


REPORT Published July 23, 2025 • 8 minute read

An AI Agenda for a New Era: Advance, Protect, & Implement



Mike Sexton, Senior Policy Advisor for Artificial Intelligence and Digital Technology

Lawmakers face a political dilemma when it comes to artificial intelligence (AI): align with AI pessimists or AI optimists? Once the party of innovation and technology, Democrats have morphed over the years to one of skepticism and hostility when it comes to technology. Republicans, on the other hand, too often have a fanatical allegiance to deregulation without true appreciation of impacts. But those trends don't have to portend the future. As AI evolves into artificial general intelligence, with the potential to exceed the intellect of humans in new domains, there is an opportunity for the powerful and continuously improving technology to address Americans' greatest challenges—if we implement it well.

In this memo, we provide an overview of two opposing camps when it comes to AI regulation. Then, we chart a new course and call on lawmakers to embrace a bold and opportunistic AI agenda built on

three pillars: advance, protect, and implement (API). The greatest risk is not failing to regulate AI but failing to enable AI to help Americans and society thrive. We need a new approach—it’s time for API.

Two Opposing Visions on AI

The Doomers

Four months after ChatGPT’s public release, a group of prominent AI experts signed a public letter calling for a six-month pause in AI development.¹ AI pause supporters on X (formerly Twitter) expressed their views by updating their names to include the pause emoji (⏸).² Some AI pessimists who think even a six-month pause would not suffice made their views known with a stop emoji (🛑).³

The fringe of AI pessimism is colloquially known as **AI doomerism**.⁴ The most famous AI doomer is Eliezer Yudkowsky, founder of the Machine Intelligence Research Institute (MIRI), who says that “the most likely result of building a superhumanly smart AI... is that literally everyone on Earth will die.”⁵ Yudkowsky supports shutting down data centers training AI—no exceptions for governments or militaries—and even suggests airstrikes on rogue AI training sites.

AI pessimists and doomers have not achieved their stated goals, but they have influenced AI discourse and development considerably. Yudkowsky’s logically rigorous arguments about AI’s existential risks have resonated with **effective altruists**, the philanthropic philosophical movement that aims to maximize the benefits of charitable donations.⁶

Many effective altruists were signatories to the six-month pause petition.⁷ Others have funded or founded AI risk-focused think tanks, including the Center for Security and Emerging Technology, the Center for the Governance of AI, and Yudkowsky’s MIRI.⁸ While not all effective altruists believe in AI doom, most take it seriously enough to support research into it.

Many AI developers have absorbed Yudkowsky’s arguments. The challenge of developing AI with built-in guardrails that prevent it from harming humans has become known as the **alignment problem**. OpenAI’s mission is a nod to the alignment problem: “to ensure that artificial general intelligence benefits all of humanity.”⁹ Likewise, its competitor, Anthropic, says it builds “AI to serve humanity’s long-term well-being.”¹⁰ Paradoxically, by articulating AI’s existential risks so well, OpenAI CEO Sam Altman believes Yudkowsky has “done more to accelerate AGI than anyone else.”¹¹

The Accelerationists

While AI pessimists and effective altruists have converged politically on the priority of addressing existential AI risk, their partnership has spawned its own evil twin: **effective accelerationism**, or **e/acc**

for short. Like their altruist twins, effective accelerationists aim to maximally improve human life—but rather than view AI as a threat to that mission, they see it as its ultimate engine.

The most prominent effective accelerationist is venture capitalist (VC) Marc Andreessen, who believes AI “is quite possibly the best thing human beings have ever created.”¹² In “The Techno-Optimist Manifesto,” Andreessen argues that “any deceleration of AI will cost lives,” and that delaying the emergence of potentially life-saving AI through regulation “is a form of murder.”¹³ Marc Andreessen has a friend in a high place: before his Senate run, JD Vance co-founded the VC firm known as Narya Capital from 2019-2022, backed by Andreessen’s investment.¹⁴

While fringe, e/acc’s political potency is that it is both prescriptive and descriptive: it sees accelerating and compounding technological advancement as both a goal and a naturally occurring phenomenon. On the second point, it is hard to dispute. There is a broad recognition in Washington that if the United States decided on principle to slow AI development, China will readily fill the void—as DeepSeek showed.

