



Phase Transition in AI

Dr. Ece Kamar

VP of Research, AI Frontiers, Microsoft Research
Affiliated Faculty, University of Washington

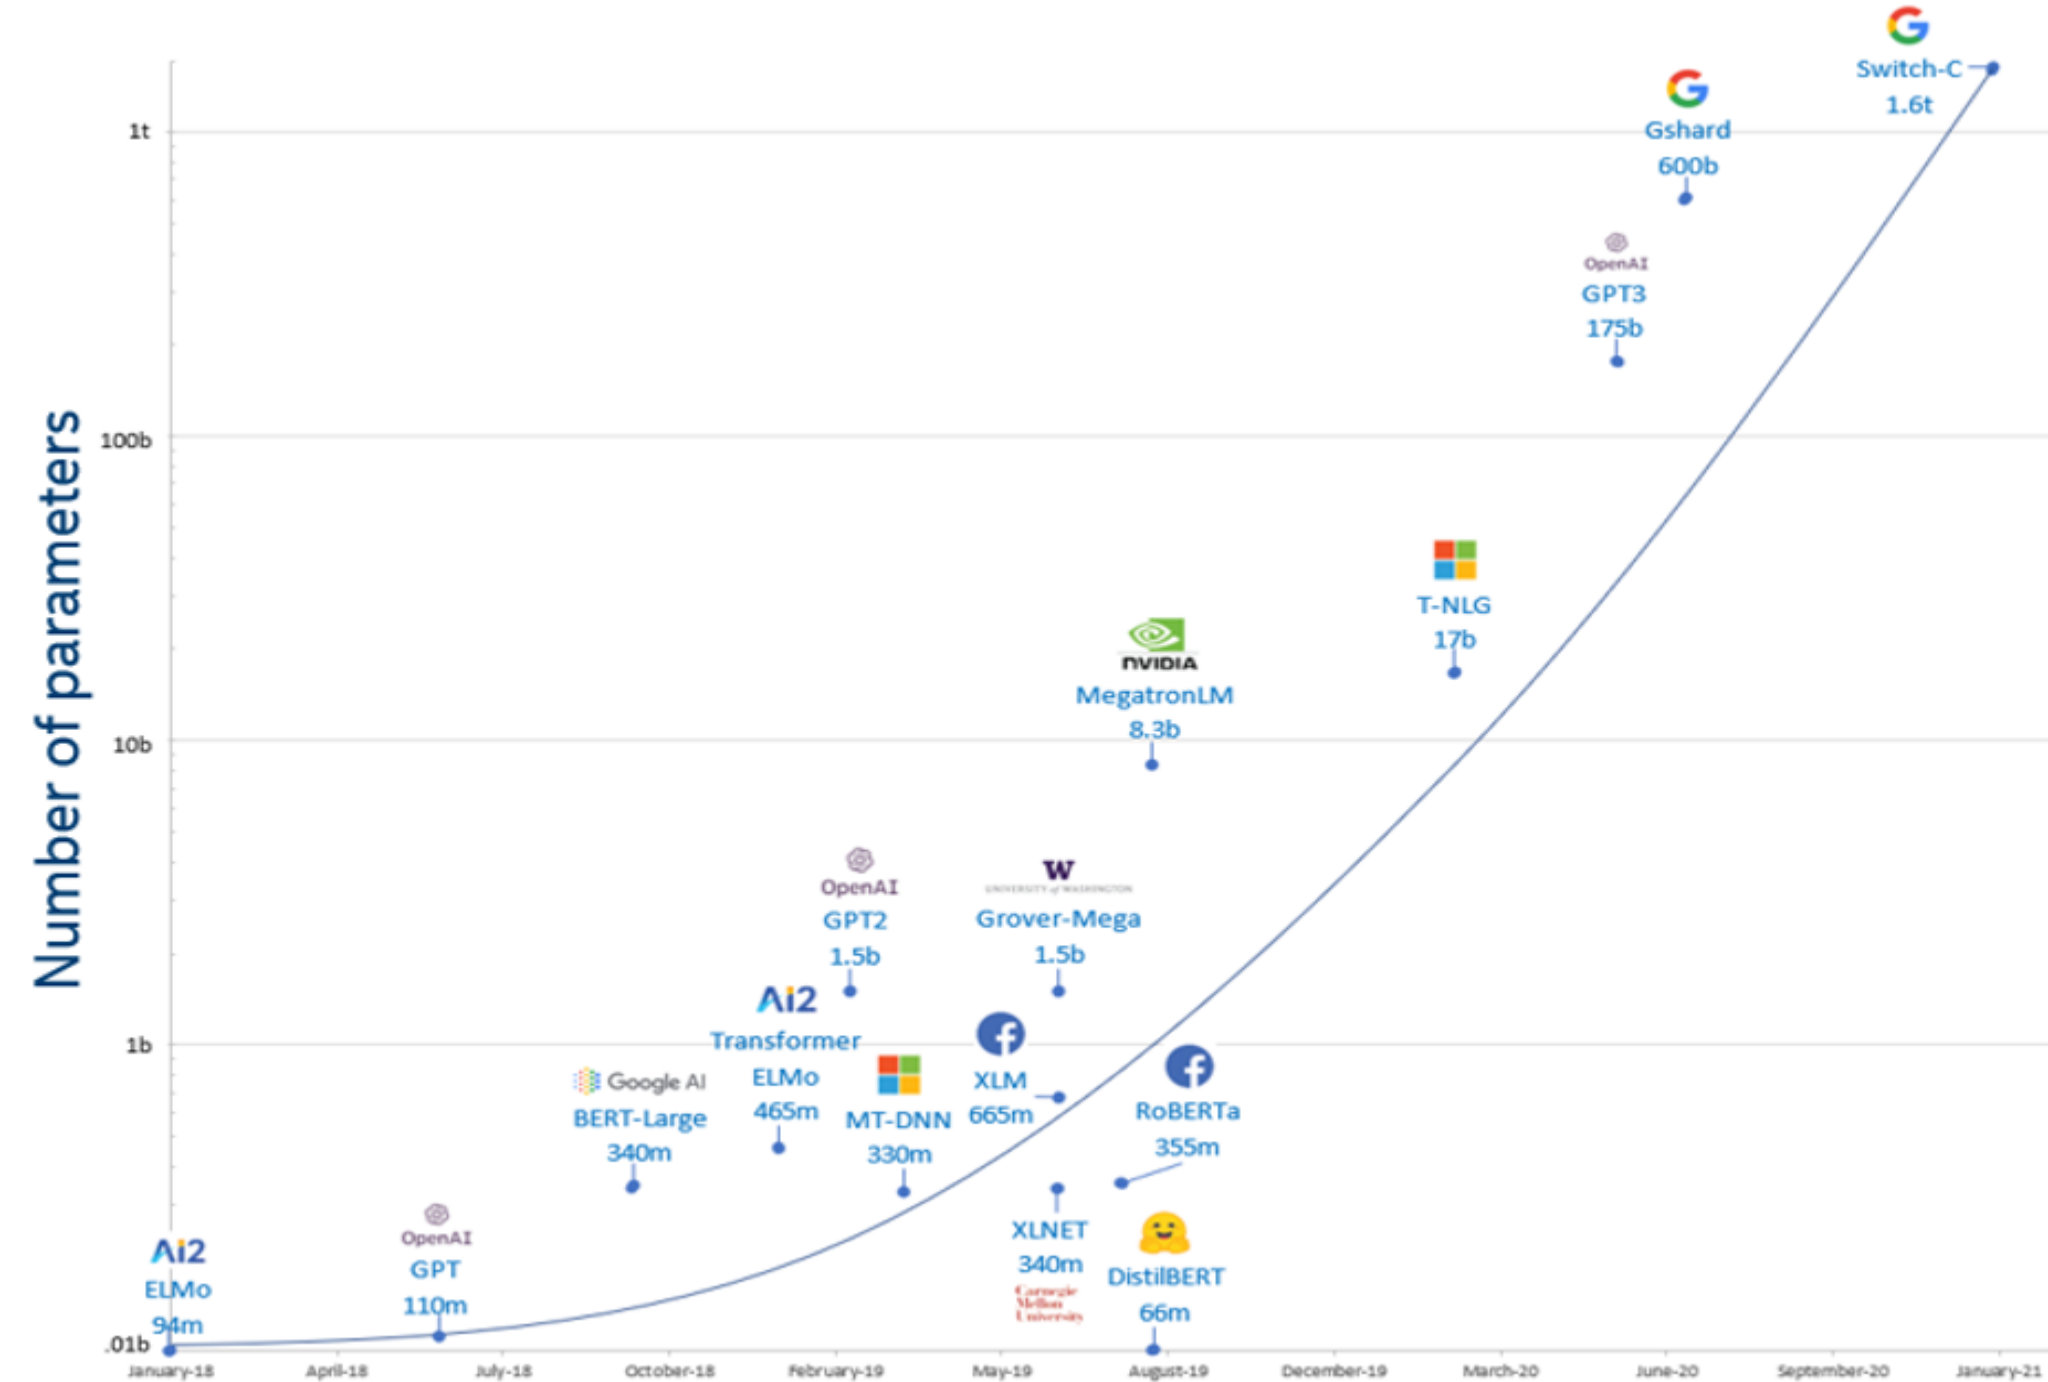


Figure 1: Exponential growth of number of parameters in DL models

[The Rise of Cognitive AI. On the role of knowledge that is...](#) | by Gadi Singer | [Towards Data Science](#)

Sparks of Artificial General Intelligence: Early experiments with GPT-4

Sébastien Bubeck Varun Chandrasekaran Ronen Eldan Johannes Gehrke
Eric Horvitz Ece Kamar Peter Lee Yin Tat Lee Yuanzhi Li Scott Lundberg
Harsha Nori Hamid Palangi Marco Tulio Ribeiro Yi Zhang

Microsoft Research

Abstract

Artificial intelligence (AI) researchers have been developing and refining large language models (LLMs) that exhibit remarkable capabilities across a variety of domains and tasks, challenging our understanding of learning and cognition. The latest model developed by OpenAI, GPT-4 [Ope23], was trained using an unprecedented scale of compute and data. In