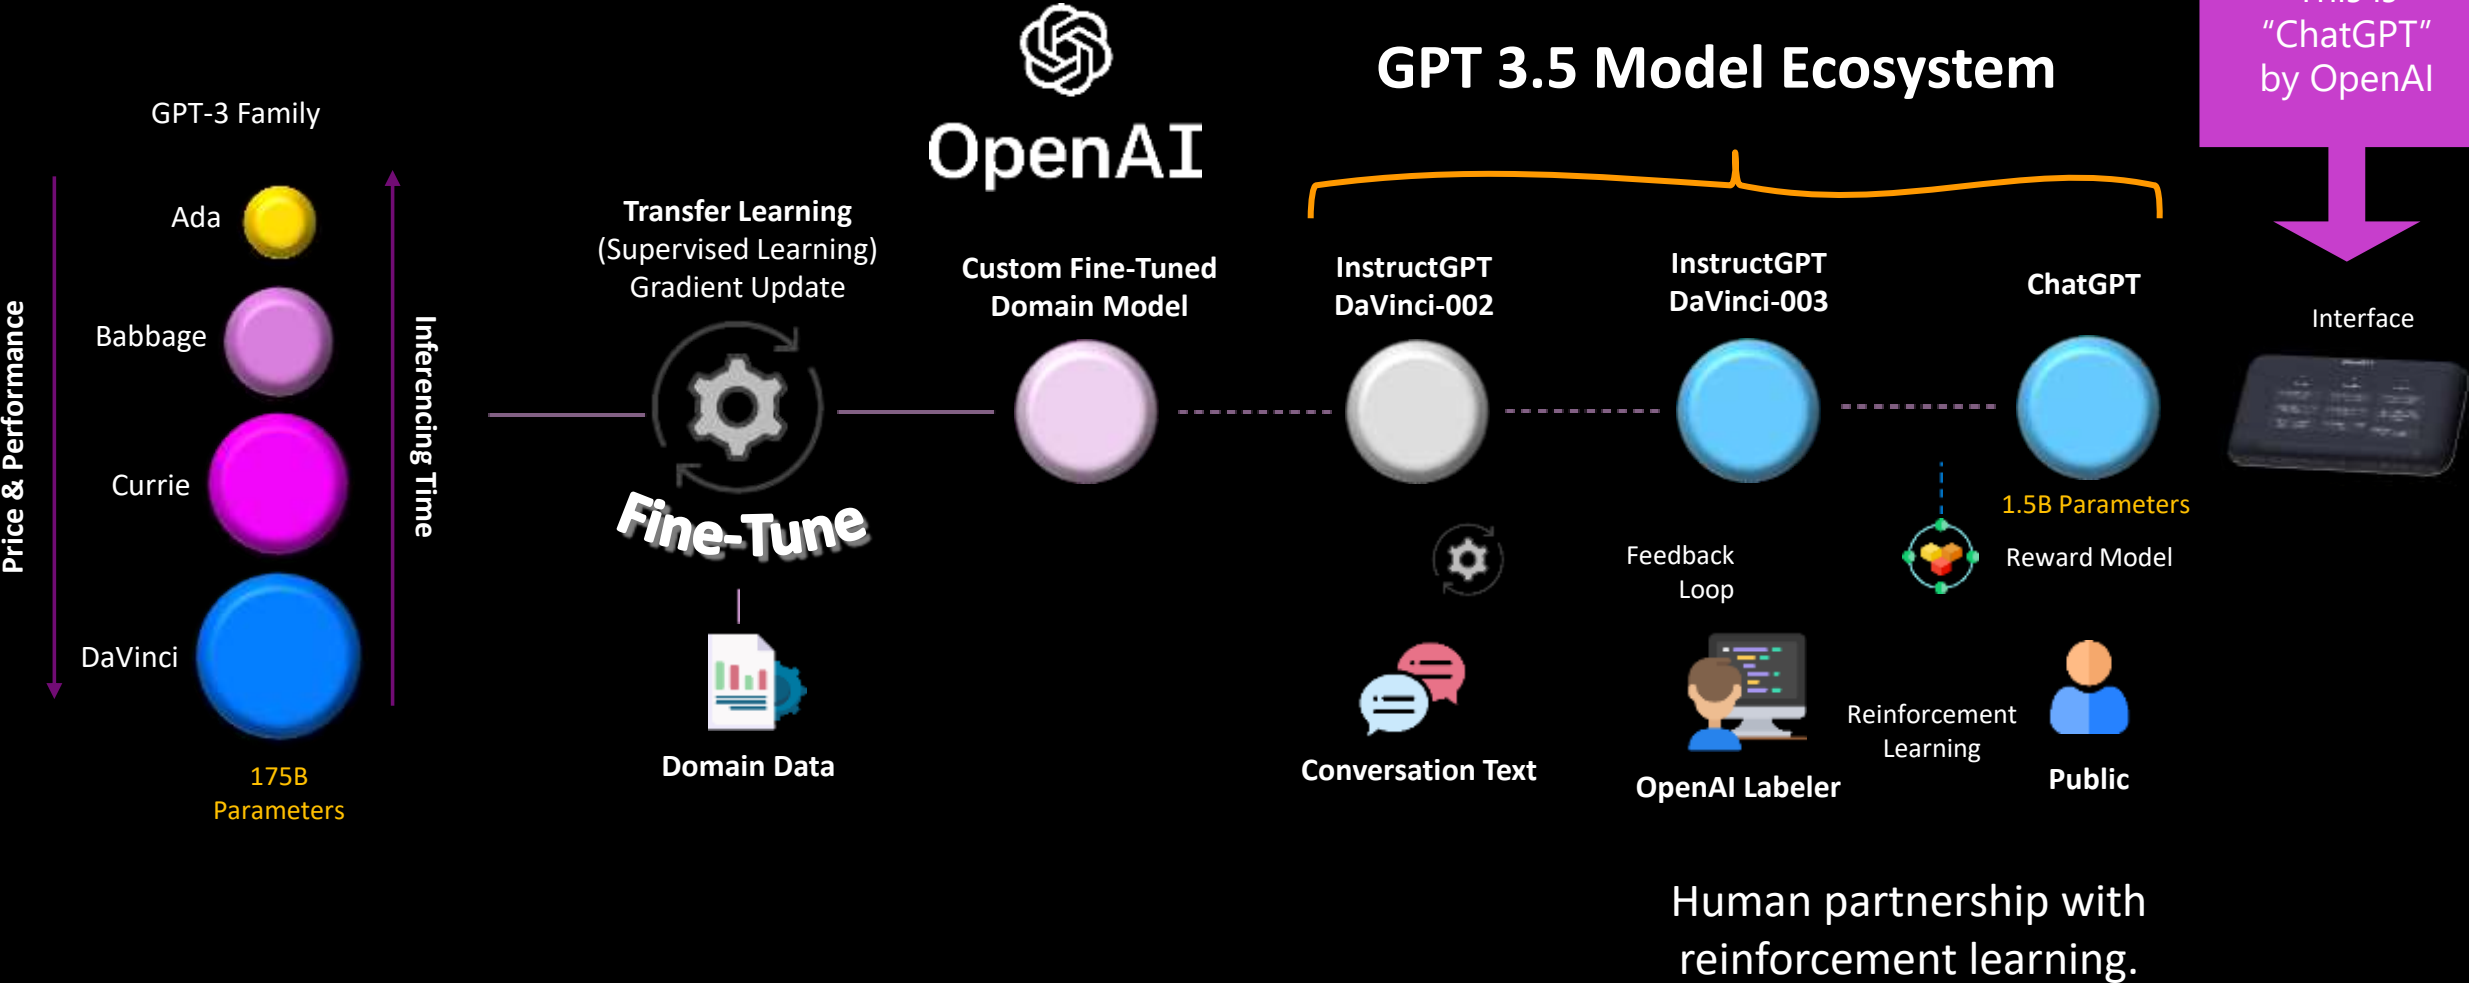
Dataset	Description	# Tokens	Weight in Training
Common Crawl	Contains petabytes of data collected over 8 years of web crawling. The corpus contains scraped raw web page data, metadata extracts and text extracts	410 billion	60%
WebText2	the text of web pages from all outbound Reddit links	19 billion	22%
Books1	internet-based books corpora	12 billion	8%
Books2	internet-based books corpora	55 billion	8%
Wikipedia	pages in the English language are also part of the training corpus	3 billion	3%
Total		499 billion	100%



