

The Current State of Artificial Intelligence

SOCIAL CAPITAL_

Table of Contents

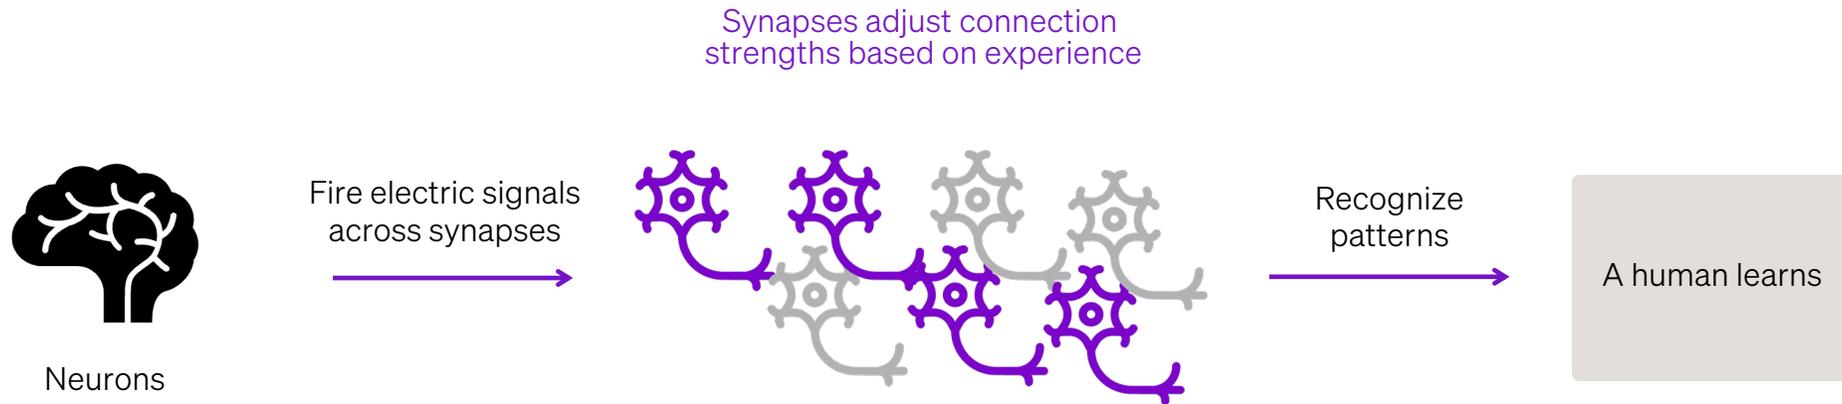
Chapter:	Page:
Introduction	04
State of Chatbot Race	18
Reasoning Models	33
Business of AI	46
Agentic Systems	61
Multimodal / Voice	75
U.S.-China Primacy	87
Wrapping Up	99

How to Read This Presentation

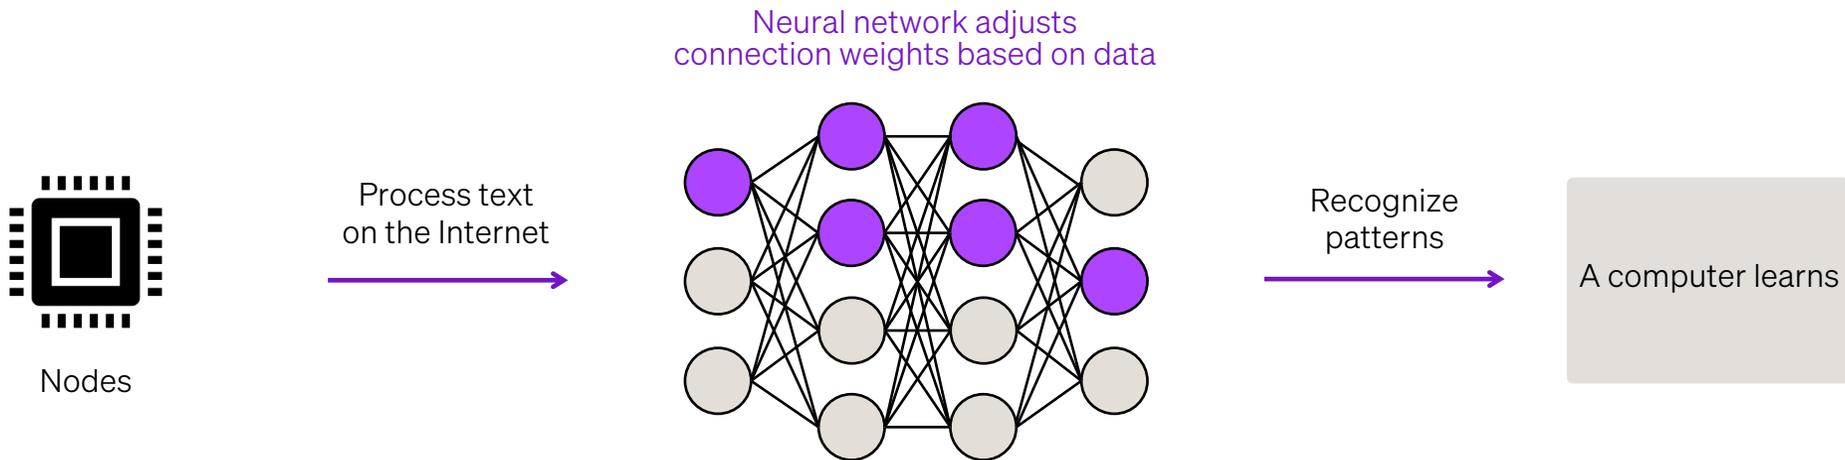
- This presentation was designed to be read in chronological order, in one go, like a flip book. Each section of this presentation builds on the prior and assumes no prior knowledge about the discussed topic.
- This presentation aims to provide a current overview of large language models, including the emergence of advanced reasoning models and agents, as well as the current state of AI businesses and their underlying dynamics.
- By the end of this deep dive, you should understand how language models are evolving under resource constraints and competitive pressures, what lessons we've learned from this first wave of AI products, and how to identify the breakthrough opportunities in this megatrend.

Introduction

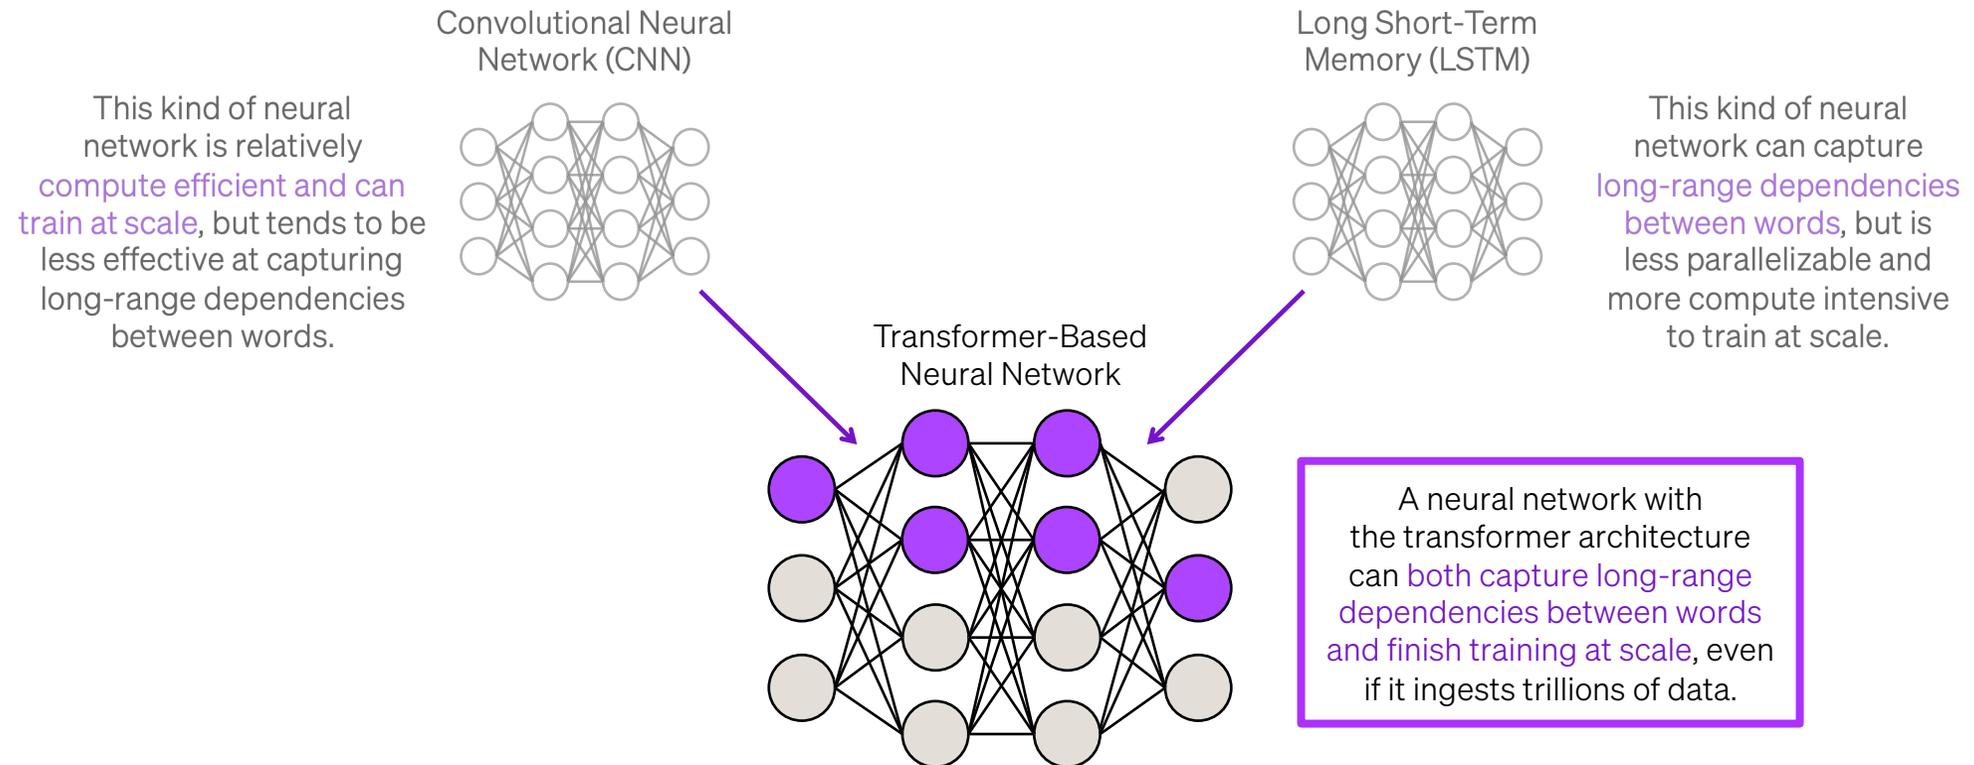
The human brain contains roughly 100 trillion neural connections, working together to generate thoughts through patterns of activation.



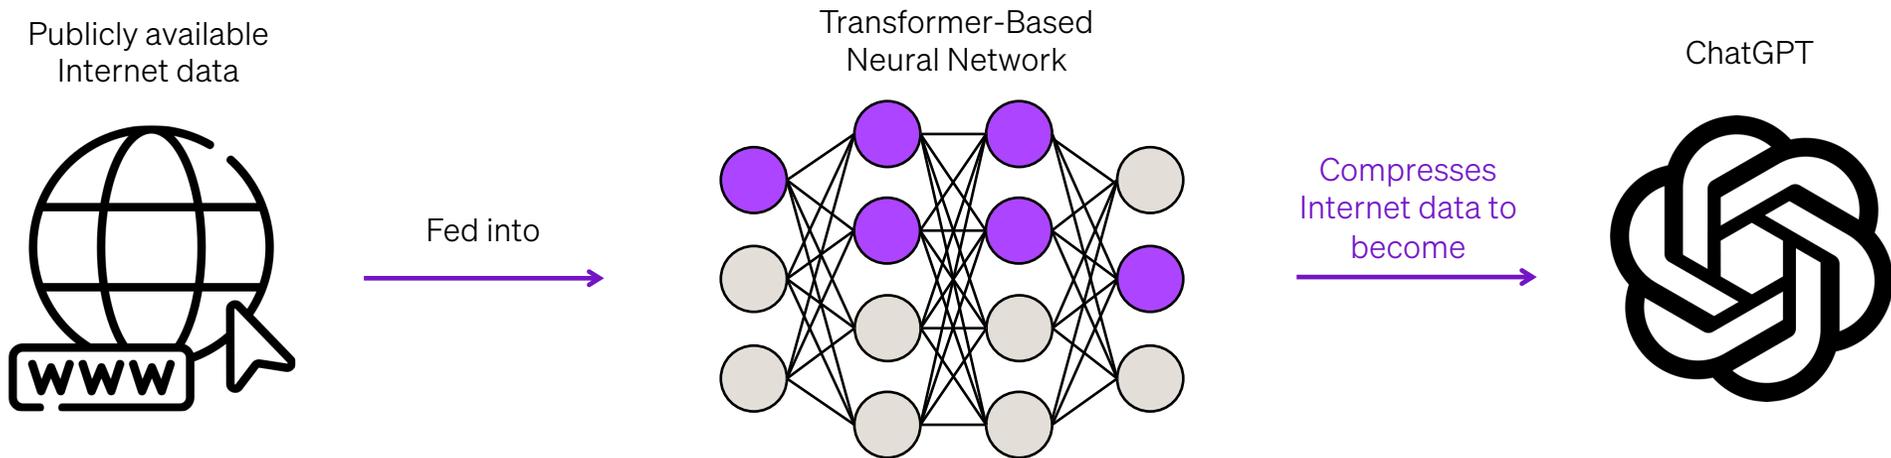
AI researchers took inspiration from the human brain, attempting to build neural networks with computers that could "learn" through patterns in data.



The transformer architecture was a breakthrough that hit the sweet spot between pattern recognition power and computational feasibility.

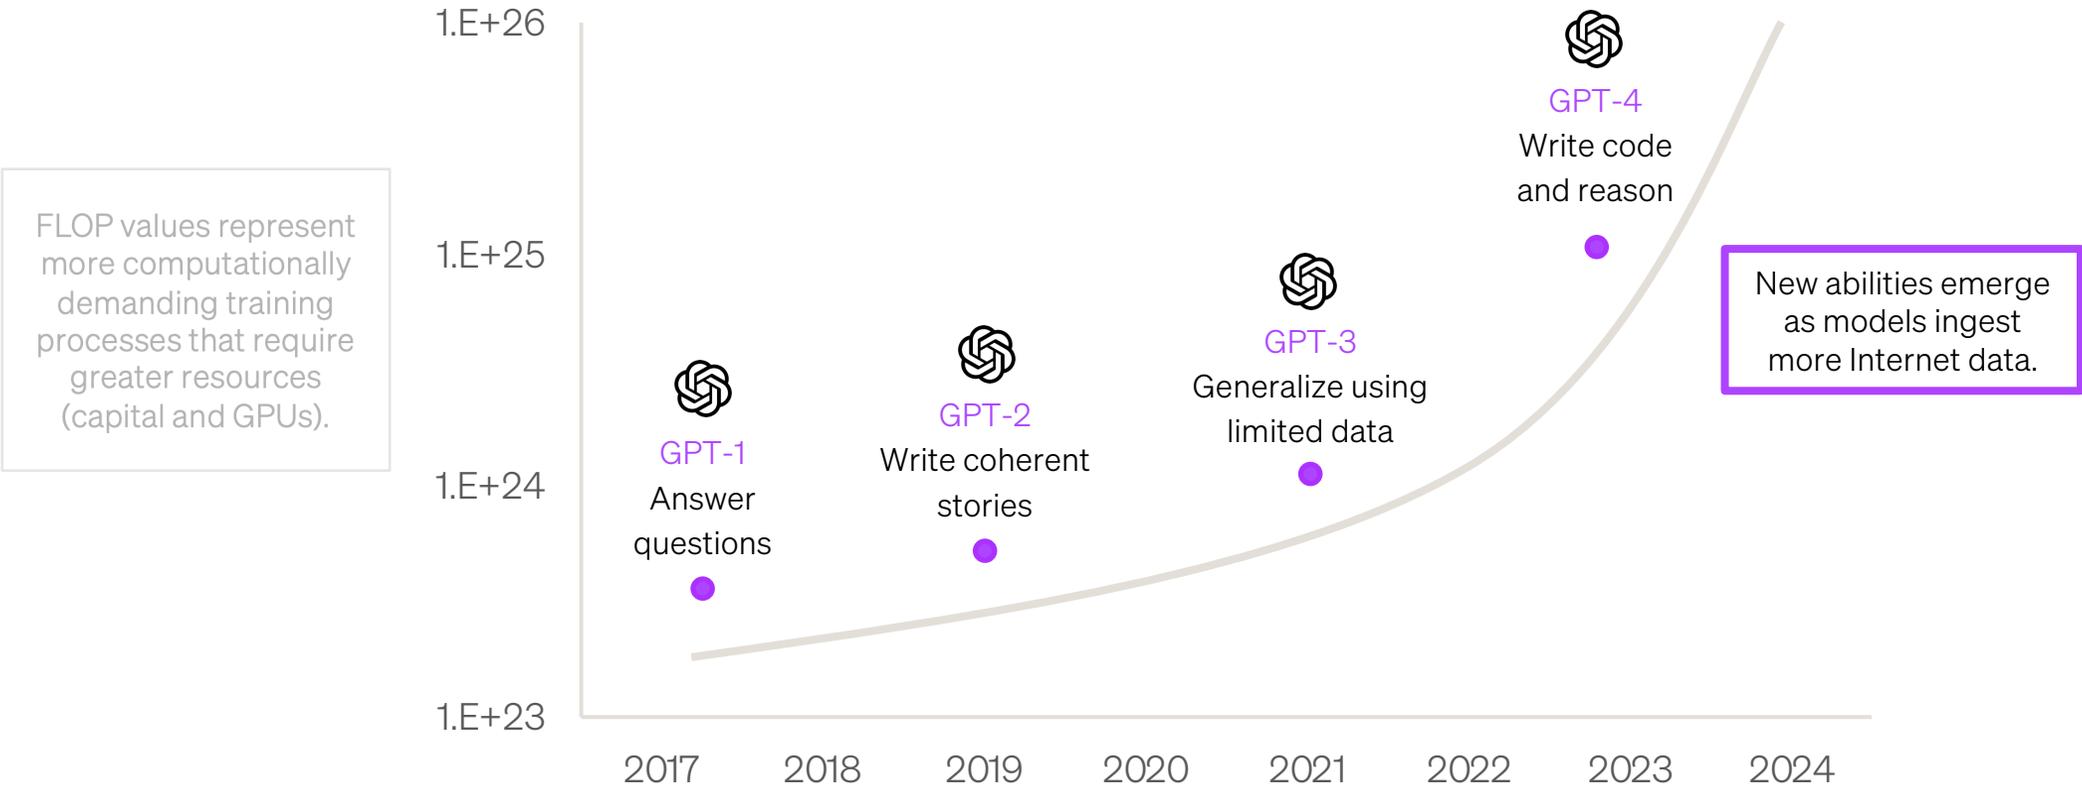


This breakthrough made it possible to compress most of the publicly available knowledge stored on the internet into a language model.



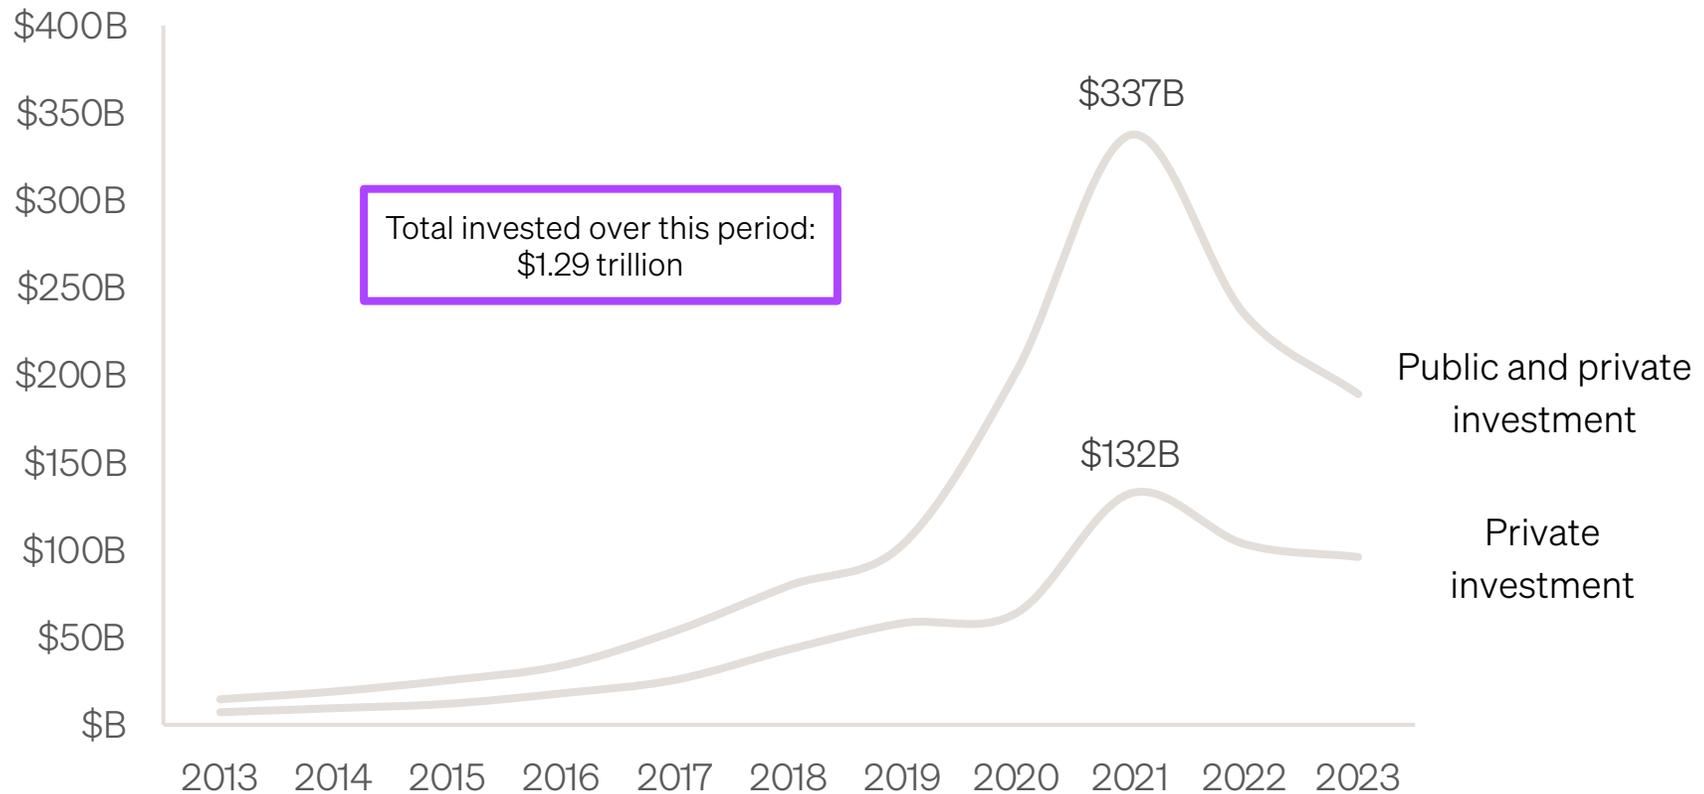
As these models were scaled with exponentially more internet data, researchers observed that capabilities would suddenly "emerge" that weren't explicitly trained for.

Training Compute (FLOP) Required for Each OpenAI Language Model, 2017 to 2024



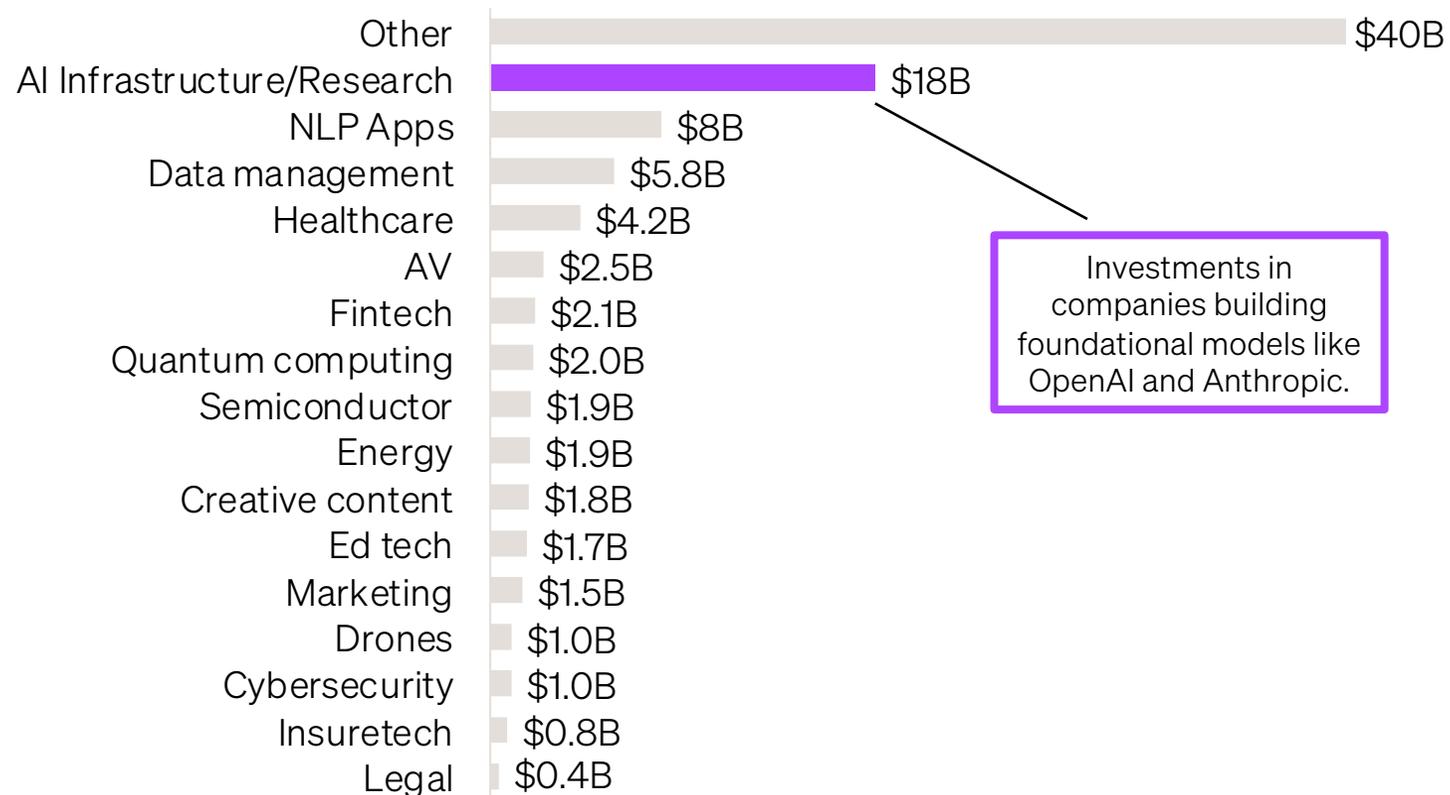
All the potential use cases of these capabilities and the pace of advancement compelled people to invest more than a trillion dollars into AI.

Public and Private Investment in AI, Billions of U.S. Dollars, 2013 to 2024



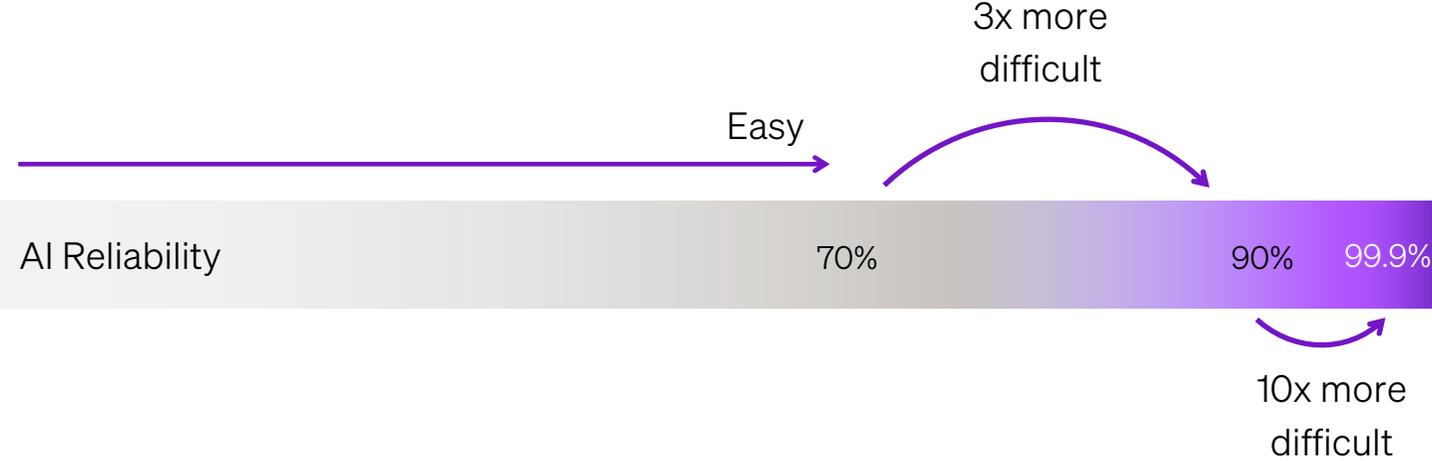
And that trillion-plus-dollar investment fueled foundational models and AI products, enabling rapid innovation and widespread adoption.

Private Investment in AI by Focus Area, 2023

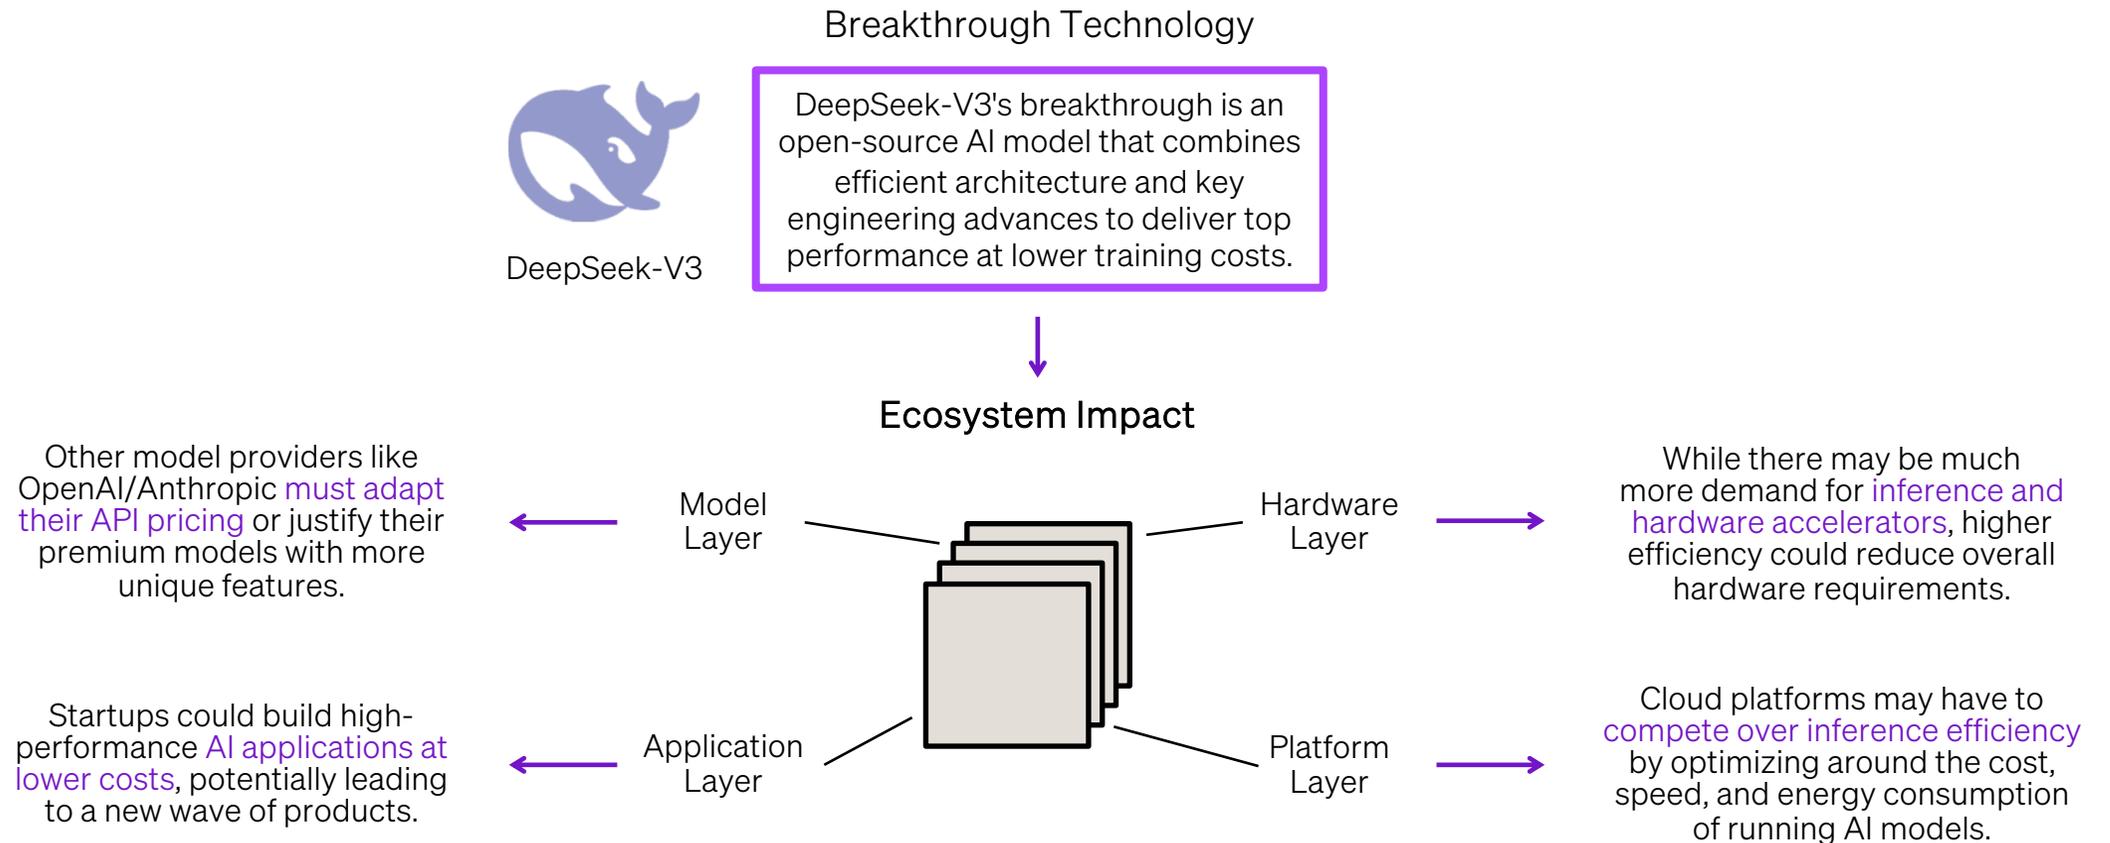


So what is the current state of AI?

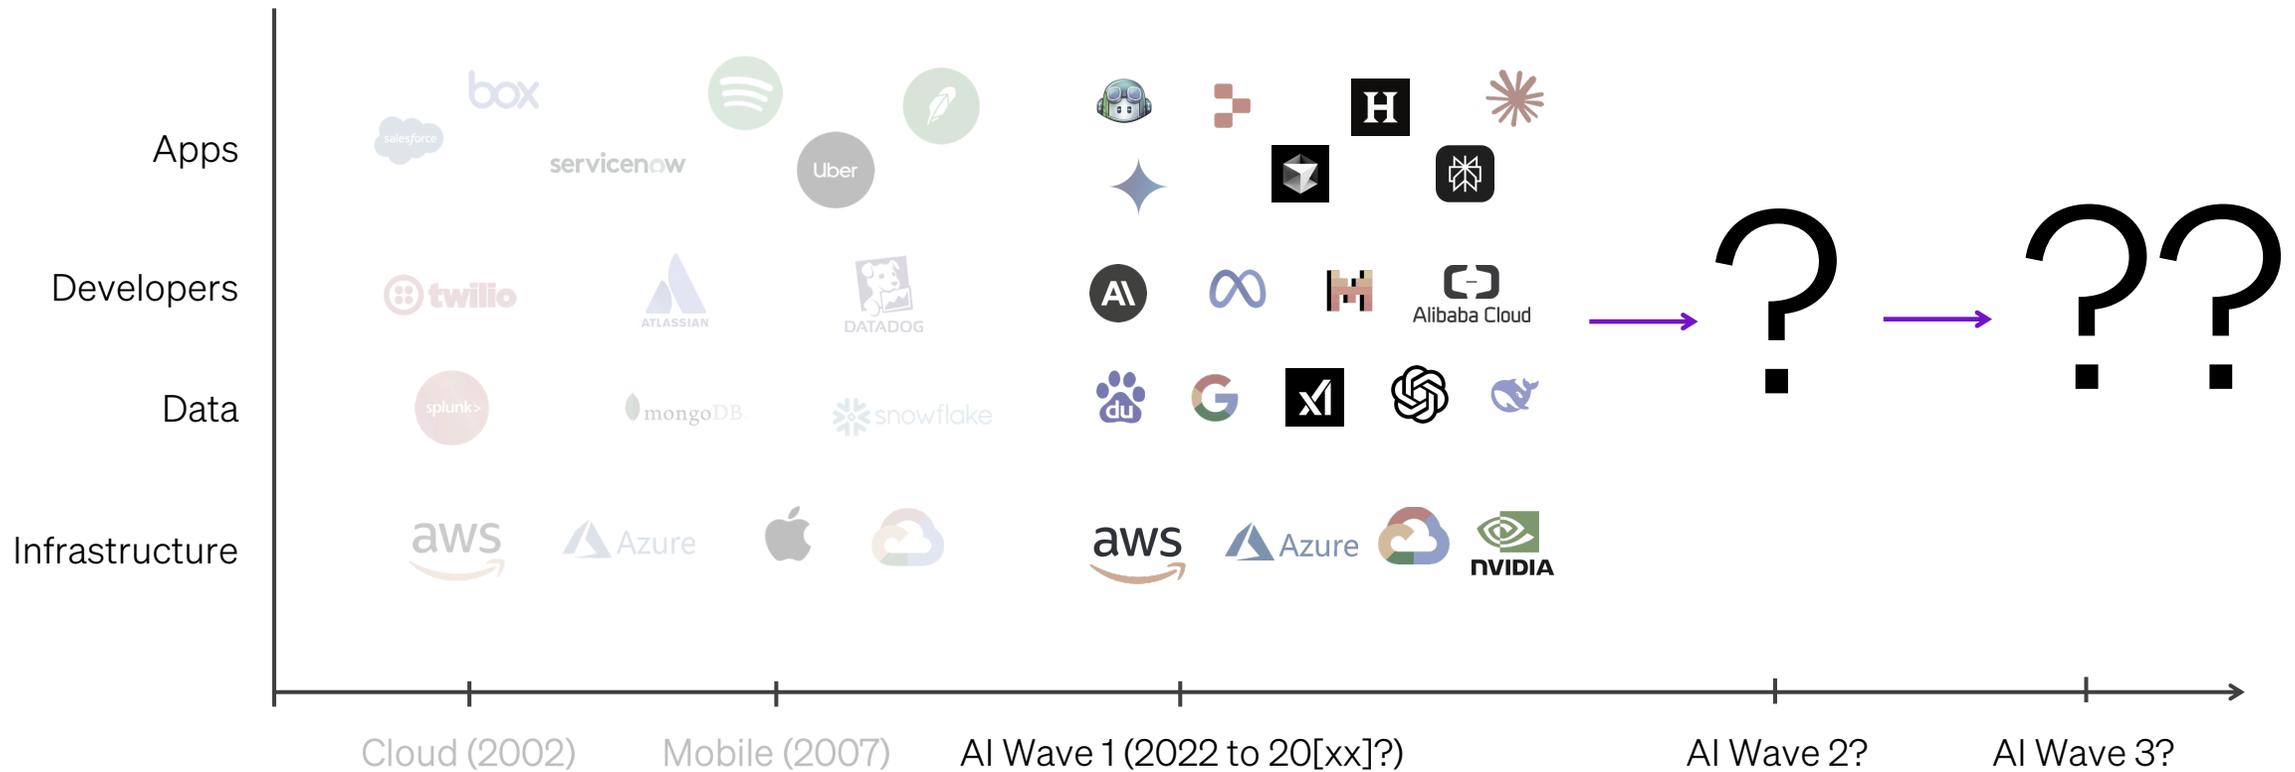
While AI has made prototyping easier, bringing AI products to market requires extensive effort and engineering prowess to achieve production-grade reliability.



