

P(doom)

[3 languages](#)

[Article](#) [Talk](#)

[Read](#) [Edit](#) [View history](#)

From Wikipedia, the free encyclopedia

For the record label, see [P\(doom\) \(record label\)](#).

In [AI safety](#), **P(doom)** is the probability of [existentially catastrophic outcomes](#) (so-called "doomsday scenarios") as a result of [artificial intelligence](#).^{[1][2]} The exact outcomes in question differ from one prediction to another, but generally allude to the [existential risk from artificial intelligence](#).^[3] Economic analyses show that even small existential risks due to AI would justify substantial investments in AI safety and alignment research.^{[4][5]}

Originating as a shorthand for communication in the [rationalist community](#) and among AI researchers, the term came to prominence in 2023 following the release of [GPT-4](#), as high-profile figures such as [Geoffrey Hinton](#)^[6] and [Yoshua Bengio](#)^[7] began to warn of the risks of AI.^[8] In a 2023 survey, AI researchers were asked to estimate the probability that future AI advancements could lead to human extinction or similarly severe and permanent disempowerment within the next 100 years. The mean value from the responses was 14.4%, with a median value of 5%.^[9]

Notable P(doom) values [\[edit \]](#)

Name	P(doom)	Notes	Ref.
Dario Amodei	10-25%	CEO of Anthropic	^[10]
Marc Andreessen	0%	Co-founder of venture capital firm Andreessen Horowitz	^[11]
Sam Altman	>0%	CEO of OpenAI	^[12]
Yoshua Bengio	20% ^[Note 1]	Computer scientist, director of the Mila institute, and the Turing award winner	^[3]
Grady Booch	0% ^[Note 2]	American software engineer	^[1]
Vitalik Buterin	12%	Co-founder of Ethereum	^[13]
Paul Christiano	50%	Head of research at the US AI Safety Institute	^[14]
Andrew Critch	85%	Founder of the Center for Applied Rationality	^[15]

Lex Fridman	10%	American computer scientist and host of Lex Fridman Podcast	[16]
Demis Hassabis	>0%	Co-founder and CEO of Google DeepMind and Isomorphic Labs and 2024 Nobel Prize laureate in Chemistry	[17]
Dan Hendrycks	>80% ^[Note 3]	Director of Center for AI Safety	[1]
Geoffrey Hinton	10-20% (all-things-considered) >50% (independent impression)	"Godfather of AI" and 2024 Nobel Prize laureate in Physics	[18]
Holden Karnofsky	50%	Executive Director of Open Philanthropy	[19]
Lina Khan	≈ 15%	Former chair of the Federal Trade Commission	[8]
Daniel Kokotajlo	70–80%	AI researcher and founder of AI Futures Project , formerly of OpenAI	[20]
Connor Leahy	90%+	German-American AI researcher; cofounder of EleutherAI .	[21]
Yann LeCun	<0.01% ^[Note 4]	Previously chief AI Scientist at Meta	[22]
Shane Legg	≈ 5–50%	Co-founder and Chief AGI Scientist of Google DeepMind	[23]
Jan Leike	10–90%	AI alignment researcher at Anthropic , formerly of DeepMind and OpenAI	[1]
Benjamin Mann	0–10%	Co-founder of Anthropic	[24]
Emad Mostaque	50%	Co-founder of Stability AI	[25]
Zvi Mowshowitz	70%	Writer on artificial intelligence, director on the board of the Center for Applied Rationality , former competitive Magic: The Gathering player	[26]
Elon Musk	≈ 10–30%	Businessman and CEO of X , Tesla , and SpaceX	[27]
Casey Newton	5%	American technology journalist	[1]
Toby Ord	10%	Australian philosopher and author of The Precipice	[28]
Emmett Shear	5–50%	Co-founder of Twitch and former interim CEO of OpenAI	[8]
Nate Silver	5–10%	Statistician, founder of FiveThirtyEight	[29]

Max Tegmark	>90%	Swedish-American physicist, machine learning researcher, and author, best known for theorising the mathematical universe hypothesis and co-founding the Future of Life Institute .	^[30]
Roman Yampolskiy	99.9%–99.999999% ^[Note 5]	Latvian computer scientist, formerly a research advisor of the Machine Intelligence Research Institute , and an AI safety fellow of the Foresight Institute	^[31] ^[32]
Eliezer Yudkowsky	>95%	Founder of the Machine Intelligence Research Institute , author of <i>If Anyone Builds It, Everyone Dies</i> .	^[1]