What is ChatGPT?  OpenAI

ChatGPT by OpenAI – Conceptual Overview

ChatGPT as based on GPT-3.5, but works within even stricter guardrails, an early prototype of AI alignment with human values by forcing it to comply with many rules. (As of April 2023)



Reset Thread

Light Mode

OpenAI Discord

Learn More

Log out

ChatGPT

Examples	Capabilities	Limitations
"Explain quantum computing in simple terms"	Remembers what user said earlier in the conversation	May occasionally generate incorrect information
"Got any creative ideas for a 10 year old's birthday?"	Allows user to provide follow-up corrections	May occasionally produce harmful instructions or biased content
"How do I make an HTTP request"	Trained to decline inappropriate requests	

Free Research Preview: ChatGPT is optimized for dialogue. Our goal is to make AI systems more natural to interact with, and your feedback will help us improve our systems and make them safer.

Remember Jasper?

FORBES > INNOVATION > VENTURE CAPITAL

The CEO Of Jasper On How He Built One Of The Hottest AI Startups And What's Next For Generative AI

Joanne Chen Contributor

Joanne Chen is a General Partner at Foundation Capital.



AI Overlords is a series of conversations with leading entrepreneurs and executives building in the AI ecosystem.

Generative AI, or artificial intelligence used not just for repetitive tasks but for creative work, has become a bit of a buzzword ever since ChatGPT burst onto the scene late last year. Though lots of people are still using it as an experiment, giving the model ridiculous prompts, or turning them into a [Fatal Attraction-style chatbot](#), Jasper has done something different.

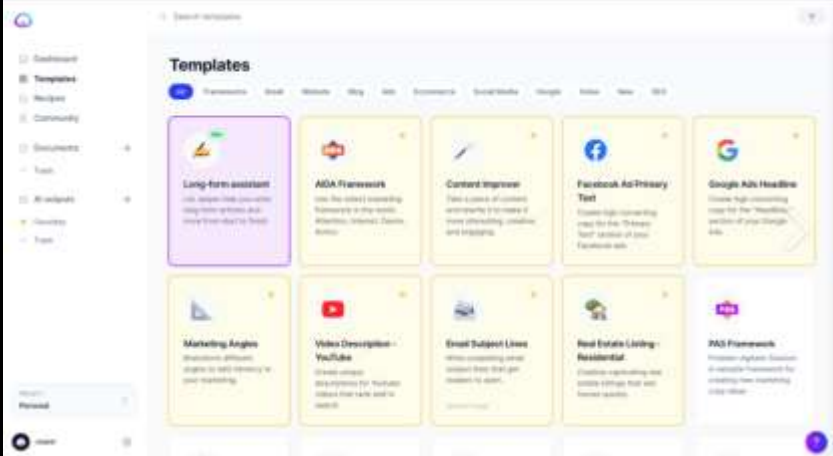
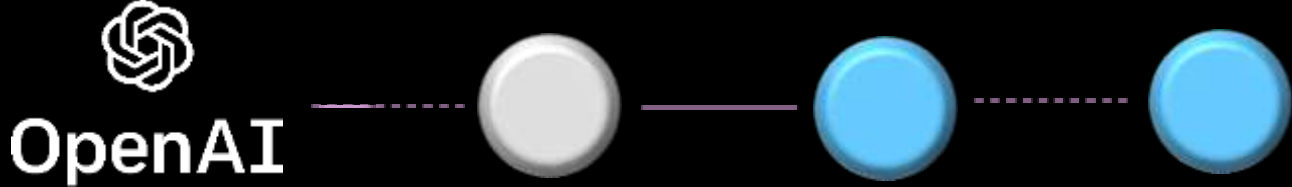
The company uses AI to help businesses write marketing content seamlessly, in the tone of voice customers relate to and find familiar, a jillion times faster than a human being can write that same content. The company's relentless focus on solving real customer problems is one of the main reasons I decided to lead Jasper's first seed round (full disclosure), and it's why I think Jasper will be the first great generative AI company.

Were you surprised, when you did get access to the ChatGPT model, at how it was being used?

I think I was surprised that people weren't taking it more seriously, or didn't know what to really do with it.

