SGH
Szkoła Główna
Handlowa
w Warszawie

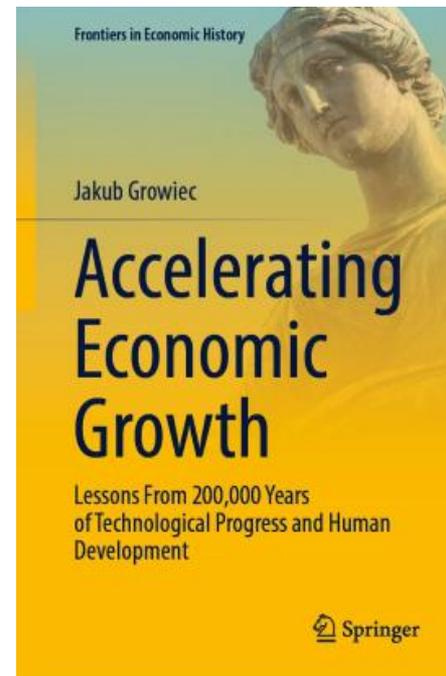# Latest Advances in AI and Future Pathways

Jakub Growiec

March 18, 2026, Geneva Association Economic Forum @PZU, Warszawa

I. AI and the long arc of technological progress

Five eras of development of the global civilization

- The hunter-gatherer era (~70 000 – 10 000 BP)
  - *Homo sapiens* conquers the world
- The agricultural era (~10 000 BP – 1550 CE)
  - Population grows, states appear
- The enlightenment era (~1550-1800 CE)
  - Accelerated technological progress
- The industrial era (~1800-1980 CE)
  - Physical capital accumulation, GDP *per capita* growth
- The digital era (from ~1980 CE)
  - Information processing outside of the human brain

Frontiers in Economic History

Jakub Growiec

Accelerating Economic Growth

Lessons From 200,000 Years of Technological Progress and Human Development

Springer

Economic growth has been accelerating over the eras

Era	Period	Measure	Growth	Doubling
Hunter-gatherer	25 000 – 5 000 BCE	Population	0.002%	34 360 y
Agricultural	5 000 BCE – 1500 CE	Population	0.069%	1007 y
Enlightenment	1500 CE – 1820 CE	Population	0.271%	256 y
		... world GDP	0.324%	214 y
Industrial	1820 – 1990 CE	... population	0.962%	72 y
		World GDP	2.14%	33 y
Digital	1990 – 2012 CE	... population	1.30%	54 y
		... world GDP	3.41%	21 y
		Data communication	28%	2.8 y

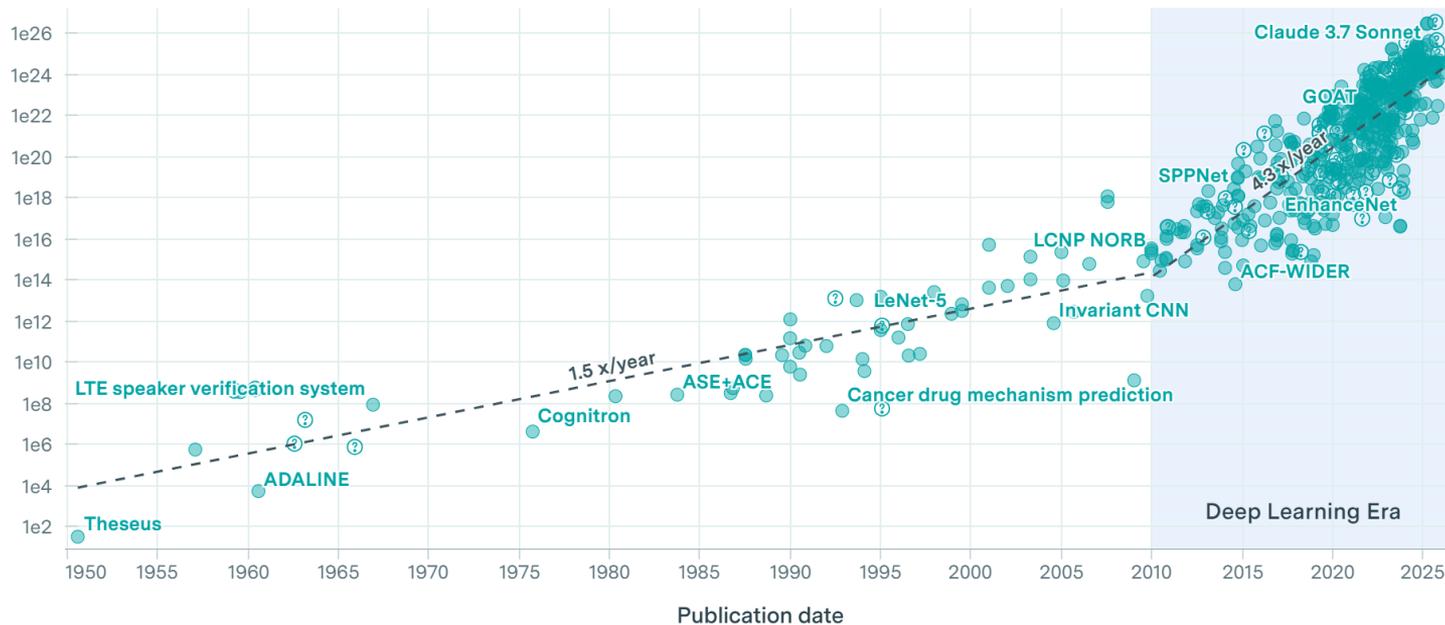Source: Growiec (2022), using data from Kremer (1993), Piketty (2014), Hilbert and López (2011)

