

Superintelligence Has Authors, the US Government Is Not One of Them

President Trump renamed artificial intelligence after a mind no one has built, and the minds who defined that word were not in the room.

Basil C. Puglisi, MPA

A Human-AI Collaboration

basilpuglisi.com

Superintelligence has authors, and so does the name it now replaces. In 1955, John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon wrote the proposal that gave a new field its name, artificial intelligence (McCarthy et al., 1955/2006). Ten years later, the statistician I. J. Good, who had worked with Alan Turing, imagined an ultraintelligent machine that could outdo the brightest person at every intellectual task. He called it the last invention people would ever need, provided the machine was “docile enough to tell us how to keep it under control” (Good, 1965). Turing himself had warned in 1951 that “we should have to expect the machines to take control,” pointing to Samuel Butler’s novel *Erewhon* (Gonçalves, 2023).

In 1998, Nick Bostrom defined superintelligence, and his 2014 book refined that definition into the one most of the field still cites (Bostrom, 1998, 2014). In those later terms, a superintelligence is an intellect that greatly exceeds human cognitive performance in virtually all domains (Bostrom, 2014).

On September 29, 2026, President Trump signed Executive Order 14434, and it tramples both records. Its opening section credits American scientists with giving the field its name, and its next section orders every executive agency to replace that name with Super Intelligence (Exec. Order No. 14434, 2026). The White House fact sheet adds that the executive branch will “no longer acknowledge the terms ‘Artificial Intelligence’ or ‘AI’” (The White House, 2026a). The replacement is the word Bostrom defined for the kind of mind Good imagined, one beyond ours. President Trump attached it to every system the law already calls AI and left Good’s condition behind. These scholars own the idea in the only way ideas are owned, through the published record that defines them. The order sends the President’s science adviser to consult agency heads on a federal definition, and it names no researcher at all.

The Room on September 29

I went looking for those minds in the room. When I set out to learn AI, I learned it from the minds that bend the machine, the researchers, ethicists, economists, and philosophers whose arguments built the field. My forthcoming book, *The Minds That Bend the Machine*, profiles

twenty-four of them. Six executives signed a voluntary pledge under the new name that day: Sundar Pichai, Mark Zuckerberg, Jensen Huang, Elon Musk, Greg Brockman, and Dario Amodei (CNN, 2026). Jeff Bezos and Palantir’s Alex Karp attended as well. My book names that table in the coda, the chapter I added last. There I define the AI industrialist as “a figure whose influence over AI runs through ownership rather than argument” (Puglisi, 2026f).

Amodei is the one person at that table with a chapter among my minds, and he signed as something else. His company raised money at a \$965 billion valuation in May and filed confidentially in June to go public (CNBC, 2026a; NPR, 2026). He remains one of my minds, and on September 29 he sat at that table in his industrial role. The pledge carries the name White House Accord on Super Intelligence (CNBC, 2026b). It asks each company to build internal controls, bring in outside auditors, set up board oversight committees, and meet on shared safety standards (Fortune, 2026). An analyst quoted by Al Jazeera described the arrangement in one line: “The companies drafted the principles, they hire the auditor, and the commitment is voluntary” (Al Jazeera, 2026).

The minds who define superintelligence were missing, Bostrom among them. Geoffrey Hinton, Yoshua Bengio, and Stuart Russell had already signed the Future of Life Institute statement calling for a prohibition on developing superintelligence (Future of Life Institute, 2026; Gold, 2025; House of Lords Library, 2026). Their prohibition holds until scientists broadly agree it can be built safely and controllably and the public supports building it. One week before the order, Timnit Gebru and Emily Bender warned that describing products as superintelligence “ascribes agency to products rather than to the companies building them” (Gebru & Bender, 2026). The people who study the word were on record, and the people who sell the technology signed under the name.

What Superintelligence Means

Bostrom’s definition sets a high bar, since a superintelligence must exceed human performance across virtually every domain that matters. His own example excludes a chess program that outplays people, because its skill stops at chess (Bostrom, 2014). The technology industry usually writes the term as one word, and the order splits it into two. After the order, AI professor David Krueger described superintelligence as AI so far beyond human ability that controlling it becomes very hard, perhaps impossible (The San Francisco Standard, 2026). Good’s 1965 machine carried that warning from the start, in the condition he attached about control. Each of these definitions measures the machine against people across the full range of thought, and none of them treats a single strong skill as enough.