this paper, we report on our investigation of an early version of GPT-4, when it was still in active development by OpenAI. We contend that (this early version of) GPT-4 is part of a new cohort of LLMs (along with ChatGPT and Google’s PaLM for example) that exhibit more general intelligence than previous AI models. We discuss the rising capabilities and implications of these models. We demonstrate that, beyond its mastery of language, GPT-4 can solve novel and difficult tasks that span mathematics, coding, vision, medicine, law, psychology and more, without needing any special prompting. Moreover, in all of these tasks, GPT-4’s performance is strikingly close to human-level performance, and often vastly surpasses prior models such as ChatGPT. Given the breadth and depth of GPT-4’s capabilities, we believe that it could reasonably be viewed as an early (yet still incomplete) version of an artificial general intelligence (AGI) system. In our exploration of GPT-4, we put special emphasis on discovering its limitations, and we discuss the challenges ahead for advancing towards deeper and more comprehensive versions of AGI, including the possible need for pursuing a new paradigm that moves beyond next-word prediction. We conclude with reflections on societal influences of the recent technological leap and future research directions.

Why Phase Transition in AI?

Can you write a proof of infinitude of primes, with every line that rhymes?



Yes, I think I can,
Though it might take a clever plan.
I'll start by noting Euclid's proof,
Which shows that primes aren't just aloof.
Assume we have a finite list,
Of primes, and that none have been missed.
Multiply them all together,
And add one, just to be clever.
The result will be a number,
That has no prime factors, oh wonder!
But every number has a prime,
That divides it, it's just a matter of time.
So we've found a contradiction,
And our finite list needs eviction.
There must be infinitely many primes,
And that's the end of my rhyming lines.