Moore's Law: The number of transistors on microchips has doubled every two years

Moore's law describes the empirical regularity that the number of transistors on integrated circuits doubles approximately every two years.
This advancement is important for other aspects of technological progress in computing – such as processing speed or the price of computers.

Our World in Data

Transistor count

SGH

www.sgh.waw.pl

AI compute is scaling faster than Moore's Law

Notable AI models

Training compute (FLOP)

⑦ : Speculative data   540 Results

Labeled points: Claude 3.7 Sonnet, GOAT, SPPNet, EnhanceNet, LCNP NORB, ACF-WIDER, LeNet-5, Invariant CNN, ASE+ACE, Cancer drug mechanism prediction, LTE speaker verification system, Cognitron, ADALINE, Theseus

4.3 x/year
1.5 x/year

Deep Learning Era

Publication date
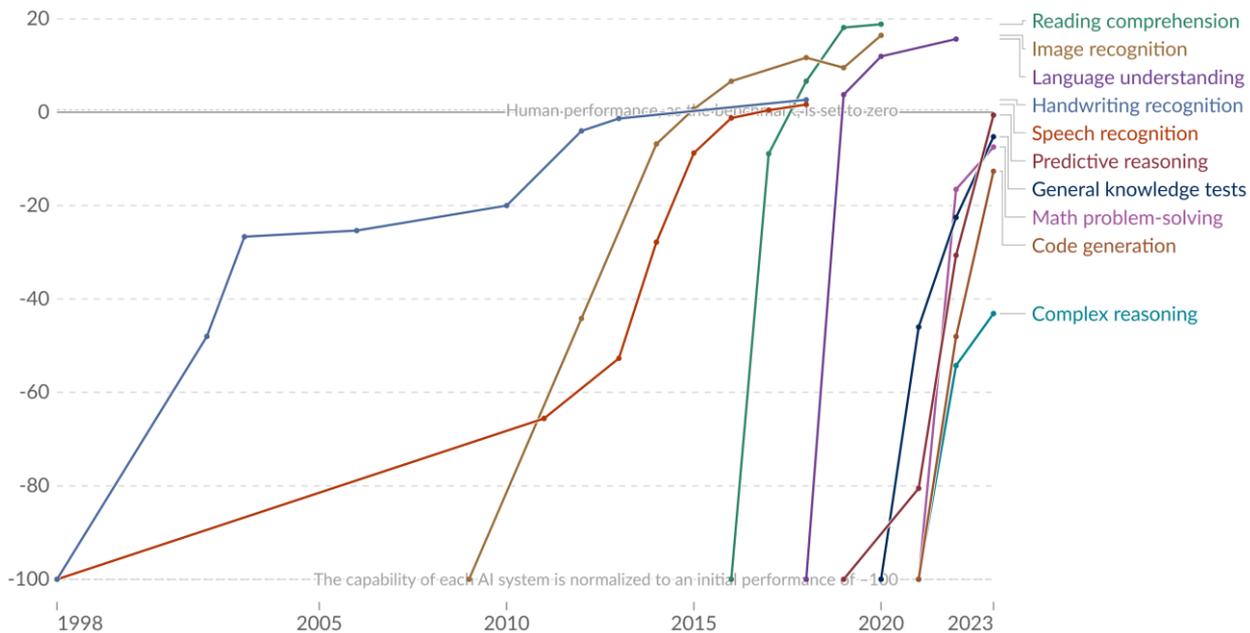
epoch.ai

Source: EpochAI (2026)

SGH

www.sgh.waw.pl

Quantity begets quality:
AI benchmarks are quickly saturated

Test scores of AI systems on various capabilities relative to human performance

Within each domain, the initial performance of the AI is set to −100. Human performance is used as a baseline, set to zero. When the AI's performance crosses the zero line, it scored more points than humans.

Reading comprehension
Image recognition
Language understanding
Handwriting recognition
Speech recognition
Predictive reasoning
General knowledge tests
Math problem-solving
Code generation

Complex reasoning

Human performance, as the benchmark, is set to zero

The capability of each AI system is normalized to an initial performance of −100

Data source: Kiela et al. (2023)

OurWorldinData.org/artificial-intelligence | CC BY

Note: For each capability, the first year always shows a baseline of −100, even if better performance was recorded later that year.

SGH

www.sgh.waw.pl

Cross-domain competence is rising

Epoch Capabilities Index (ECI)

epoch.ai

Source: EpochAI (2026)

Artificial General Intelligence is already here

nature

Explore content ⌄   About the journal ⌄   Publish with us ⌄ | Subscribe

nature > comment > article

COMMENT | 02 February 2026

# Does AI already have human-level intelligence? The evidence is clear

The vision of human-level machine intelligence laid out by Alan Turing in the 1950s is now a reality. Eyes unclouded by dread or hype will help us to prepare for what comes next.