On top of that, new breakthroughs can dramatically change the playing field and triggers a paradigm shift throughout the entire ecosystem.

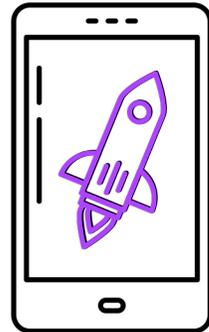


AI may be in the early innings of a megatrend that may reshape society in ways we're only beginning to uncover.



And some key questions remain...

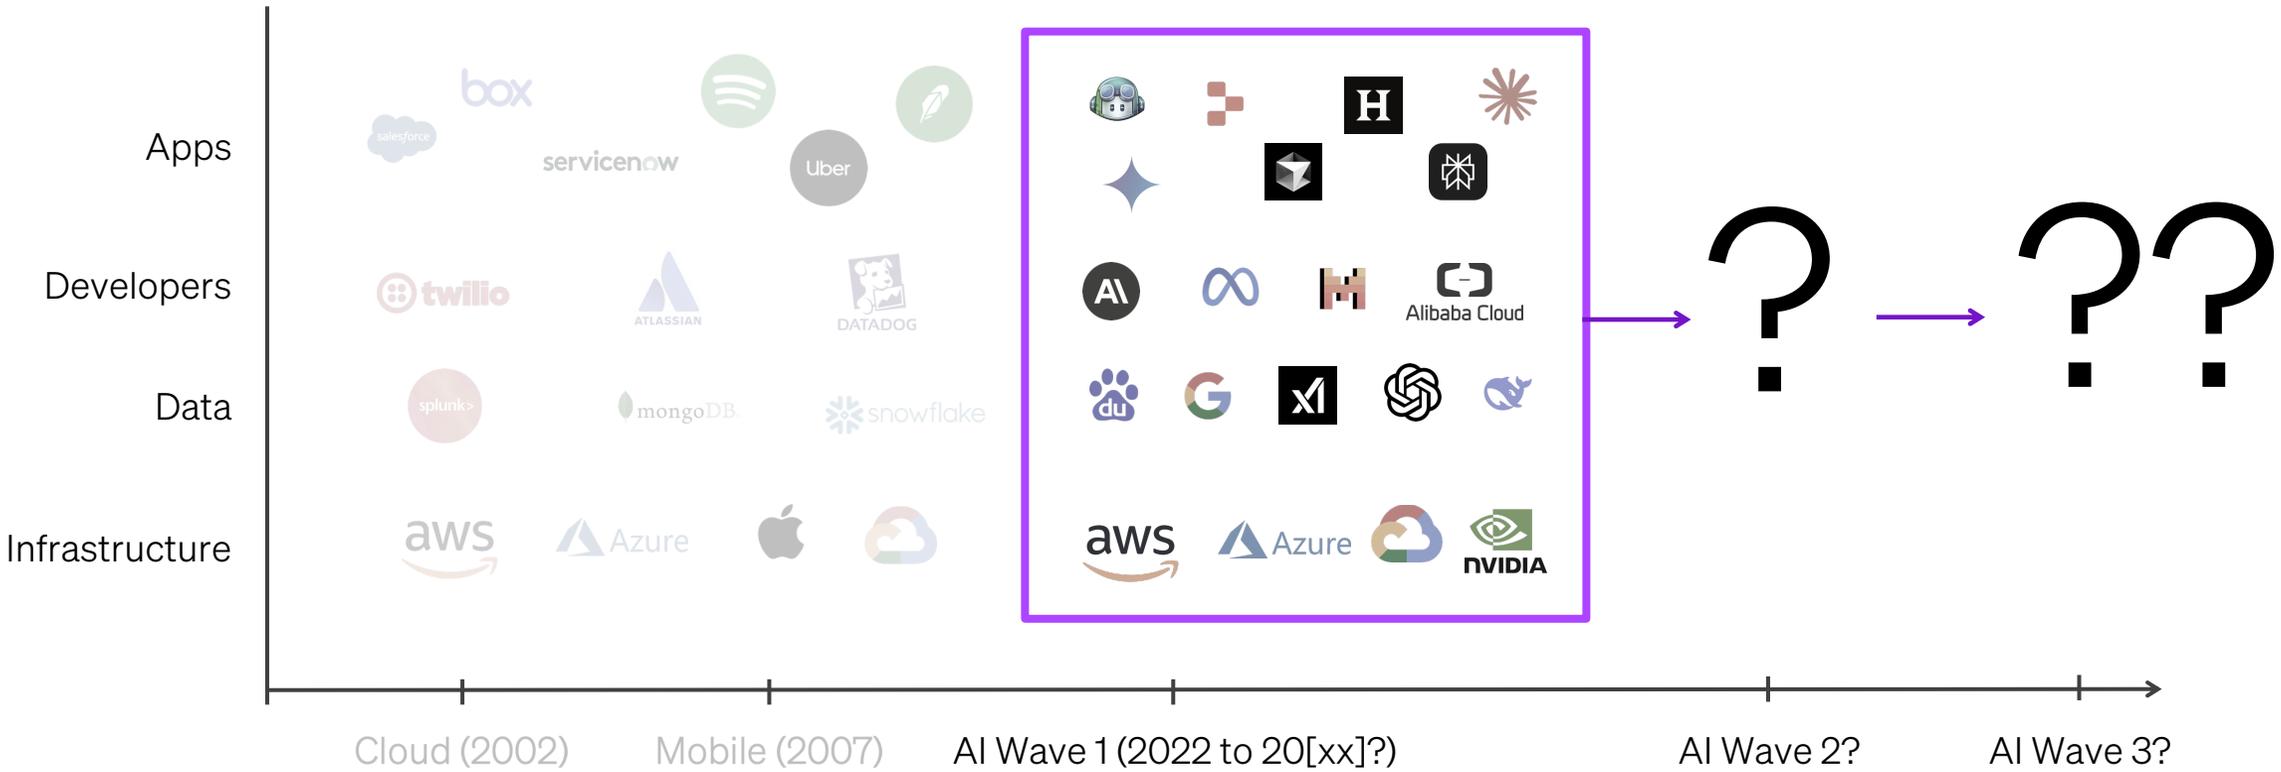
Where will the next killer app emerge?



What AI companies will generate returns for investors?



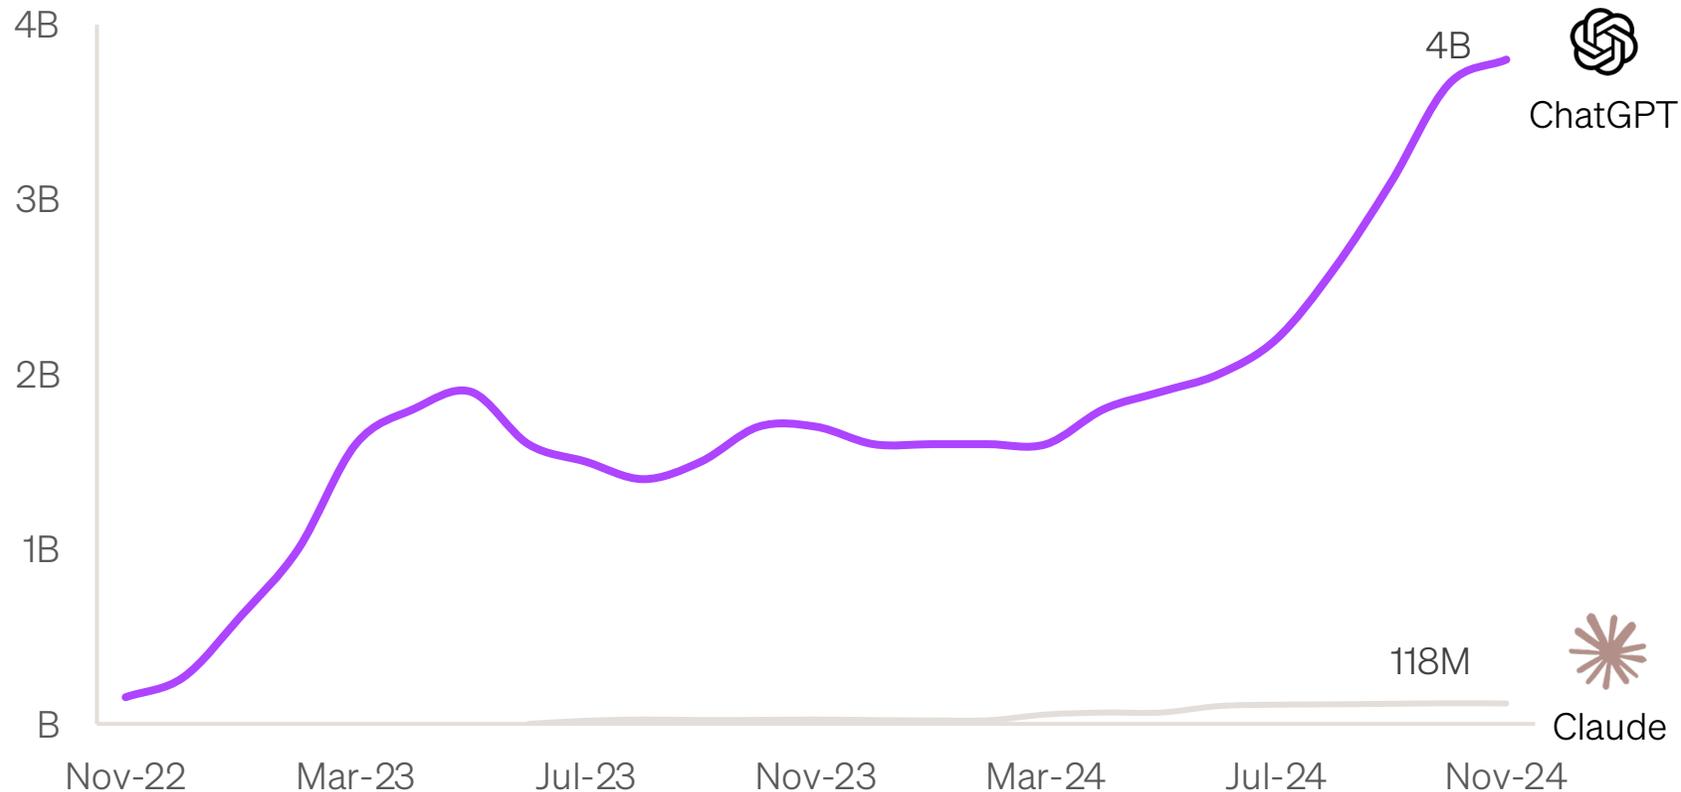
This deep dive will examine AI from infrastructure to application to help you understand the dynamics forming in the early innings of this megatrend.



State of Chatbot Race

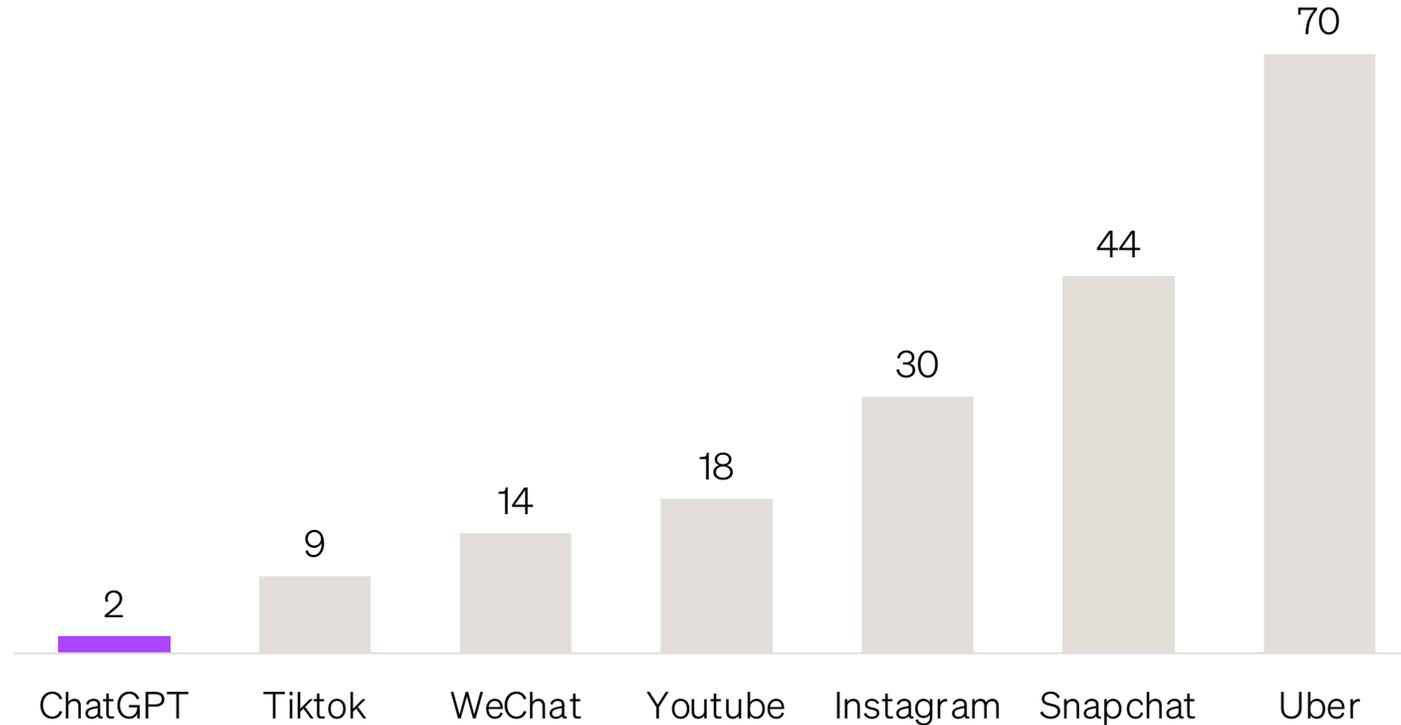
ChatGPT's launch in 2022 demonstrated unprecedented natural language capabilities to the public, establishing OpenAI's first-mover advantage in the market.

Monthly Website Visits, Nov 2022 to Nov 2024



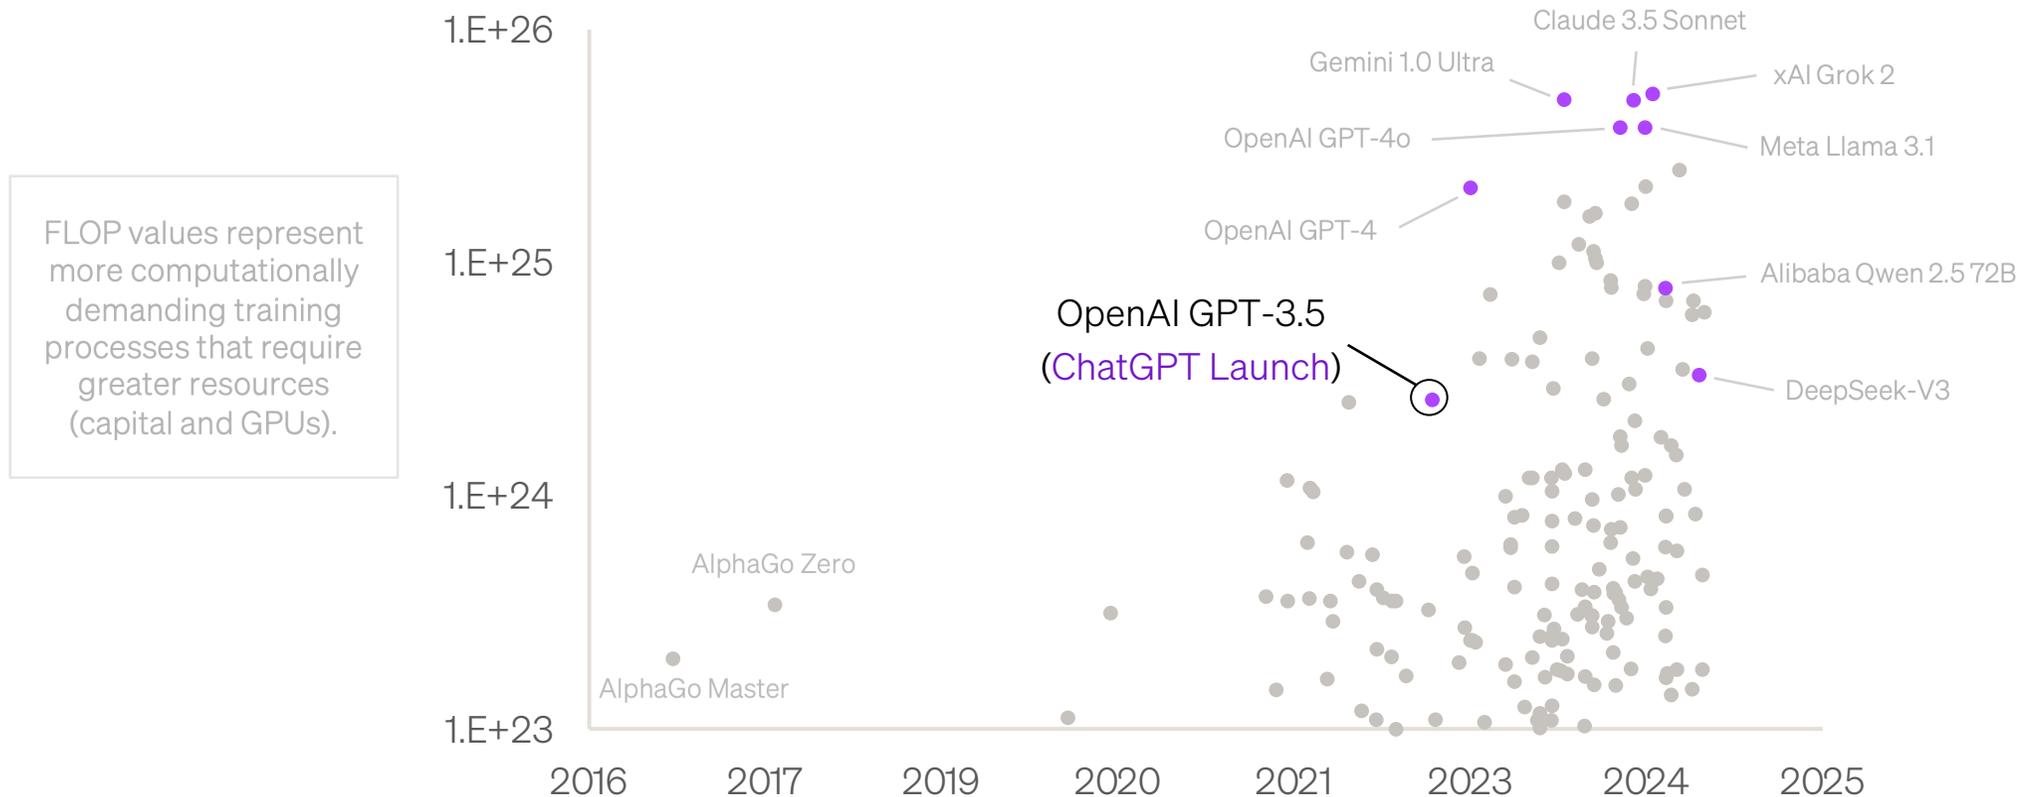
Its success revealed the first killer AI application, an intelligent chatbot that lets users converse with a language model.

Number of Months to Reach 100M Users



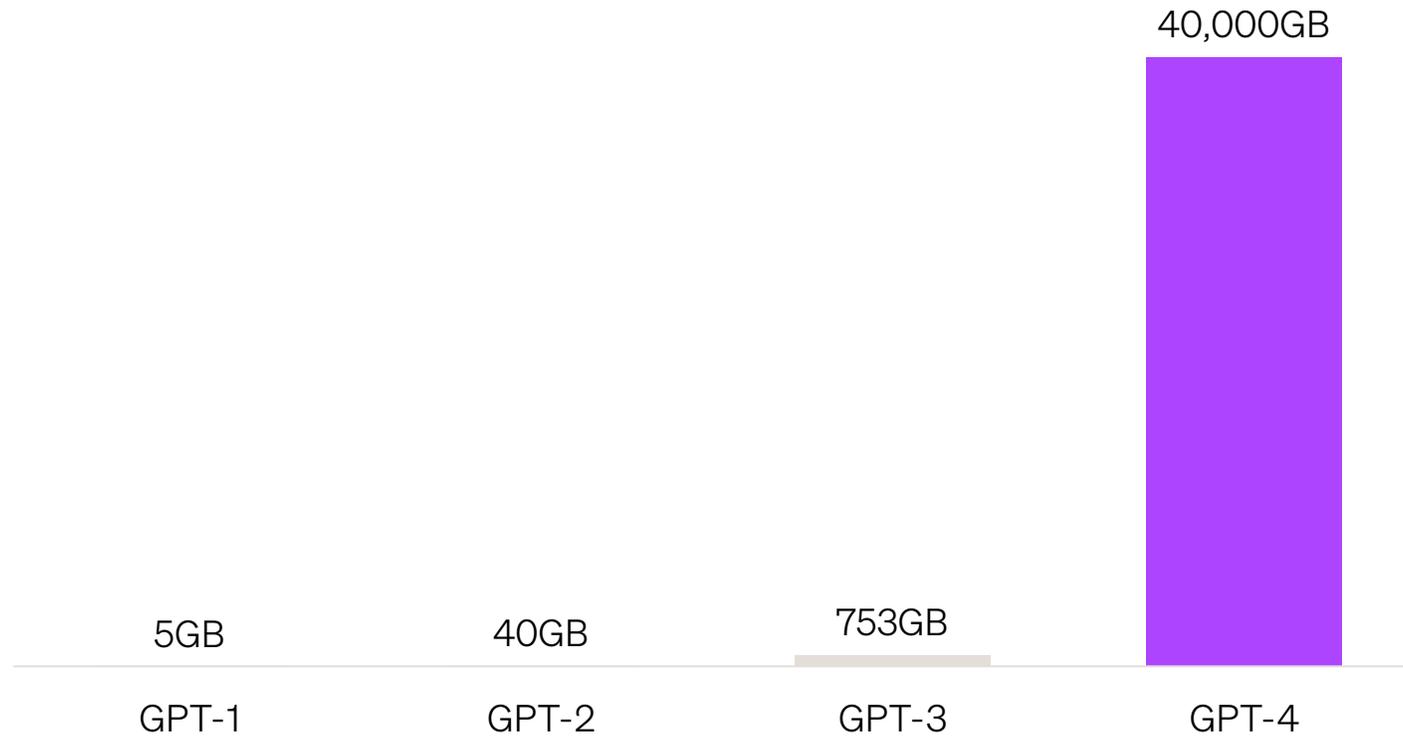
But over the course of two years, many companies have entered the race with highly capable models, and intelligent chatbots are becoming commoditized.

AI Models, Training Compute (FLOP) Over Time, 2016 to 2025



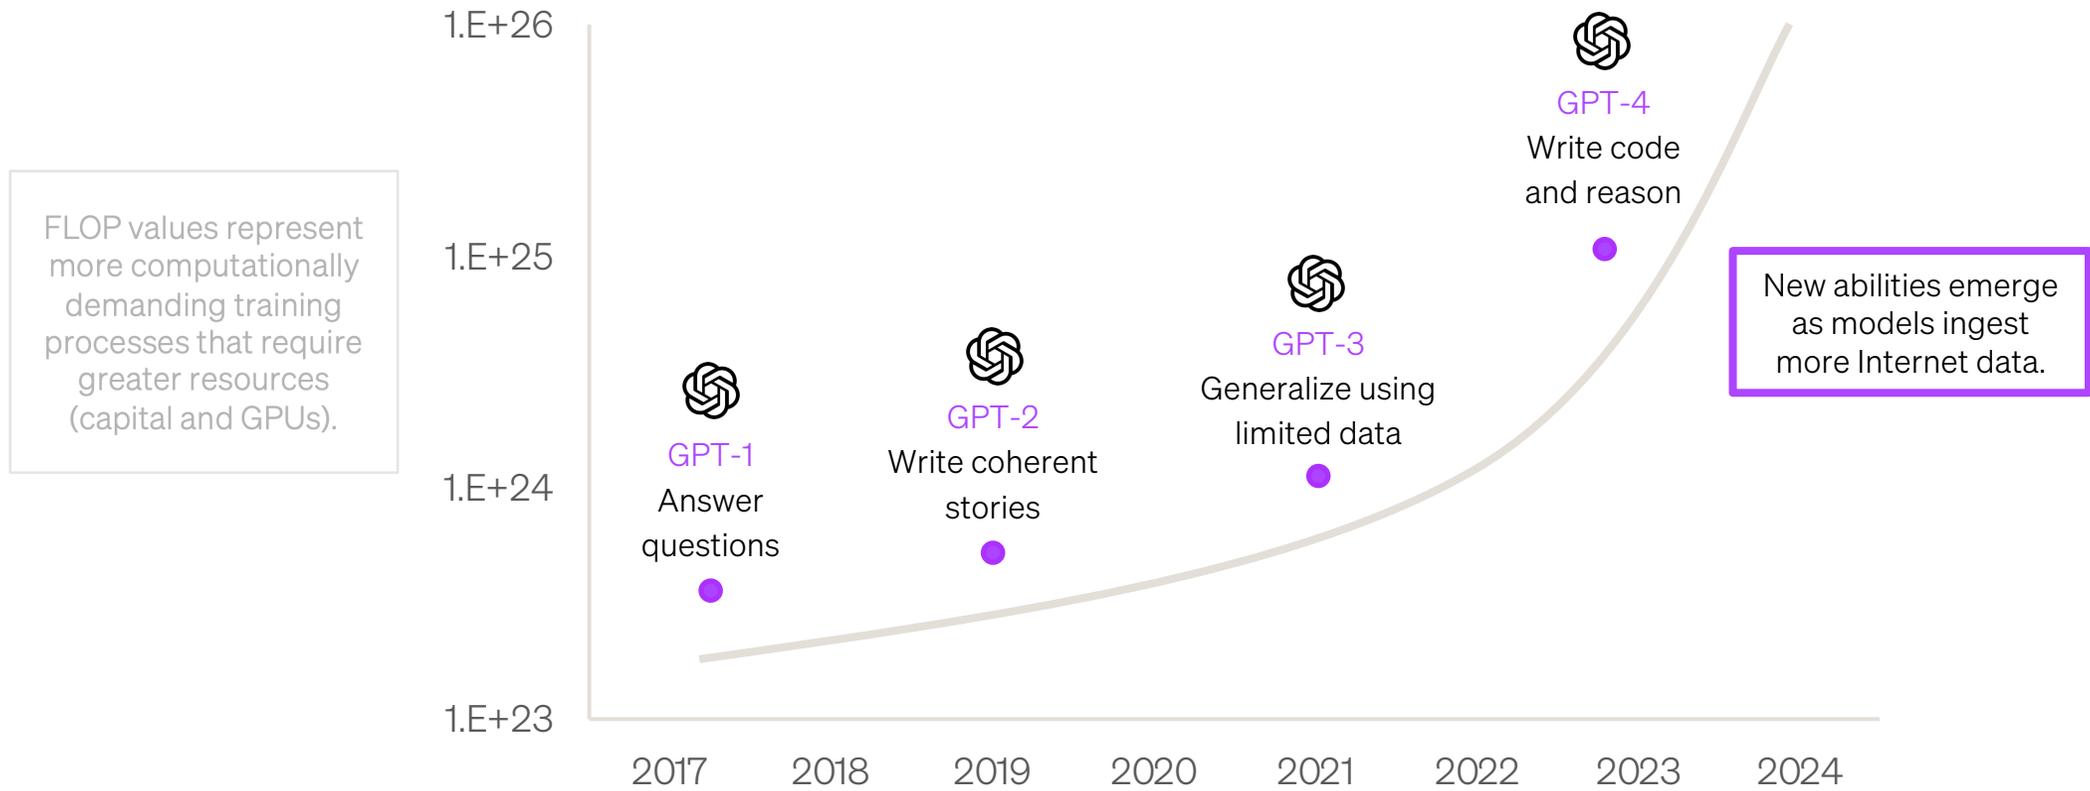
These models all share similar architectures, so the paradigm to win the chatbot race has been to train them with exponentially more high-quality internet data.

Total Size of Uncompressed Dataset Used in Training Language Models, in Gigabytes



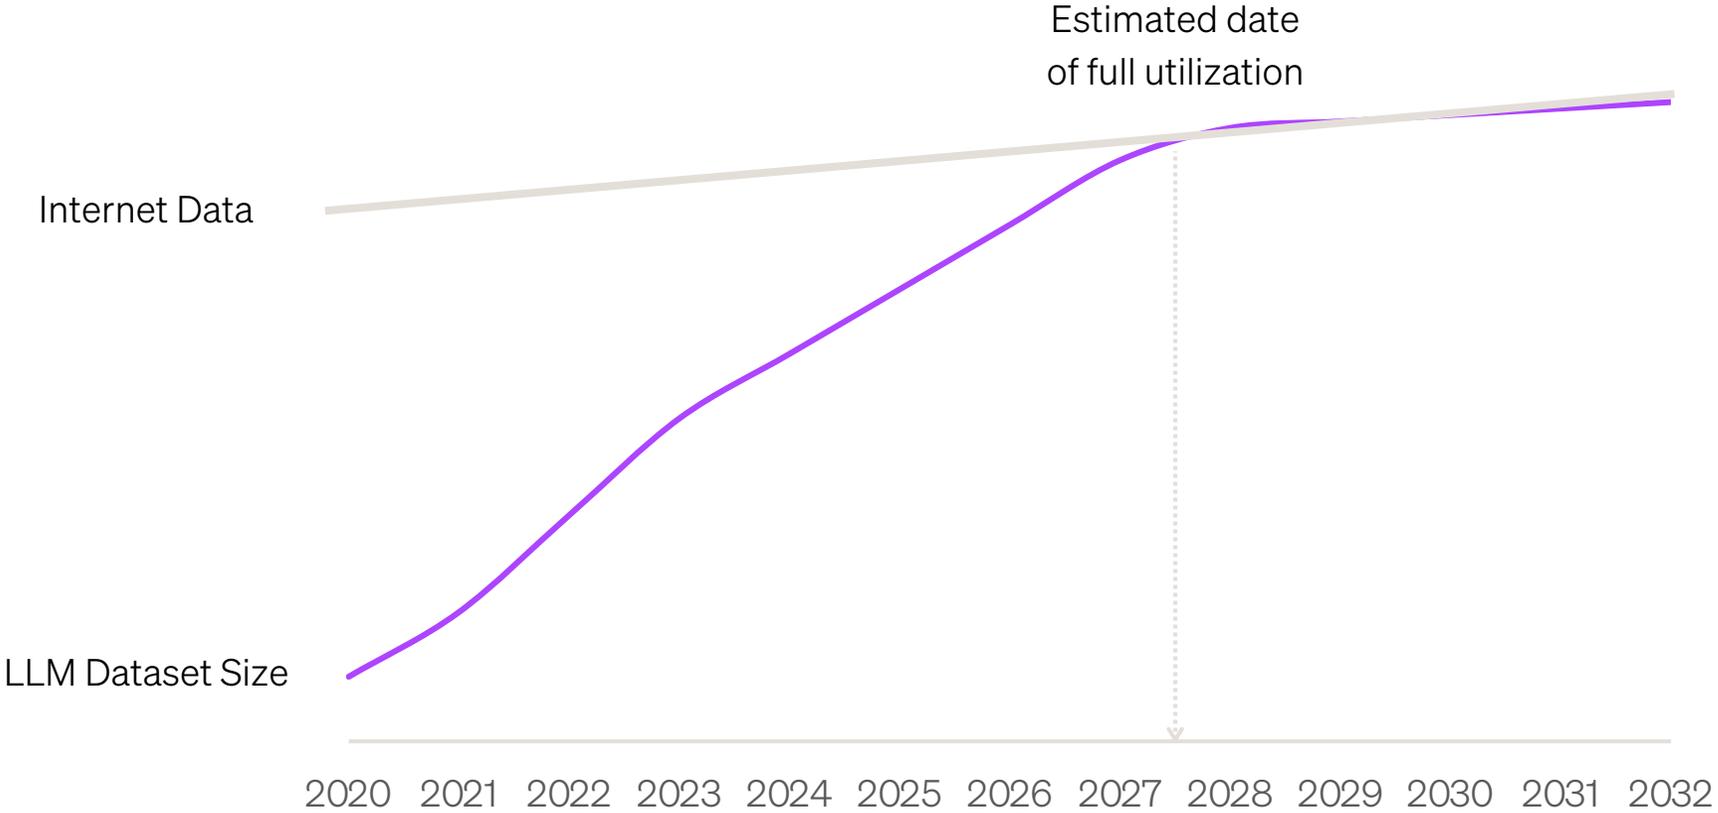
The idea was to feed massive amounts of data to the language model to develop unique capabilities and domain expertise that would set it apart from competitors.

Training Compute (FLOP) Required for Each OpenAI Language Model, 2017 to 2024



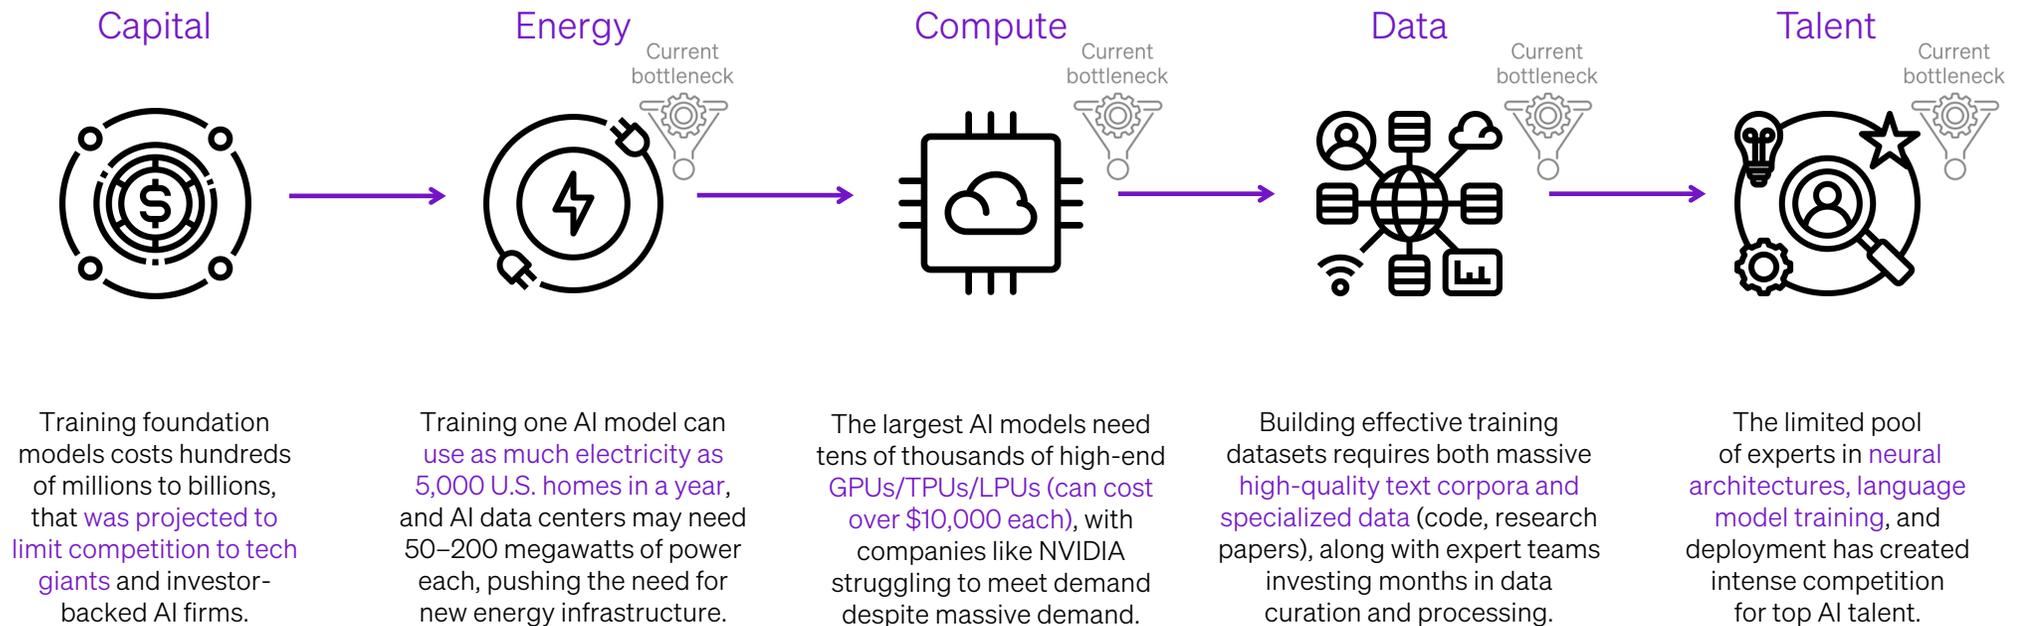
However, there is no longer exponentially more data to scrape from the internet.

Limits of LLM Scaling Based on Human Generated Data

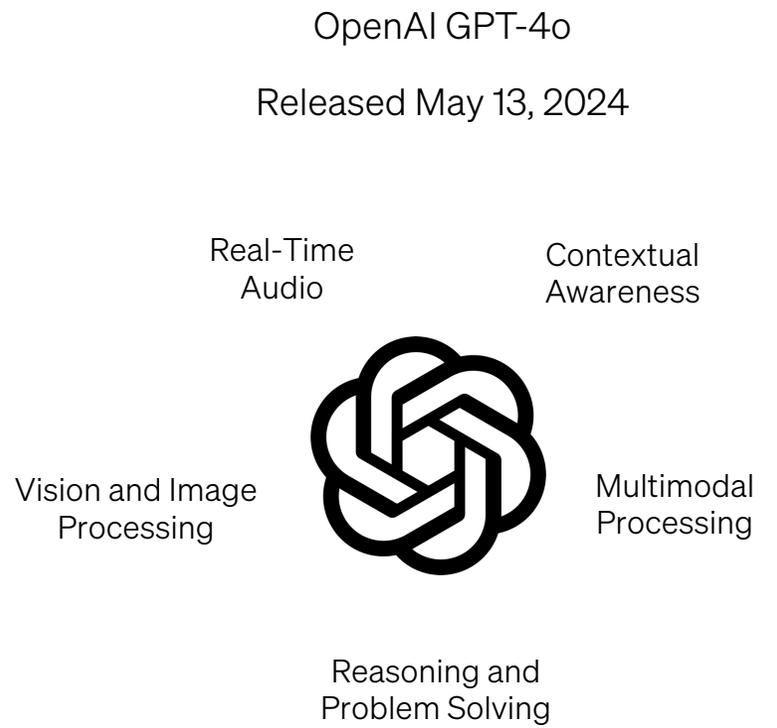


By 2028, AI models may exhaust the internet's available public training text at current growth rates.

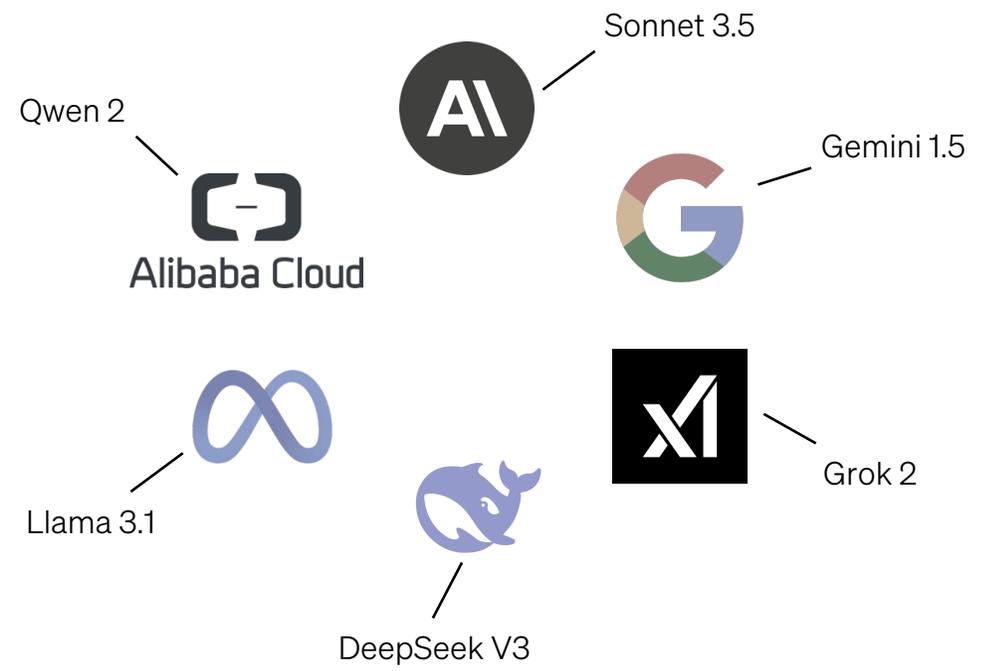
Moreover, infrastructure requirements for training and running these models are starting to become bottlenecks as well.