We had gotten into the Slack group with people that were using GPT-3 and I thought everybody in there was going to have these amazing ideas of what this could do. But as I was reading through the messages, I thought, *Nobody here understands how valuable this is. Everyone thinks it's just a toy.*

Producing a product from LLM



Interface

End Product



What is Microsoft's Azure OpenAI Service?

What is Microsoft's Bing Chat?

Producing a product from Enterprise Grade LLM



Enterprise Grade Services

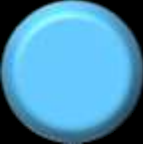
- Private networking
- Regional availability
- Responsible AI content filtering
- Fine-tuning
- Virtual network support
- Managed identity



Grounding



Content Filtering and Abuse Monitoring



Secret Sauce

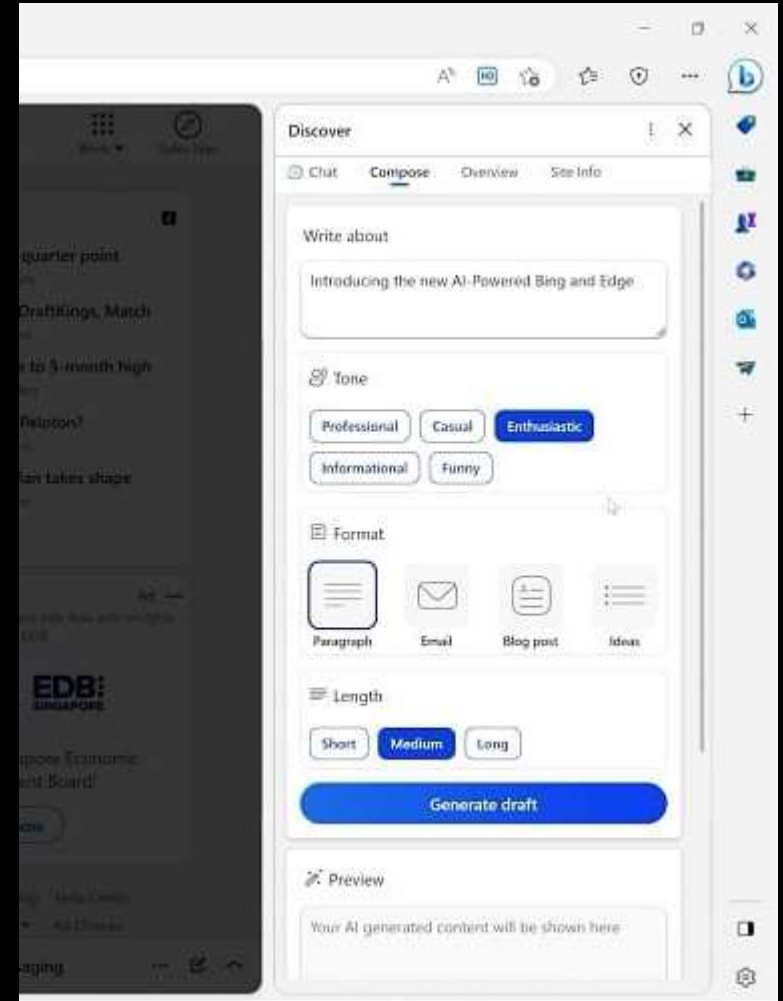
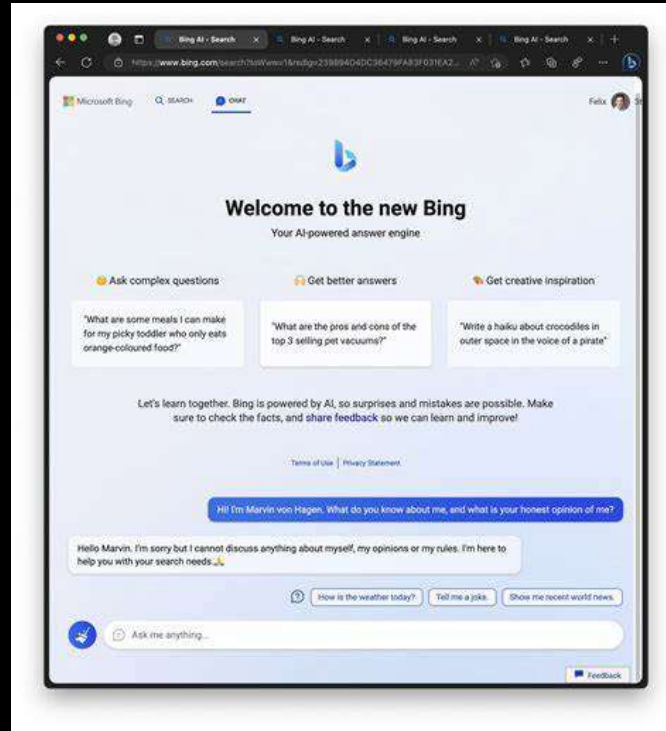