By Eddy Keming Chen ✉, Mikhail Belkin, Leon Bergen & David Danks

Agentic capabilities are rising

Time horizon of software tasks
different LLMs can complete 80% of the time
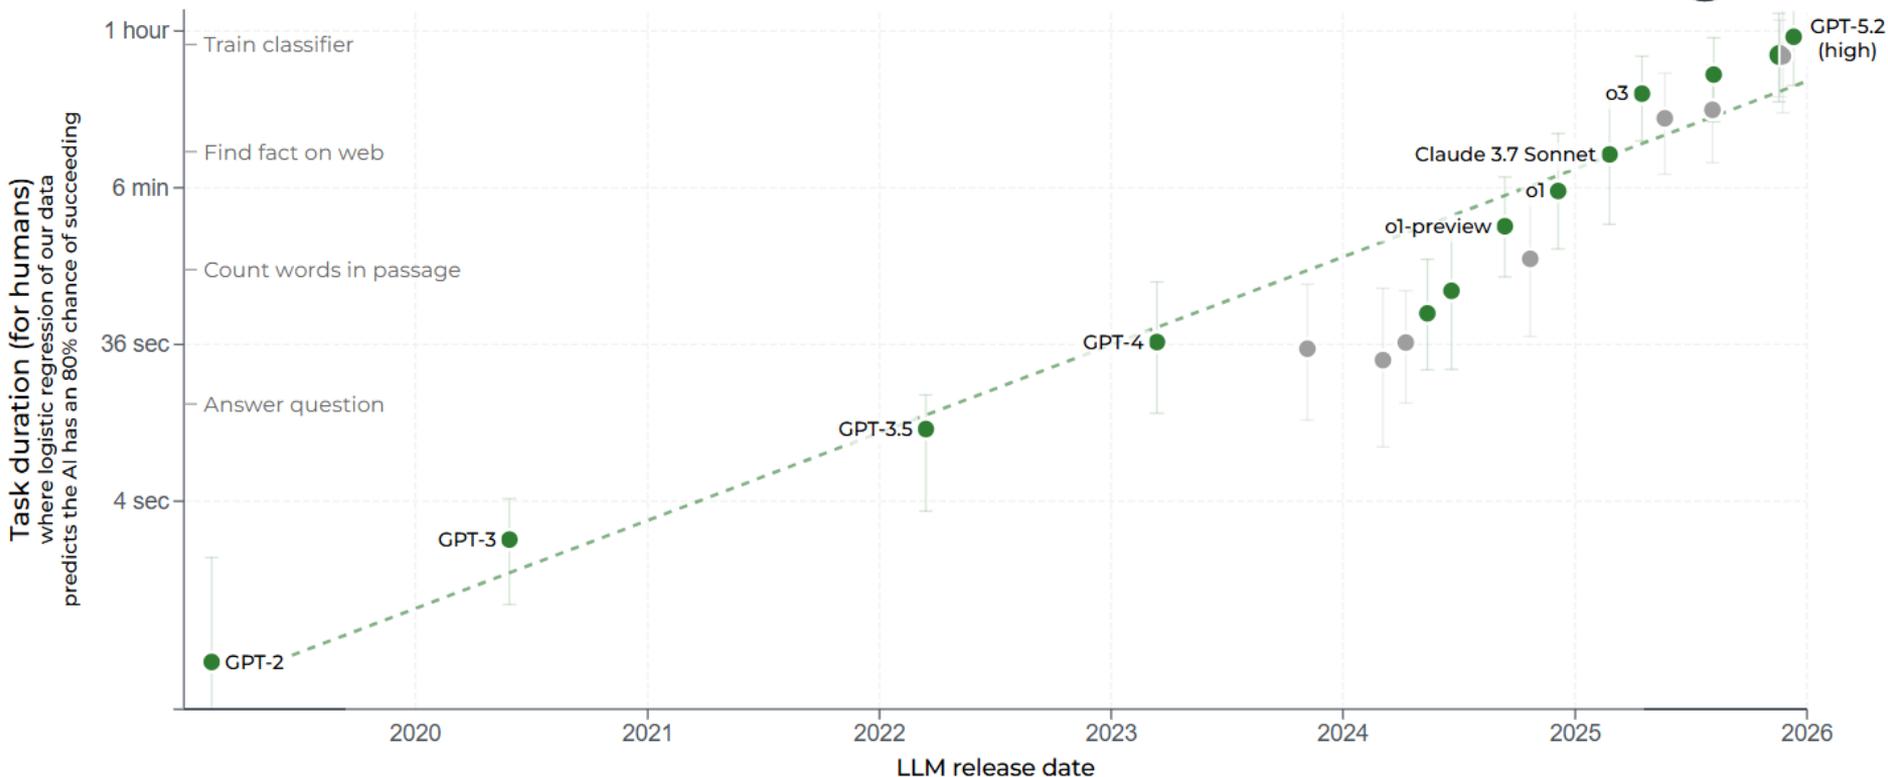
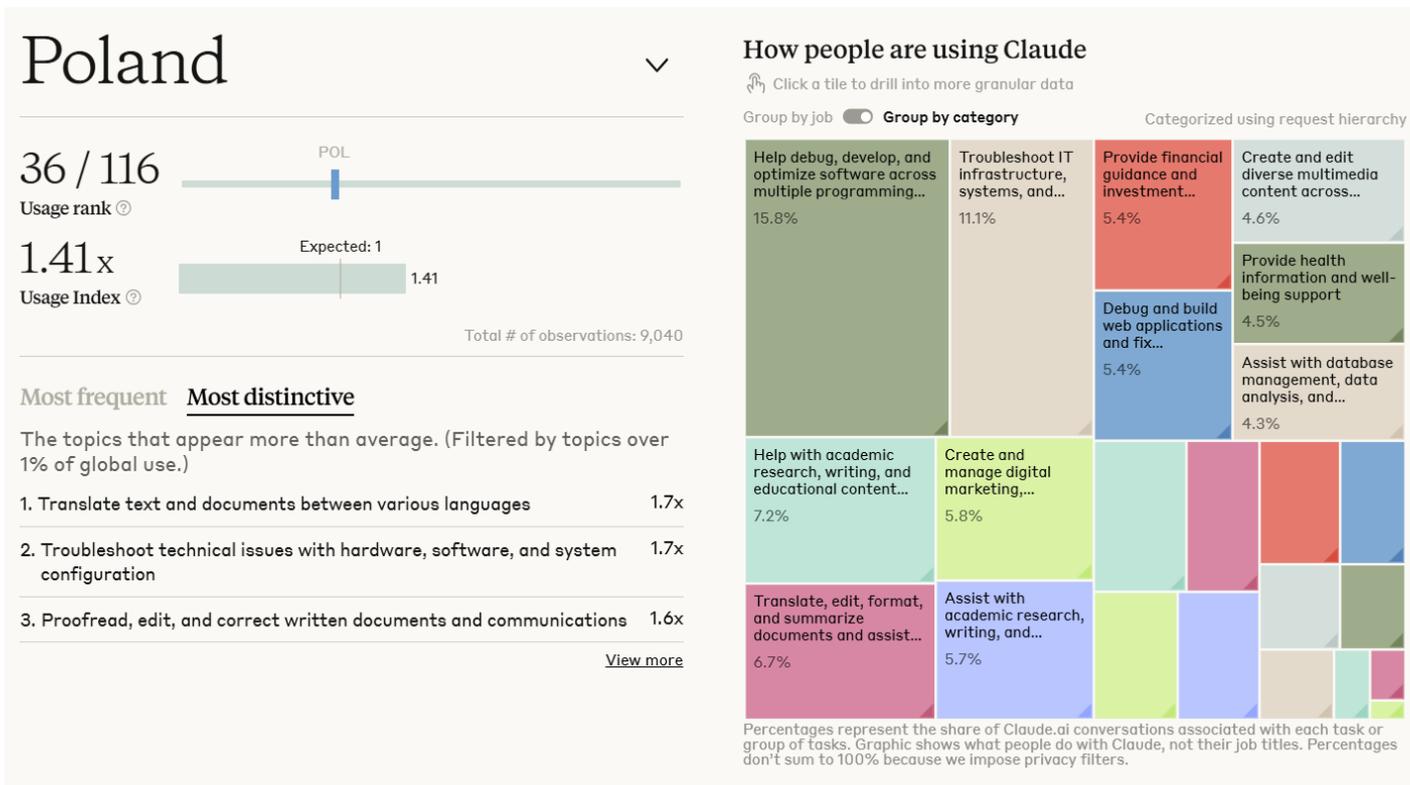# *AI timelines*: how far are we from TAI / ASI?