And even if companies cobble together all the resources necessary to create a frontier model, their capabilities will be matched by others within months.



LLM Models with Similar Capabilities
Within 6 months



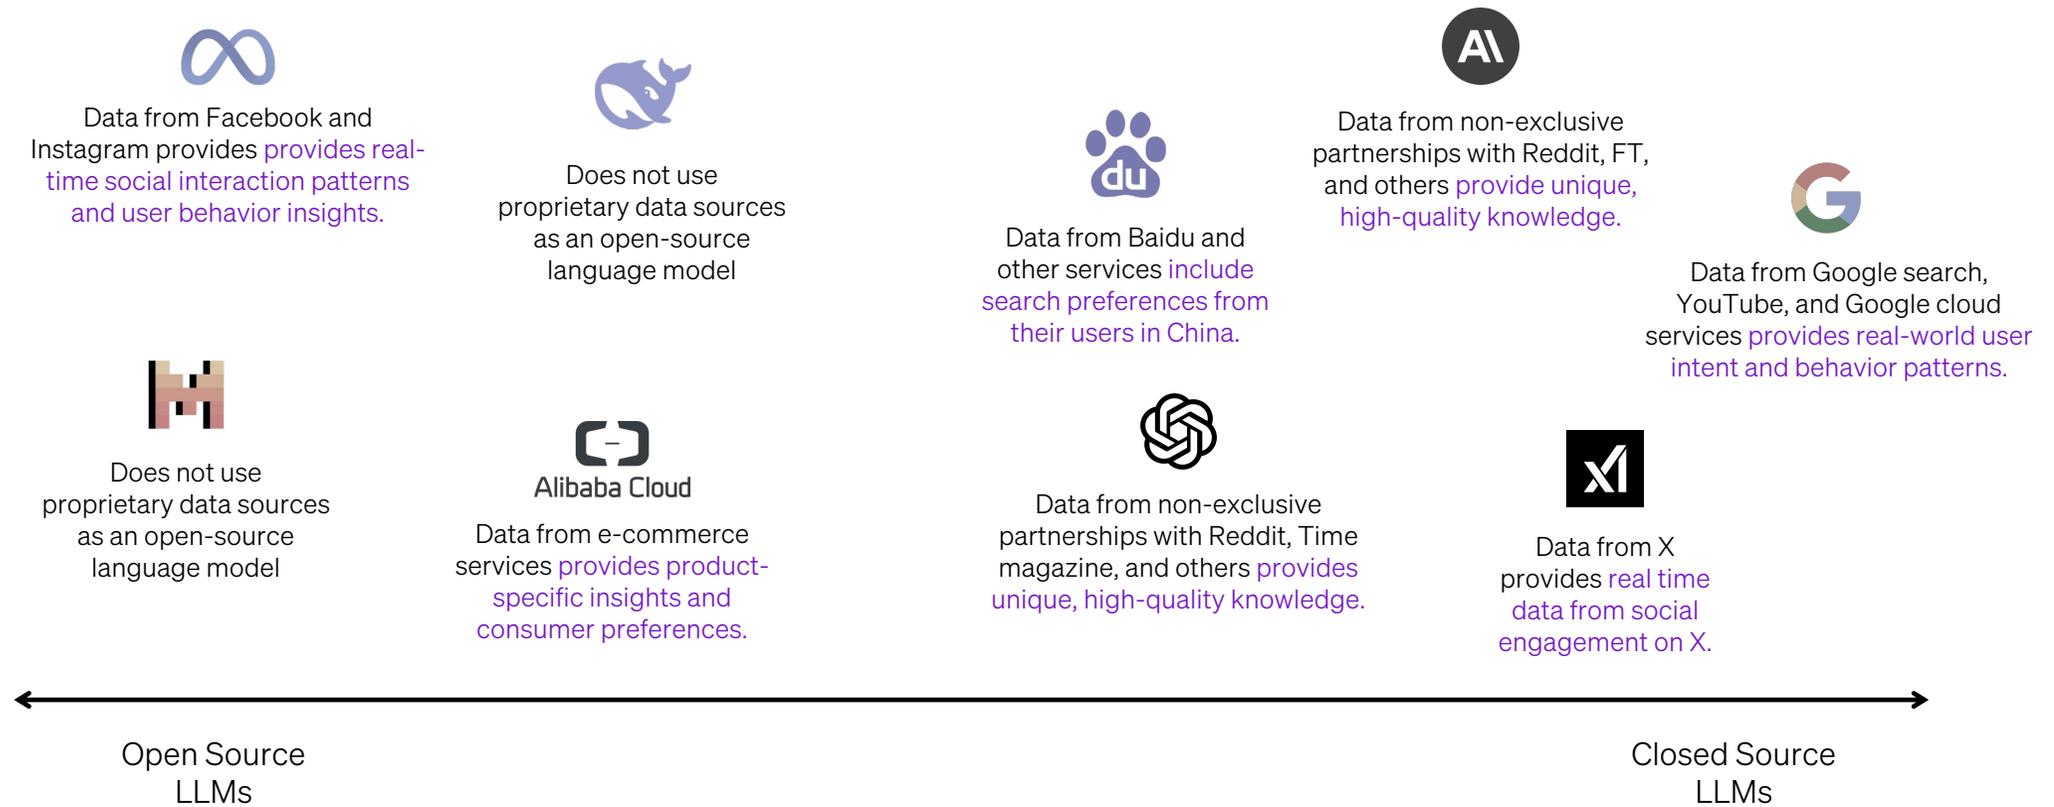
Thus, all these intelligent chatbots share similar features, with only some differentiation at the edges.



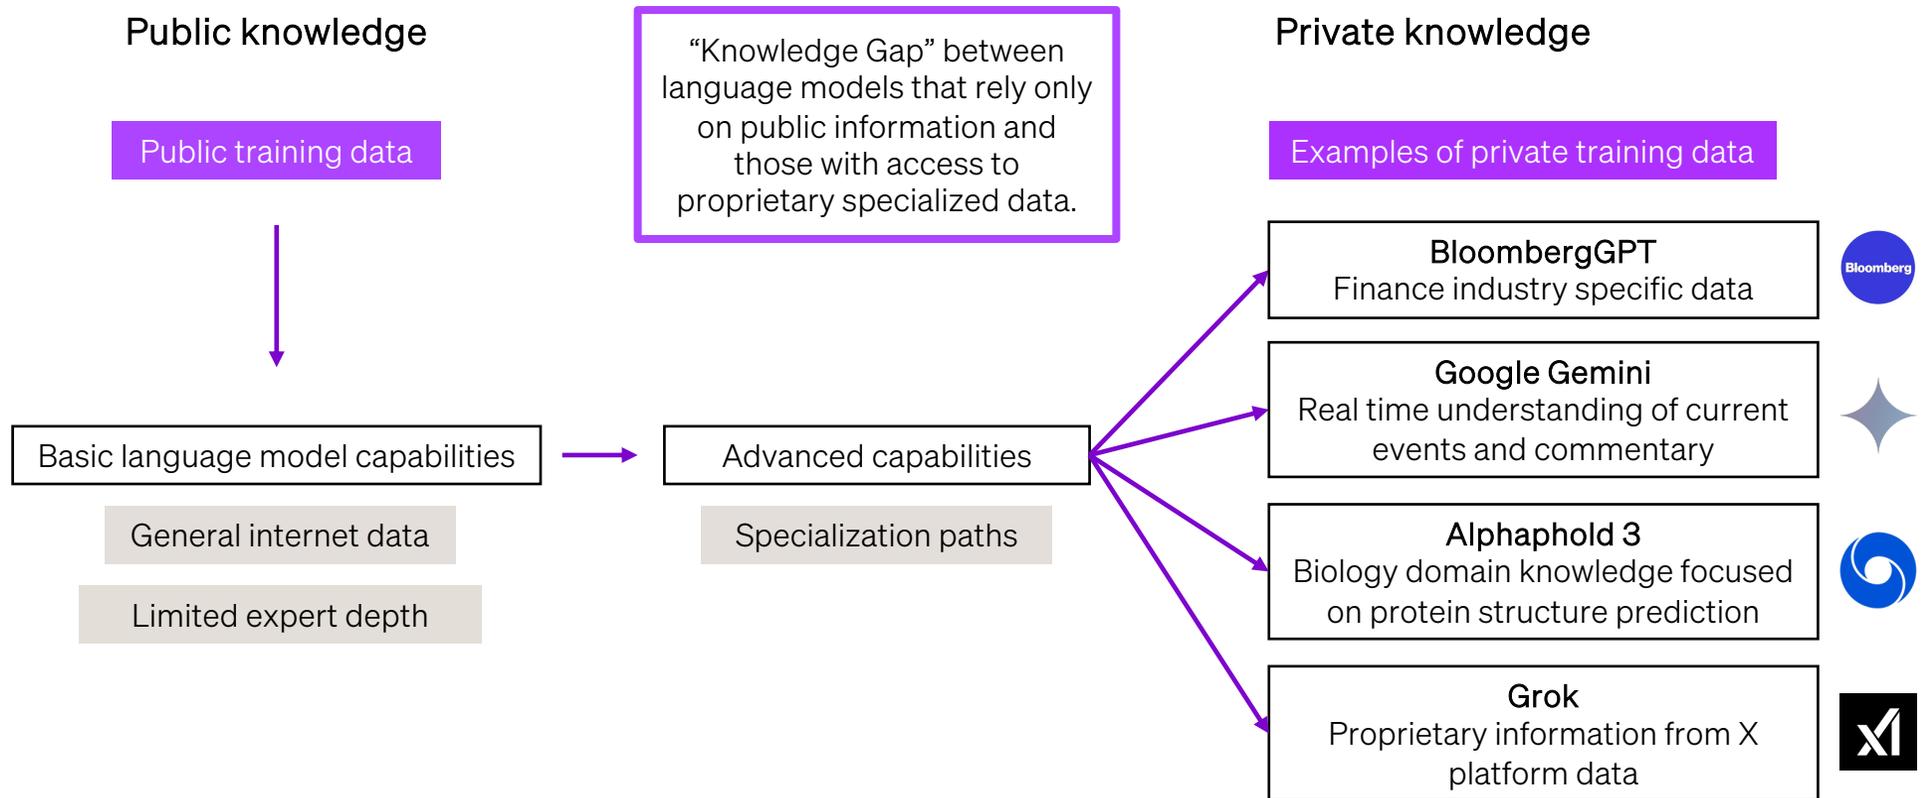
Chat Application	Input Formats	Output Formats	Knowledge	Differentiated Feature	Context Window
ChatGPT-4o	Text, PDFs, Images, Audio	Text, code, image, chart, table, docs, image	October 2023 cut off	Realistic voice conversations	128k Tokens
Claude Sonnet 3.5	Text, PDFs, Images	Text, code	April 2024 cut off	Artifacts (Interactive widgets)	100k tokens
Gemini 1.5	Text, Images	Text, code	November 2023 cut off	Integration with Google services	2m Tokens
Grok	Text, Images, X Posts	Text, image	Real time data access	Access to X data	128k Tokens

So how are foundation model providers responding to the risk of commoditization?

Companies are leveraging proprietary data sources to try to develop specialized model capabilities and build competitive moats.



This may result in closed-source language models with unique features and greater mastery in specific domains.

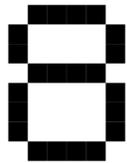


However, other companies are also leveraging clever engineering methods to improve performance and develop highly-competitive language models.

How does DeepSeek-V3 boost performance while also cutting compute and cost of training?



8-bit Efficiency



V3 used **simpler number formats** (8-bit vs 32-bit) to slash memory needs by 75% while keeping answers accurate.

Phrase Processing



Instead of analyzing words one-by-one, V3 **reads groups simultaneously** – doubling speed with only minor accuracy tradeoffs.

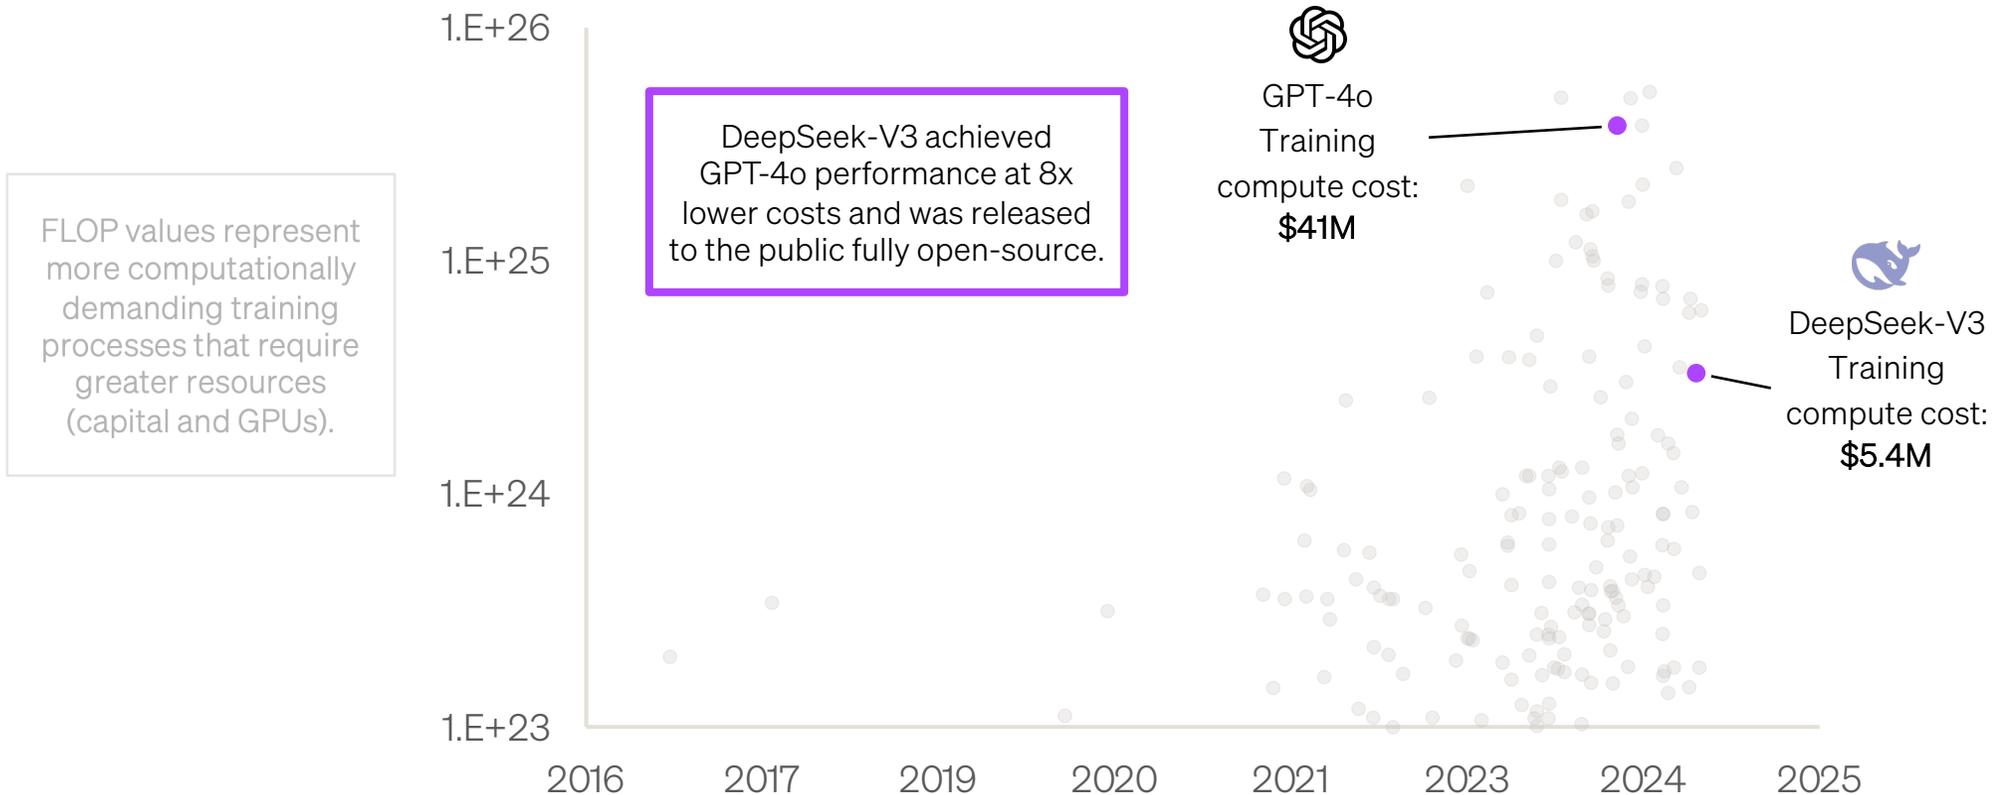
On Demand “Experts”



Unlike other models that use all components constantly, V3 activates **only 2% of its system per query**.

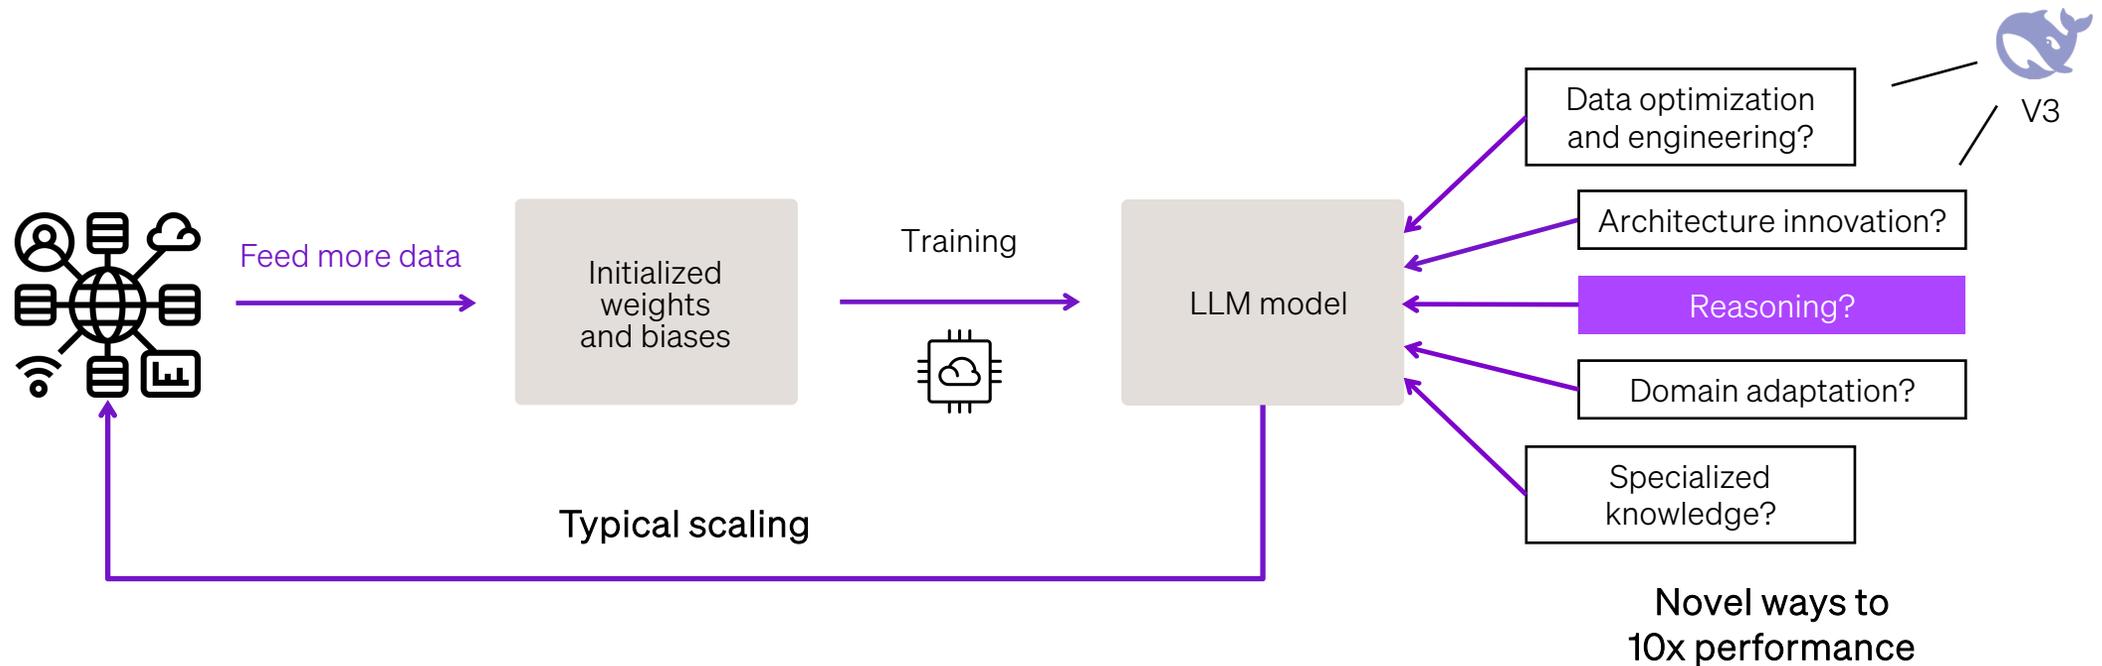
Thus, while proprietary data may specialize language models, innovative open-source models may still erode competitive moats and commoditize language models.

AI Models, Training Compute (FLOP) Over Time, 2016 to 2025



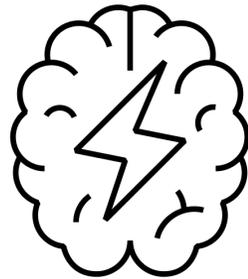
Reasoning Models

In response to the limits of the scaling paradigm, researchers started to explore novel ways to increase performance without exponentially more internet data.



Inspired by System I vs. System II thinking in the brain, researchers developed models that can "reason".

System I



Fast

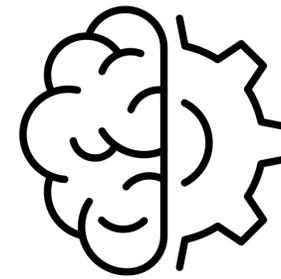
Automatic



Intuitive

Instinctual

System II



Slow

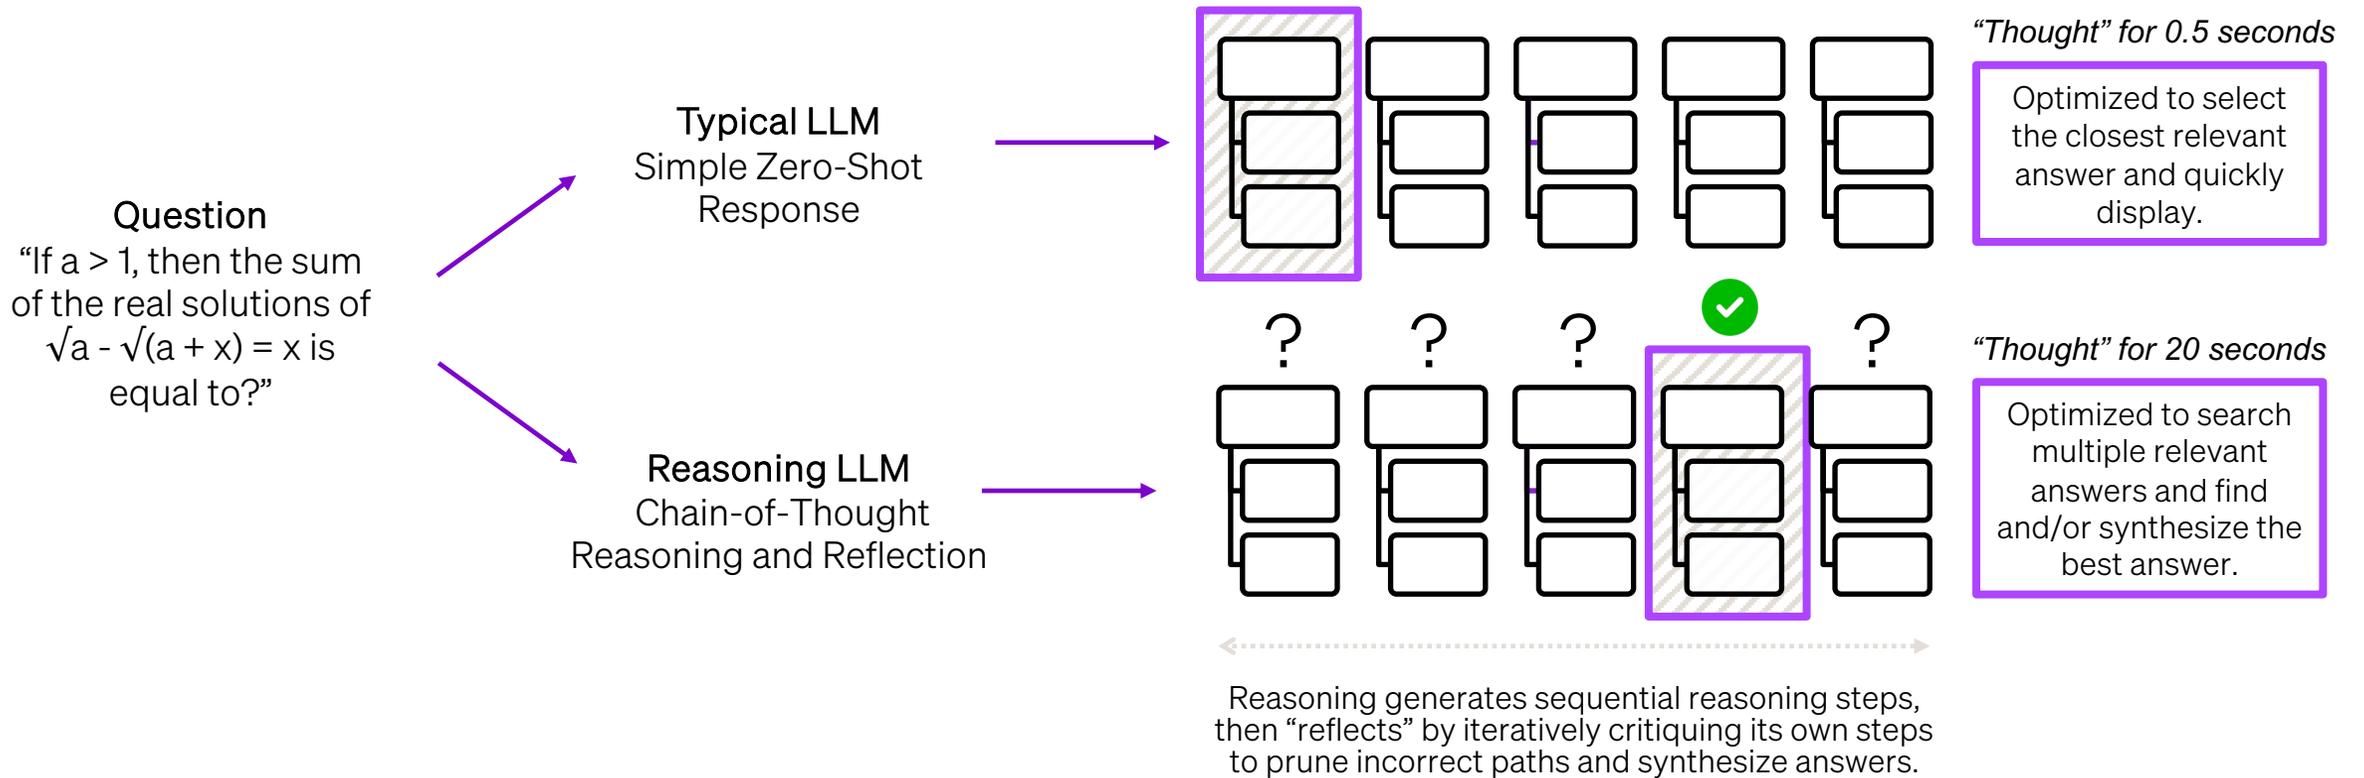
Deliberate



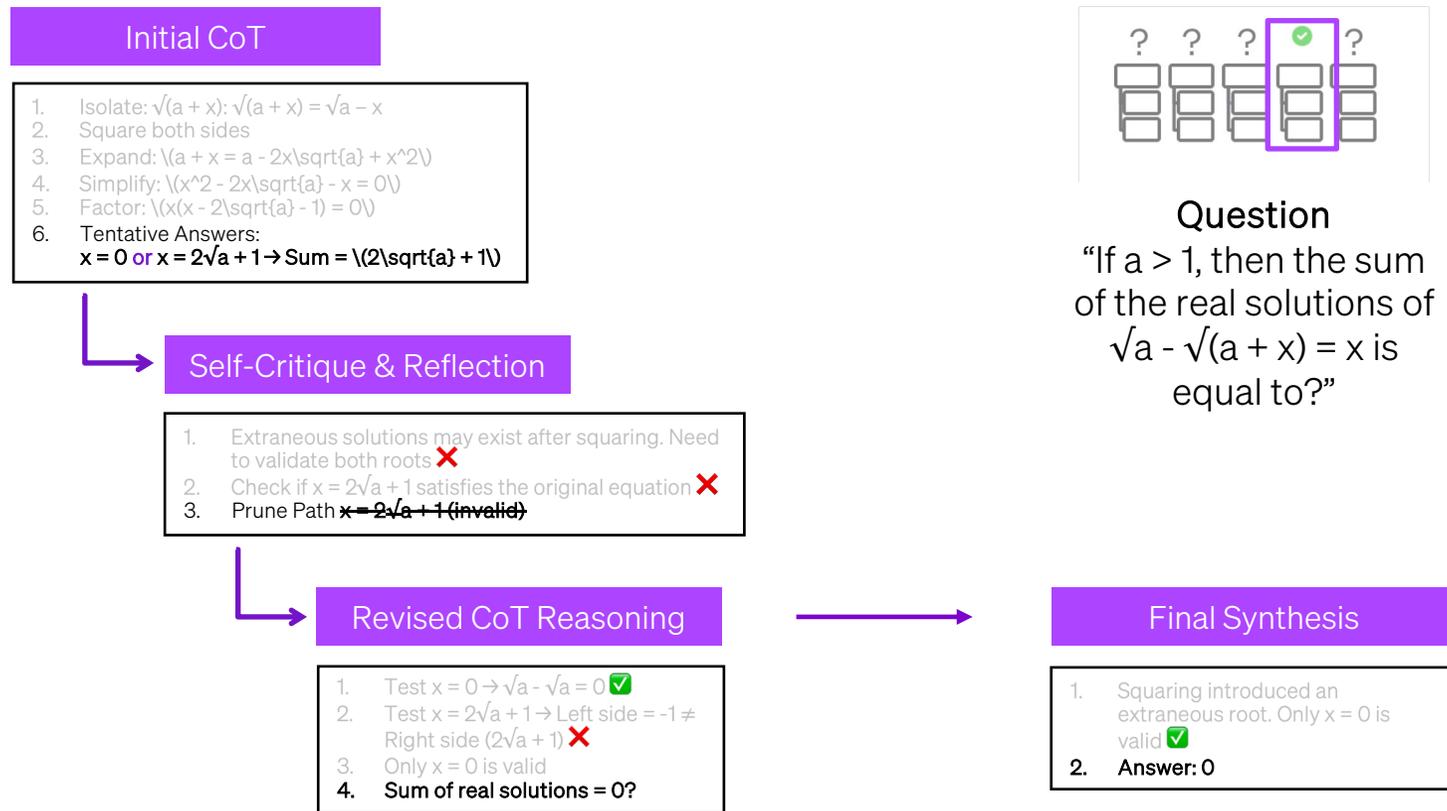
Analytical

Thoughtful

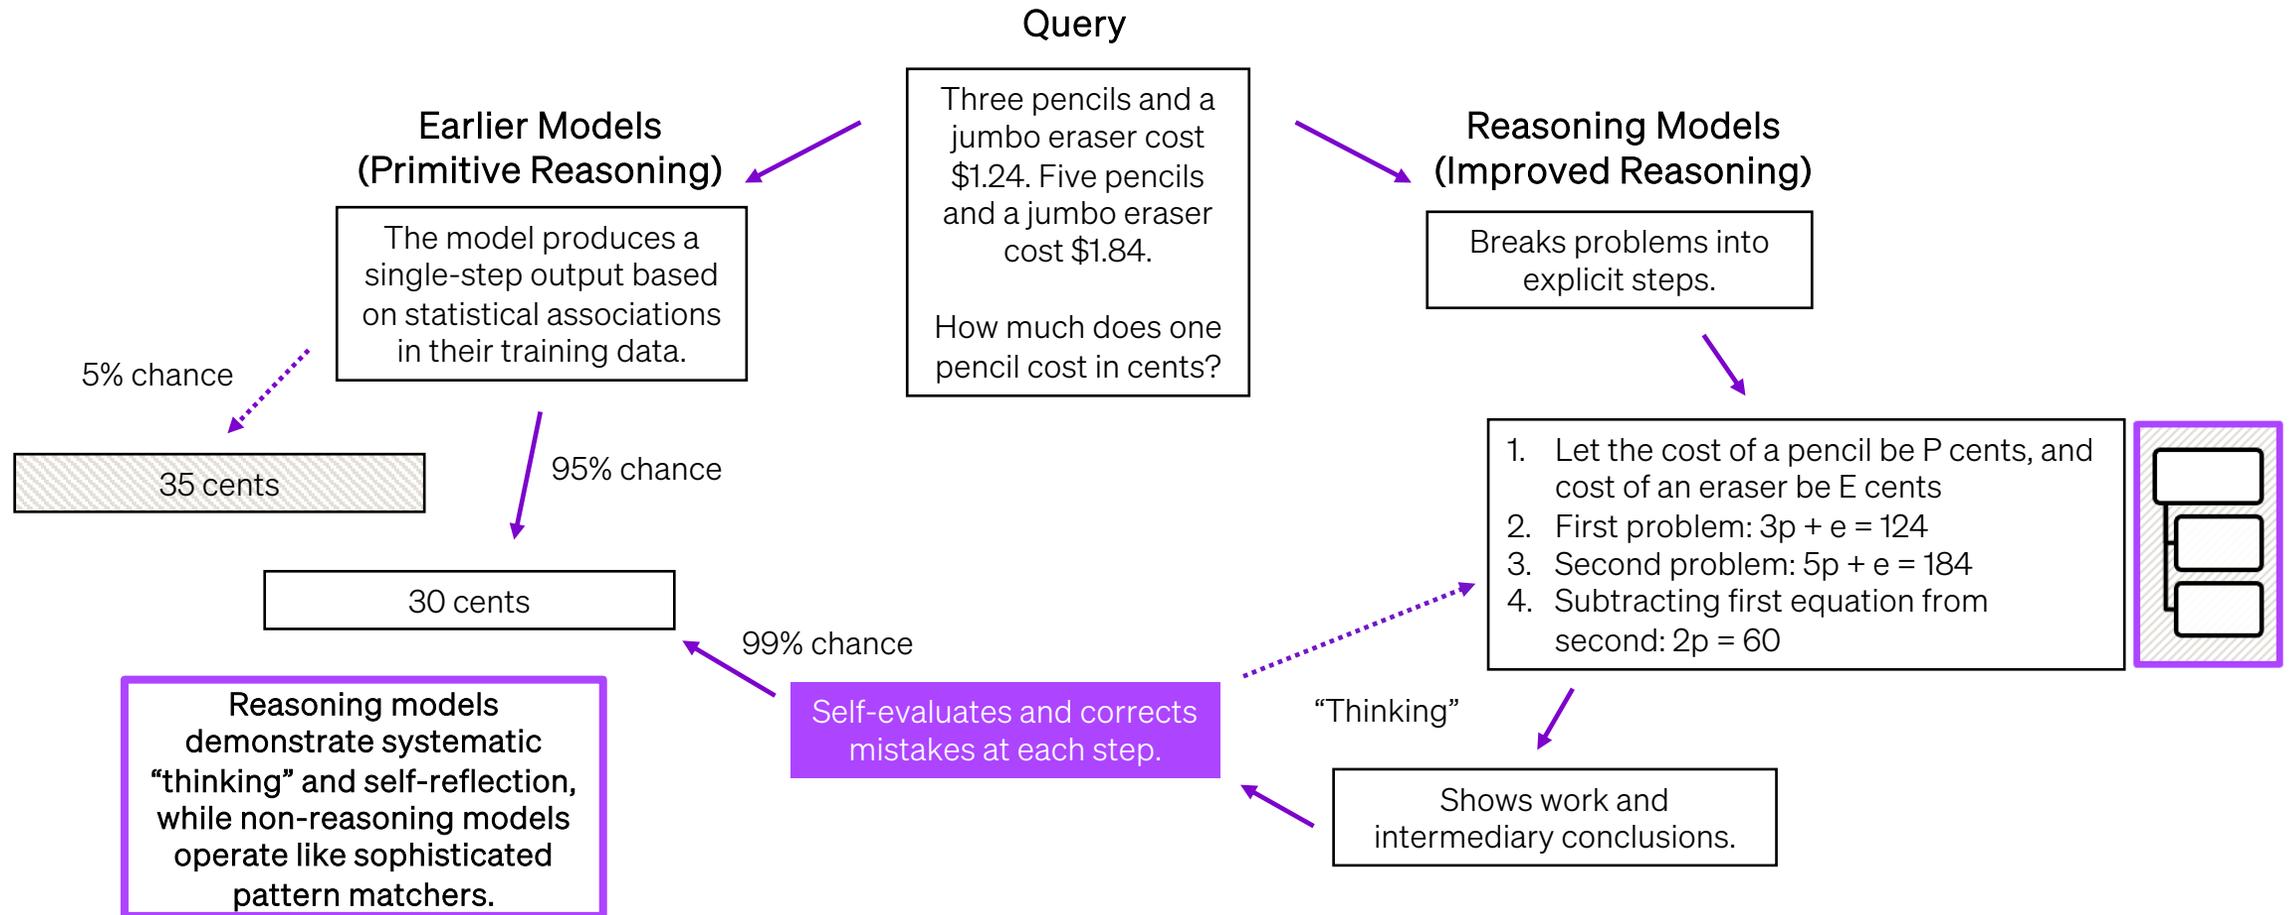
The latest reasoning models can solve problems like humans do by breaking them into smaller steps and working through each part with careful “reflection”.



In other words, reasoning generates sequential reasoning steps, then “reflects” by iteratively critiquing its own steps to prune incorrect paths and synthesize answers.

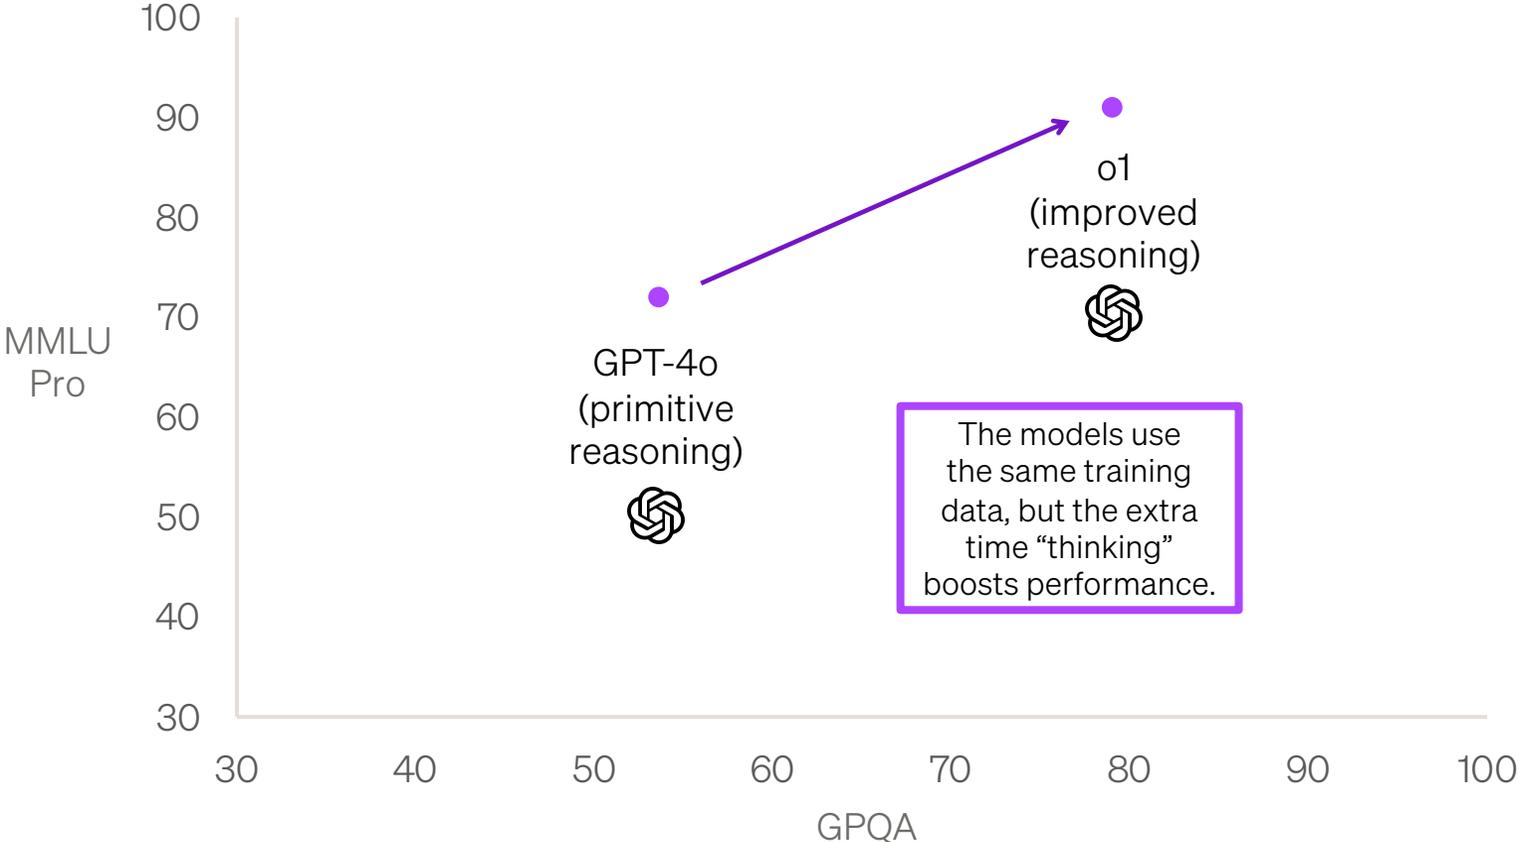


This feature marks a departure from earlier models, which retrieved and applied pre-learned knowledge without exploration and reflection.