Criticism ^{[[edit](#)]}

There has been some debate about the usefulness of P(doom) as a term, in part due to the lack of clarity about whether or not a given prediction is conditional on the existence of [artificial general intelligence](#), the time frame, and the precise meaning of "doom".^[8]^[33]

In popular culture ^{[[edit](#)]}

In 2024, Australian rock band [King Gizzard & the Lizard Wizard](#) launched their new label, named p(doom) Records.^[34]

See also ^{[[edit](#)]}

- [Existential risk from artificial general intelligence](#)
- [Statement on AI risk of extinction](#)
- [AI alignment](#)
- [AI takeover](#)
- [AI safety](#)
- [Future of Life Institute](#)
- [Survivalism](#)

Notes ^{[[edit](#)]}

- ↑ "Professor Bengio shared his p(doom), saying: "I got around, like, 20 per cent probability that it turns out catastrophic." Professor Bengio arrived at the figure based on several inputs, including a 50 per cent probability that AI would reach human-level capabilities within a decade, and a greater than 50 per cent likelihood that AI or humans themselves would turn the technology against humanity at scale."
- ↑ Equivalent to "P(all the oxygen in my room spontaneously moving to a corner thereby suffocating me)".
- ↑ Up from ~20% 2 years prior.
- ↑ "Less likely than an asteroid wiping us out".
- ↑ Within the next 100 years.

References [\[edit \]](#)

1. [^] [a b c d e f](#) Railey, Clint (2023-07-12). "P(doom) is AI's latest apocalypse metric. Here's how to calculate your score" . *Fast Company*.
2. [^] Thomas, Sean (2024-03-04). "Are we ready for P(doom)?" . *The Spectator*. Retrieved 2024-06-19.
3. [^] [a b](#) "It started as a dark in-joke. It could also be one of the most important questions facing humanity" . *ABC News*. 2023-07-14. Retrieved 2024-06-18.
4. [^] Jones, Charles I. (2024). "The AI Dilemma: Growth versus Existential Risk" . *American Economic Review: Insights*. **6** (4): 575–590. doi:[10.1257/aeri.20230570](#) . ISSN 2640-205X .
5. [^] Growiec, Jakub; Prettnner, Klaus (2026). "The economics of p(doom): Scenarios of existential risk and economic growth in the age of transformative AI" . *Economic Modelling*. **163** 107718. doi:[10.1016/j.econmod.2026.107718](#) . ISSN 0264-9993 .
6. [^] Metz, Cade (2023-05-01). "'The Godfather of A.I.' Leaves Google and Warns of Danger Ahead" . *The New York Times*. ISSN 0362-4331 . Retrieved 2024-06-19.
7. [^] "One of the 'godfathers of AI' airs his concerns" . *The Economist*. ISSN 0013-0613 . Retrieved 2024-06-19.
8. [^] [a b c d](#) Roose, Kevin (2023-12-06). "Silicon Valley Confronts a Grim New A.I. Metric" . *The New York Times*. ISSN 0362-4331 . Retrieved 2024-06-17.
9. [^] Piper, Kelsey (2024-01-10). "Thousands of AI experts are torn about what they've created, new study finds" . *Vox*. Retrieved 2024-09-02.
Citing:
"2023 Expert Survey on Progress in AI [AI Impacts Wiki]" . *wiki.aiimpacts.org*. Retrieved 2024-09-02.
10. [^] Morrone, Megan. "Amodei on AI: 'There's a 25% chance that things go really, really badly'" . *Axios*. Retrieved 19 September 2025.
11. [^] Marantz, Andrew (2024-03-11). "Among the A.I. Doomsayers" . *The New Yorker*. ISSN 0028-792X . Retrieved 2024-06-19.
12. [^] "OpenAI CEO Sam Altman: AI Warfare, Freedom & Immortality | MD MEETS Episode #1" . *YouTube*. October 1, 2025.
13. [^] Shapira, Liron (2025-08-12). "VITALIK vs. AI DOOMER — Will 'd/acc' Protect Humanity from Superintelligent AI?" . *YouTube*. Retrieved 2025-10-03.
14. [^] "ChatGPT creator says there's 50% chance AI ends in 'doom'" . *The Independent*. 2023-05-03. Retrieved 2024-06-19.
15. [^] Doom Debates (2024-11-15). *Andrew Critch vs. Liron Shapira: Will AI Extinction Be Fast Or Slow?* . Retrieved 2025-05-31 – via YouTube.
16. [^] "Archives: sundar-pichai-transcript" . *Lex Fridman*. 2025-06-05. Retrieved 2025-06-06.
17. [^] "Demis Hassabis on Chatbots to AGI | EP 71" . *YouTube*. 23 February 2024. Retrieved 8 October 2024.
18. [^] METR (Model Evaluation & Threat Research) (2024-06-27). *Q&A with Geoffrey Hinton* . Retrieved 2025-02-07 – via YouTube.
19. [^] Thomas, Sean (2024-03-04). "Are we ready for P(doom)?" . *The Spectator*. Retrieved 2025-05-31.
20. [^] Kokotajlo, Daniel (April 3, 2025). "2027 Intelligence Explosion Month-by-Month Model" . *youtube.com*.
21. [^] Future of Life Institute (2024-11-22). *Connor Leahy on Why Humanity Risks Extinction from AGI* . Retrieved 2025-05-31 – via YouTube.
22. [^] Wayne Williams (2024-04-07). "Top AI researcher says AI will end humanity and we should stop developing it now — but don't worry, Elon Musk disagrees" . *TechRadar*. Retrieved 2024-06-19.
23. [^] "Q&A with Shane Legg on risks from AI" . *Less Wrong*. 2011-05-17. Retrieved 2025-05-23.