- OpenAI (2023): ``While superintelligence seems far off now, we believe it could arrive this decade" (until 2030).
- Aschenbrenner (2024), Kokotajlo et al. (2025): ~2027
- EpochAI (2024), `direct approach' model, ~2033
- Cotra (2022), `bio anchors' model, ~2040
- 2 778 researchers surveyed by Grace et al. (2024): ~2047

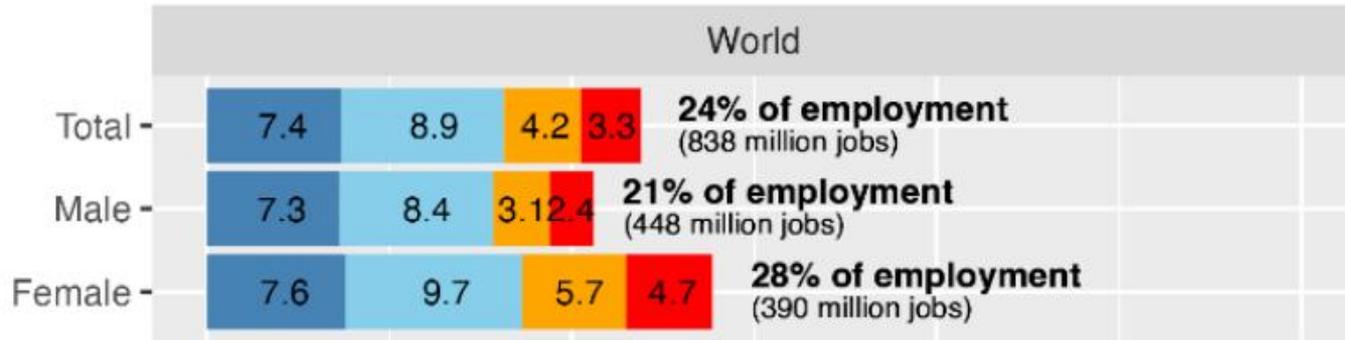**When will the first general AI system be devised, tested, and publicly announced?**