The Order Skips a Rung

My own position has held for a year, and it starts with a ladder. In September 2025, I wrote, “Today’s systems remain narrow AI, designed for specific tasks like drafting text or processing data. AGI imagines cross-domain adaptation. ASI imagines surpassing human capability” (Puglisi, 2025a). In the same piece I wrote that artificial general intelligence “is not something that exists publicly today, nor can it be demonstrated in any credible research.” I went further then and said superintelligence would stay out of reach. My position now concerns timing: superintelligence comes after AGI, and AGI is not here.

The record holds across the year. Eight days after that piece, I wrote on Medium that when people call it AGI, I cringe, because what follows is usually about money and not reality (Puglisi, 2025b). When Mo Gawdat predicted AGI in 2026, I answered that “no current consensus holds that any deployed system meets that threshold” (Puglisi, 2026a). On September 26 of this year, one year after the Medium piece, I wrote that AGI has not arrived and that superintelligence would come well after it (Puglisi, 2026e).


The strongest counterclaim came this September. OpenAI president Greg Brockman said GPT-6 Astra had opened an AGI era, and Nvidia’s Jensen Huang posted that AGI had arrived (Security Boulevard, 2026; Williams, 2026). Toby Walsh, chief scientist of the UNSW AI Institute, answered, “I’d be amazed if it really has matched all human cognitive capabilities” (Williams, 2026). Rebecca Johnson of the University of Sydney said she doesn’t believe Astra has reached AGI, because no stable definition exists against which to judge it (Williams, 2026).

A sharper dissent comes from an author of this vocabulary. Mark Gubrud says he coined the term artificial general intelligence in 1997 (Gubrud, 2025), and in March he wrote that “we have achieved AGI” (Gubrud, 2026). He pointed to models that command language and general knowledge at roughly a high-human level, and he granted that major deficiencies remain.

Even granting the AGI claim, superintelligence sits on the rung above it, and the order takes that top rung’s name without ever naming the middle one. Anyone who reads the law can see the gap. The order defines Super Intelligence as whatever the National Artificial Intelligence Initiative Act of 2020 already calls artificial intelligence (Exec. Order No. 14434, 2026). That statute covers any machine-based system that makes predictions, recommendations, or decisions “for a given set of human-defined objectives” (National Artificial Intelligence Initiative Act of 2020, 15 U.S.C. § 9401(3)). A spam filter fits that description, and so does a playlist recommender. The chess program Bostrom excluded from superintelligence now qualifies for the federal label. An order may define a word for its own purposes, and this term of art now collides with the meaning the research record already holds.

The order names the top rung and defines the bottom one

Executive Order 14434 and the ladder from narrow AI to superintelligence



Sources: Puglisi (2025a, 2026e); Good (1965); Bostrom (1998, 2014); Exec. Order No. 14434 (2026); 15 U.S.C. § 9401(3)

basilpuglisi.com

Figure 1. Executive Order 14434 takes its name from the top rung of the ladder and its legal meaning from the bottom rung, and it never names the rung between them.

If It Is Far Off, the Name Is a Bad Joke

If my reading of the timeline holds, the rename changes the label and leaves the machine exactly as it was. The order gives the new term the same legal meaning as the old one. No system became more capable or better governed on September 30 than it was on September 28. The order also limits its own reach, applying the new name to non-statutory documents and requiring no change to earlier regulations, contracts, grants, or other historical documents (Exec. Order No. 14434, 2026). The statutes still say artificial intelligence, so federal agencies will now describe the same systems in two vocabularies, one for the law and one for everything else.

Axios reported the new label as part of a wider effort in Washington and Silicon Valley to change how the public sees AI (Axios, 2026). That effort follows a public backlash against the technology and its data centers. The President said the word artificial made the technology sound like a hoax (Axios, 2026). A week earlier, at the United Nations, he had compared AI safety warnings to climate warnings he has also called a hoax (Agence France-Presse, 2026). I examined that pattern in *The AI HOAX Distraction* (Puglisi, 2026d).

Simon Coghlan lectures in digital ethics at the University of Melbourne. He said the new name “borrows the prestige of a term researchers reserve for a much bigger leap” (Forbes India, 2026). That is the joke, and it carries a cost. Federal agencies must now use it across websites, reports, policy documents, and correspondence. The public, meanwhile, hears its government describe a spam filter in words researchers reserve for the most capable system they can imagine. On my timeline, the President has promoted ordinary software to the top of a ladder whose middle rung no one has reached.