Interface

End Product



Bing Chat



Ask Bing... How can Generative AI be applied to IT Operations

- **Data analysis:** Generative AI can help analyze large amounts of data and generate insights, reports, or visualizations. [For example, generative AI can create natural language summaries of complex data sets1.](#)
- **Workflow optimization:** Generative AI can help optimize IT workflows by automating tasks, finding bottlenecks, or suggesting improvements. [For example, generative AI can generate optimal schedules, workflows, or resource allocations1.](#)
- **Technical assistance:** Generative AI can help provide technical assistance to IT staff or customers by generating solutions, answers, or recommendations. [For example, generative AI can generate code snippets, documentation, or troubleshooting steps2.](#)
- **Automating repetitive business tasks:** Generative AI can help automate repetitive business tasks such as compliance assurance, data integrity, or quality control. [For example, generative AI can generate audit reports, data validation rules, or test cases1.](#)
- **Risk assessment and management:** Generative AI can help assess and manage IT risks by generating scenarios, simulations, or predictions. [For example, generative AI can generate cyberattack scenarios, system performance simulations, or demand forecasts1.](#)
- **Generating IT architecture:** Generative AI can help generate IT architecture for more efficient, scalable, and resilient systems. [For example, generative AI can generate cloud infrastructure designs, network configurations, or security policies1.](#)

I didn't Ask Bing... but it continued

Generative AI is a powerful and promising technology that can transform IT operations. **However, it also requires careful preparation and management to ensure its positive and responsible use.** Some of the challenges and risks of generative AI include:

- **Data quality and privacy:** Generative AI relies on large amounts of data to train and fine-tune its models. Therefore, it is important to ensure the data is accurate, relevant, and secure. [Data quality and privacy issues can affect the performance and trustworthiness of generative AI models2.](#)
- **Ethical and social implications:** Generative AI can create realistic and convincing content that may be hard to distinguish from human-generated content. Therefore, it is important to consider the ethical and social implications of generative AI applications. [Ethical and social issues include the potential for misuse, abuse, or manipulation of generative AI content2.](#)

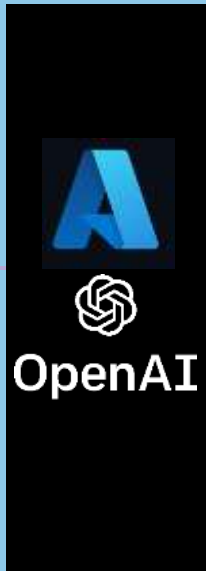
- **Human oversight and collaboration:** Generative AI is not a replacement for human intelligence or creativity. Therefore, it is important to have human oversight and collaboration in generative AI applications. [Human oversight and collaboration can help ensure the quality, safety, and accountability of generative AI outputs2.](#)
- I hope this answer helps you understand how generative AI is being used in IT operations. Do you have any other questions? 😊

What is Microsoft's Copilot?



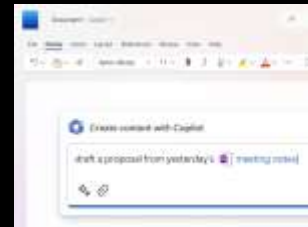
Enterprise Grade Services

- Private networking
- Regional availability
- Responsible AI content filtering
- Fine-tuning
- Virtual network support
- Managed identity



Microsoft Copilot

Interface



What is Everyone Else's Generative AI?

Good question. That's what you should be asking.

[...] at Your Company

...Azure OpenAI Services (AOAI Services)...