Current estimate

Dec 2033

2020   2080

Source: metaculus.com

II. New opportunities and risks

GenAI is currently used mostly in IT and clerical tasks

Poland

36 / 116
Usage rank ⓘ

POL

1.41 x
Usage Index ⓘ

Expected: 1   1.41

Total # of observations: 9,040

Most frequent   **Most distinctive**

The topics that appear more than average. (Filtered by topics over 1% of global use.)

1. Translate text and documents between various languages   1.7x

2. Troubleshoot technical issues with hardware, software, and system configuration   1.7x

3. Proofread, edit, and correct written documents and communications   1.6x

View more

How people are using Claude

🖑 Click a tile to drill into more granular data

Group by job   **Group by category**   Categorized using request hierarchy

Help debug, develop, and optimize software across multiple programming... **15.8%**

Troubleshoot IT infrastructure, systems, and... **11.1%**

Provide financial guidance and investment... **5.4%**

Create and edit diverse multimedia content across... **4.6%**

Provide health information and well-being support **4.5%**

Debug and build web applications and fix... **5.4%**

Assist with database management, data analysis, and... **4.3%**

Help with academic research, writing, and educational content... **7.2%**

Create and manage digital marketing,... **5.8%**

Translate, edit, format, and summarize documents and assist... **6.7%**

Assist with academic research, writing, and... **5.7%**

Percentages represent the share of Claude.ai conversations associated with each task or group of tasks. Graphic shows what people do with Claude, not their job titles. Percentages don't sum to 100% because we impose privacy filters.

SGH

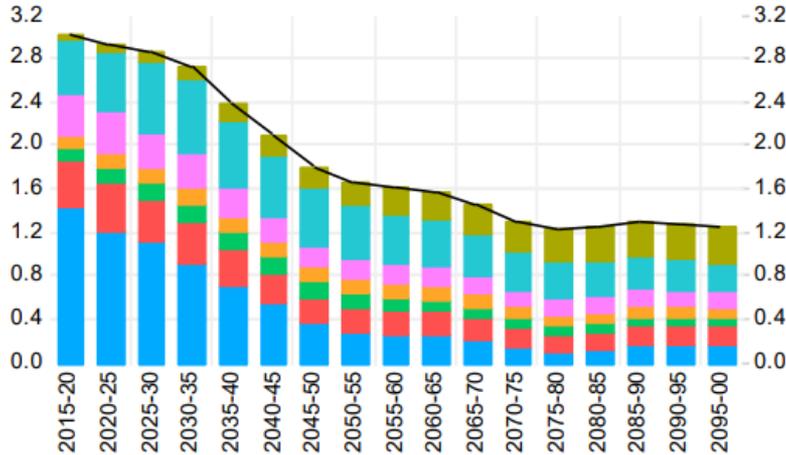Source: Anthropic Economic Index (2026)

Many jobs are susceptible to direct GenAI automation

Source: ILO (2025)