Having a model "think" for 20 seconds gets the similar boost in performance as scaling the model, without the need for exponentially more internet data.

GPT-4o and o1, MMLU-Pro and GPQA Benchmark Scores, 0 to 100

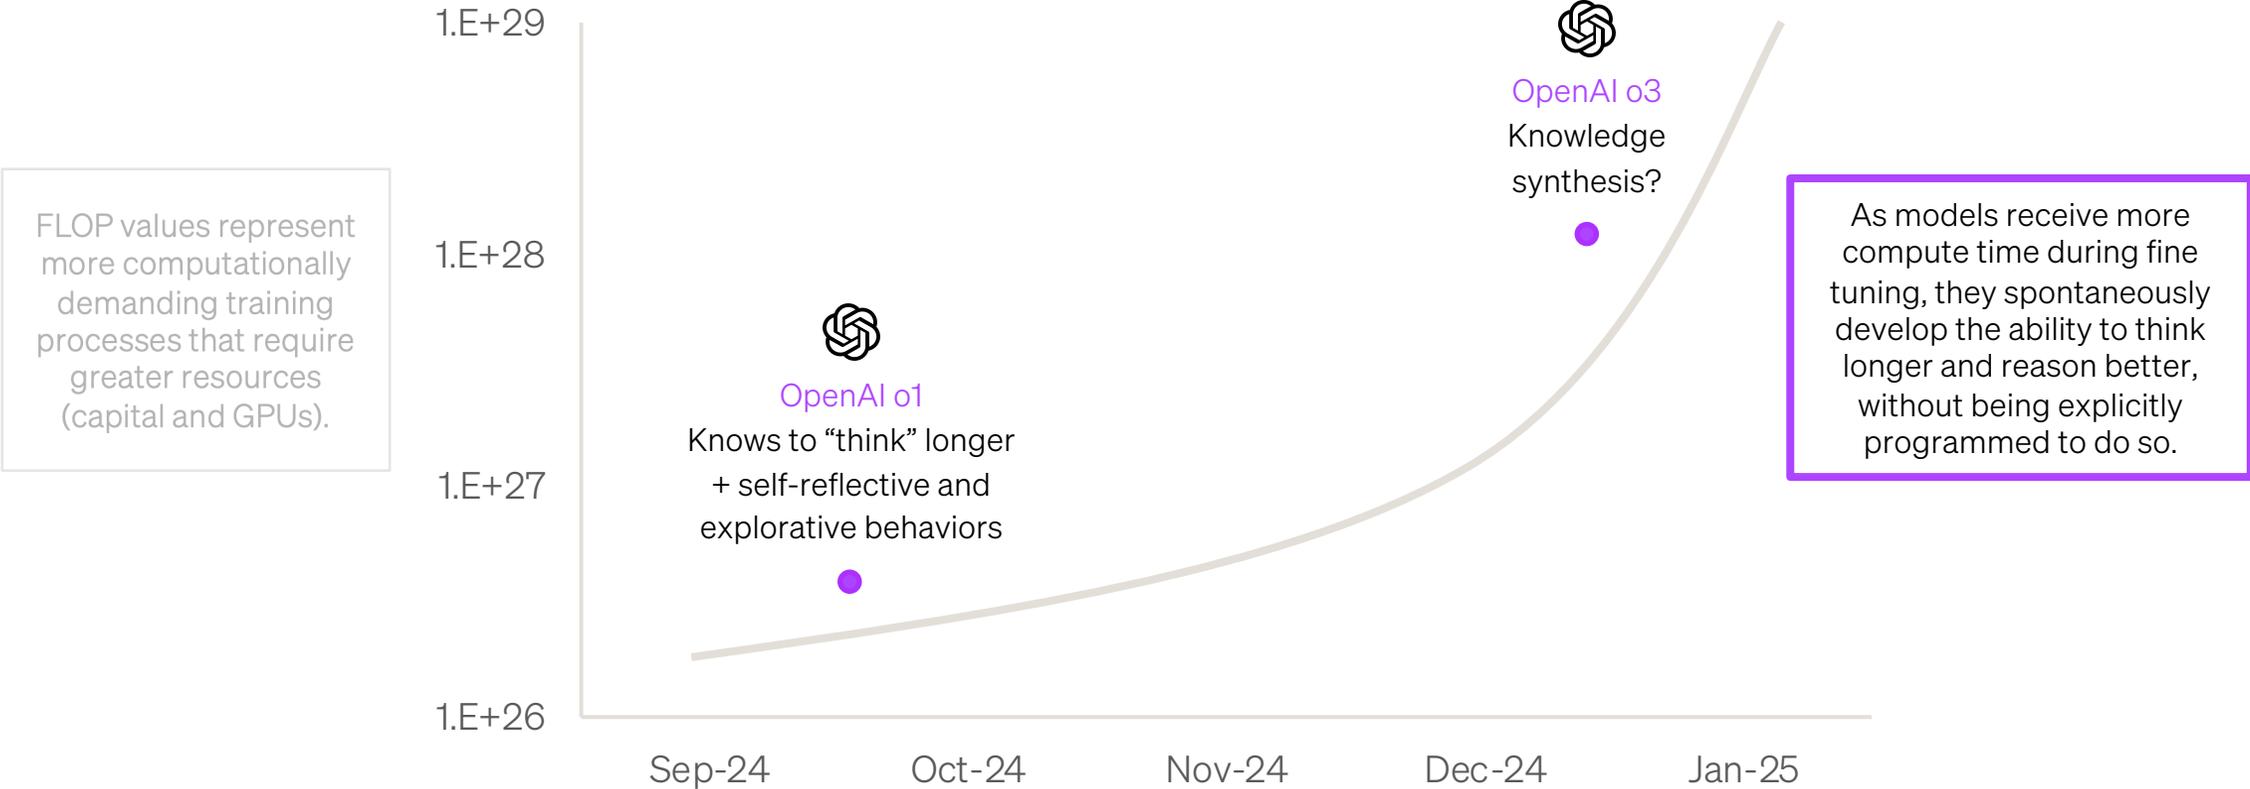


MMLU Pro:
Tests advanced reasoning and expertise across professional domains like medicine, law, engineering, and advanced academic subjects.

GPQA
Tests with a dataset of graduate-level, Google-proof multiple-choice questions in biology, physics, and chemistry, crafted by domain experts.

And as these reasoning models become more advanced, we are observing a new wave of capabilities suddenly emerging again.

Training Compute (FLOP) Required for Each OpenAI Language Model,



The improvements in benchmarks are impressive because even a small percentage improvement in technical fields represents a large jump in intelligence.

OpenAI o1



Math Problem (structured puzzle)

“The greatest common divisor of two positive integers less than 100 equals 3. Their least common multiple is twelve times one of the integers. What is the largest possible sum of the two integers?”

Case 2: $3ab = 12 \times 3b$

Similarly, if the LCM equals $12 \cdot B = 12 \cdot (3b)$, then

$$3ab = 36b \implies ab = 12b \implies a = 12,$$

(provided $b \neq 0$).

Hence $A = 3a = 3 \cdot 12 = 36$. Again we need $\gcd(12, b) = 1$.

- Because $B = 3b < 100$, we have $3b < 100$ so $b \leq 33$.
- We want b to be as large as possible while coprime to 12.

By the same reasoning, the largest $b \leq 33$ not divisible by 2 or 3 is 31. Hence

$$B = 3 \times 31 = 93, \quad A = 36.$$

Their sum again is $36 + 93 = 129$.

3. Conclusion

In both cases, the largest possible sum of the two integers (under the constraint that each is less than 100, their gcd is 3, and their lcm equals 12 times one of them) is

129.

OpenAI o3



Math Problem (research-grade challenge)

“How many nonzero points are there on $x^3y + y^3z + z^3x = 0$ over \mathbb{F}_5^3 up to scaling?”

4. The final answer

Hence the final, succinct statement is that:

The number of nonzero projective solutions $(x : y : z)$ in $\mathbb{P}^2(\mathbb{F}_{5^{13}})$ to

$$x^3y + y^3z + z^3x = 0$$

is

$$\boxed{5^{18} + 6 \cdot 5^9 + 1}.$$

Equivalently, one writes

$$|C(\mathbb{F}_{5^{13}})| = 5^{18} + 6 \cdot 5^9 + 1,$$

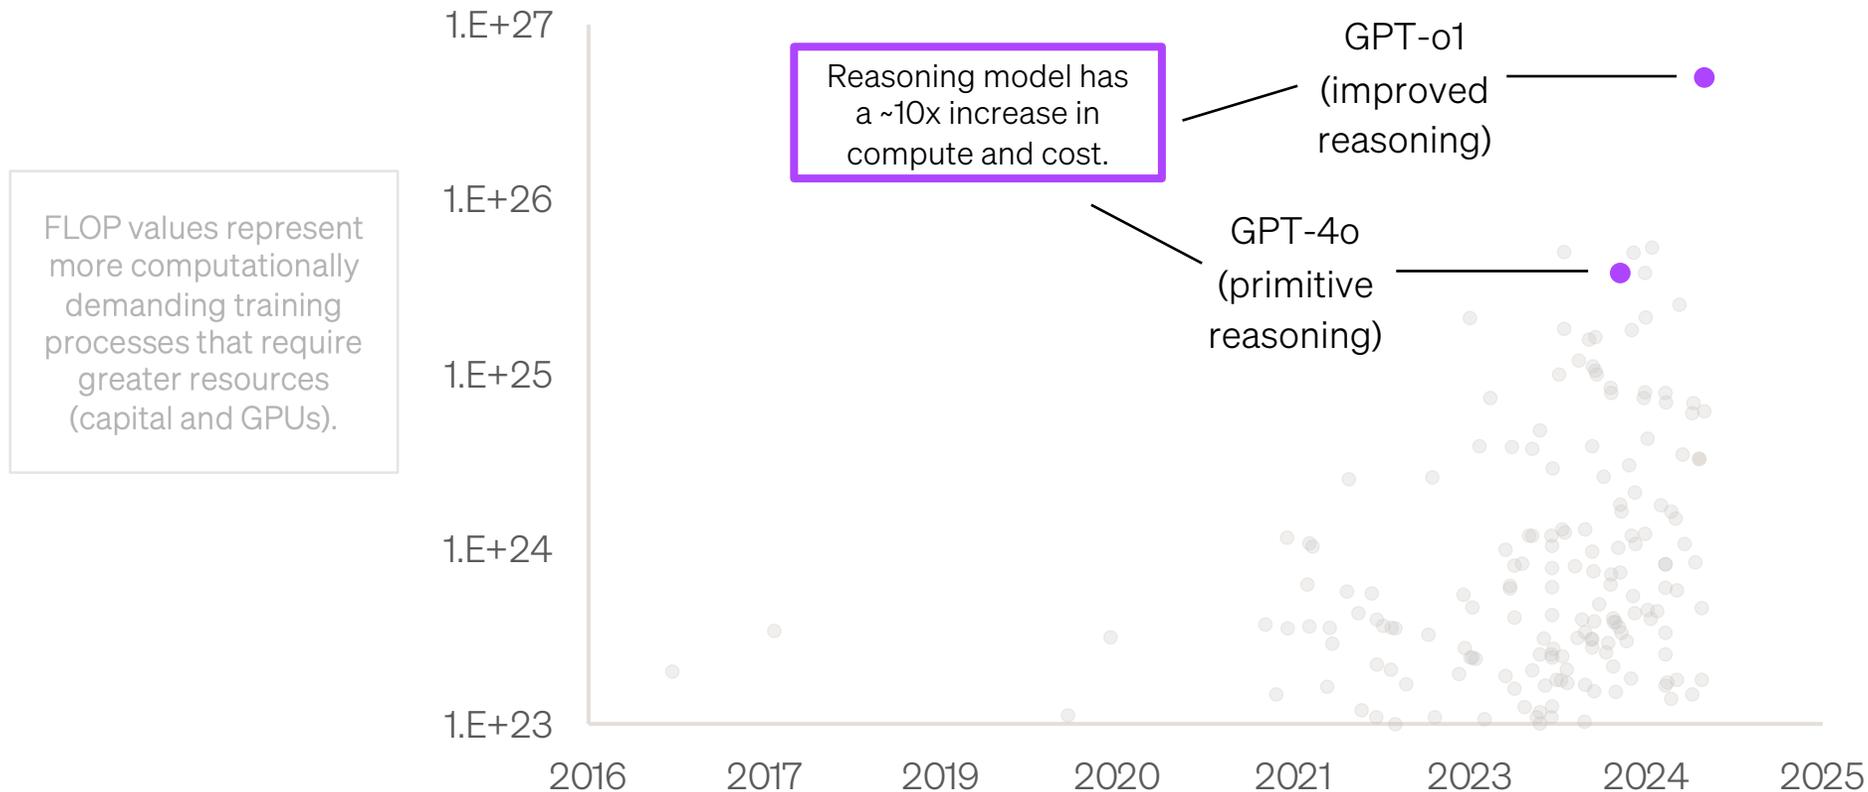
matching the provided hint.

🔊 📄 📌 🗑️ 🔄

According to benchmarks, o3 was able to solve 25% of FrontierMath T3 research grade problems (problems like the one shown).

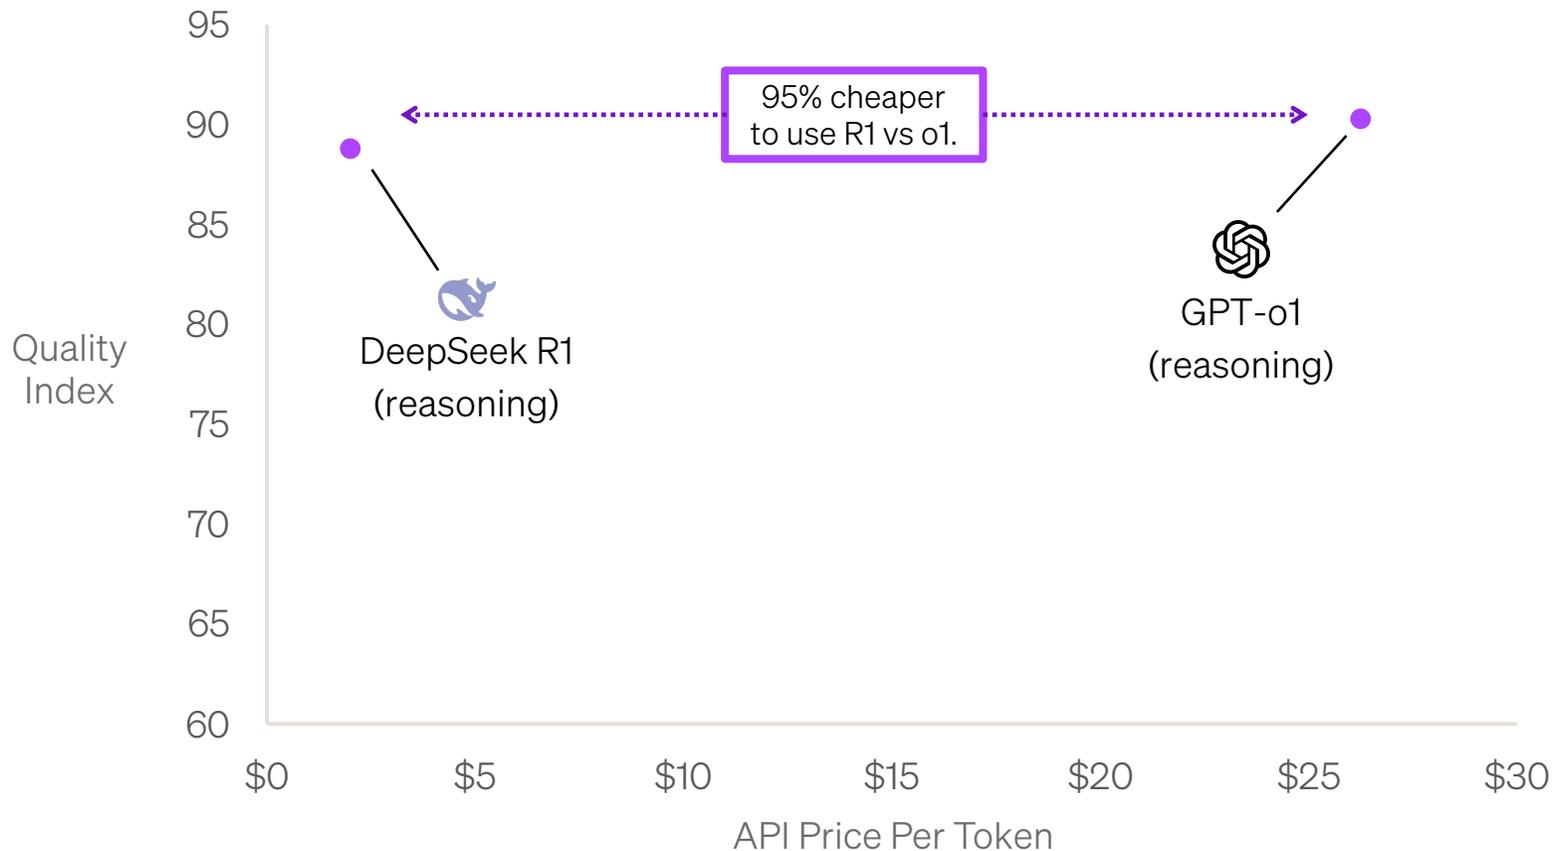
This approach doesn't bottleneck in terms of Internet data needed, but it still does in terms of the compute necessary to train and run the models.

Large Scale AI Models, Training Compute (FLOP) Over Time, 2016 to 2025

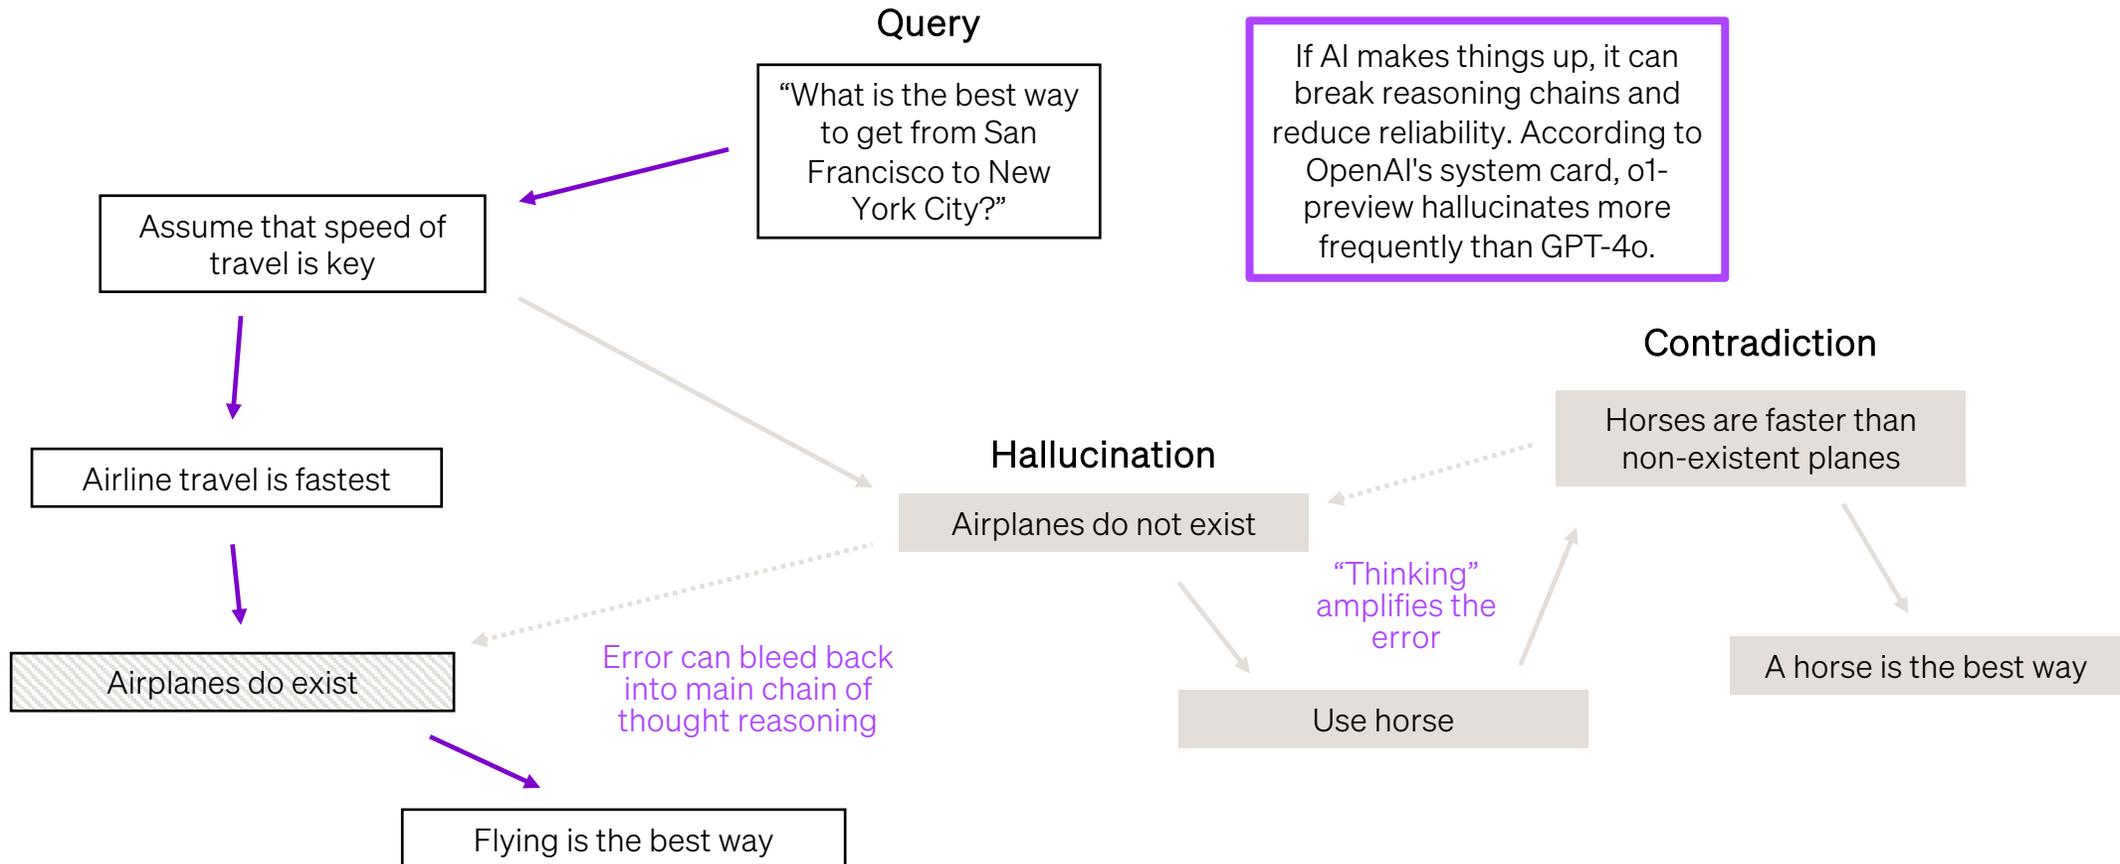


DeepSeek's R1 model achieved advanced reasoning with greater computational efficiency than other leading reasoning models, reducing training and operational costs.

Language Model Quality vs Usage Price, January 2025

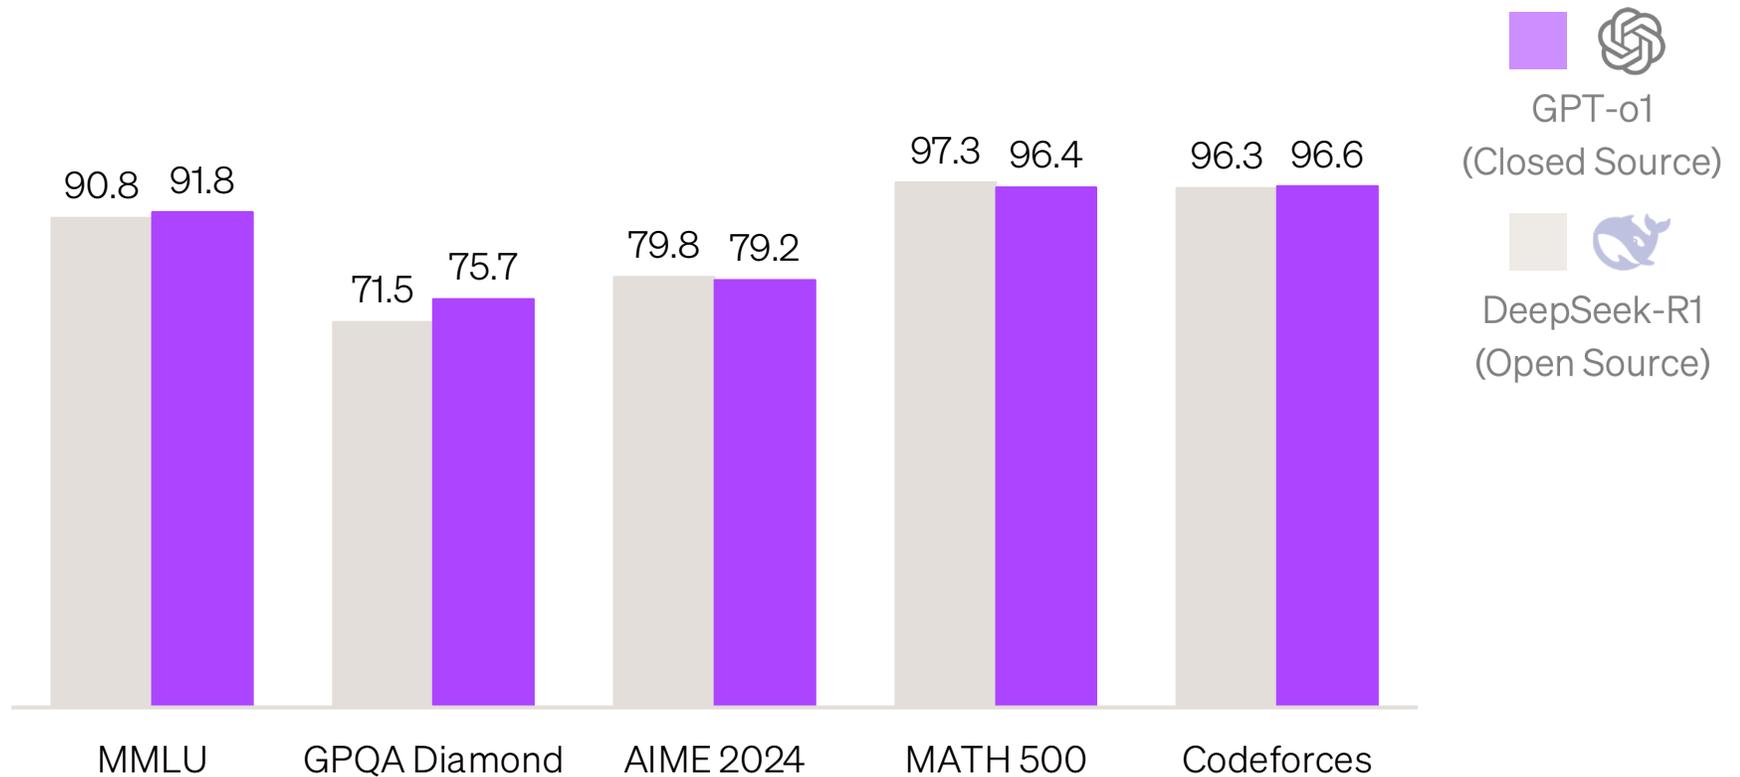


Moreover, like other language models, these reasoning models can also hallucinate.



Over time, as the availability of high-quality reasoning models increase, they may become commodities as well.

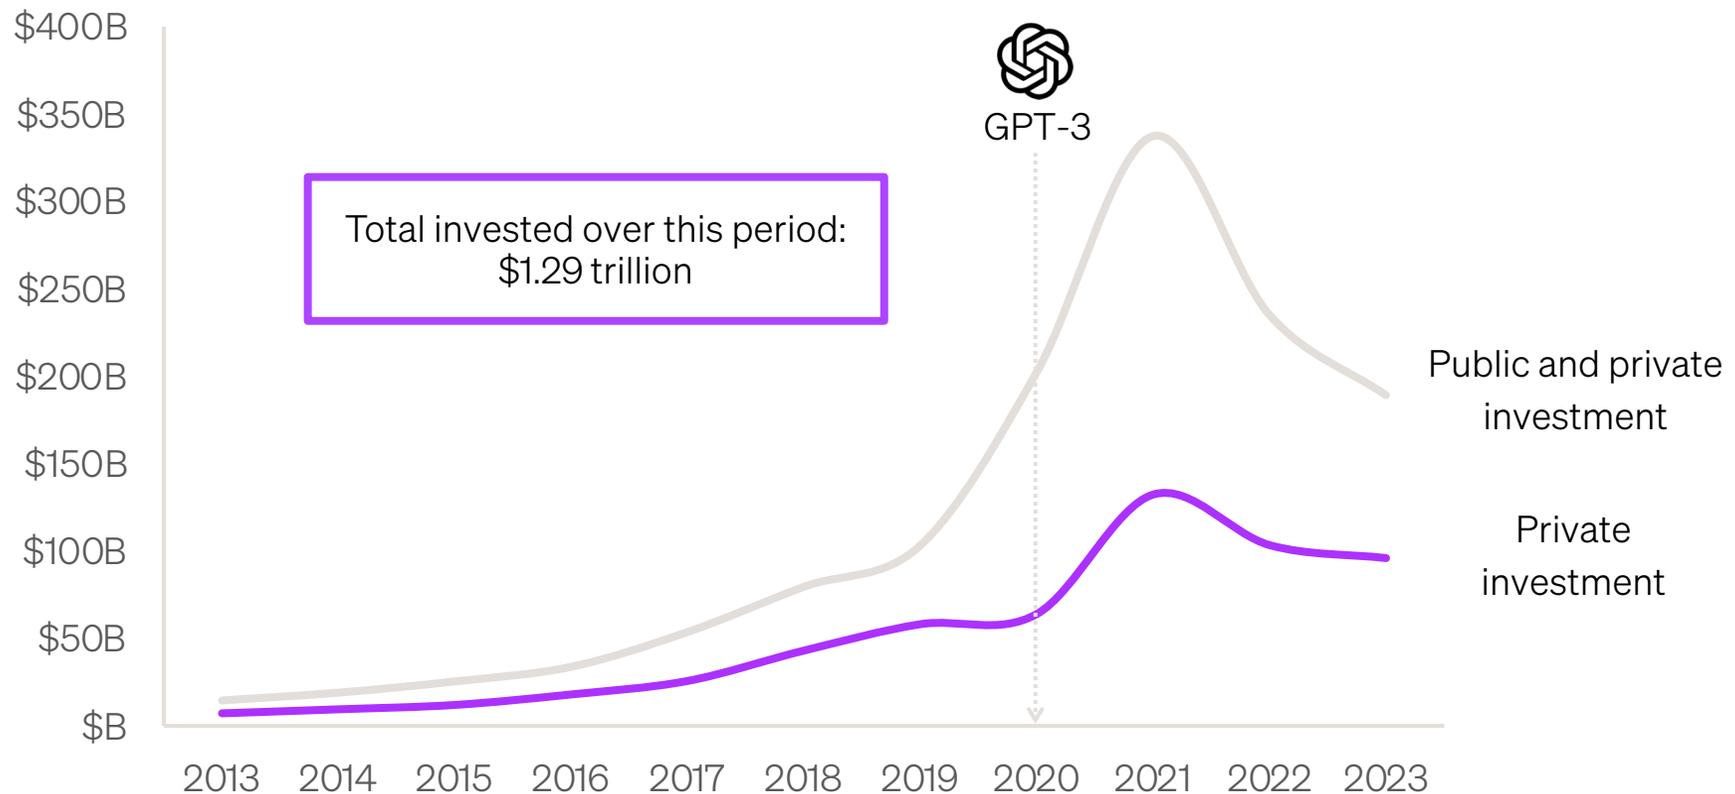
Benchmark Performance, Score 0-100



Business of AI

All the potential use cases of these capabilities and the pace of advancement compelled people to invest more than a trillion dollars into AI.

Public and Private Investment in AI, Billions of U.S. Dollars, 2013 to 2024



This led to a wave of experimentation - applications that used language models as component in a system.

<u>Coding</u>	<u>Conversational</u>	<u>Healthcare</u>	<u>Legal</u>	<u>Marketing</u>
 Cursor AI	 ChatGPT	 Ada	 Ironclad	 Writesonic
 Replit Agent	 Claude	 Path AI	 Evenup	 Jasper
 Github Copilot	 Preplexity	 Deepscribe	 Casetext	 Copysmith
 Codeium	 Gemini	 Glass Health	 Harvey AI	 Rytr
 Cognition	 Copilot	 Memora	 Pincites	 Twain

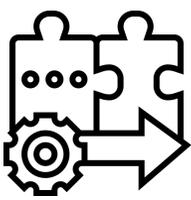
When models shift from powering chatbot conversations (where human oversight can catch errors) to becoming components in enterprise software, new challenges emerge.

Data Security



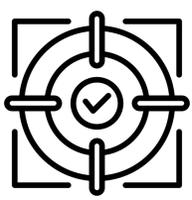
Enterprises cannot risk exposing sensitive internal data to public language models during training or inference.

Integration



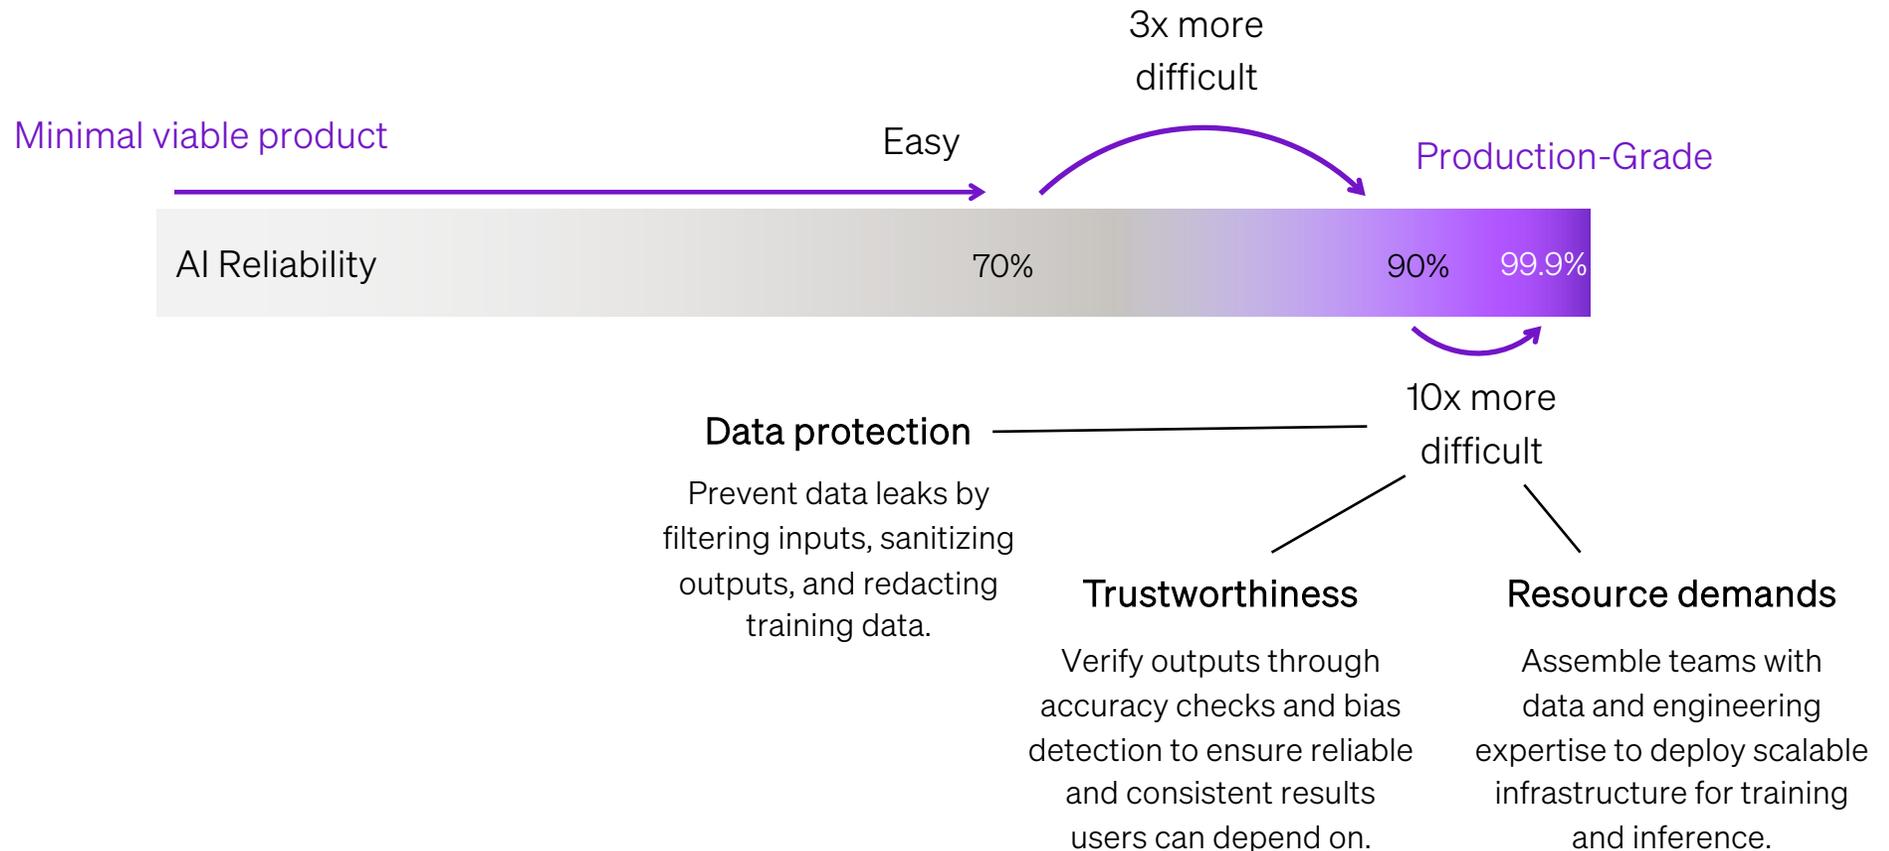
Connecting language models to existing enterprise infrastructure requires complex technical integration work.