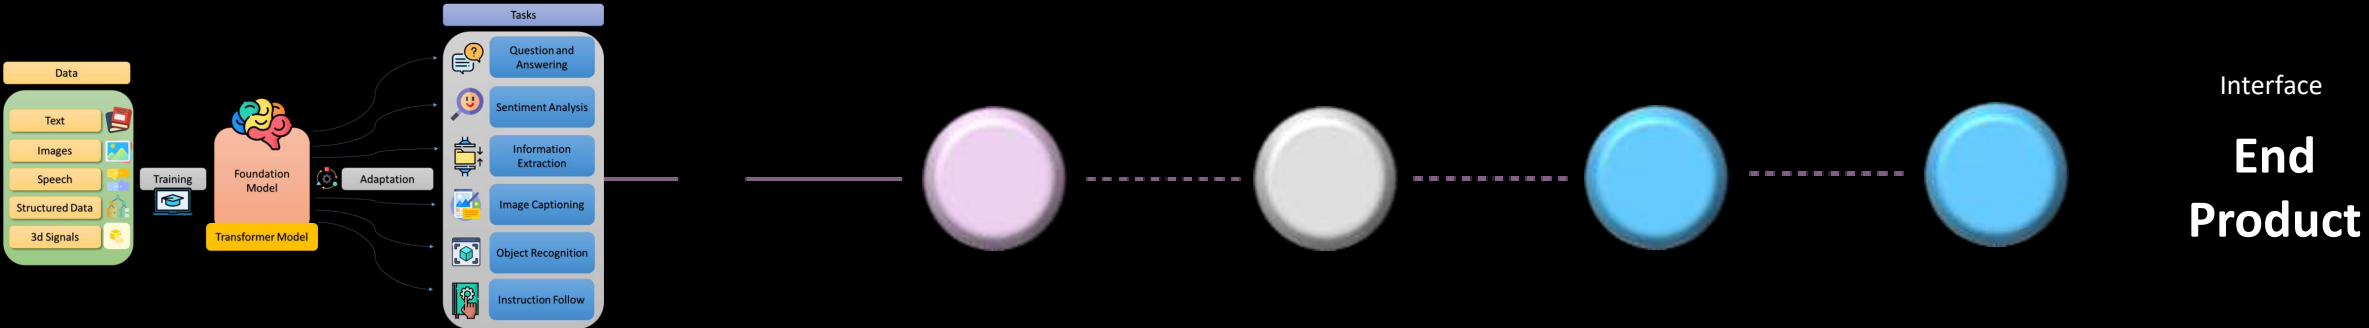
...ChatGPT technology...

...Copilot...