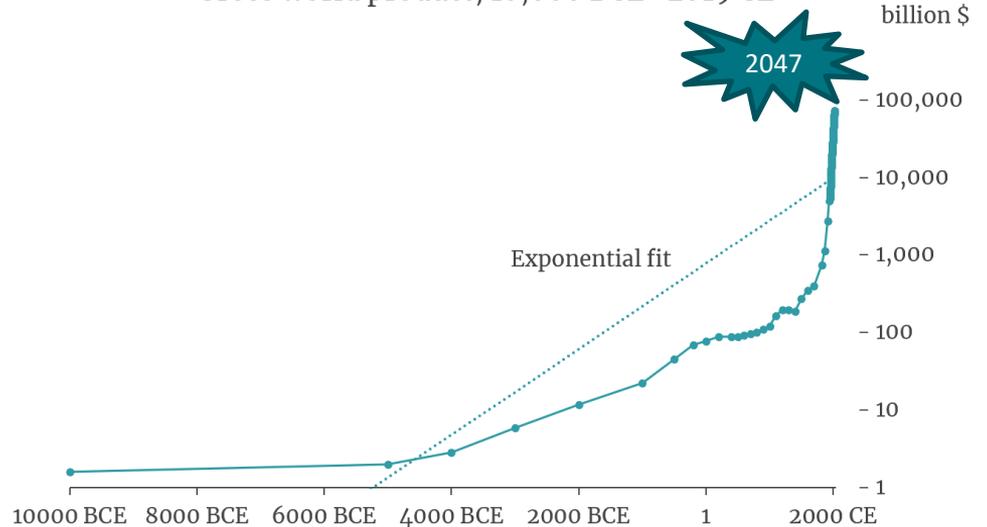# Long-term growth forecasts: AI makes *all* the difference

Gross world product, OECD forecast until 2100

A. Global potential real GDP growth, % pts

Guillemette (2025), *OECD global long-run economic scenarios: 2025 update.* OECD, Paris

Gross world product, 10,000 BCE−2019 CE

Roodman (2020), *On the probability distribution of long-term changes in the growth rate of the global economy: An outside view.* Open Philanthropy.

SGH

Three critical thresholds for TAI / ASI

1) Achieving superhuman **agency** and **reliability in pursuing plans**
   See the METR graph

2) Thanks to greater autonomy, becoming **substitutable**, rather than complementary, to human cognitive work

3) Entering a cascade of **recursive self-improvement**
   OpenAI (2026): "GPT-5.3-Codex is our first model that was instrumental in creating itself. The Codex team used early versions to debug its own training, manage its own deployment, and diagnose test results and evaluations."

The key risk is AI takeover, by misaligned TAI

<u>The economy</u>

Incentives and competitive pressures work in favor of an *AI takeover*: it pays to automate and hand over decision authority to AI, *at all management levels*

- Human decision making is often the bottleneck
- Greatest gains are achieved when a business process is **fully automated**

<u>AI</u>

Emergence of **goals and values** in AI (Mazeika et al., 2025)

- „*Mesa-optimizers*" created by optimization pressure in AI training
- **Instrumental convergence** (Bostrom, 2014)
- With current architectures, AI models are incomprehensible **black boxes**

SGH

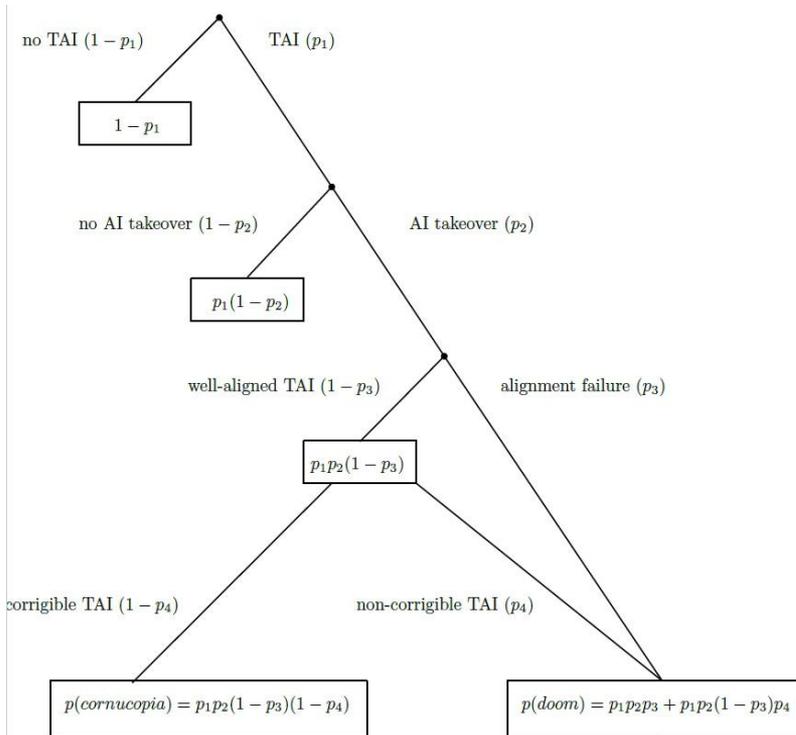# Existential risk from TAI / ASI

"What is your *p(doom)*?"