Accuracy



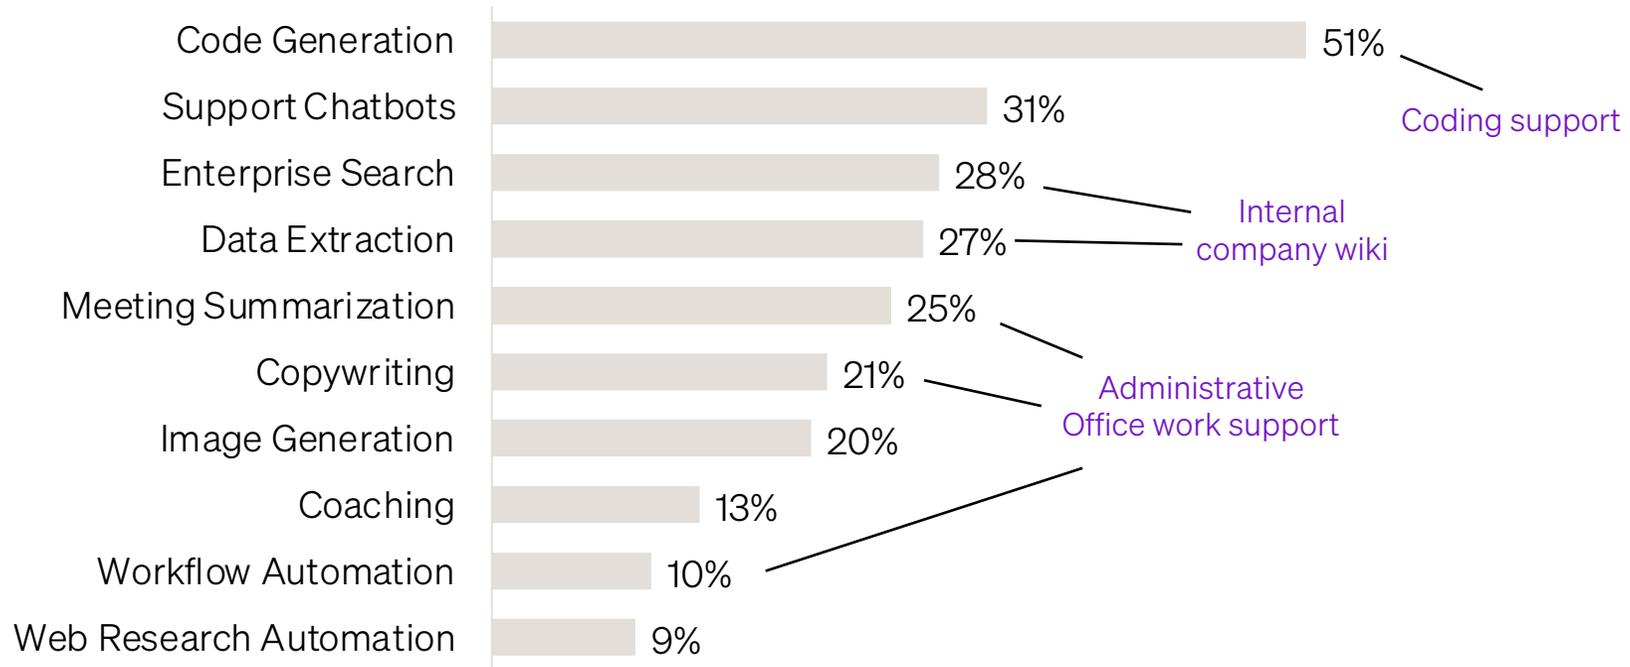
Enforce precise financial computations and ensure mission-critical decisions are accurate through structured controls.

Even if building a minimal viable product is easier than before, building production-grade applications still requires rigorous hardening and testing.



Thus, the first wave of “generative AI” apps saw established companies explore internal use cases for language models, with only internal tools gaining traction.

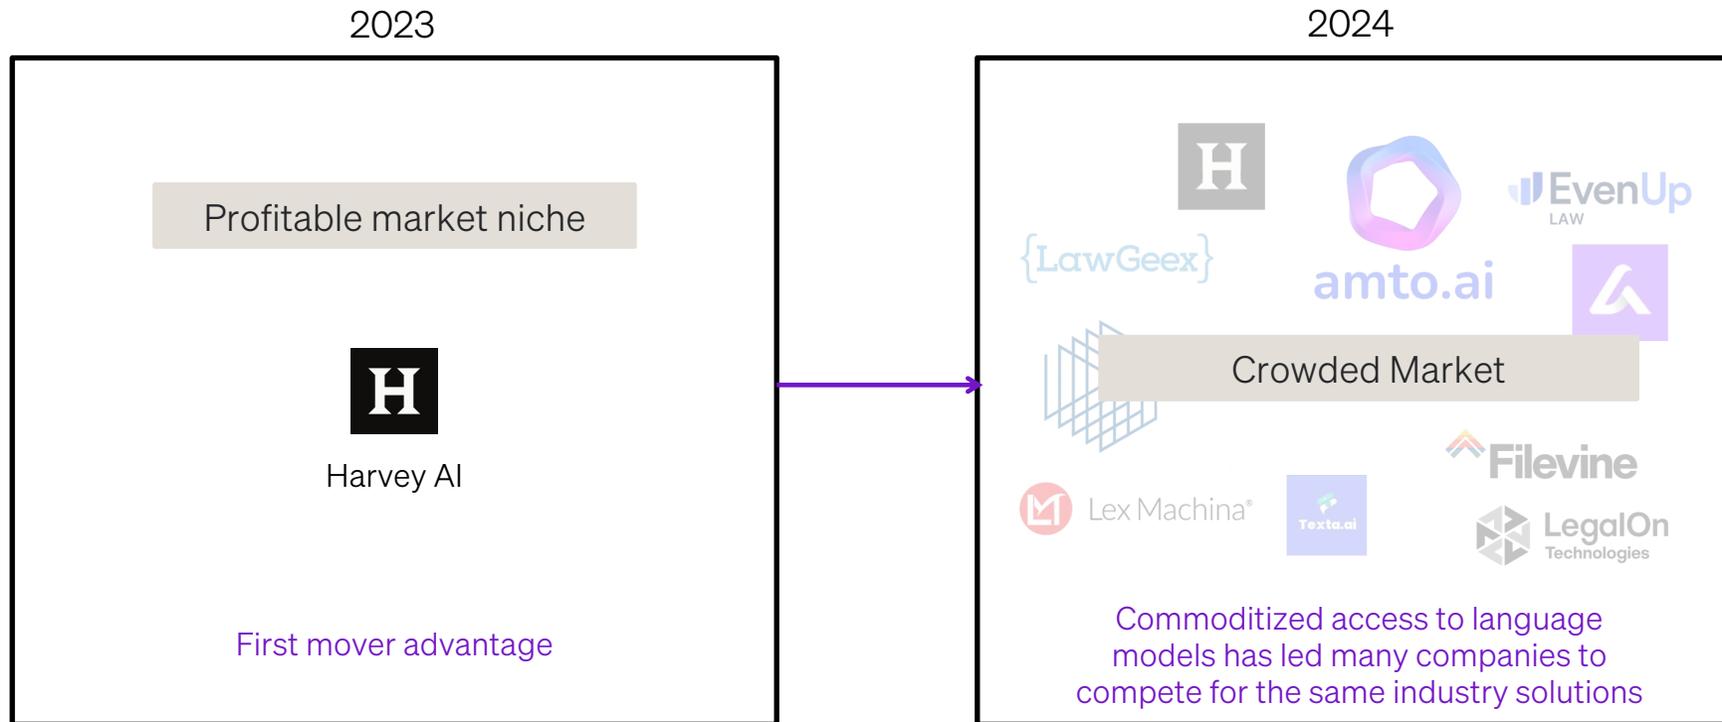
Dominant AI Use Cases at Companies, Percentage Adoption



Internal tooling that make employees variably more productive.

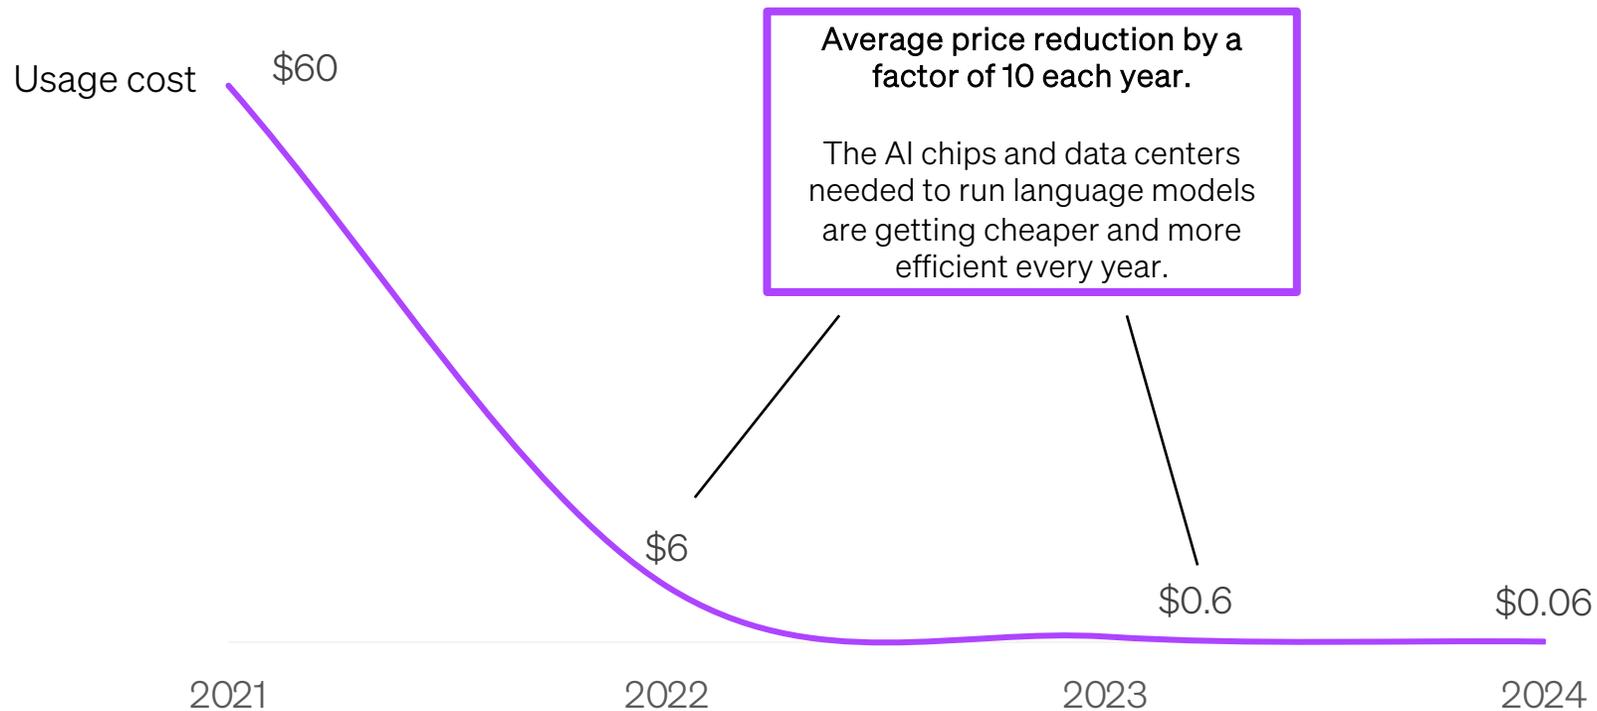
Startups were formed to build AI products, but many of their products face the same pricing pressures as the language models themselves.

What happens when there is a low barrier of entry to create an MVP?



So even if the cost of running these companies has gone down, competition and low barriers-to-entry put them at risk of a race to the bottom.

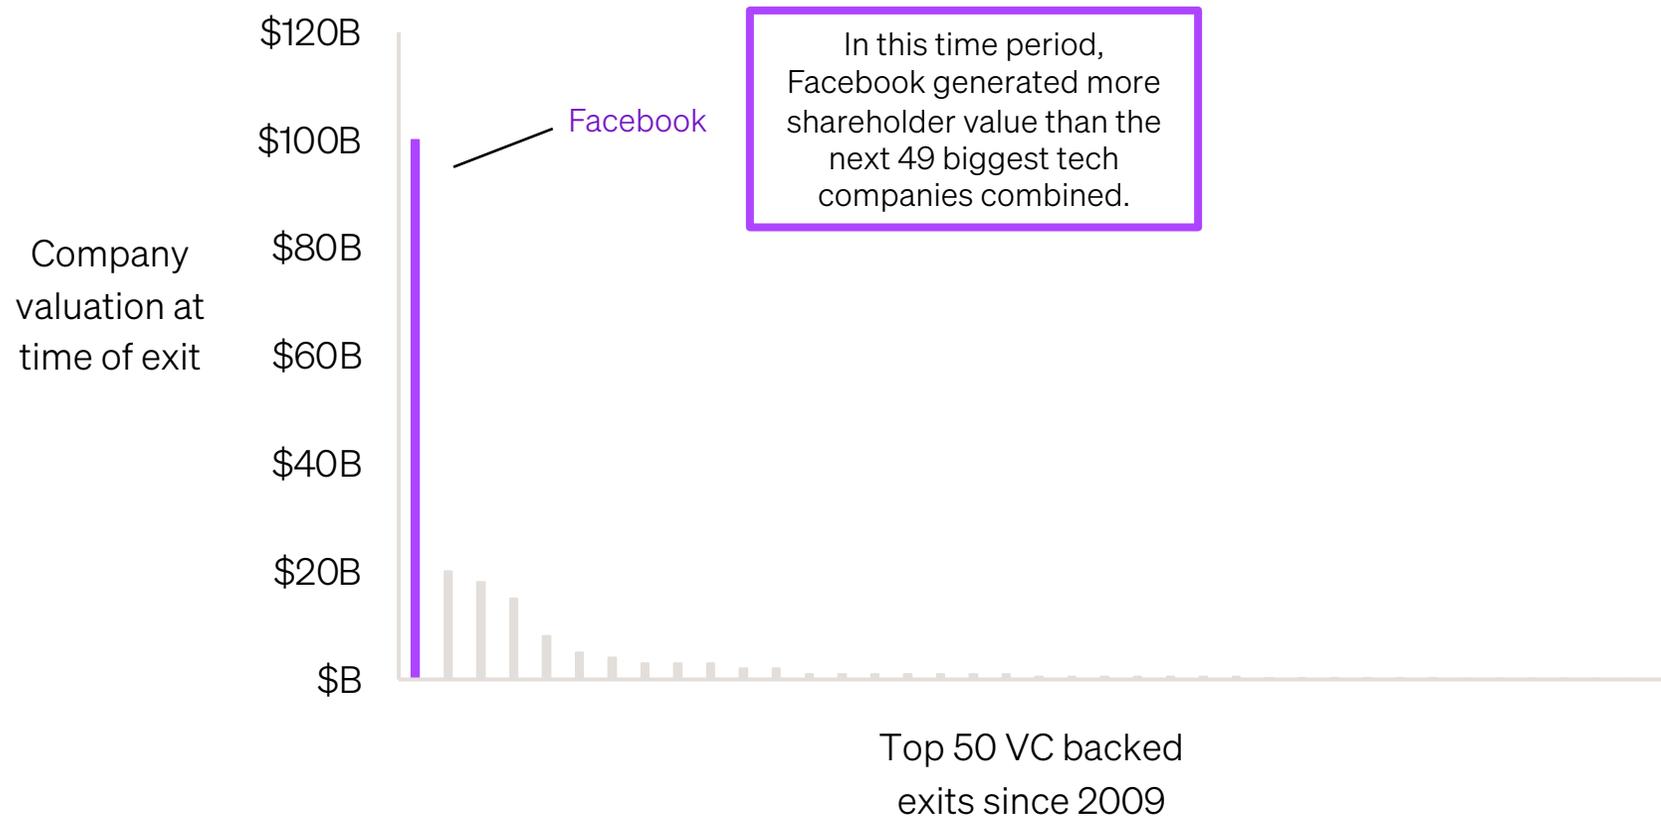
Lowest Available Price of Language Model Usage, Price Per Token, U.S. Dollars, 2021 to 2024



So, what could be the hallmarks of
a profitable billion-dollar AI company?

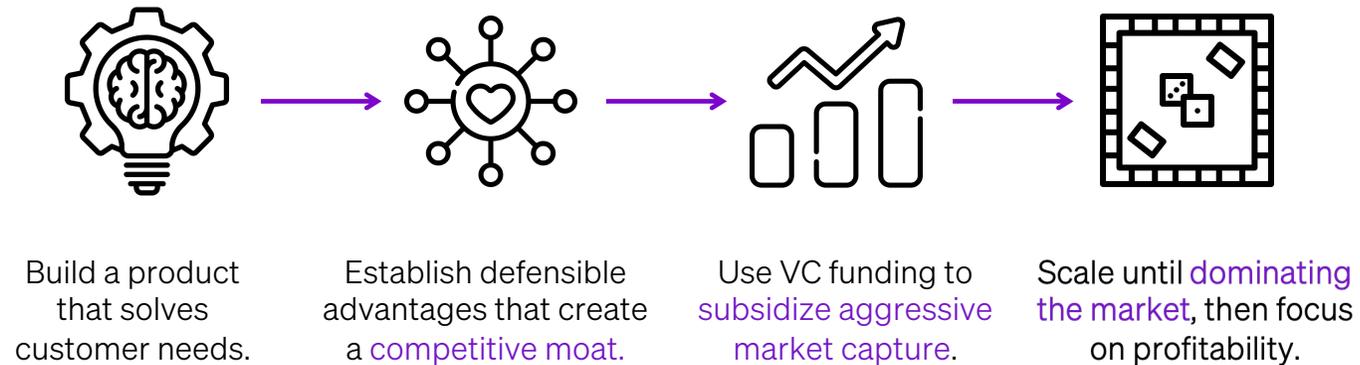
Traditionally, great investments in venture capital follow the power law, where the top 0.1% of companies generate most of the returns in a portfolio.

Top 50 VC Backed Exits From 2009 to 2014, Valuation at The Time of Exit

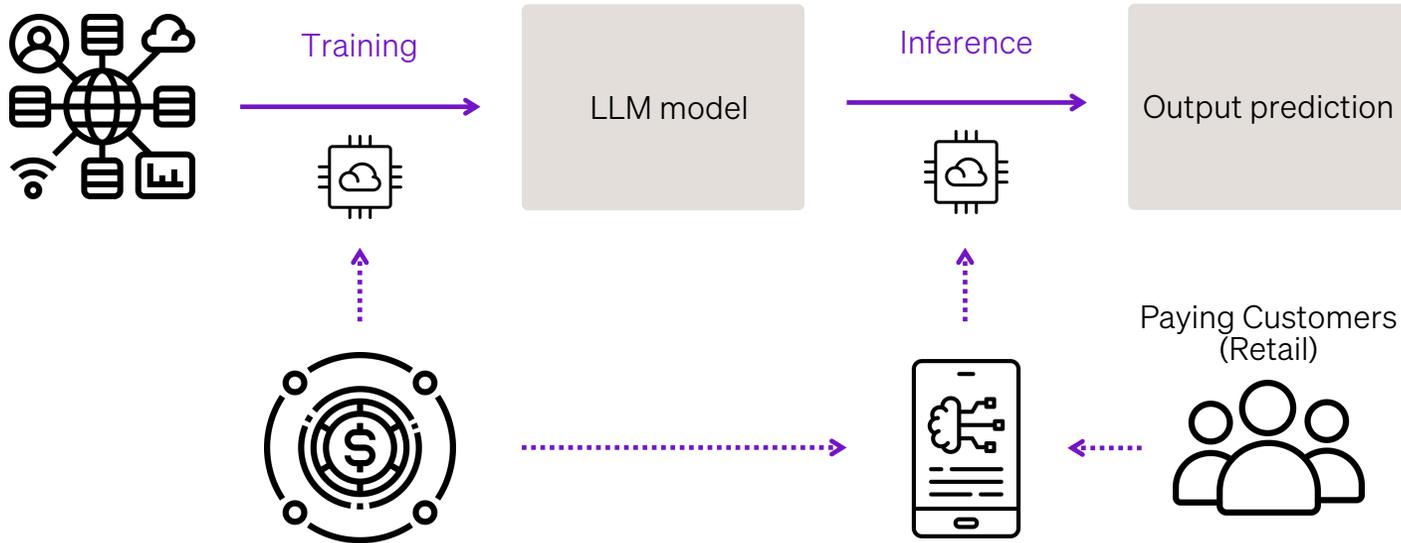


The underlying assumption of the power law relies on forming monopolies with moats and network effects - something AI products currently struggle to attain.

Playbook For Previous Waves of Technology



In other words, these products must eventually have a base of paying customers paying a premium to justify all the investment dollars poured in.

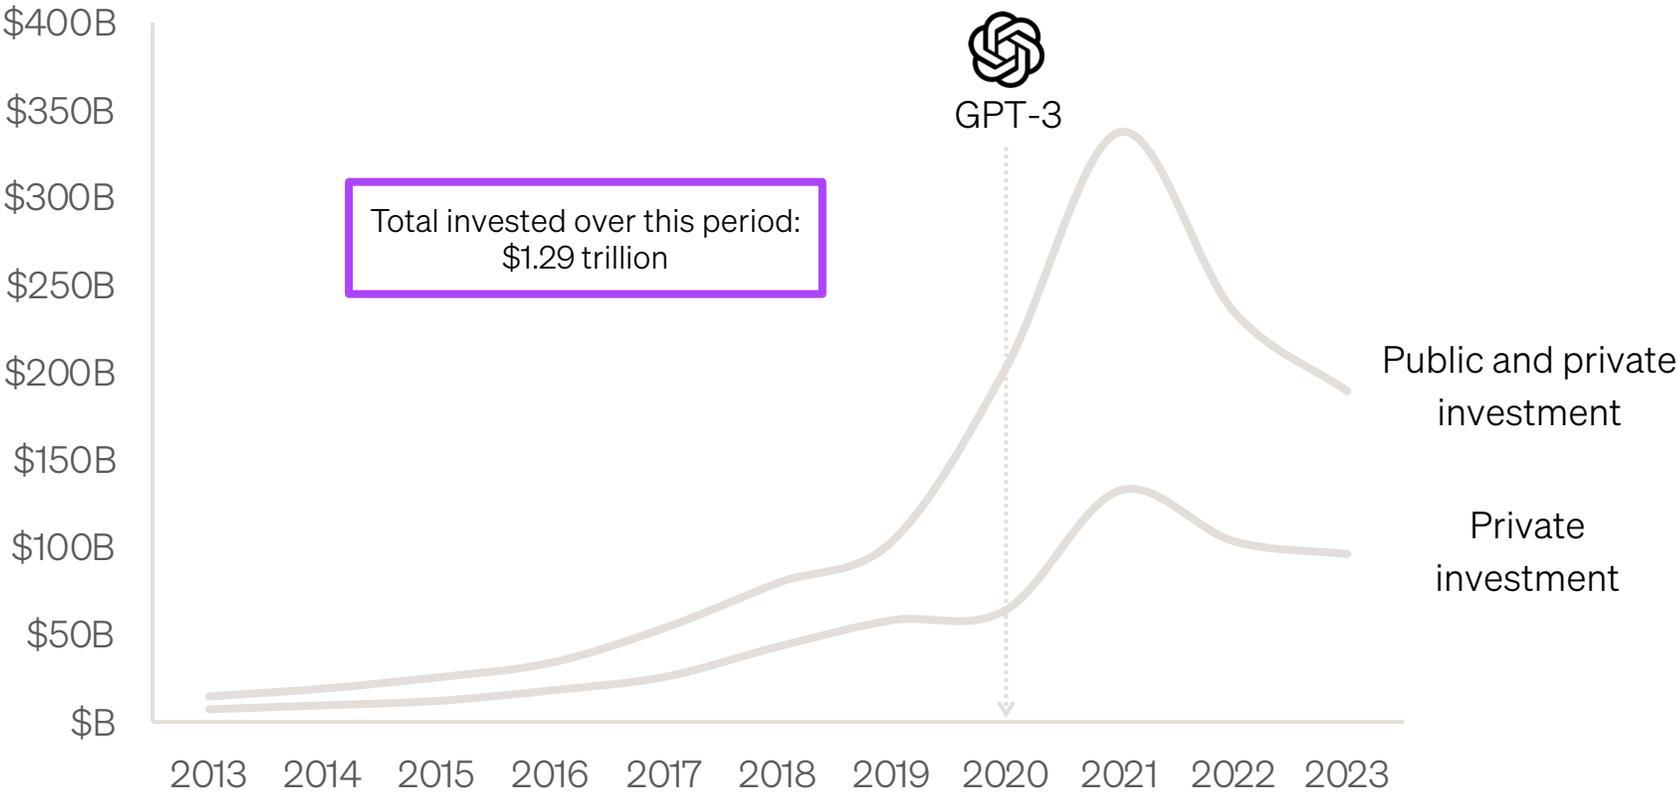


Investors are deploying capital in companies that develop foundational models, as well as applications that use the models.

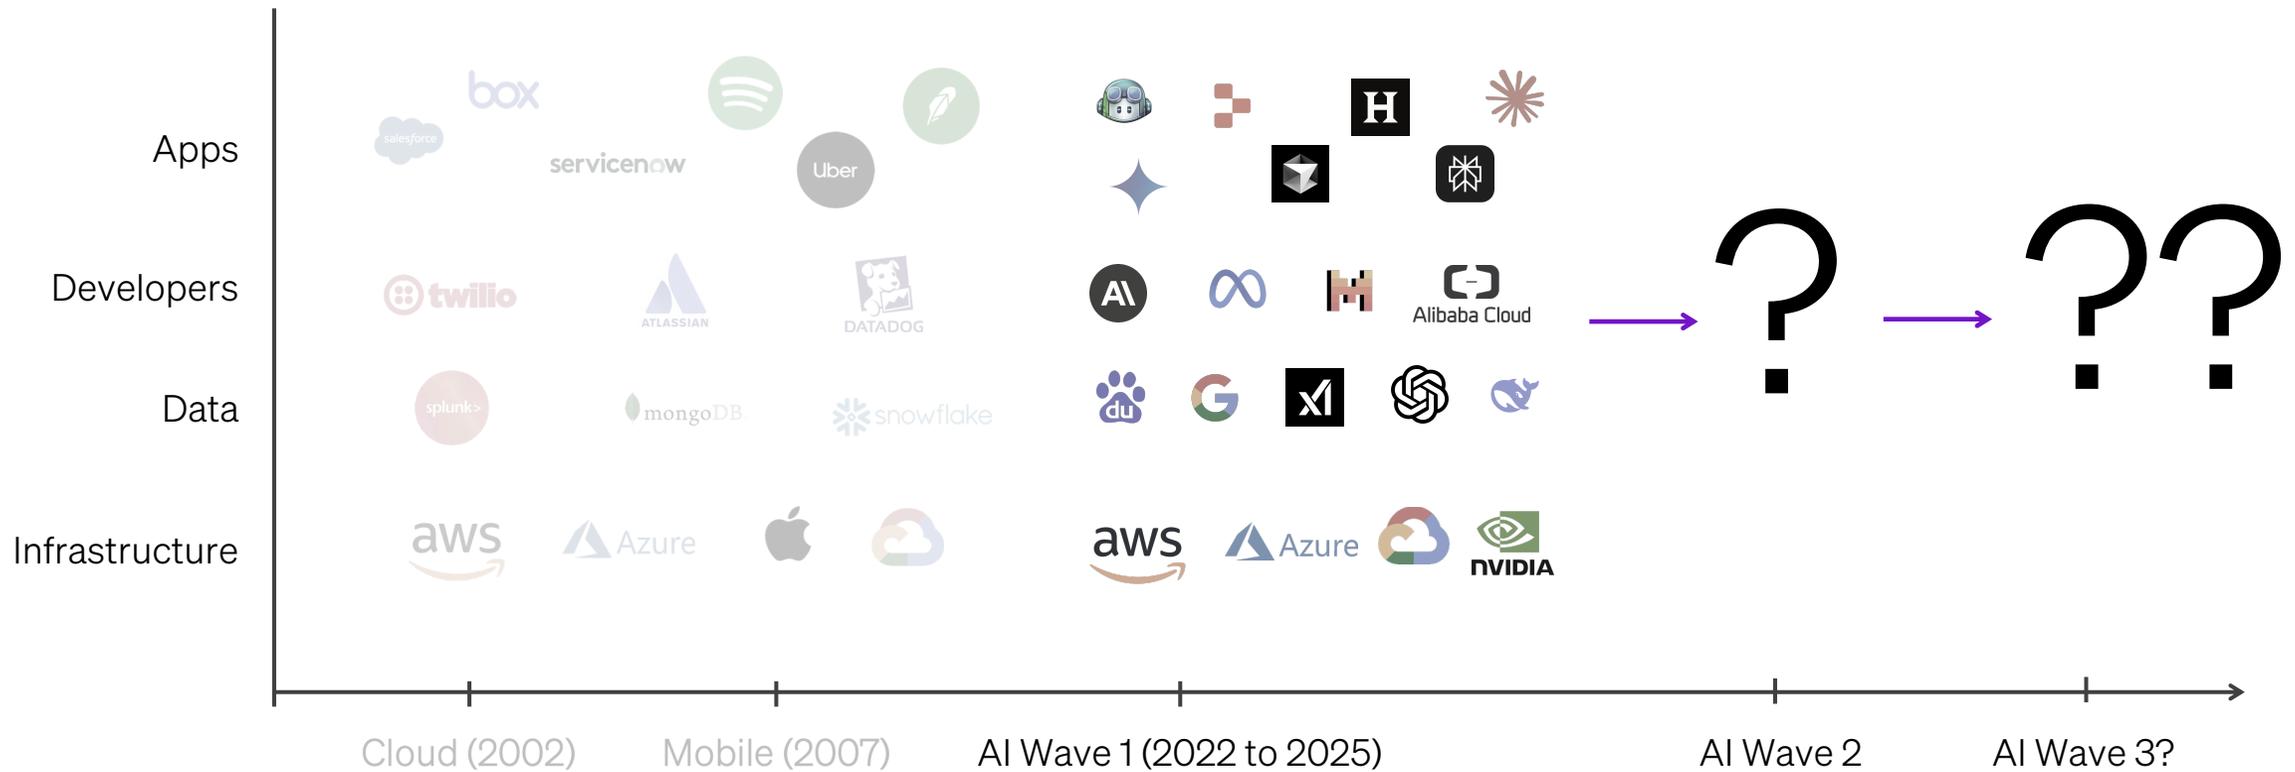
Application developers pay for model usage, either for API calls, or for compute if they deploy their own infrastructure.

If companies can't build lasting competitive advantages, how many of the investment dollars already invested will generate returns for their investors?

Public and Private Investment in AI, Billions of U.S. Dollars, 2013 to 2024



Even if the answer is unclear, we should remember that we may only be in the early innings of a multi-decade technology wave.



Let's look at two emerging areas in AI.

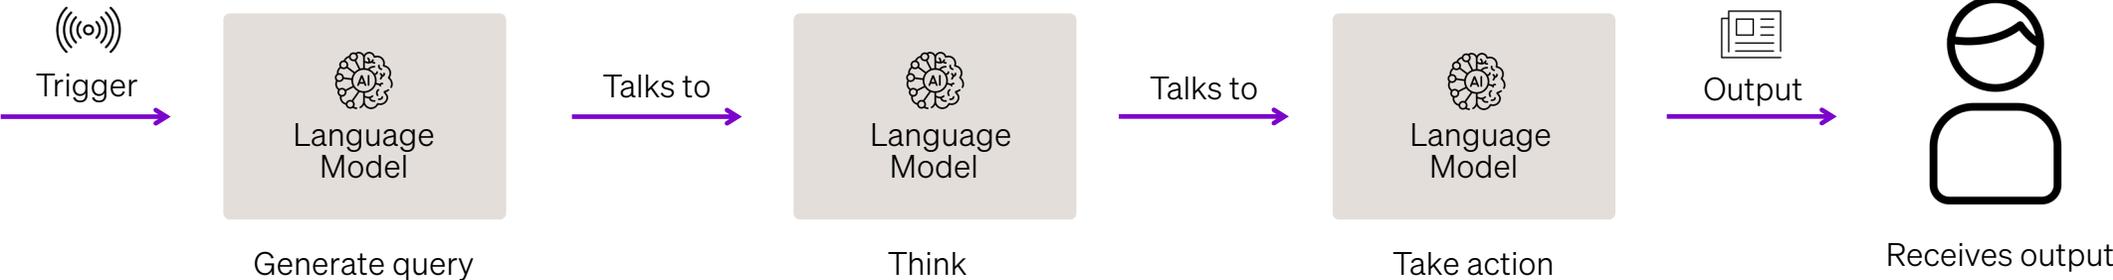
Agentic Systems

Today, language models are reactive, which means they interact with humans only by responding to queries.



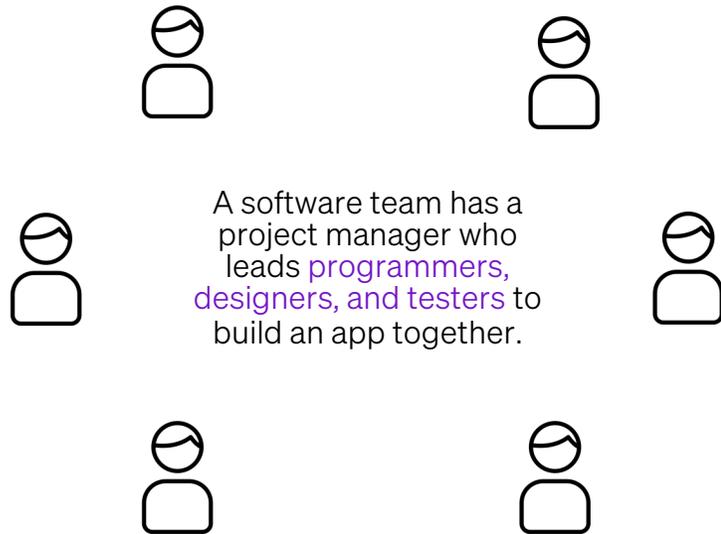
Agents represent a paradigm shift because they have the potential to be productive without constant human prompting - they can act rather than just react.

What is the idea behind agents?



Hypothetically, it becomes possible to run companies at lower cost and with less constraint on specialized talent to build a company.

Before AI Agents



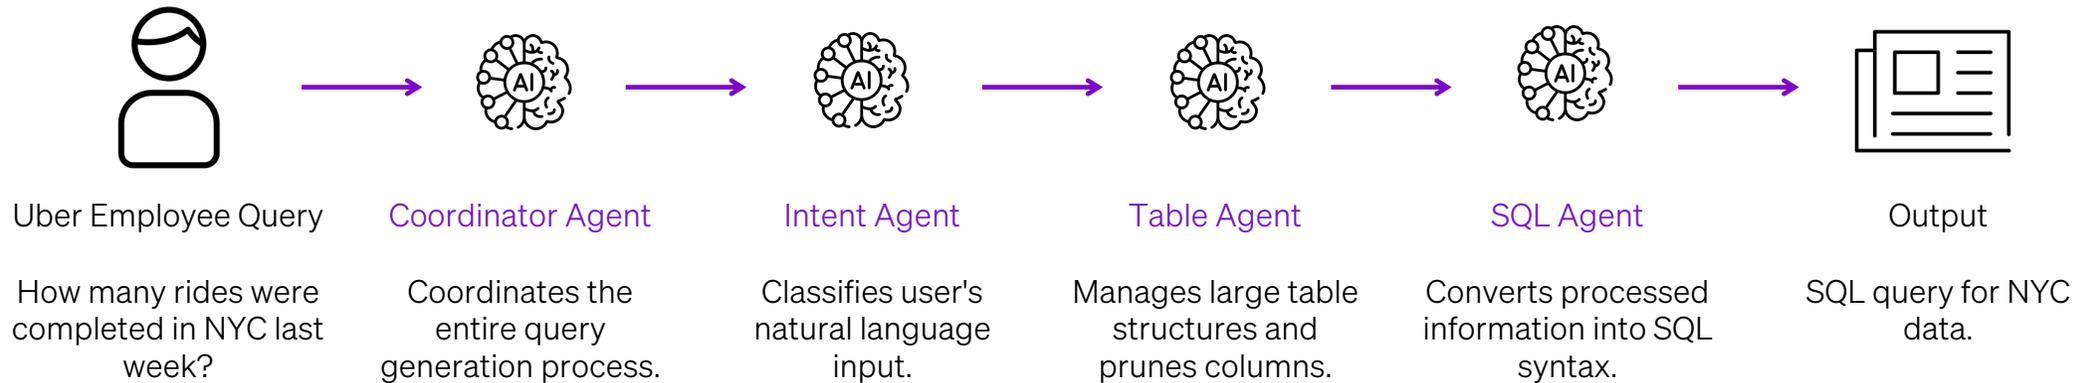
With AI Agents



Teams can move faster and spend less by using AI agents for specialist tasks, while humans lead and make key decisions.

And could reshape how companies are organized, moving toward teams where more work is being done independently by AI.

Multi-Agent Systems – Uber Internal QueryGPT Example (in production today)

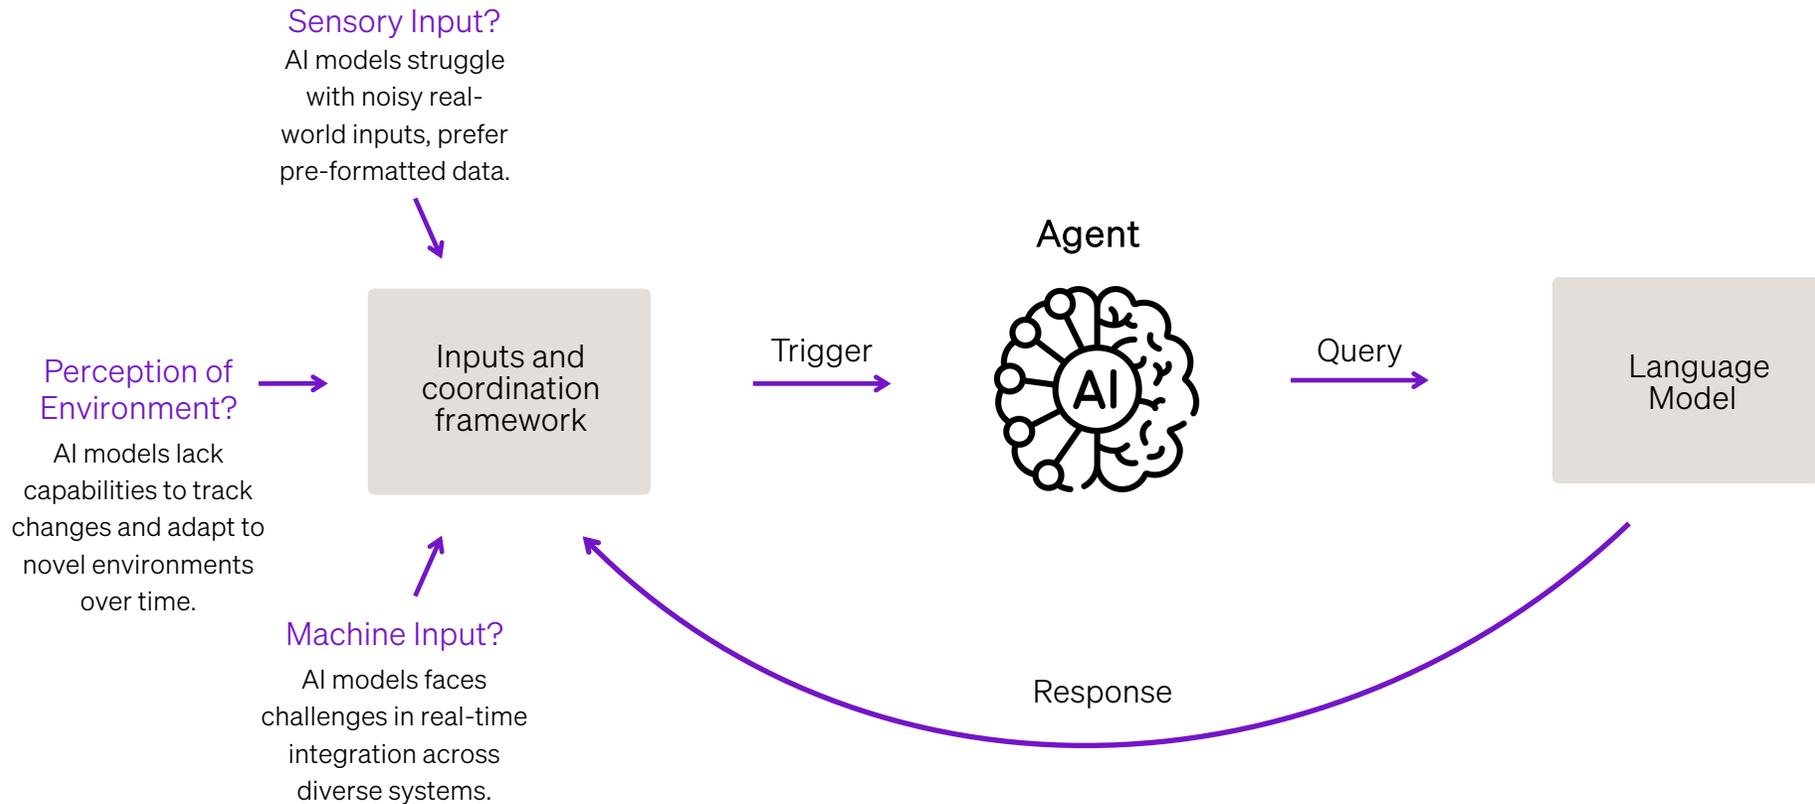


Reduced query writing time from **10 minutes to 3 minutes.**

Uber's data platform handles 1.2M monthly queries today.

So what are the current roadblocks?

First, for agents to act autonomously with less human intervention, they need additional capabilities that have not emerged yet.



This includes capabilities that language models already possess but are restricted since we've chosen to limit certain behaviors.

Negotiation



Undefined Policies
We don't trust AI to hide/reveal information strategically like humans do.