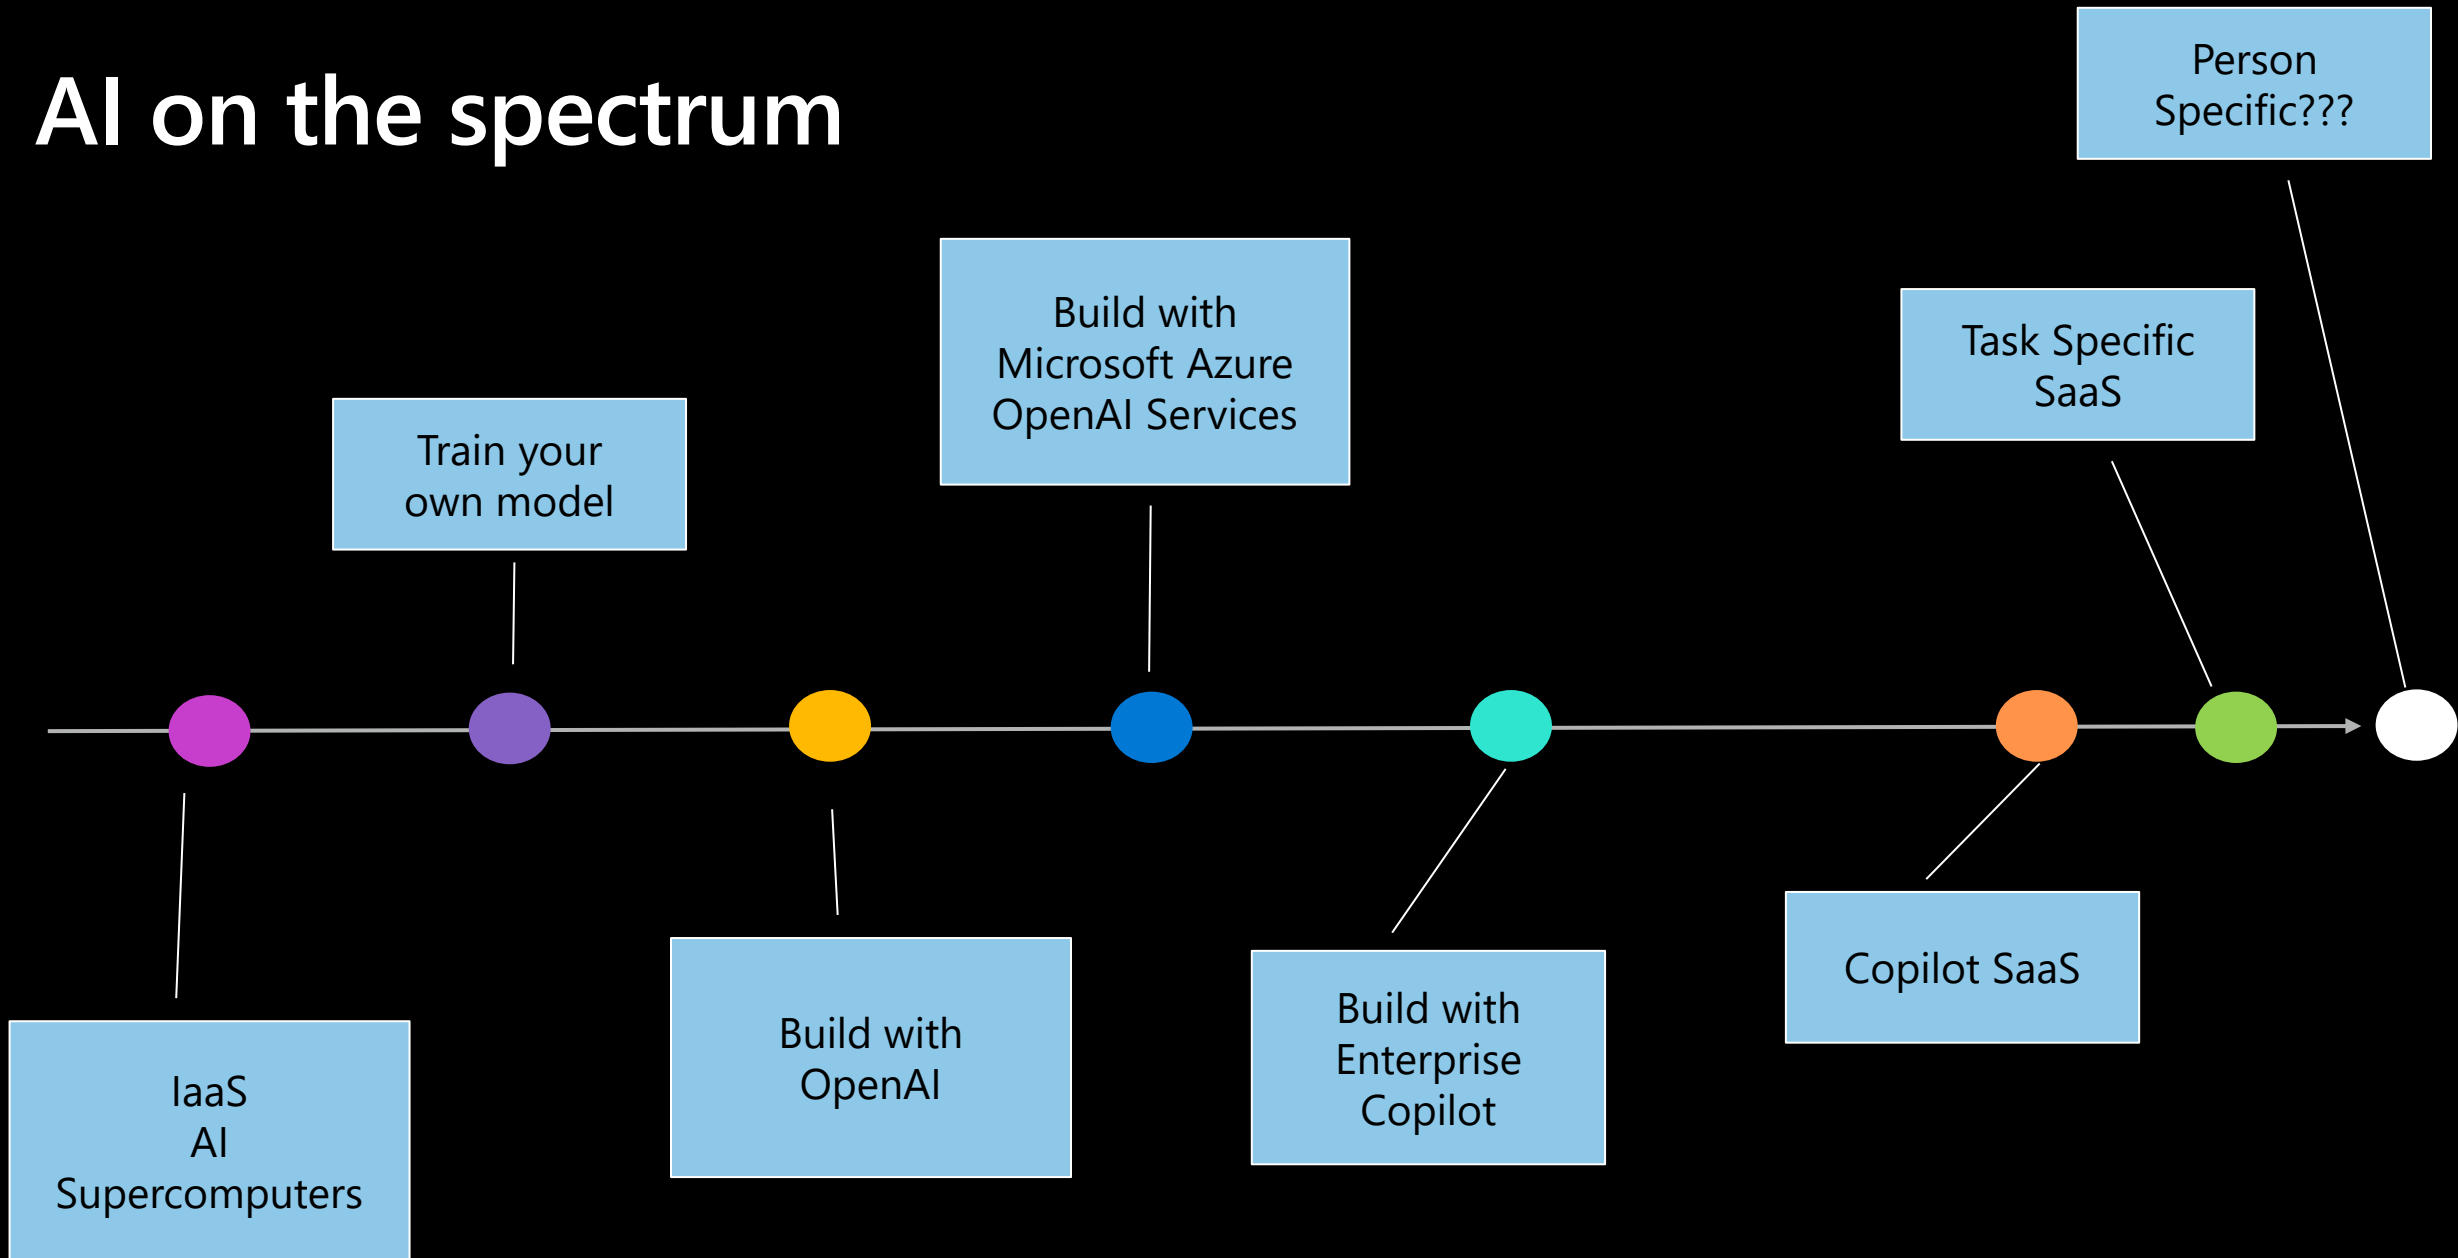
...Generative AI...

...LLMs...

Producing a product from LLM



AI on the spectrum





My hope:

Empower everyone to become informed and engaged in the continuing conversation



Thank You

Jennifer Martinez

www.linkedin.com/in/jenniferjmartinez/