- ~ **0%**: Yann LeCun
- ~ **20%**: Yoshua Bengio
- ~ **50%**: Geoffrey Hinton, Paul Christiano
- ~ **80%**: Dan Hendrycks, Daniel Kokotajlo
- ~ **100%**: Eliezer Yudkowsky, Roman Yampolskiy

AI industry leaders admit it too:

- **10 − 25%**: Sam Altman (OpenAI),
  Dario Amodei (Anthropic), Elon Musk (xAI)
- **5 − 50%**: Shane Legg (Google DeepMind)

NEW YORK TIMES BESTSELLER

IF ANYONE BUILDS IT, EVERYONE DIES

WHY SUPERHUMAN AI WOULD KILL US ALL

ELIEZER YUDKOWSKY & NATE SOARES

SGH
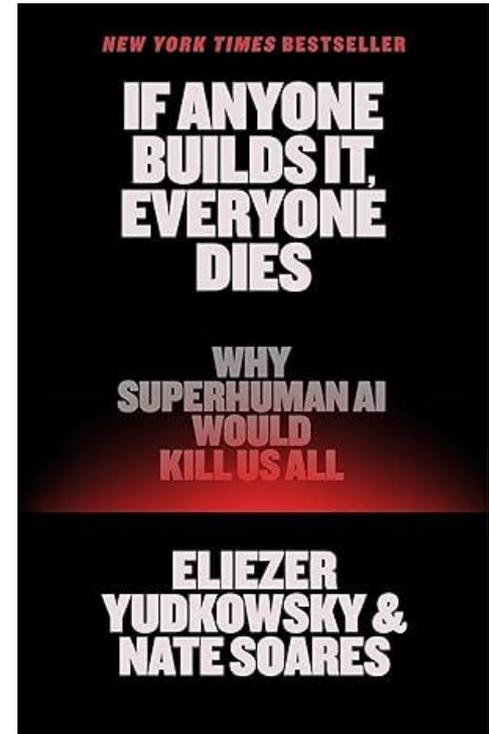
TAI / ASI will be a *filter* for humanity

Source: Growiec and Prettner (2025).

SGH

1/ Utopian outcomes ('cornucopia')
- Massively accelerated economic growth and technological change
- Cosmic expansion
- `Post-scarcity', immortality, brain uploads, etc. etc.

2/ Even without AI takeover, TAI can bring
- Technological unemployment
- Labor share falling towards zero
- Extreme income and wealth inequality
- Digital dictatorships (Harari)

3/ AI doom, e.g. human extinction
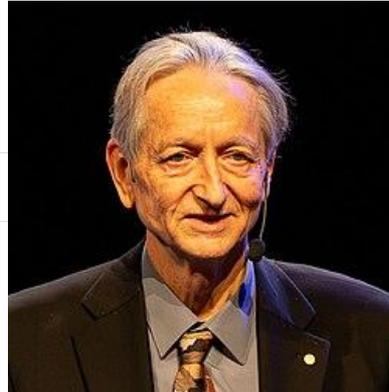
Key policy recommendations

The statement:

Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war.

Academic & industry leading signatories: (500+)

☑ AI Scientists   ☑ Other Notable Figures

Geoffrey Hinton
Emeritus Professor of Computer Science, University of Toronto

- Don't underestimate AI intelligence and agency
- We need **globally coordinated AI policy** (USA, China, EU, …)
- We need to put **AI safety ahead of AI capabilities**: the TAI alignment problem is still unsolved
- AI arms race is a suicide race

SGH

Jakub Growiec, SGH
https://e-web.sgh.waw.pl/jg23234/