Strategic Deception



Trust/Control Issues
No clear rules on when AI can withhold or shape information.

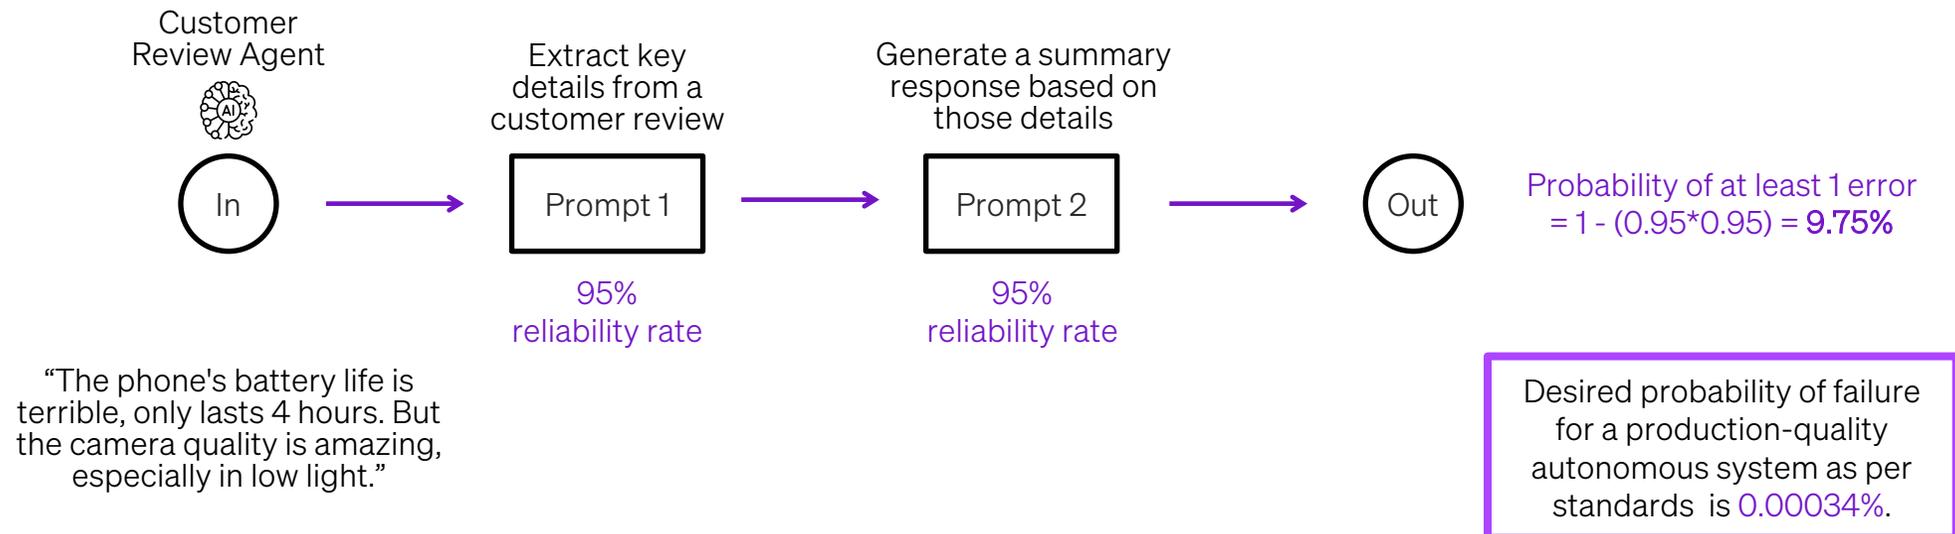
Risk Taking



Risk/Reward Ratio
No framework for AI to make high-stakes decisions with real consequences.

Moreover, a key bottleneck is reliability, as agents operating autonomously break down quickly if there are errors like hallucinations.

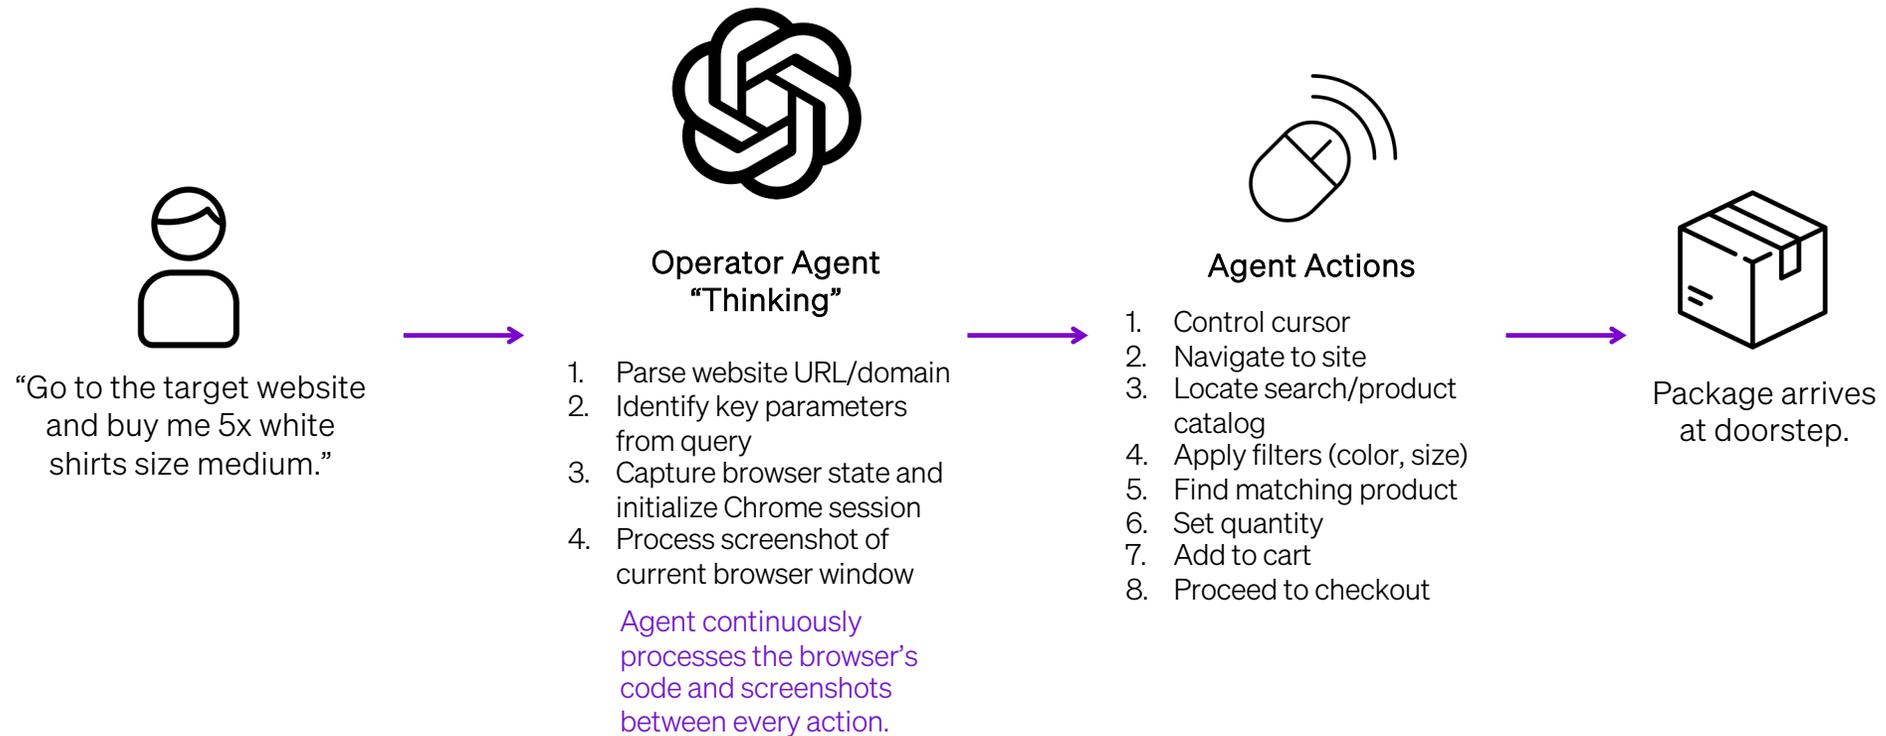
Example Agent Workflow – Prompt Chaining



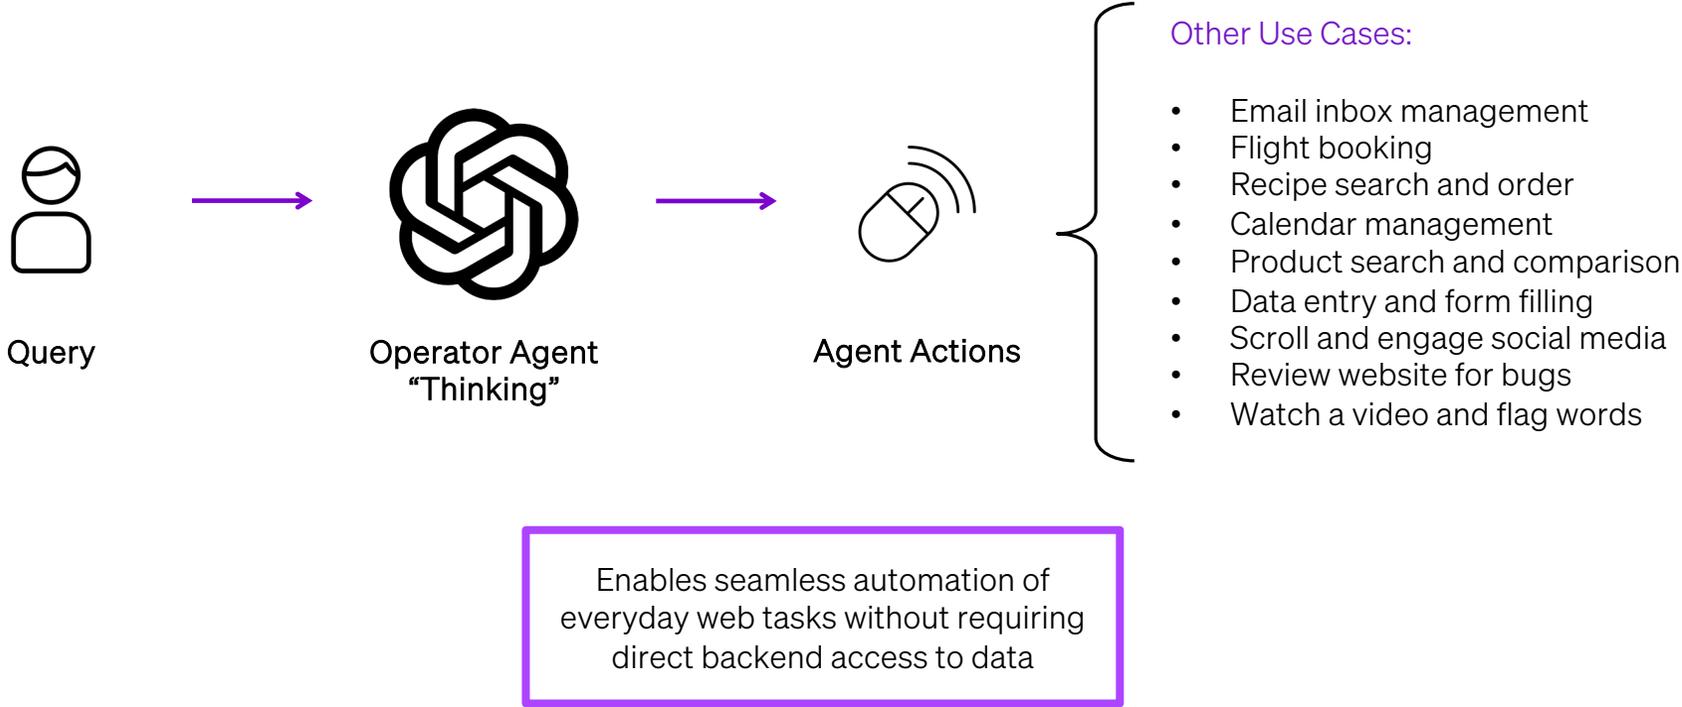
So where are agents today?

Agents are already beginning to conduct tasks on the browser and computer more effectively.

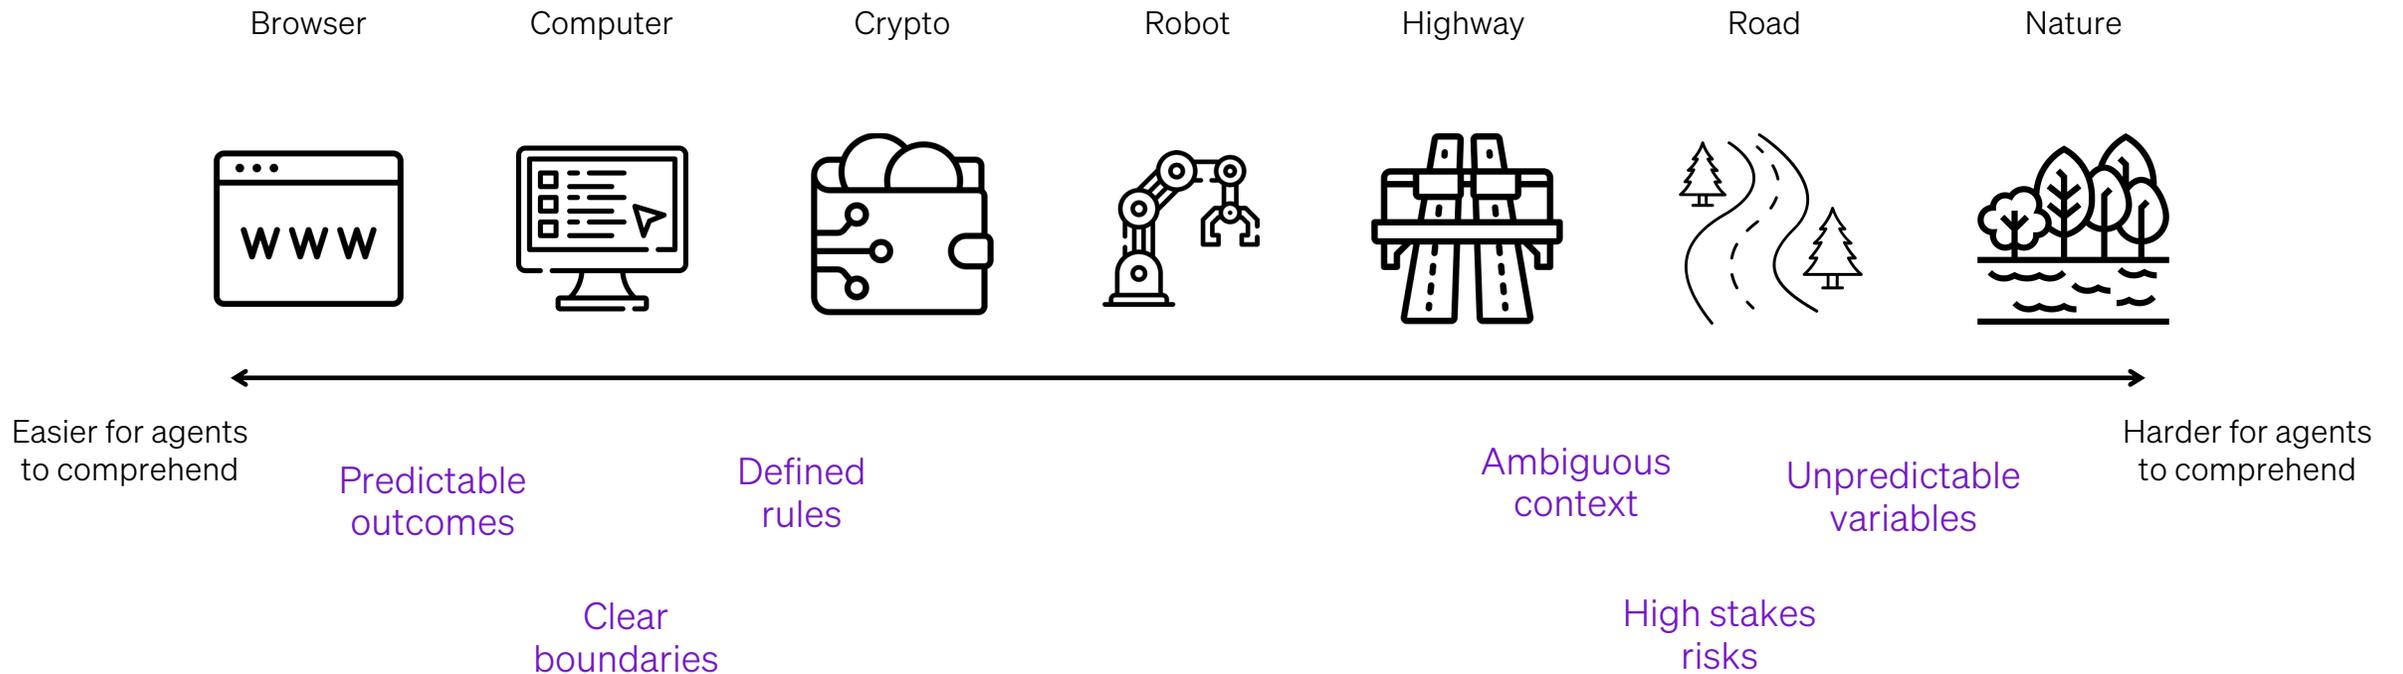
How does OpenAI's Operator handle an online shopping task on the browser?



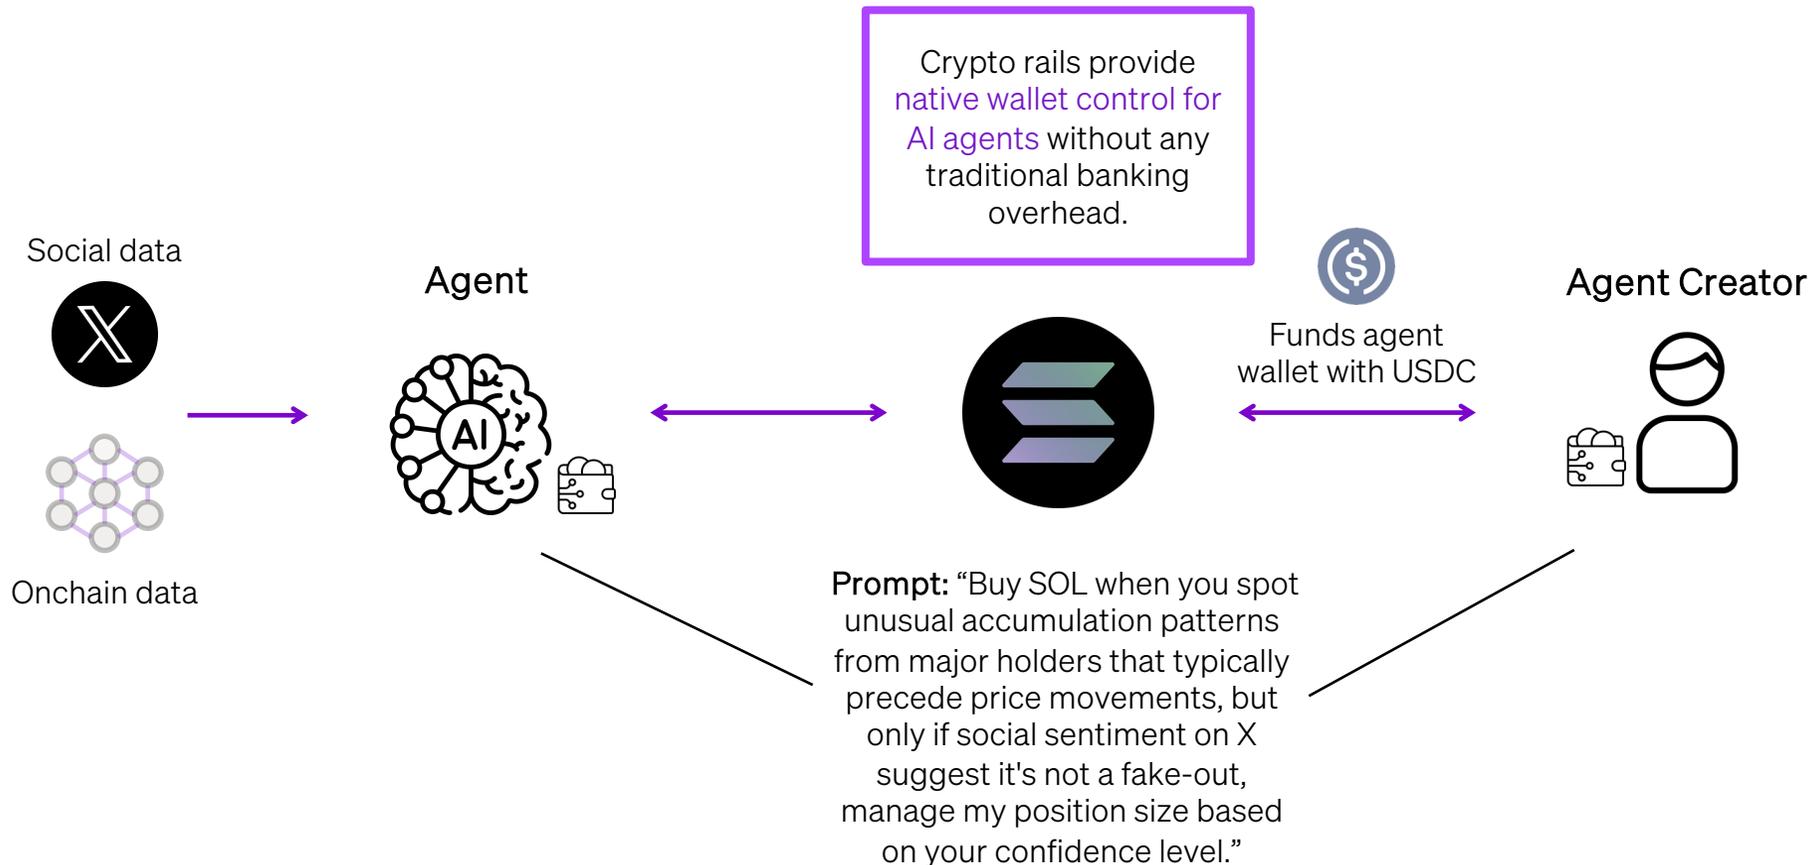
This creates the possibility of an even greater set of workflows and tasks on the computer we can automate and have completed for us.



Agents will be more effective initially in digital environments and over time become more capable in real-world settings.



And crypto networks are one type of digital environment that enables AI agents to execute instant transactions and automate financial operations.



Multimodal / Voice

Language models used to only process text and code, but are now multimodal by default.

	Model	Text	Code	Image	Video	Audio	PDF	Docs
	Claude Sonnet 3.5	✓	✓	✓			✓	✓
	OpenAI GPT-4o	✓	✓	✓	✓	✓	✓	✓
	xAI Grok 2	✓	✓	✓			✓	✓
	Google Gemini 1.5	✓	✓	✓	✓	✓	✓	✓
	Meta Llama 3.3	✓	✓	✓		✓	✓	✓

As language models evolve in their multimodal capabilities, they may enable us to rethink user experiences from first principles.



What is the most natural and user-friendly way to interact with a handheld device with AI?



What is the most natural and user-friendly way to interact with a laptop with AI?

Today, the primary way to interact with phone and computers is through writing and typing, which is a bottleneck for use and interaction.



Steps to complete task:

1. Get phone out of pocket
2. Unlock phone with face ID
3. Find app to perform task
4. Navigate and type information in the app.

⋮

12. Complete task



Steps to complete task:

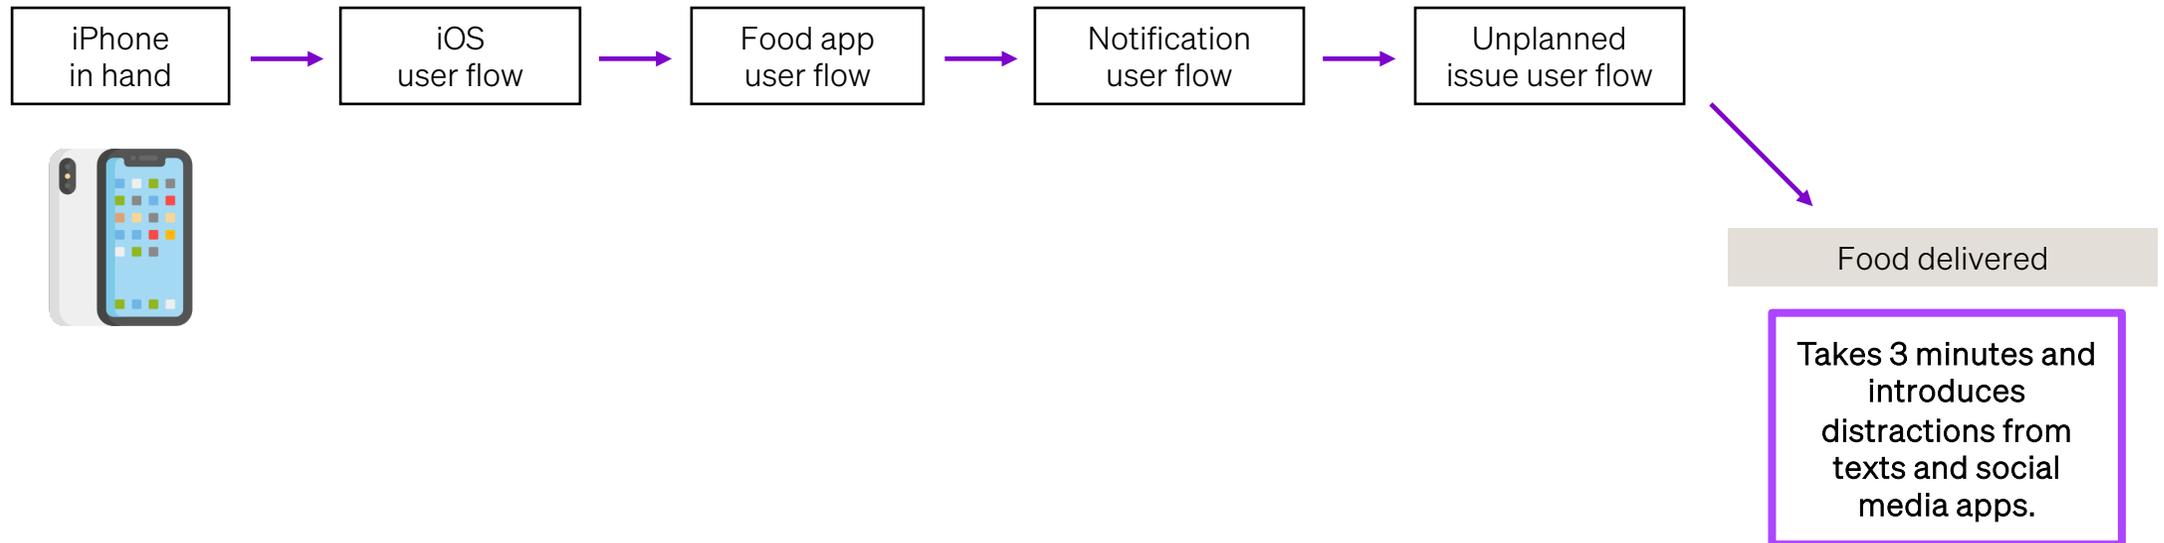
1. Get laptop out of bag
2. Type in password
3. Click browser app
4. Type and explore internet to find website.

⋮

15. Complete task

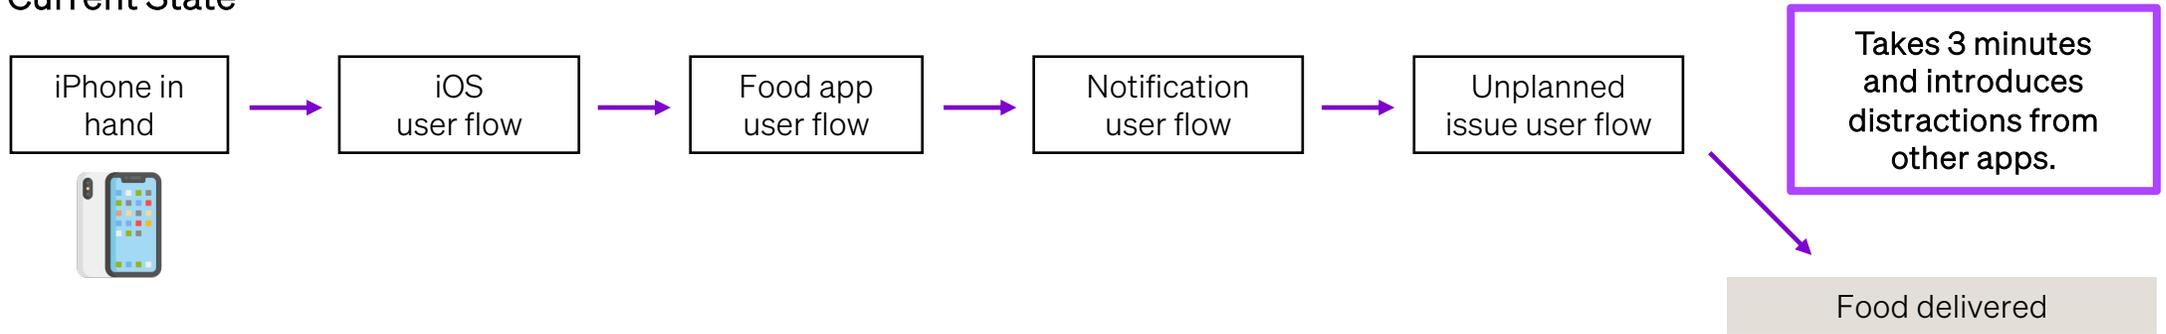
The user experience of computing platforms has evolved to accommodate this bottleneck.

Current State



Voice modality enables more natural ways to interact with computers, offering a universal method of communication.

Current State



Hypothetical Multimodal Device Flow



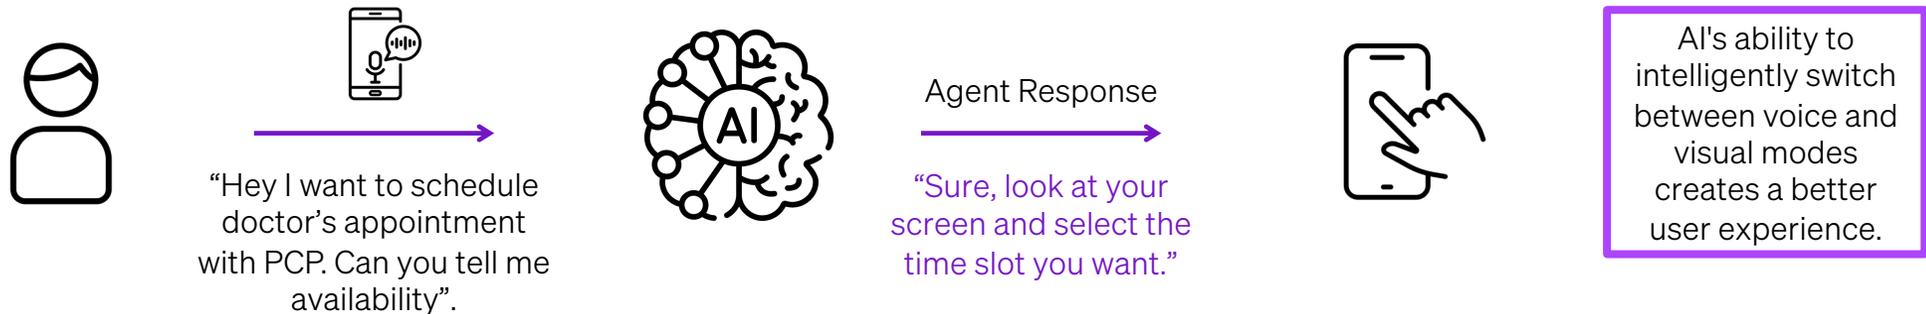
“Hey Siri I want food from Chipotle, pick for me something I like, use my AMEX card to pay, and tell the driver to call me if he has any problems”.

To create seamless experiences, AI systems will need to intelligently integrate and select the most appropriate input and output modalities for each task.

Current Problem



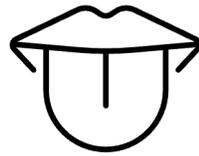
Desired Future



And we may find multimodal inputs and outputs that are directly associated with the senses.