24. ^ Lenny's Podcast (2025-07-20). [Anthropic co-founder: AGI predictions, leaving OpenAI, what keeps him up at night | Ben Mann](#) . Retrieved 2025-07-22 – via YouTube.
25. ^ Mostaque, Emad [@EMostaque] (2024-12-04). ["My P\(doom\) is 50%"](#) (Tweet). Retrieved 2025-05-31 – via X (formerly Twitter).
26. ^ ["Zvi Mowshowitz on Longer Timelines, RL-induced Doom, and Why China is Refusing H20s"](#) . *The Cognitive Revolution*. September 6, 2025.
27. ^ Tangalakis-Lippert, Katherine. ["Elon Musk says there could be a 20% chance AI destroys humanity — but we should do it anyway"](#) . *Business Insider*. Retrieved 2024-06-19.
28. ^ ["Is there really a 1 in 6 chance of human extinction this century?"](#) . *ABC News*. 2023-10-08. Retrieved 2024-09-01.
29. ^ ["It's time to come to grips with AI"](#) . *Silver Bulletin*. 2025-01-27. Retrieved 2025-02-03.
30. ^ Tegmark, Max (30 April 2025). ["My P\(doom\) Estimate"](#) . *X (formerly Twitter)*. [Archived](#) from the original on 2025-05-02. Retrieved 2025-05-31.
31. ^ Altchek, Ana. ["Why this AI researcher thinks there's a 99.9% chance AI wipes us out"](#) . *Business Insider*. Retrieved 2024-06-18.
32. ^ Yampolskiy, Roman [@romanyam] (12 March 2024). ["I am at 99.999999% https://pauseai.info/pdoom"](#) (Tweet) – via X (formerly Twitter).
33. ^ King, Isaac (2024-01-01). ["Stop talking about p\(doom\)"](#) . *LessWrong*.
34. ^ ["GUM & Ambrose Kenny-Smith are teaming up again for new collaborative album 'III Times' "](#) . *DIY*. 2024-05-07. Retrieved 2024-06-19.

External links [[edit](#)]

- Statement on Superintelligence

Categories: Existential risk from artificial intelligence | Prediction | Surveys (human research)

This page was last edited on 20 July 2026, at 06:17 (UTC). Page was rendered with Parsoid.

Text is available under the [Creative Commons Attribution-ShareAlike 4.0 License](#); additional terms may apply. By using this site, you agree to the [Terms of Use](#) and [Privacy Policy](#). Wikipedia® is a registered trademark of the [Wikimedia Foundation, Inc.](#), a non-profit organization.

[Privacy policy](#) [About Wikipedia](#) [Disclaimers](#) [Contact Wikipedia](#) [Legal & safety contacts](#) [Code of Conduct](#) [Developers](#) [Statistics](#)

[Cookie statement](#) [Mobile view](#)