If the Fears Are Right, the Name Is Ignorance at Best

I don’t believe superintelligence will end us, and I said so on September 26 (Puglisi, 2026e). One of the idea’s authors came to fear exactly that. Good’s 1965 paper opened by tying human survival to building the machine, and in 1998 he wrote that he now suspected “survival” should be replaced by “extinction” (Dvorsky, 2013). The order still has to answer to the people who hold those fears, and on their terms it fails worse. Taken at its word, the order declares that superintelligence has arrived, and the White House met that arrival with a pledge its signers enforce on themselves. Good’s condition was control. At the United Nations, the President said the United States “totally rejects any attempt to construct a globalist scheme of control” for AI (The White House, 2026b). In *Governing AI*, I named this the Economic Override Pattern: corporate incentives systematically prioritize capability advancement over safety validation across all risk domains (Puglisi, 2025c). The pattern predicts that voluntary governance commitments fail under competitive pressure, and a pledge whose signers hire their own auditors sits squarely inside it.

The name also spends the alarm before anyone needs it. If a system that truly exceeds human ability ever arrives, the federal government will already have used the word on chatbots and spam filters. The warning that should signal something new will sound like the brochure.

The definition fight makes the problem worse. By November 28, 60 days after the order, the President’s science adviser must submit proposed legislative language defining Super Intelligence. The proposal must also assess whether the new definition should modify, expand upon, or supersede the statutory definition of artificial intelligence (Exec. Order No. 14434, 2026). Until then, the order’s own definition is interim, and it governs only until a later Presidential action or an Act of Congress supersedes it (Exec. Order No. 14434, 2026). A proposal that reserves the word for systems beyond human ability would answer part of this argument. If Congress instead writes SI into law to mean today’s AI, every future law aimed at the real thing inherits the wrong word.

Legislation already before Congress uses the word the research way. Senator Bernie Sanders and Representative Greg Casar have introduced the Ban Artificial Superintelligence Act. It defines

its target as AI that “exceeds human cognitive performance and capabilities across most domains” (Casar, 2026). If the executive branch calls today’s AI Super Intelligence while Congress debates banning artificial superintelligence, the federal government will use one word for two very different things.

Geburu and Bender saw the deeper problem a week early. The superintelligence label moves agency from the companies to their products, and it helps those companies escape accountability for what they build (Geburu & Bender, 2026). A government that calls the product super while its makers sign a pledge to police themselves has written the companies’ marketing for them. It has also left unasked the question of who answers for the product. At best, that’s ignorance of what the word means, and if the fears are to be believed, it is dangerous.

What the Order Actually Describes

Read past the title, and the order describes something else. Its own justification says today’s systems extend human ingenuity and creativity and open work that people could not do before (Exec. Order No. 14434, 2026). That is augmentation, and its lineage is older than Good’s ultraintelligent machine. J. C. R. Licklider proposed man-computer symbiosis in 1960, and Douglas Engelbart laid out a framework for augmenting human intellect in 1962 (Engelbart, 1962; Licklider, 1960). The statute the order borrows points the same way, since every system it covers works toward objectives that people set.

I call this Augmented Intelligence, the other AI, in which the human governs and the machine extends what the human can do (Puglisi, 2026b). The order’s description and its chosen name point in opposite directions. One puts a person in charge of a stronger tool, and the other places the machine above the person.

The Question the Name Does Not Answer

Whatever the government calls the technology, the governance question stays the same: who decides, by what authority, at what checkpoint? My answer, in Checkpoint-Based Governance, is a named human with binding authority at a defined point in the work, accountable through moral, professional, civil, and criminal channels (Puglisi, 2026c). At that checkpoint, the person can accept, modify, or reject what the machine produced, and the decision carries that person’s name. A rename doesn’t supply that person, and neither does a pledge whose signers police themselves.

The science adviser has until November 28 to propose a federal definition. The published record already gives superintelligence an established technical meaning and a documented lineage, and

the researchers who carry that record forward are available to ask. If the definition that reaches Congress keeps their meaning, the word stays ready for the day it is needed.