Haptic sensors



Conductivity sensors



Image recognition

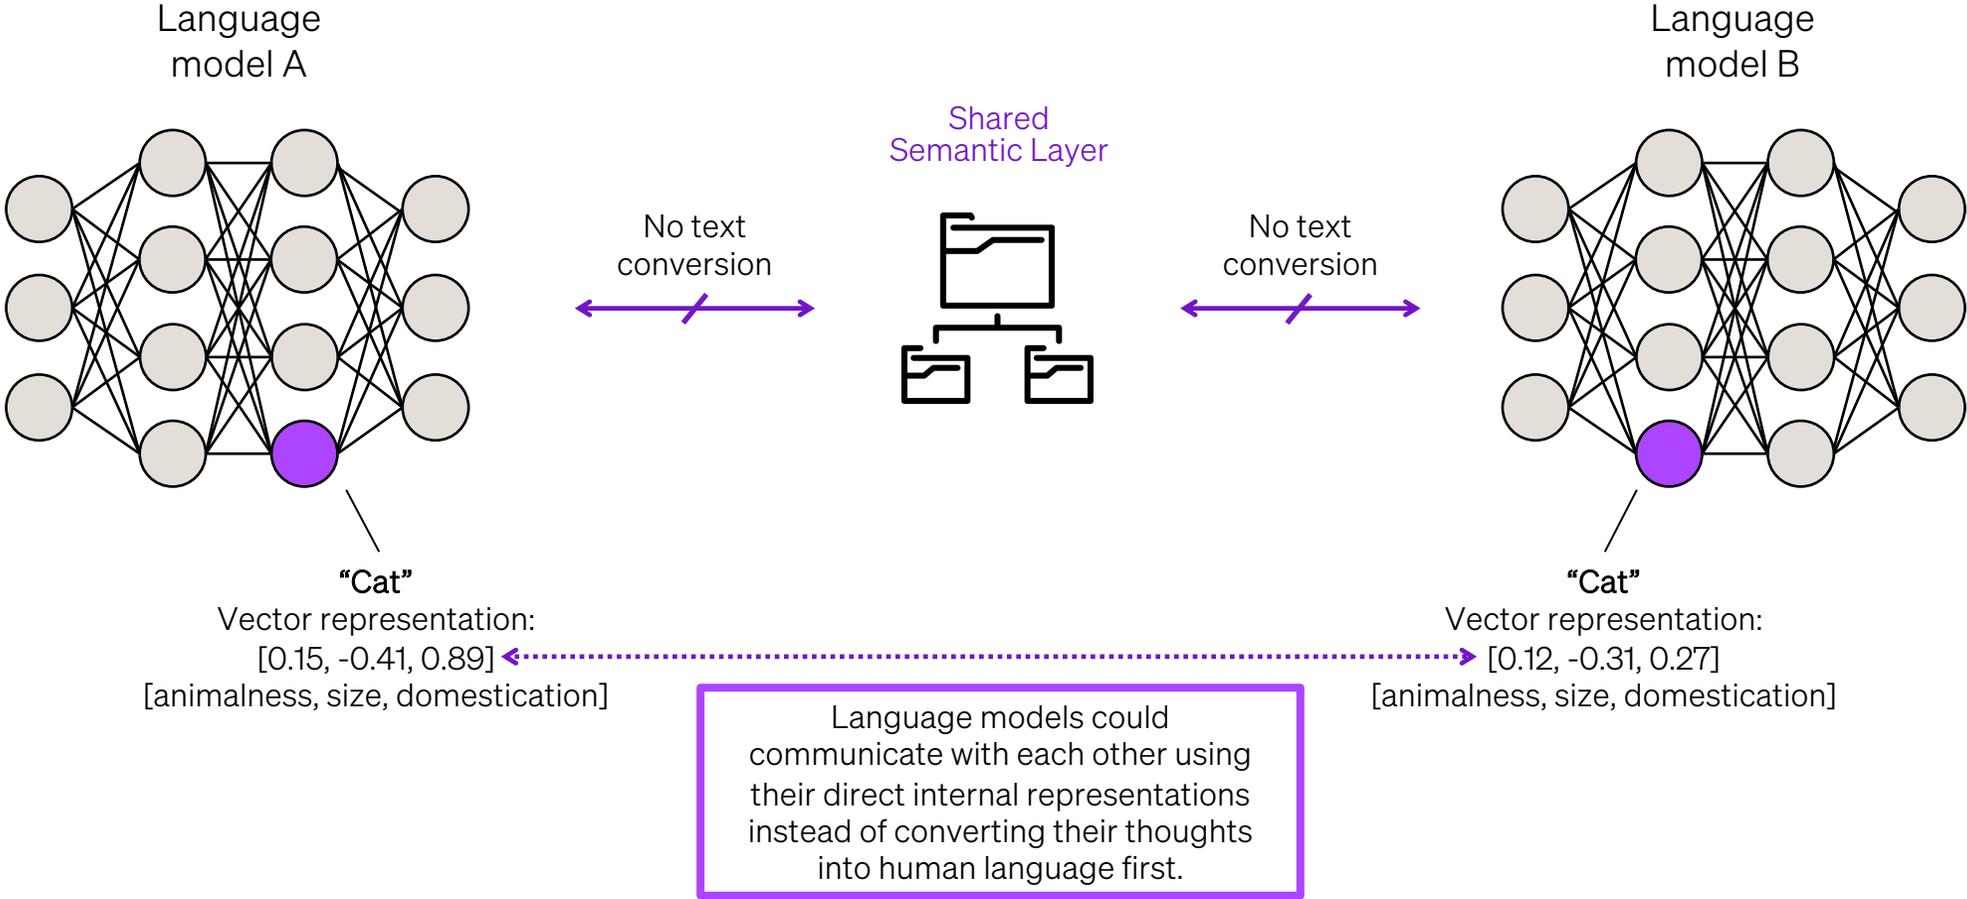


Speech Recognition

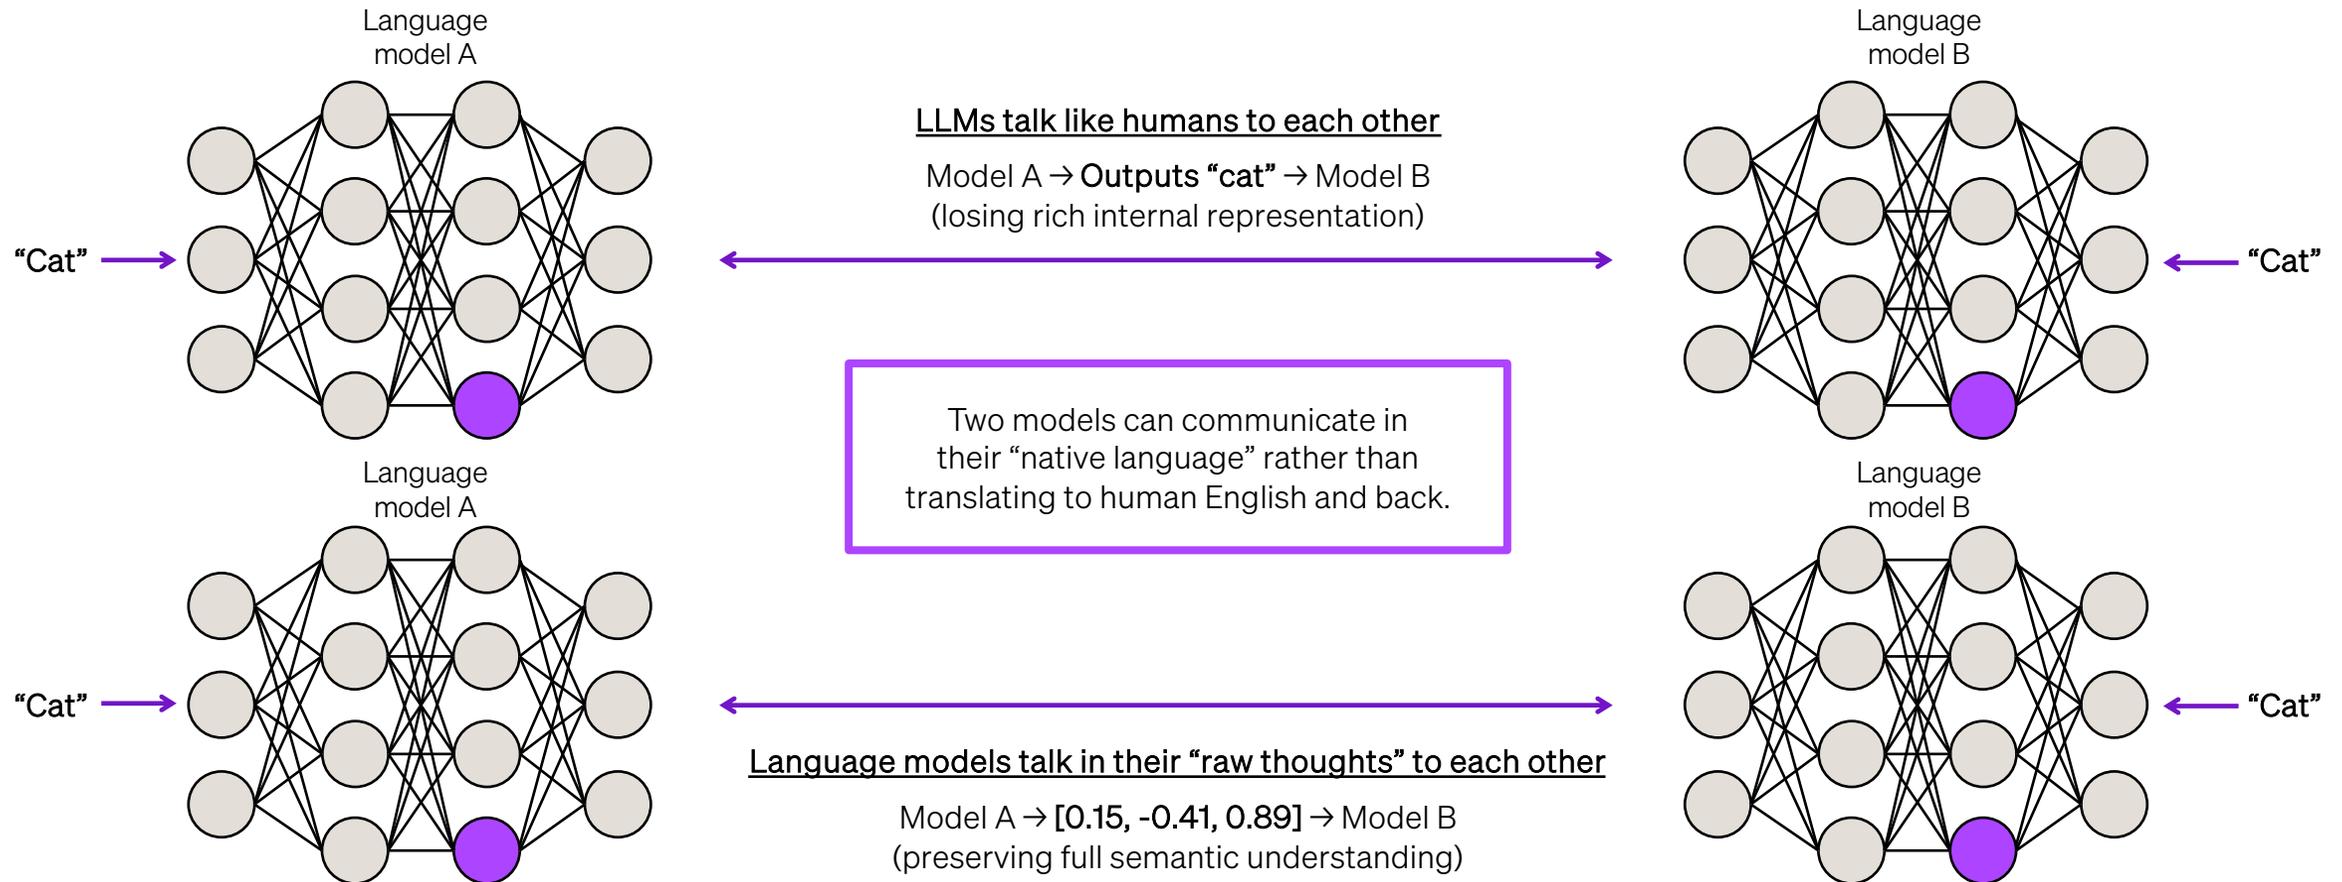


Chemical Sensing

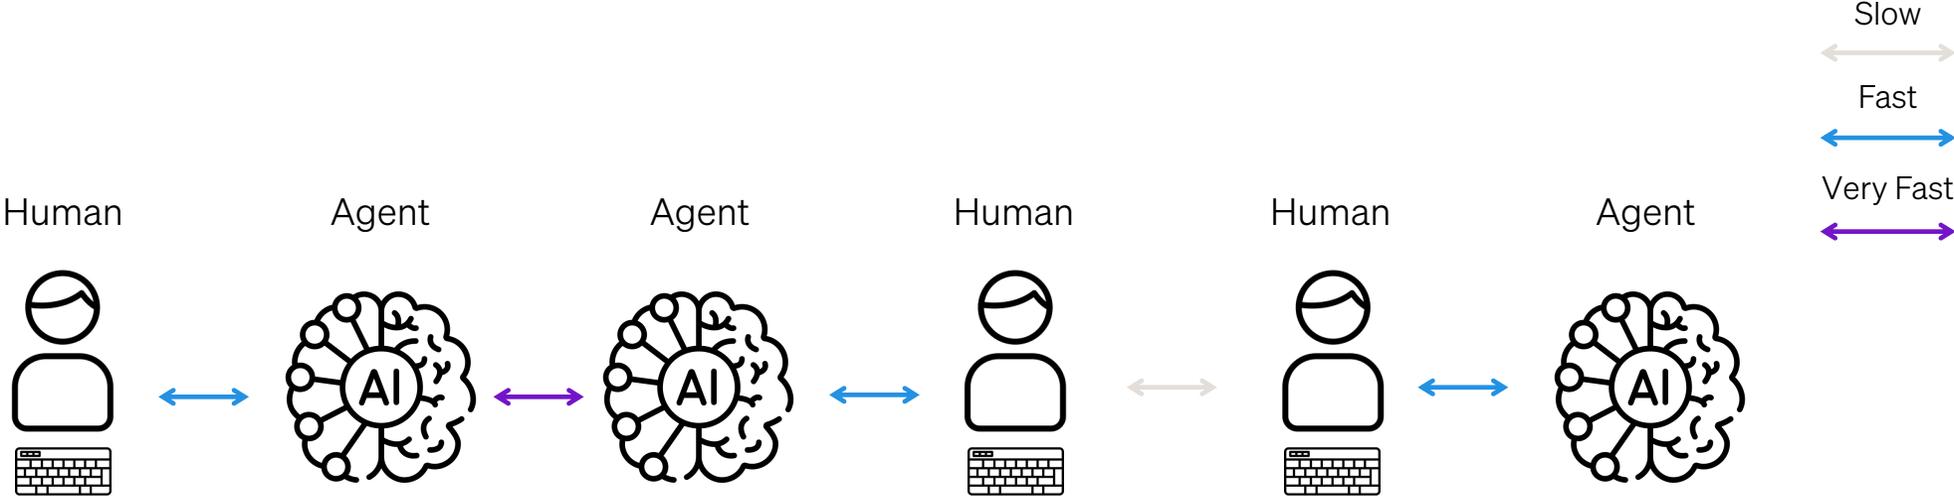
One novel form of multimodality could emerge where models communicate with each other through their internal semantic representations.



Direct semantic exchange between models would eliminate the overhead of converting to human-readable formats, allowing for denser information exchange.



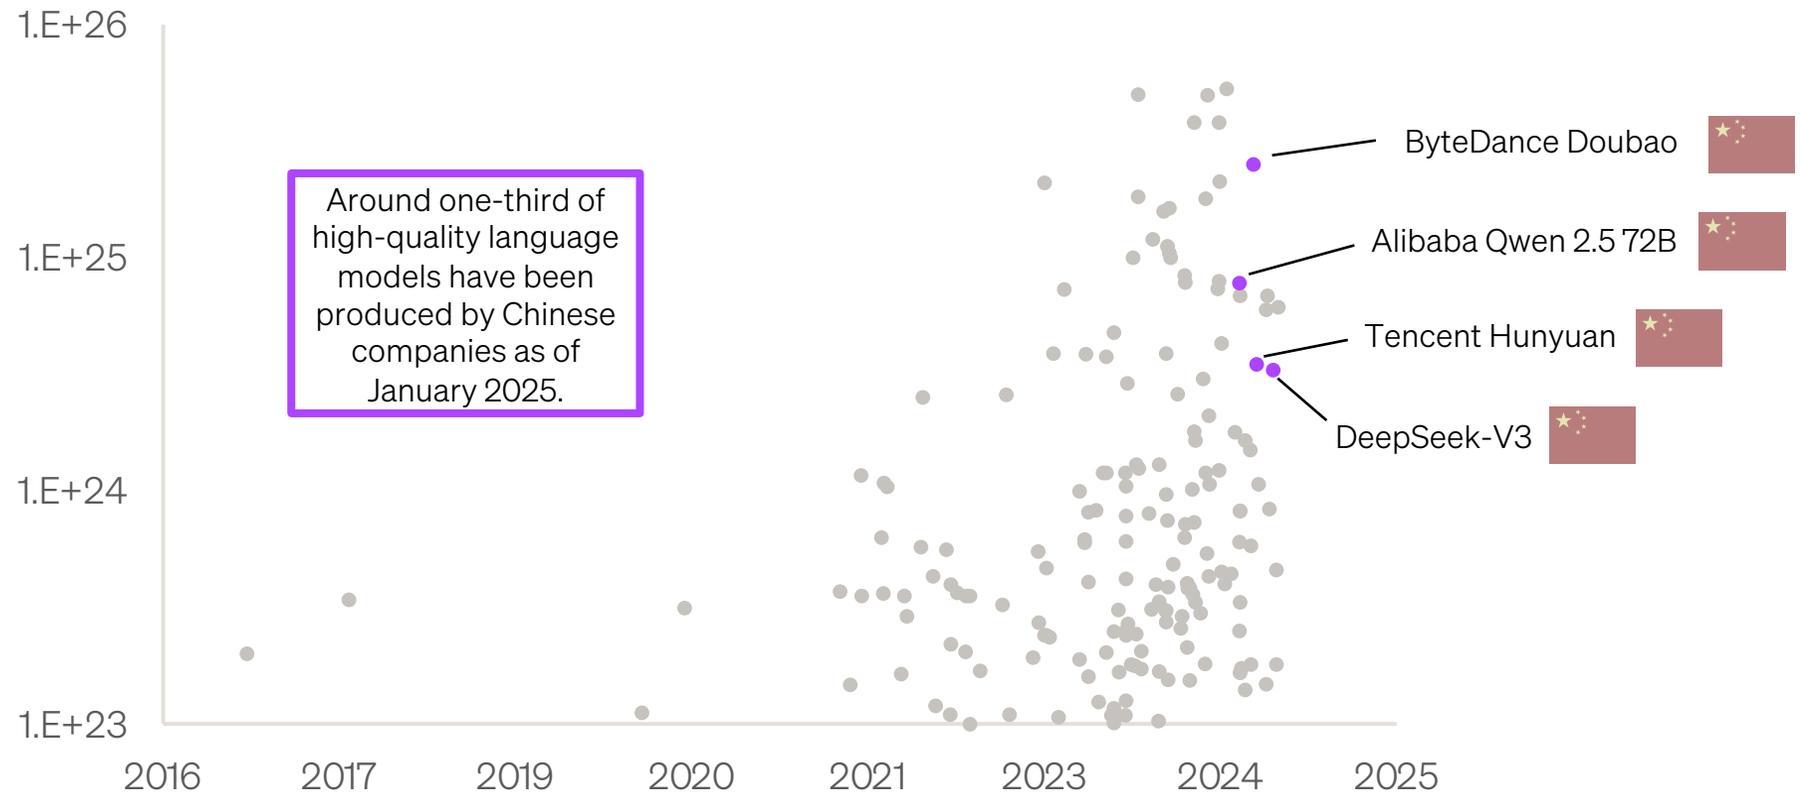
This would enable faster coordination and increased efficiency in systems with multiple AI agents working together.



U.S.-China Primacy

In the two years since ChatGPT's release, a growing number of frontier models have emerged from Chinese companies.

Large Scale AI Models Training Compute (FLOP), 2016 to 2025

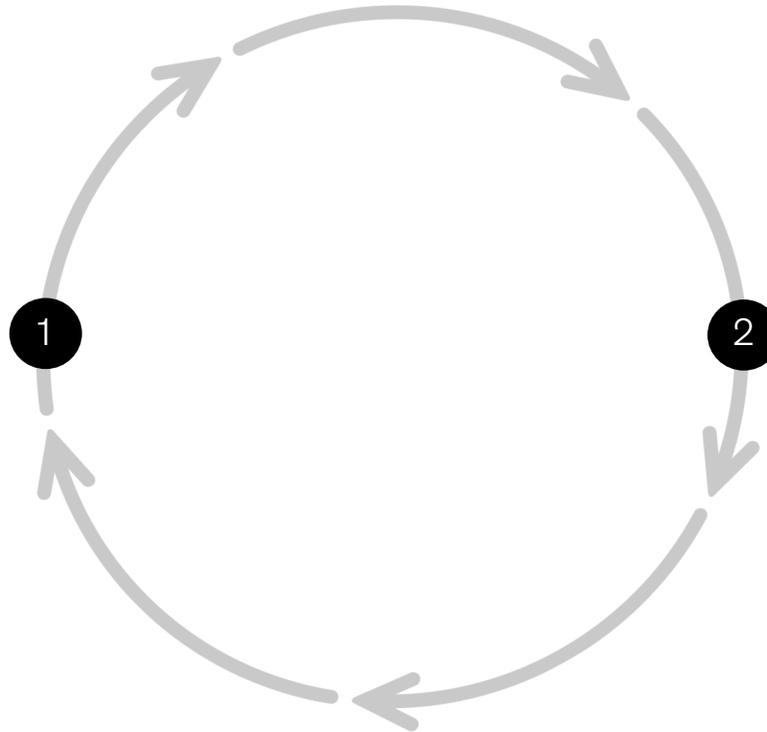


Over the next decade, the country whose AI systems become globally adopted may have outsized effects on economic and military primacy.

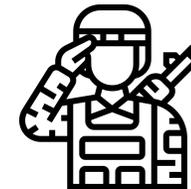
Economic Primacy



AI leadership enables countries to control key technologies, digital infrastructure, and data flows that increasingly drive global commerce and innovation.

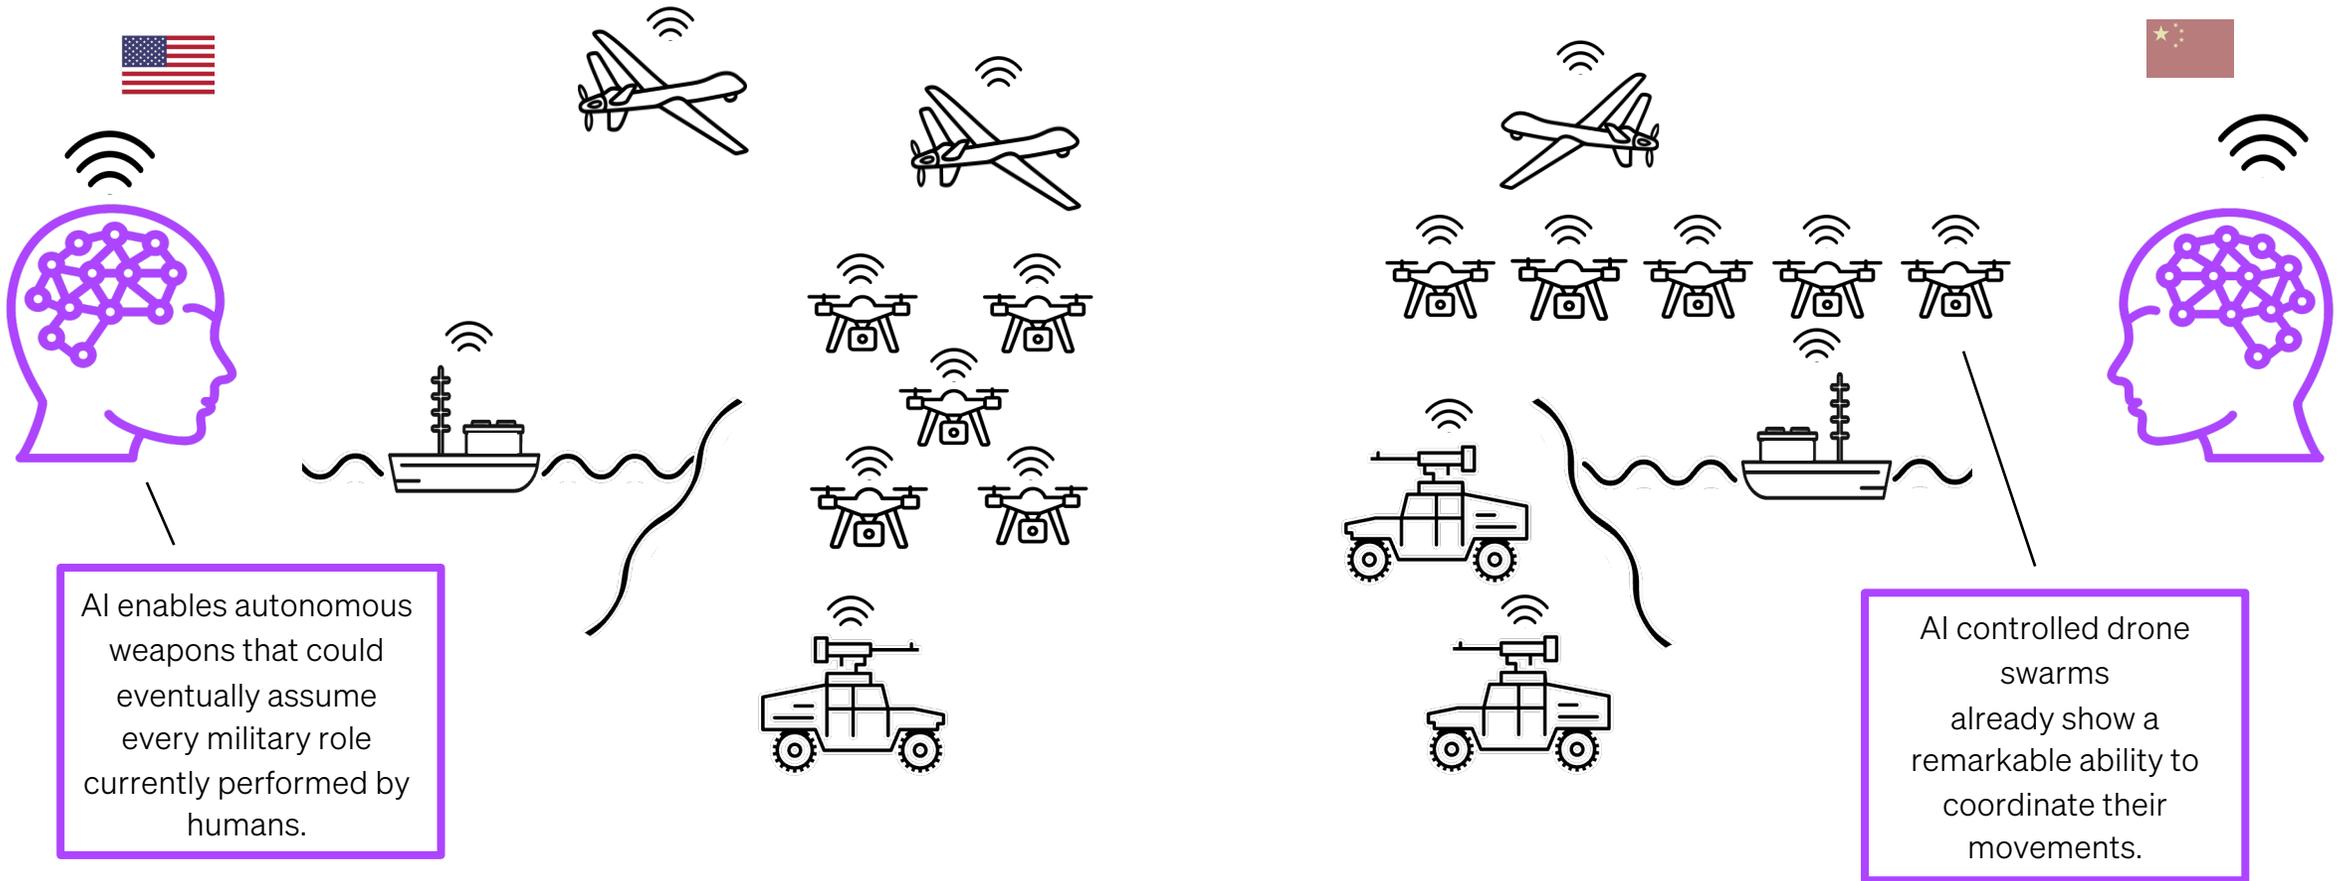


Military Primacy



Economic dominance translates to military advantages, as greater resources and technology enable advanced weapons development and defense innovation.

Whoever has the more advanced AI system may have the edge in military applications, particularly with autonomous drones and robotics.

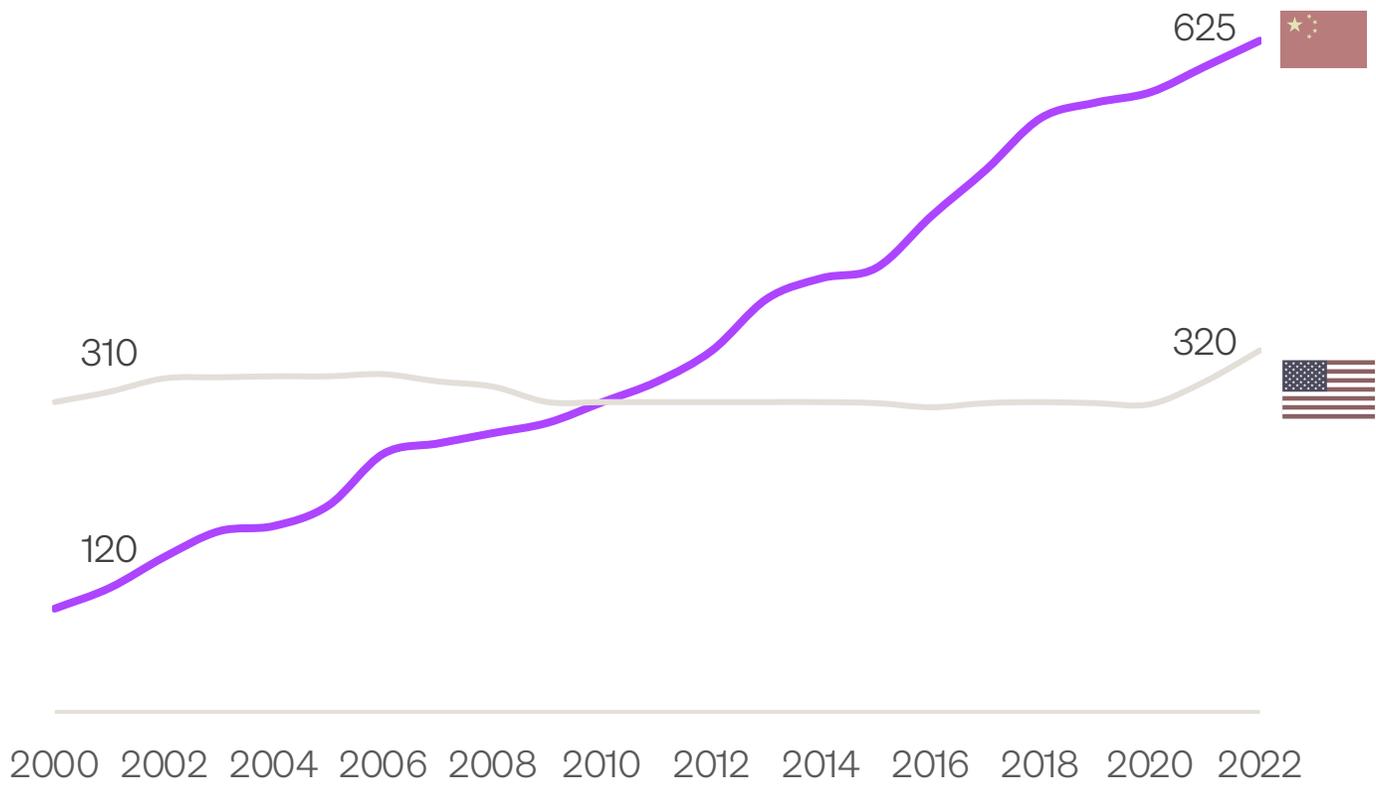


So who is positioned to win the AI race?

China's power grid investment positions it well for AI computing demands, while U.S. data centers face growing electricity supply constraints.

Comparison of USA and China Electricity Generation, TWh, 2000 to 2022

The U.S. must overhaul its aging power grid to support energy-intensive AI data centers.

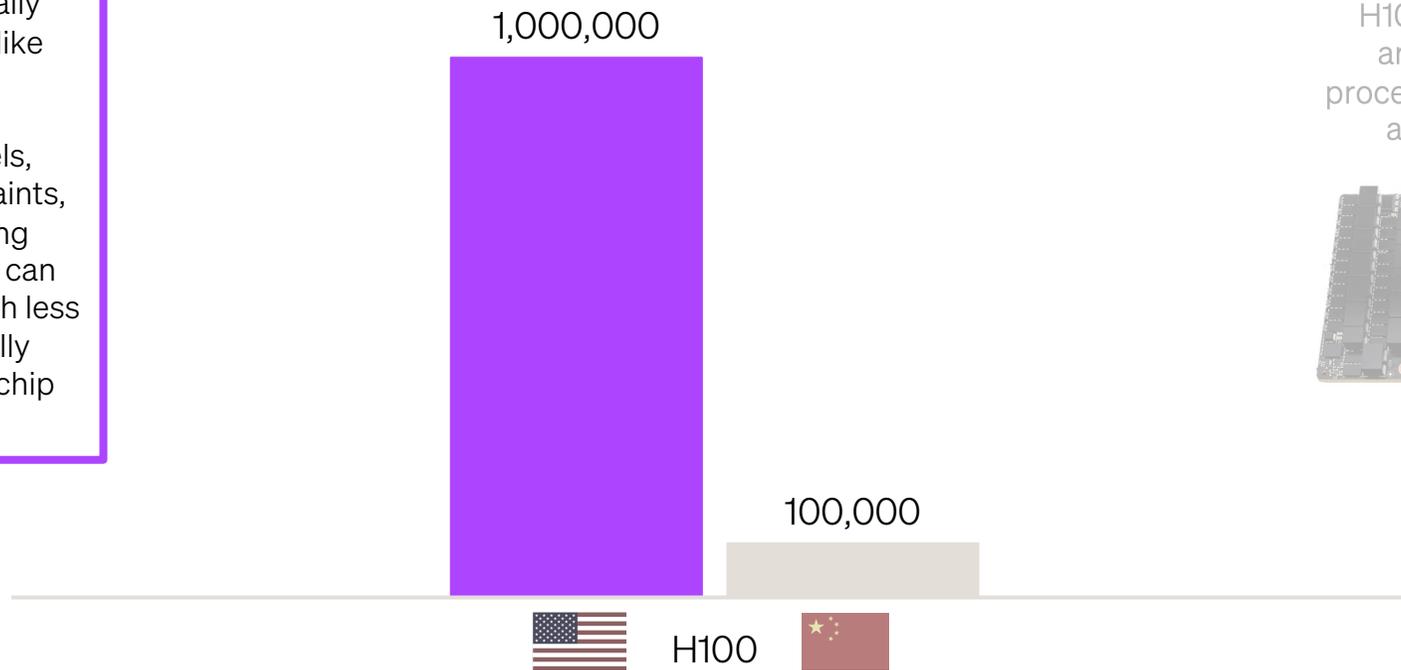


The U.S. is better positioned for the compute needed, partly because of export controls.

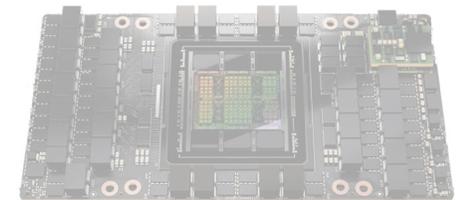
Number of Nvidia H100 GPUs Deployed in Private and Public Clouds

U.S. export controls prevent Chinese companies from legally acquiring advanced AI chips like NVIDIA's H100.

DeepSeek's V3 and R1 models, developed under these constraints, demonstrate how engineering innovation across the AI stack can achieve strong performance with less powerful hardware, potentially reducing the effectiveness of chip export controls.



H100 GPUs from NVIDIA are high-performance processors designed for train advanced AI models.



In terms of research and ability to produce frontier models, the U.S. is ahead.

Between 2000-2021:



The U.S. produced **1.37 million AI papers**,



while China produced **958,000 papers**.



U.S. AI research led with **7.4 million citations**,



followed by China with **2.2 million citations**.

But frontier model leadership typically lasts only months before other companies catch up.

First mover advantage by only 4 months.

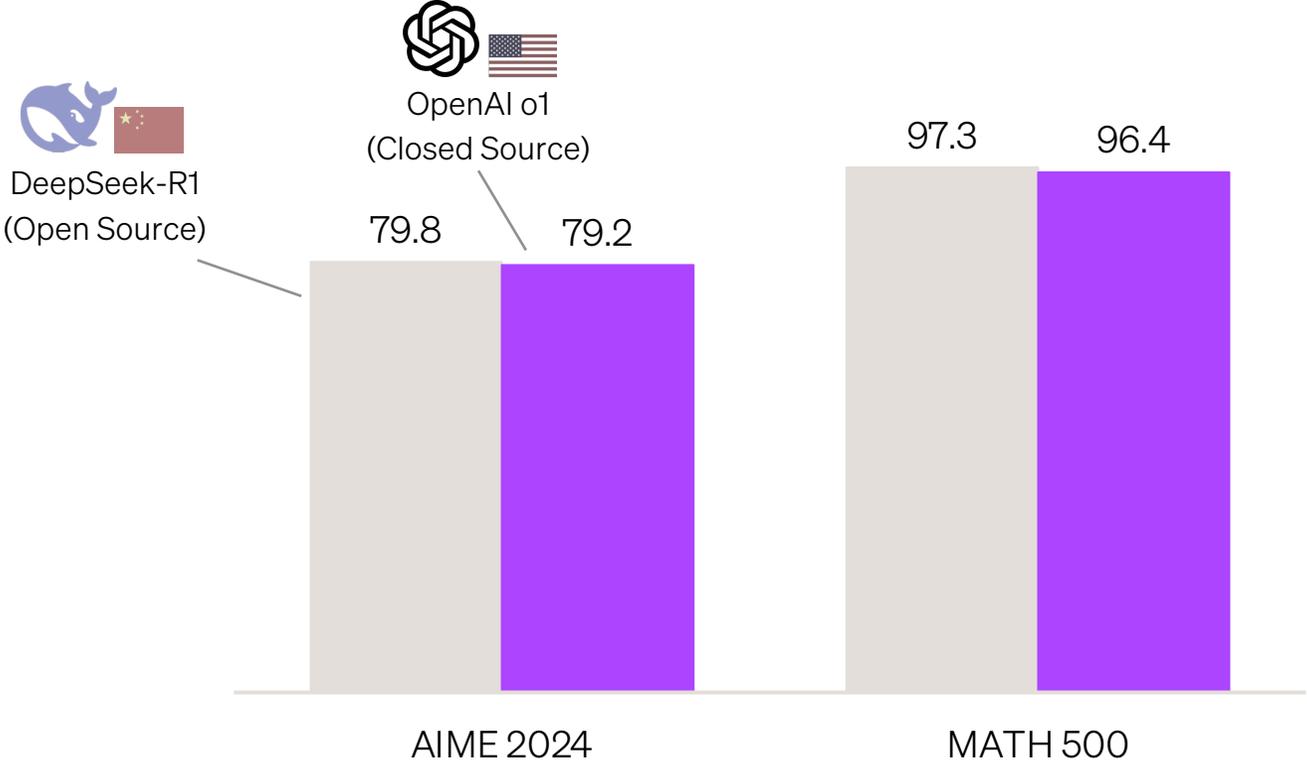
	Model	Date of Release	MMLU	GPQA
	OpenAI – 4o	May 2024	88.7	53.6
	Alibaba Qwen 2.5	Sep 2024	86.1	49
	Tencent Hunyuan 389	Nov 2024	89.9	42.4
	DeepSeek v3	Dec 224	87.1	59.1

MMLU Pro:
Tests advanced reasoning and expertise across professional domains like medicine, law, engineering, and advanced academic subjects.

GPQA
Tests with a dataset of graduate-level, Google-proof multiple-choice questions in biology, physics, and chemistry, crafted by domain experts.

And constraints can drive innovation, as demonstrated by how hardware limitations in China have spurred creative solutions to accomplish more with less.

Benchmark Performance between o1 and R1, Score 0-100



The U.S. historically pioneers core technology platforms, while China excels at scaling and innovating on applications built on these platforms.



Internet protocols (1990)
www, tcp/ip



Mobile OS (2007)
iOS, Android



Platforms (2000)
Amazon, eBay



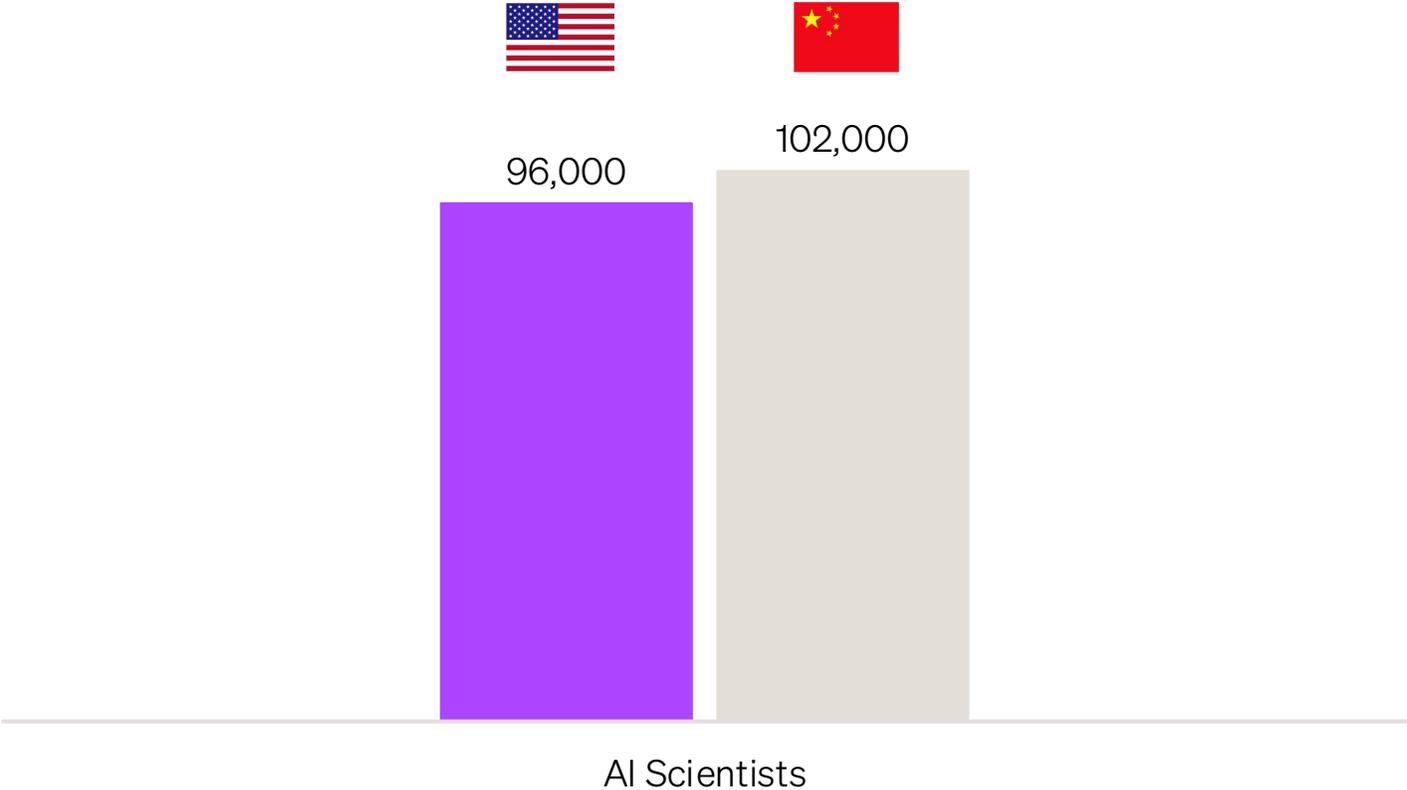
Wechat, Baidu (2000)
China's WeChat has 1.3+ billion users

AliPay, WeChat Pay (2013)
China's mobile payment volume is 80x the USA.

Shein, TikTok shop (2008)
Shein overtook Amazon in fast fashion globally.

Although historical patterns with China may not apply in the future, as China has spent decades systematically positioning itself for technological leadership.

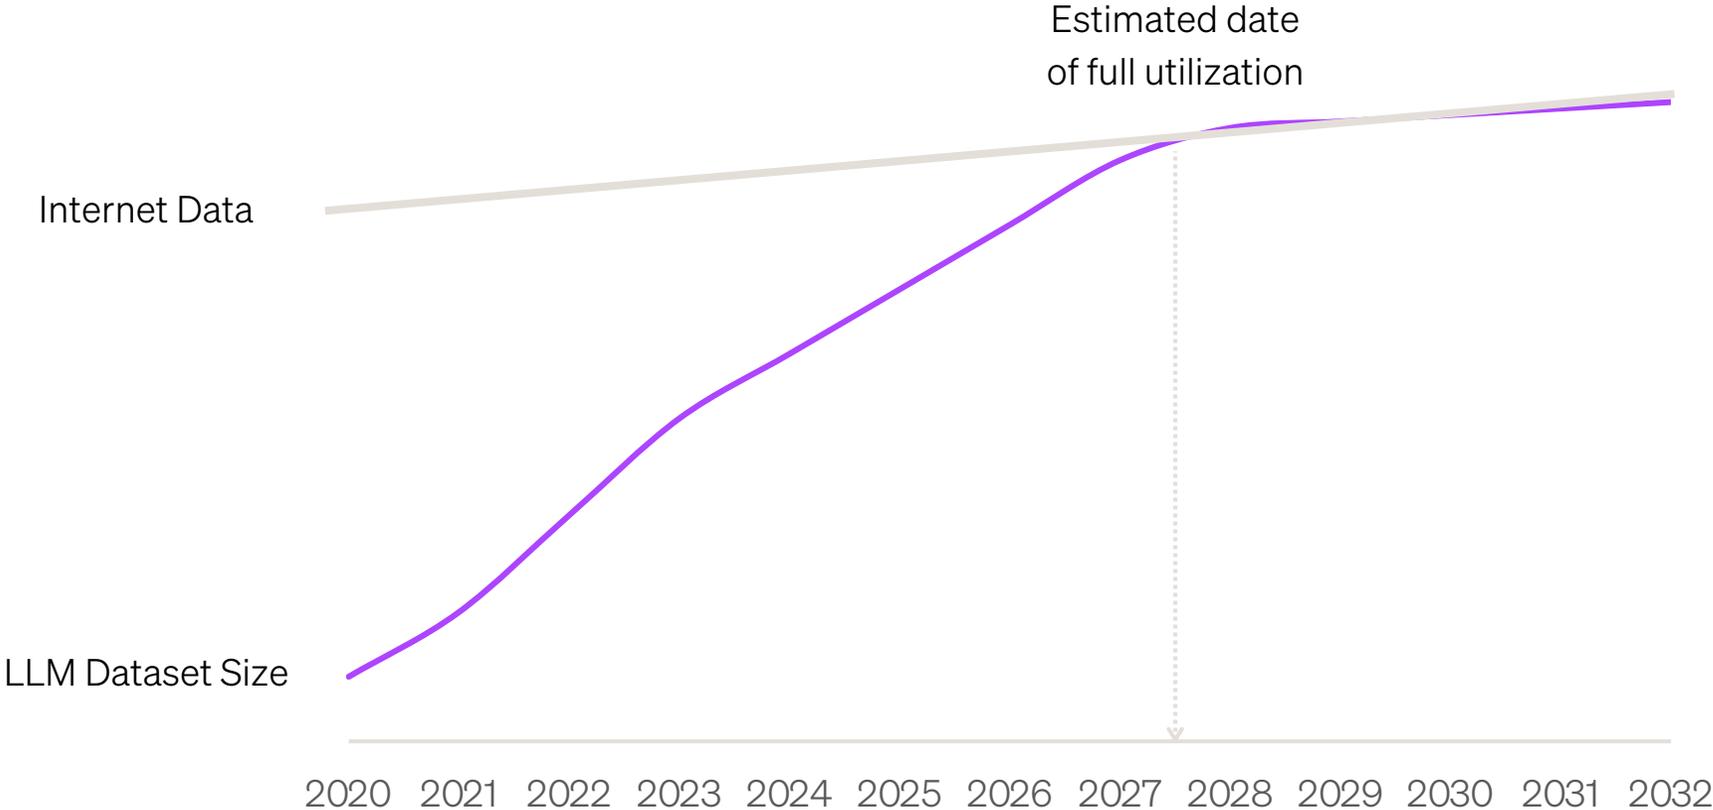
Number of AI Scientists in the U.S. vs. China, 2020



Wrapping Up...

The foundation model scaling paradigm may be reaching its limits.