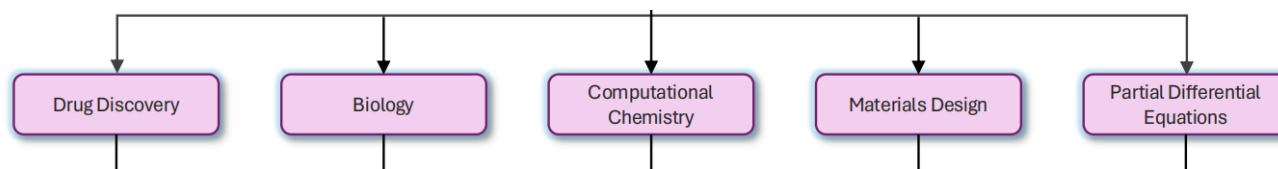
Significant progress observed in:

- Reasoning, planning
- Multi-modal translation
- Mental models, theory of mind
- Coding, tool usage
- Interactivity
- General application areas
- And more ...

Will LLMs Suffice?

The Impact of Large Language Models on Scientific Discovery: a Preliminary Study using GPT-4

Microsoft Research AI4Science
Microsoft Azure Quantum
July, 2023



- Weakness in precision and computation
- Expert on language of people, not on language of nature
- Not possible to verify correctness



Why do We Need New Approaches?

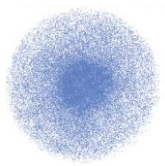
- Quantification, numerical calculations and verification is fundamental.
- Need to allow for iteration and experimentation.
- Can work with sparse & domain-specific data.
- Handle diverse data types and scientific literature.
- Domain specific knowledge and structure as inductive bias.

Good news:

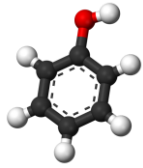
- Shared infrastructure & architectures
- Common domain-specific laws for simulation & emulation