References

- Agence France-Presse. (2026). Trump renames AI “super intelligence” as leaders jostle over dangers. *Yahoo News*. <https://www.yahoo.com/news/politics/articles/trump-rejects-international-regulation-ai-160957080.html>
- Al Jazeera. (2026, September 29). Trump, tech bosses sign voluntary pact pledging “robust” AI safeguards. <https://www.aljazeera.com/news/2026/9/29/trump-top-tech-firms-sign-accord-to-self-police-ai-development>
- Axios. (2026, October 1). Trump and CEOs try to shed AI’s toxic branding. <https://www.axios.com/2026/10/01/superintelligence-si-ai-trump>
- Bostrom, N. (1998). How long before superintelligence? *International Journal of Futures Studies*, 2. <https://nickbostrom.com/superintelligence>
- Bostrom, N. (2014). *Superintelligence: Paths, dangers, strategies*. Oxford University Press.
- Casar, G. (2026, September 23). *NEWS: Casar, Sanders introduce legislation to create new federal agency to ban artificial superintelligence, pause advanced AI development* [Press release]. <https://casar.house.gov/media/press-releases/news-casar-sanders-introduce-legislation-create-new-federal-agency-ban>
- CNBC. (2026a, June 1). Anthropic confidentially files IPO prospectus with SEC, prepping Wall Street for landmark AI deal. <https://www.cnbc.com/2026/06/01/anthropic-ipo-s1-prospectus.html>
- CNBC. (2026b, September 29). Trump tries to rename AI “super intelligence” as polls show him sinking on key issue. <https://www.cnbc.com/2026/09/29/trump-ai-super-intelligence.html>
- CNN. (2026, September 29). Top AI executives sign commitment to “self-police” after meeting at White House. <https://www.cnn.com/2026/09/29/business/amodei-huang-karp-trump>
- Dvorsky, G. (2013, October 2). Why a superintelligent machine may be the last thing we ever invent. *Gizmodo*. <https://gizmodo.com/why-a-superintelligent-machine-may-be-the-last-thing-we-1440091472>
- Engelbart, D. C. (1962). *Augmenting human intellect: A conceptual framework* (Summary Report AFOSR-3223). Stanford Research Institute.
- Exec. Order No. 14434. (2026, September 29). Inaugurating the era of super intelligence. The White House. <https://www.whitehouse.gov/presidential-actions/2026/09/inaugurating-the-era-of-super-intelligence/>
- Forbes India. (2026). Trump’s proposal to change AI to SI: Tech leaders call it “misleading.” <https://www.forbesindia.com/article/news/trumps-proposal-to-change-ai-to-si-tech-leaders-call-it-misleading/2998676/1>
- Fortune. (2026, October 1). AI’s biggest players promise to police themselves at the White House. <https://fortune.com/2026/10/01/trump-ai-regulation-luncheon-huang-amodei-pichai-brockman-zuckerberg-musk-joint-committment-frontier-responsibilities>

- Future of Life Institute. (2026, March 27). *Prominent scientists, faith leaders, policymakers and artists call for a prohibition on superintelligence, as poll shows Americans don't want it* [Press release]. <https://futureoflife.org/press-release/prominent-scientists-faith-leaders-policymakers-and-artists-call-for-a-prohibition-on-superintelligence/>
- Gebru, T., & Bender, E. M. (2026, September 22). Don't be fooled by this summer of AI hype. *MIT Technology Review*. <https://www.technologyreview.com/2026/09/22/1144867/dont-be-fooled-summer-ai-hype/>
- Gold, A. (2025, October 22). AI leaders push to pause superintelligence. *Axios*. <https://www.axios.com/2025/10/22/superintelligence-ai-pause-yoshua-bengio>
- Gonçalves, B. (2023). Irony with a point: Alan Turing and his intelligent machine utopia. *Philosophy & Technology*, 36(3), Article 50. <https://doi.org/10.1007/s13347-023-00650-7>
- Good, I. J. (1965). Speculations concerning the first ultraintelligent machine. In F. L. Alt & M. Rubinoff (Eds.), *Advances in computers* (Vol. 6, pp. 31 to 88). Academic Press. [https://doi.org/10.1016/S0065-2458\(08\)60418-0](https://doi.org/10.1016/S0065-2458(08)60418-0)
- Gubrud, M. [@mgubrud]. (2025, October 16). @hendrycks Yo. I invented the term “artificial general intelligence” in 1997 and gave the first definition, the one that has [Post]. X. <https://x.com/mgubrud/status/1978838579594825999>
- Gubrud, M. [@mgubrud]. (2026, March 23). @MaMoMVPY Well, Lars, I INVENTED THE TERM and I say we have achieved AGI. Current models perform at roughly high-human [Post]. X. <https://x.com/mgubrud/status/2036262415634153624>
- House of Lords Library. (2026). *Superintelligent AI: Should its development be stopped?* <https://lordslibrary.parliament.uk/superintelligent-ai-should-its-development-be-stopped/>
- Licklider, J. C. R. (1960). Man-computer symbiosis. *IRE Transactions on Human Factors in Electronics, HFE-1*, 4 to 11.
- McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (2006). A proposal for the Dartmouth summer research project on artificial intelligence, August 31, 1955. *AI Magazine*, 27(4), 12. <https://doi.org/10.1609/aimag.v27i4.1904> (Original work written 1955)
- National Artificial Intelligence Initiative Act of 2020, 15 U.S.C. § 9401 (2021). <https://www.law.cornell.edu/uscode/text/15/9401>
- NPR. (2026, June 1). AI giant Anthropic prepares to sell stock to the public, files preliminary IPO paperwork. <https://www.npr.org/2026/06/01/nx-s1-5843199/anthropic-ipo-filing-ai-large>
- Puglisi, B. C. (2025a, September 18). The human advantage in AI: Facts, not fantasies. basilpuglisi.com. <https://basilpuglisi.com/the-human-advantage-in-ai-facts-not-fantasies/>
- Puglisi, B. C. (2025b, September 26). When they call it AGI or Generative AI, I cringe and then realize the content to follow is about money and not reality. Medium. <https://medium.com/the-other-ai/when-they-call-it-agi-or-generative-ai-i-kringe-and-then-realize-the-content-to-follow-is-about-f8b1c4ed9383>
- Puglisi, B. C. (2025c). *Governing AI: When capability exceeds control*. Digital Ethos.