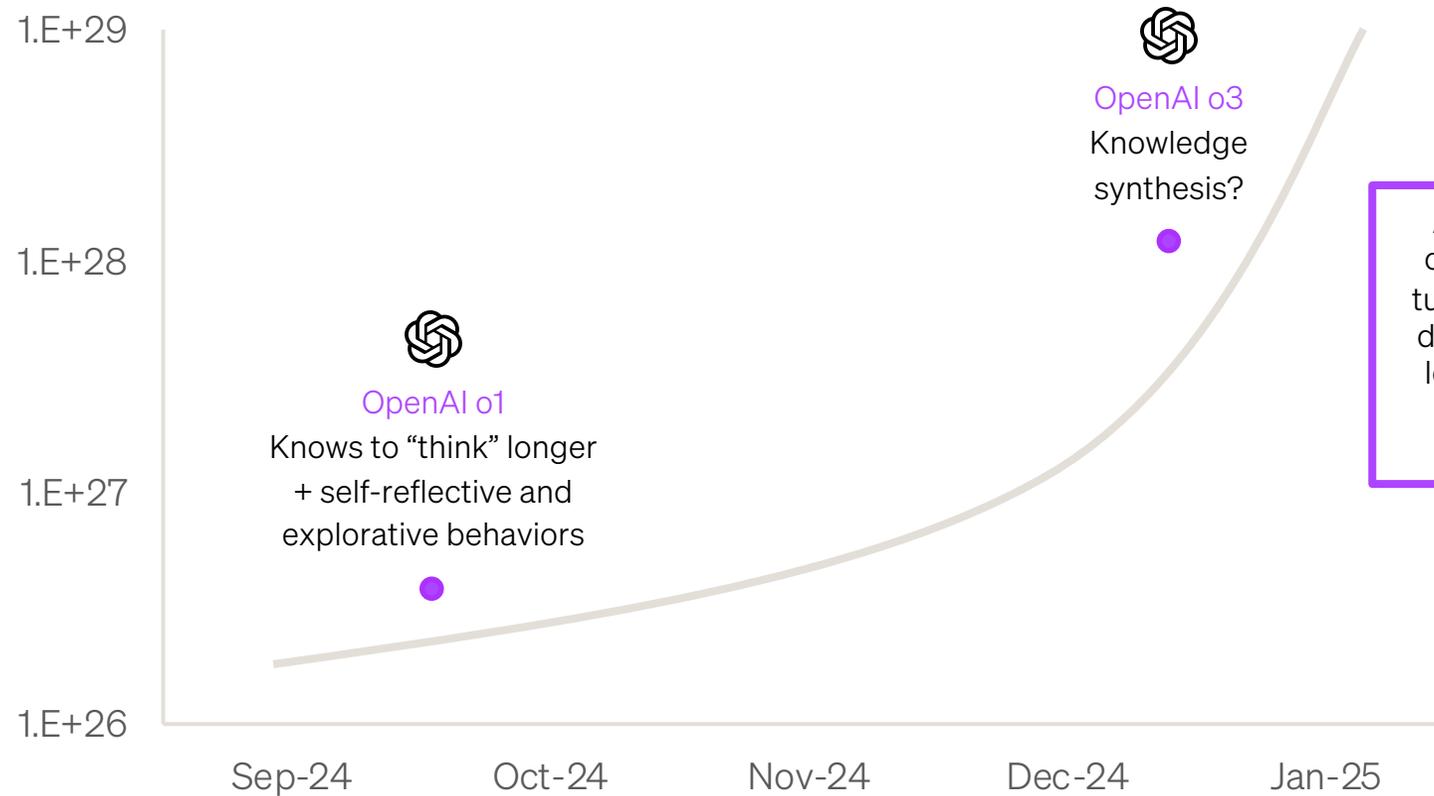
Limits of LLM Scaling Based on Human Generated Data



By 2028, AI models may exhaust the internet's available public training text at current growth rates.

New frontiers in reasoning are showing promise of emergent capabilities.

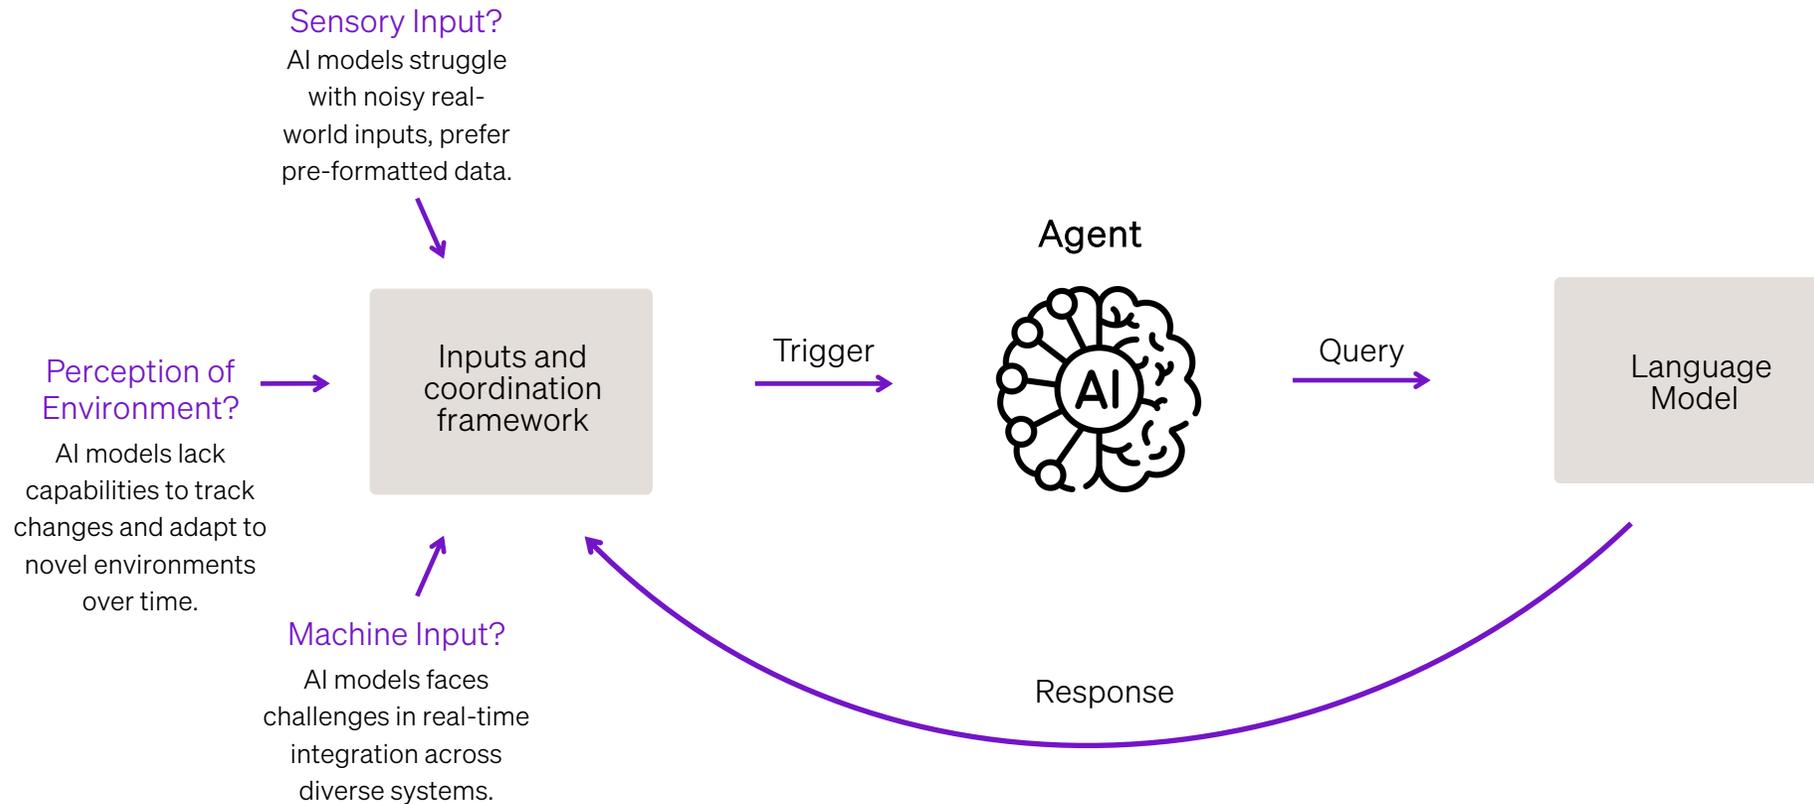
Training Compute (FLOP) Required for Each OpenAI Language Model,



FLOP values represent more computationally demanding training processes that require greater resources (capital and GPUs).

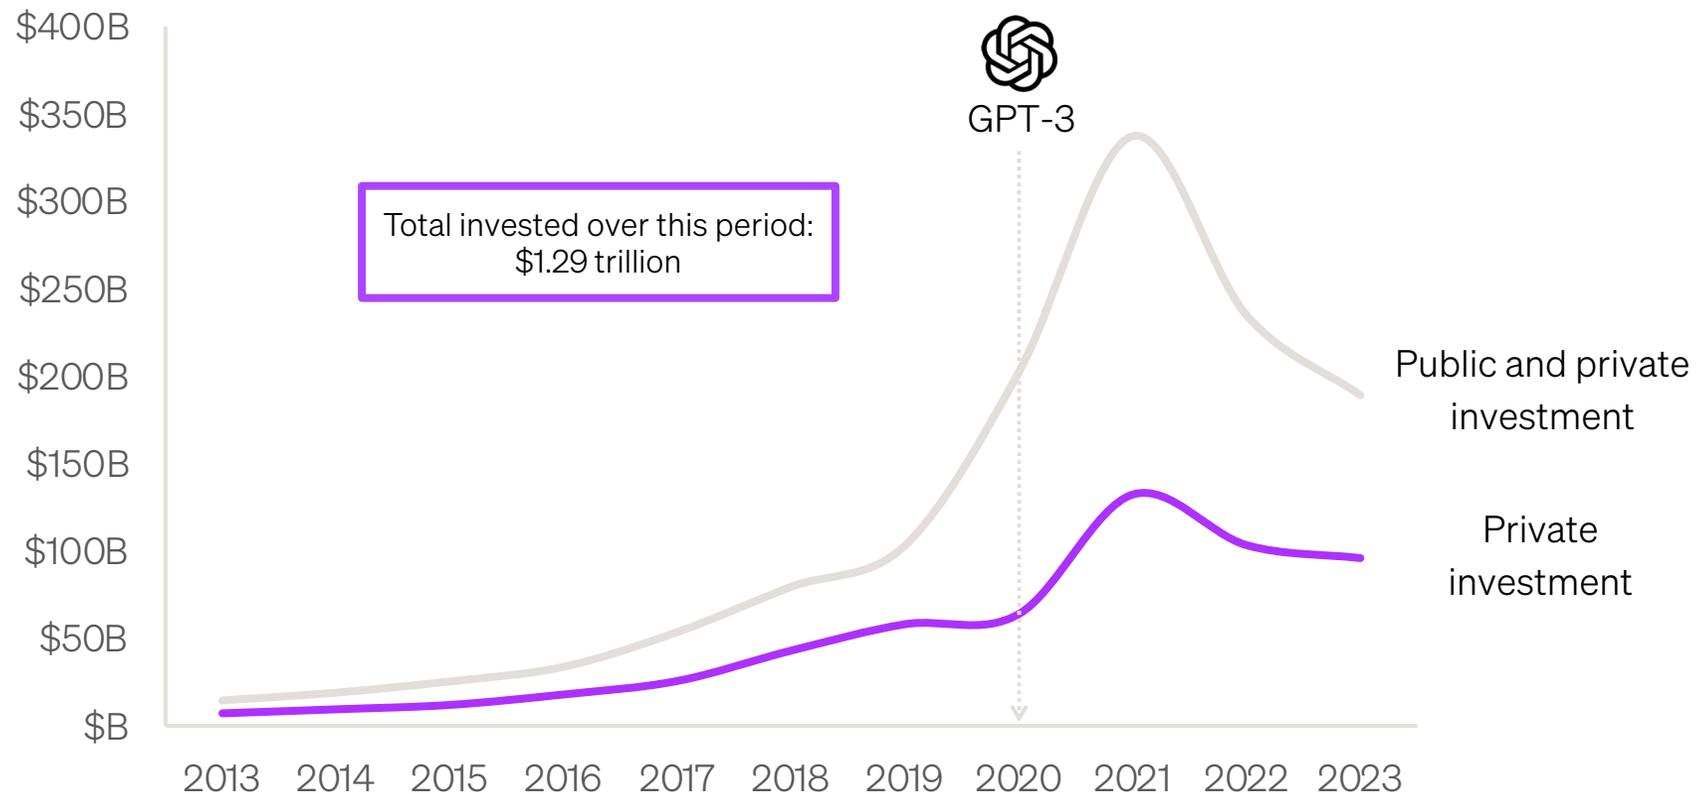
As models receive more compute time during fine tuning, they spontaneously develop the ability to think longer and reason better, without being explicitly programmed to do so.

Today's language models still have weaknesses in areas like reasoning, planning, creativity, and perception.



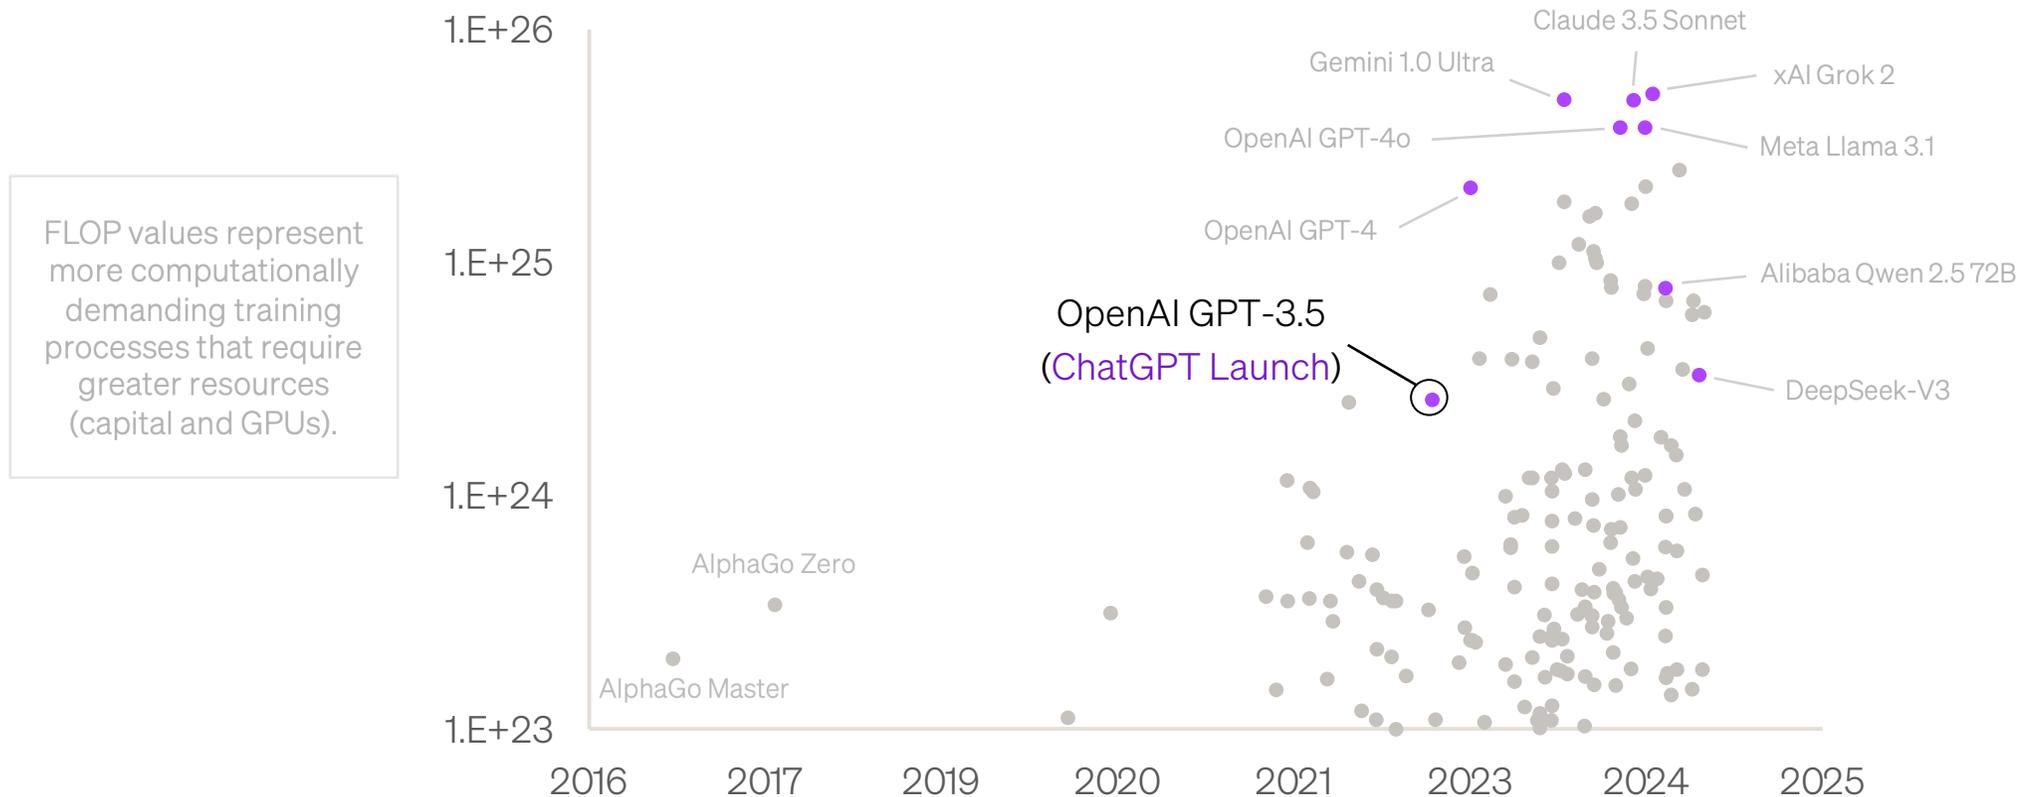
Private investment dollars have poured in, but it remains unclear where value will be captured in a defensible way.

Public and Private Investment in AI, Billions of U.S. Dollars, 2013 to 2024



Across the board, the dynamics of AI appear to be deflationary, both in foundational models and the first wave of applications.

AI Models, Training Compute (FLOP) Over Time, 2016 to 2025

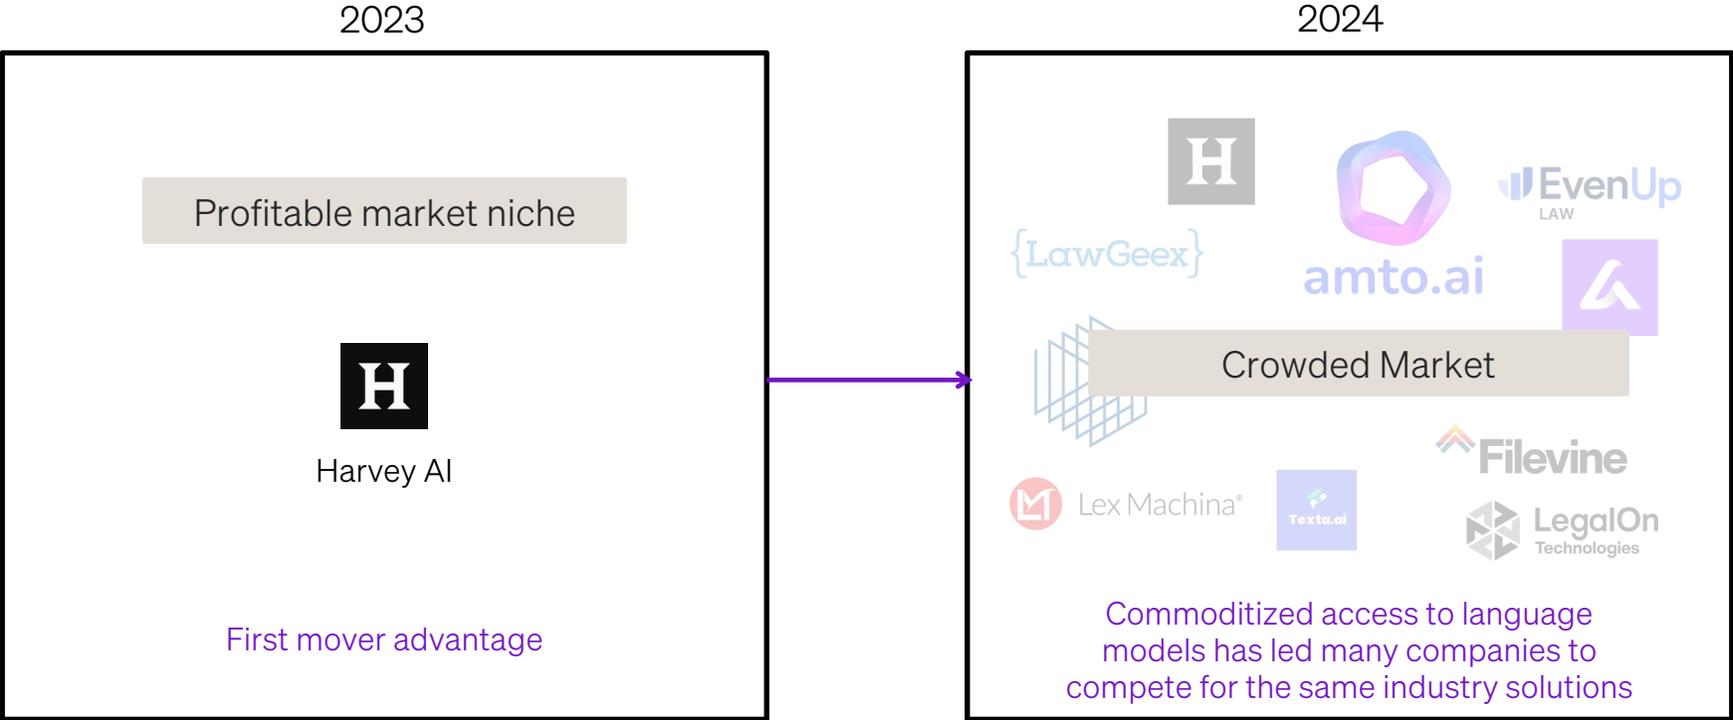


The first wave of applications are chatbots, coding, and internal tools, and there is currently a wave of agentic and voice applications emerging.

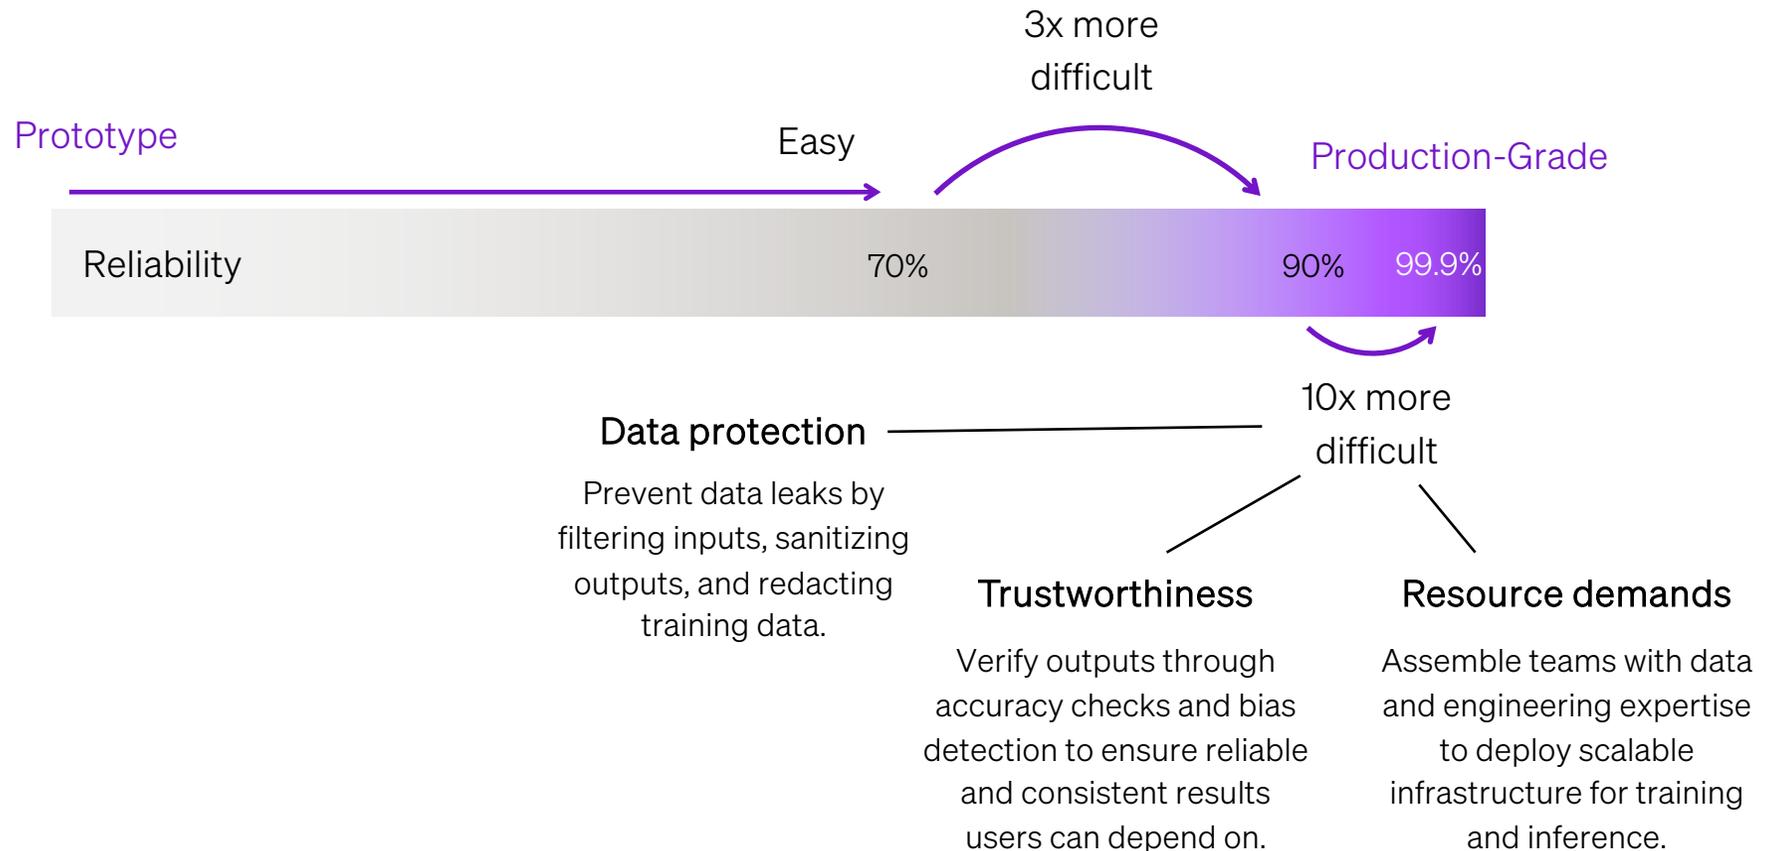
<u>Coding</u>	<u>Conversational</u>	<u>Healthcare</u>	<u>Legal</u>	<u>Marketing</u>
 Cursor AI	 ChatGPT	 Ada	 Ironclad	 Writesonic
 Replit Agent	 Claude	 Path AI	 Evenup	 Jasper
 Github Copilot	 Preplexity	 Deepscribe	 Casetext	 Copysmith
 Codeium	 Gemini	 Glass Health	 Harvey AI	 Rytr
 Cognition	 Copilot	 Memora	 Pincites	 Twain

Despite clear utility for many of these applications, value capture and creating defensibility in the company is challenging.

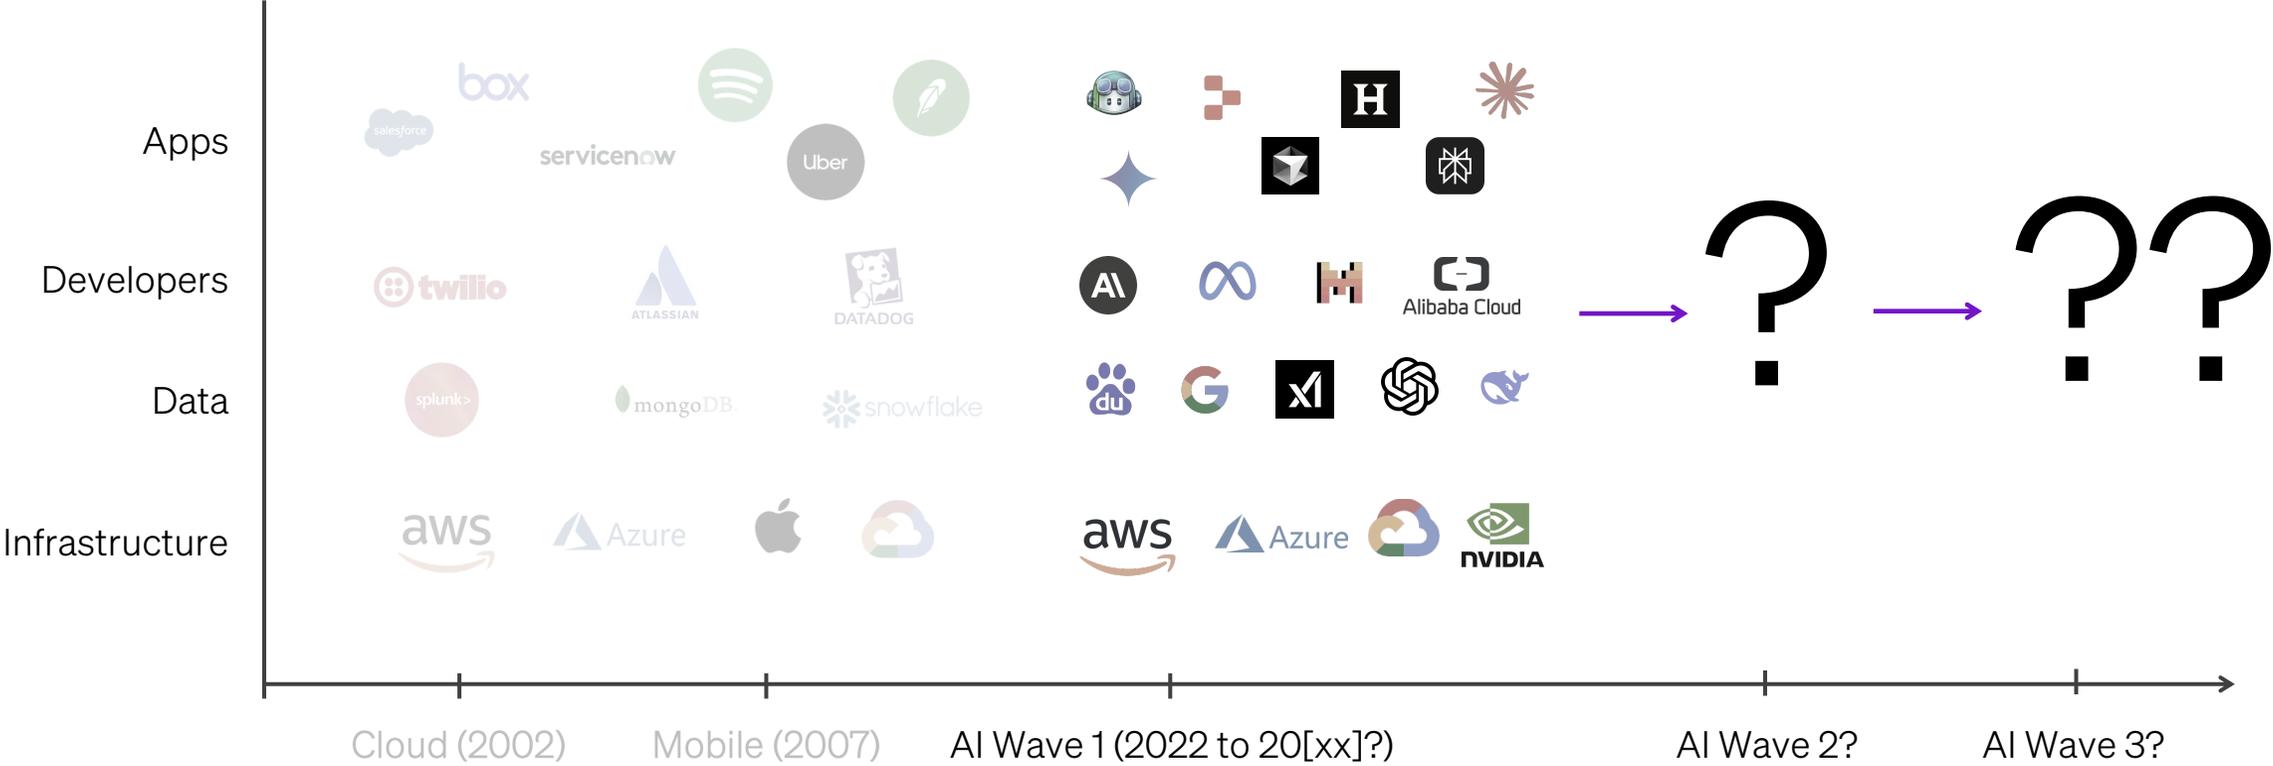
What happens when there is a low barrier of entry to create an MVP?



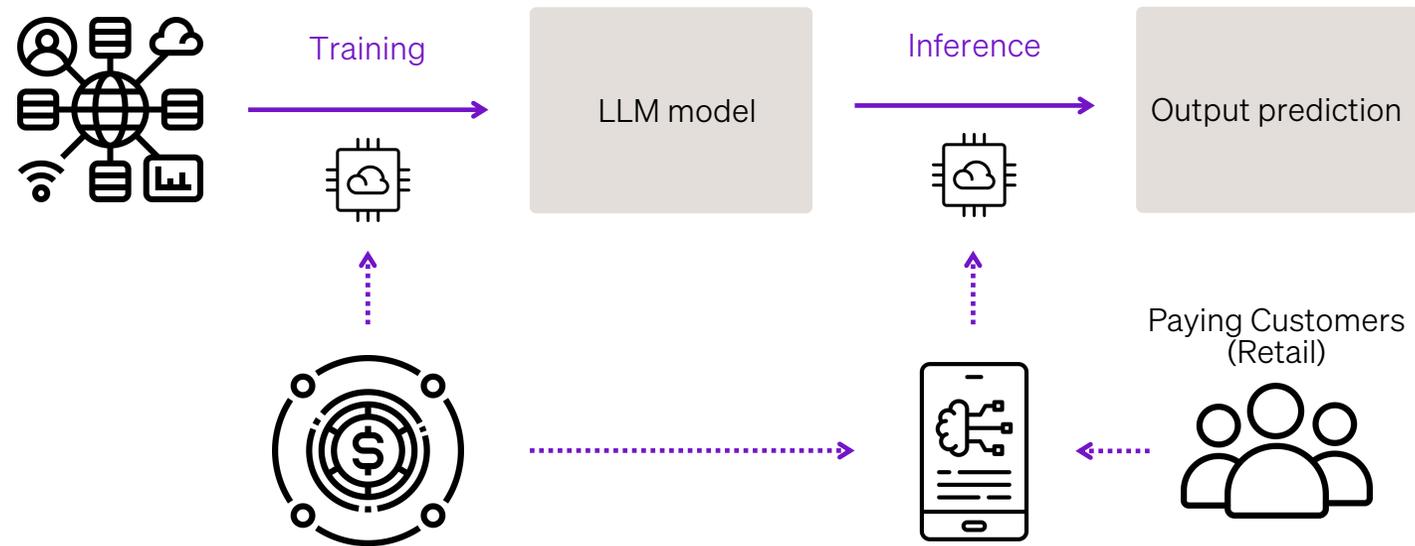
Even though prototyping with language models is easier than ever before, the real challenges emerge when engineering production-grade applications.



AI may be at the earliest innings of a multiple-decade-long megatrend.



Particular dynamics need to be in place for private investment dollars to result in outsized returns.

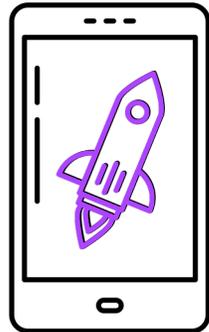


Investors are deploying capital in companies that develop foundational models, as well as applications that use the models.

Application developers pay for model usage, either for API calls, or for compute if they deploy their own infrastructure.

The questions remain: When will the next killer app emerge, and where will these billions of private investment dollars find returns?

Where will the next killer app emerge?



Where will investment dollars generate returns?



Future Deep Dives

Deep Dive Topics	Premise	Published
A Current State of Artificial Intelligence	How are language models evolving under resource constraints and competitive pressures? What lessons have we learned from this first wave of AI products? How do we identify the billion-dollar opportunities and killer apps in this megatrend?	✓
A Primer on The Magnificent Seven	How are the Magnificent Seven businesses structured and organized? What drives their revenue and bottom line? How are they allocating capital today, and are they positioned for continued outsized returns? How should we understand the Magnificent Seven in relation to the S&P 500?	
Quantum Computing and its Potential Applications	What are the different areas of quantum computing research? What is the current state of quantum computing, and how many breakthroughs away is it from practical application? From first principles, why does quantum computing represent a paradigm shift from previous computing approaches?	
Understanding Science and Evaluating Scientific Research	What is science? What is the right way to understand what science is and isn't? How has science evolved over the centuries as society has shifted and different forms of reasoning have emerged? How has the production of scientific research changed as universities have proliferated? How should we evaluate scientific research?	
China's Economic and Geopolitical Position	What has driven China's economic growth and is China positioned for continued growth? What is China's geopolitical strategy and position? What is the state of its technology industry and military capabilities? How should the U.S. position itself in response? Is conflict between the U.S. and China inevitable?	
A Primer on Marijuana and Psychedelics	What does scientific research say about marijuana and various psychedelics in terms of benefits and side effects? How do marijuana's complex interactions with anxiety, creativity, and memory vary across individuals and contexts? How is the regulatory landscape evolving to accommodate medical applications?	
Crypto Stablecoins and P2P Payments	What are the use cases for stablecoins, and why have stablecoins achieved product-market fit? Where is the high transaction volume coming from, and how are different stablecoin designs addressing technical and regulatory challenges? How should we understand and evaluate stablecoins as an asset class?	
A Primer on Materials Science	How have different atomic and molecular structures enabled different material properties? What physics explains why certain structures create specific properties? How do material choices that Apple and SpaceX make reveal about the interplay between aesthetics, engineering, and manufacturing?	
How to Invest and Allocate Capital	What is Warren Buffett's investing philosophy? How do modern portfolio theory and models like CAPM work? How do we think about investments at Social Capital and synthesize the thinking of key investors? Is today's investing climate different enough to warrant an approach different from the great investors of the past?	
A Primer on Longevity and Aging Science	What are the key areas of research when it comes to longevity and aging science? What are the most prominent longevity companies and where are they in terms of developing products? What is metabolic dysfunction and its relationship to chronic disease? Is there consensus on what we should do to be healthy and well?	
A Primer on Battery Technology	What is a battery and how do batteries work? Why do batteries need certain rare earth metals, and what is the cutting-edge battery technology and how could it be game-changing to fields like electric vehicles and energy transition? Why has battery innovation followed a linear rather than exponential improvement curve?	
The Current State of American Universities	What was the original purpose of college and how has that changed over time? Why do more people believe that college is not worthwhile? How did U.S. News & World Report and the proliferation of colleges and universities change what college is? Is it still worthwhile to go to college, and what subjects are most worthwhile to study?	

Disclaimer

This document is provided for educational purposes only. Nothing contained in this document is investment advice, a recommendation or an offer to sell, or a solicitation of an offer to buy, any securities or investment products. References herein to specific sectors are not to be considered a recommendation or solicitation for any such sector. Additionally, the contents herein are not to be construed as legal, business, or tax advice.

Statements in this document are made as of the date of this document unless stated otherwise, and there is no implication that the information contained herein is correct as of any other time. Certain information contained or linked to in this document has been obtained from sources believed to be reliable and current, but accuracy cannot be guaranteed.

This document contains statements that are not purely historical in nature but are “forward-looking statements” or statements of opinion or intention. Any projections included herein are also forward-looking statements. Forward-looking statements involve known and unknown risks, uncertainties (including those related to general economic conditions), assumptions and other factors, which may cause actual results, performance or achievements to be materially different from those expressed or implied by such forward-looking statements. Accordingly, all forward-looking statements should be evaluated with an understanding of their inherent uncertainty and recipients should not rely on such forward-looking statements. There is no obligation to update or revise these forward-looking statements for any reason.

This document also contains references to trademarks, service marks, trade names and copyrights of other companies, which are the property of their respective owners. Solely for convenience, trademarks and trade names referred to in this document may appear without the ® or ™ symbols, but such references are not intended to indicate, in any way, that such owner will not assert, to the fullest extent under applicable law, its rights or the right of the applicable licensor to these trademarks and trade names.