Moonshot

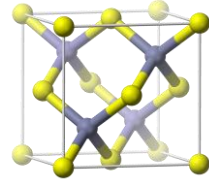
Electron



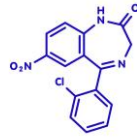
Molecule



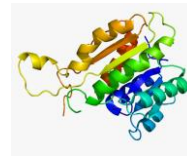
Material



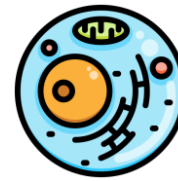
Drug



Protein



Cell



Human



Population



Foundation model for science

Getting to Scale

1

Unified data sources

2

Unified computing platforms

3

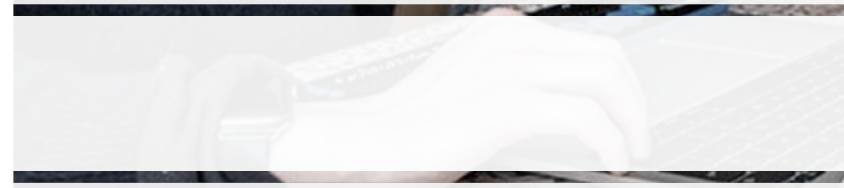
Sharing experiments, know-how

4

Tracking progress in quantitative and qualitative ways

5

Trading off experimentation and exploitation



Facial recognition for sensing
Microsoft's Project Tokyo

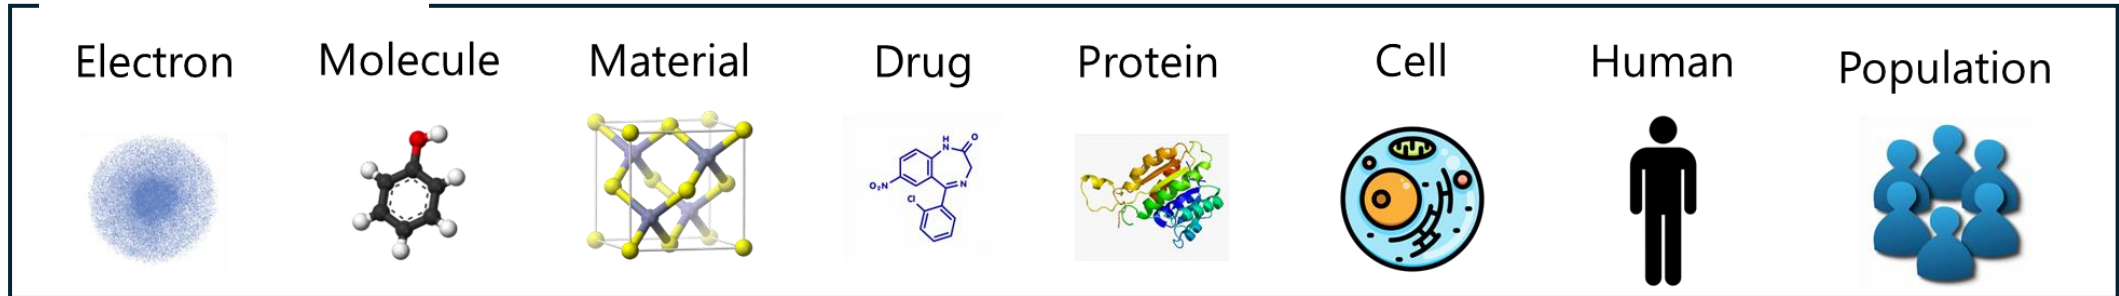


Risks: Dual Use

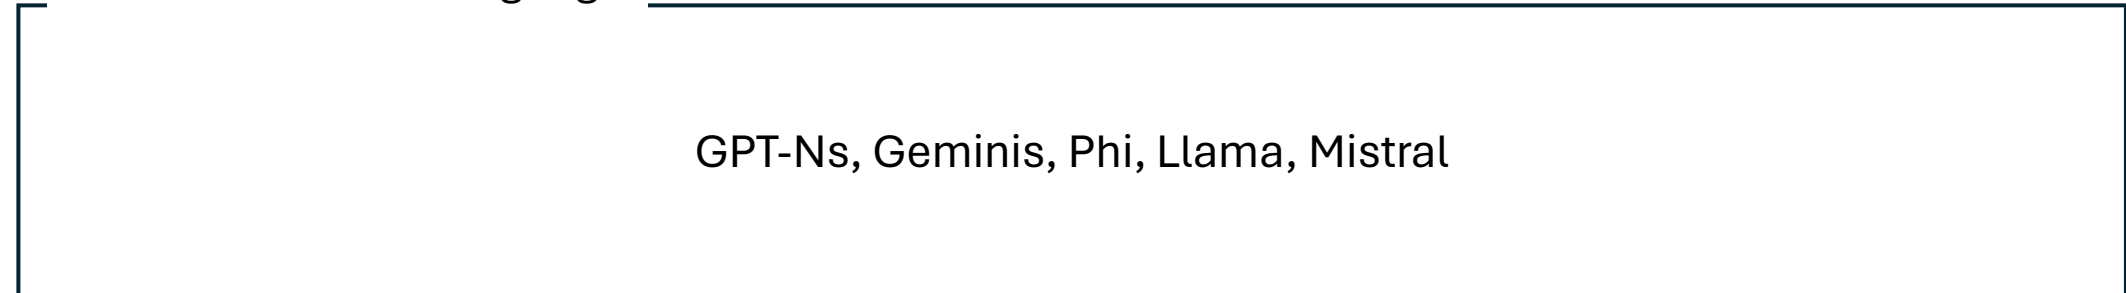
Facial recognition used to identify Lafayette Square protesters
Ira Black / Getty Images
The Washington Post, November 2, 2020

Translating Nature to Human Language

Models of nature



Models of human language



Accountability, complementarity and advancing human ingenuity