- Puglisi, B. C. (2026a, April). The inevitable is a choice: Testing Mo Gawdat's FACE RIPS forecast across two interviews against the governance architecture that could make it optional. basilpuglisi.com. <https://basilpuglisi.com/mo-gawdat-inevitable-choice/>
- Puglisi, B. C. (2026b, May). The Other AI: Augmented intelligence and the honest future of human-AI collaboration. basilpuglisi.com. <https://basilpuglisi.com/the-other-ai-augmented-intelligence-governance/>
- Puglisi, B. C. (2026c, September 11). What is Checkpoint-Based Governance? Human authority. basilpuglisi.com. <https://basilpuglisi.com/cbg/>
- Puglisi, B. C. (2026d, September 17). The AI HOAX distraction: Safety, money, and why you can't tell which is driving. basilpuglisi.com. <https://basilpuglisi.com/ai-hoax/>
- Puglisi, B. C. (2026e, September 26). A year ago I told you AGI was hype. Superintelligence is further off than sold, and it will not end us. basilpuglisi.com. <https://basilpuglisi.com/agi-hype-superintelligence/>
- Puglisi, B. C. (2026f). *The minds that bend the machine* [Unpublished manuscript].
- Security Boulevard. (2026, September). AGI has not arrived. Here is what 2026 actually proved. <https://securityboulevard.com/2026/09/agi-has-not-arrived-here-is-what-2026-actually-proved/>
- The San Francisco Standard. (2026, October 1). Trump wants to rebrand AI "Super intelligence." Will Silicon Valley play along? <https://sfstandard.com/2026/10/01/trump-rebrand-ai-superintelligence/>
- The White House. (2026a, September 29). *Fact sheet: President Donald J. Trump inaugurates the era of super intelligence*. <https://www.whitehouse.gov/fact-sheets/2026/09/fact-sheet-president-donald-j-trump-inaugurates-the-era-of-super-intelligence/>
- The White House. (2026b, September 22). *President Trump at the United Nations: "While others have talked, I have acted"* [Release]. <https://www.whitehouse.gov/releases/2026/09/president-trump-at-the-united-nations-while-others-have-talked-i-have-acted/>
- Williams, T. (2026, September 7). OpenAI says "the AGI era" is here. Experts disagree. *Information Age*. <https://ia.acs.org.au/article/2026/openai-says-the-agi-era-is-here-experts-disagree.html>

Disclaimer

I am not a lawyer, and this article does not provide legal advice. This is thought research and governance analysis based on public sources, cited materials, and human-AI review. It is intended to help executives, practitioners, insurers, and governance teams think more clearly about AI risk, liability exposure, and documentation practices. Readers should not rely on this article as a legal opinion, compliance determination, or substitute for qualified counsel. Any organization facing a legal, regulatory, contractual, or insurance question should consult its own attorney, broker, or professional adviser before acting.

The Other AI: Audio Briefings on Augmented Intelligence and AI Governance

[Spotify](#) | [Apple Podcasts](#) | [Amazon Music](#) | [YouTube Playlist](#)

#AIassisted using the HAIA Ecosystem | CC BY-NC-SA 4.0

Free for personal, educational, and noncommercial research use with attribution. Commercial exploitation, paid productization, and enterprise commercialization require separate permission and licensing.