A New Way Forward: Advance, Protect, Implement

Deliberately regulating AI to slow its development is a fool’s errand. Technology will advance regardless, and its trajectory will be worse under China’s leadership. At the same time, shunning all regulation under the theory that technology will automatically help society is a dangerous and risky assumption. Instead, lawmakers need a multifaceted and targeted agenda to address pessimists’ anxieties, harness accelerating technological growth to benefit Americans, and appeal to the entrepreneurs and innovators who will make it possible.

The answer is **API: advance, protect, and implement**. Policymakers should **advance** the state-of-the-art in AI, **protect** Americans from emerging AI risks, and **implement** AI in ways that improve Americans’ lives.

Advance

The current AI ecosystem in America has enabled Silicon Valley to lead the world in forging the future of AI development. That includes advancement in open-source AI, where AI models’ code and parameters are released publicly, furthering its proliferation. This is a vastly different approach than what has taken place in the European Union. The EU leads the world in AI regulation, with the 108-page EU AI Act adopted in May 2024.¹⁵ This has not strengthened but stifled its AI ecosystem. Most recently, Meta opted not to release its open-source multimodal LLaMa AI model in the EU.¹⁶

Lawmakers should continue cultivating an ecosystem in the United States that enables state-of-the-art AI to develop and thrive. To do that, we must:

- Increase engagement with allies and trading partners to support American interests on the front lines of the battle with China, protect American interests from international overreach, and develop an aggressive export strategy to counter China.
- Loosen permitting requirements to build clean, firm energy that can power AI data centers today and in the future.
- Free up land for data centers by allocating a portion of federal lands to the advance AI.
- Allocate more funding to federal research and development to spur and incentivize private sector research activities.
- Eliminate Trump-era tariffs on innovative products and components to reduce unnecessary costs on emerging technology.
- Encourage foreign AI talent to move to and work in the United States, by overhauling our high-skill immigration and visa system.
- Oppose legal limits on open-source software and AI as well as licensing regimes that would entrench an incumbent advantage over the future innovators.
- Establish clear rules of the road that only regulate what's necessary so as to not suffocate the innovation we're trying to grow.

Protect

Rather than constraining AI development as the primary strategy for managing risk, the United States should empower institutions and citizens to navigate and mitigate the challenges posed by AI.

Lawmakers should operate under the assumption that AI development at the international level is fundamentally anarchic and instead prioritize giving government—including law enforcement—the authority and resources to counteract harm from AI. Crucially, lawmakers must also recognize that adversary powers and terrorist organizations may be capable of causing catastrophic harm with AI in the future, and rogue AI itself poses security challenges.¹⁷

Lawmakers can pursue multiple lines of effort to minimize AI risk in the short- and long-term—including through bipartisan cooperation—to get it done:

- Work across party lines to criminalize deepfake porn and AI-generated child sexual abuse material.
- Support a “right to contest AI” when AI bias has the potential to undermine civil rights, such as financial lending.
- Cultivate strategic international partnerships to expand the United States’ influence over AI development—and limit China’s.
- Reverse the Trump administration’s job cuts at the Cybersecurity and Infrastructure Security Agency (CISA).
- Increase resources and coordination among CISA, the FBI, and the National Institute of Standards and Technology to counter rogue AI threats to critical infrastructure.
- Pursue public-private partnerships with the above agencies to protect AI models from foreign espionage and validate model alignment with human interests.
- Protect and modernize intellectual property and copyright law to support innovators.
- Standardize AI regulations at the federal level to support developers’ compliance.

Implement

Finally, lawmakers must establish plans to implement AI and other emerging technologies in ways that benefit the American people. From delivering veterans’ health care to educating the next generation, there are countless activities that can and should be optimized with artificial intelligence. Policymakers should develop AI and robotics implementation agendas to ensure these technologies are able to help all Americans. Some possibilities include:

- Revamp public education for the AI age by deploying and studying AI-based tutors and guidance counselors in public schools as well as assisting teachers with lesson planning and continuing education.
- Upskill Americans for new roles overseeing or working alongside AI and robots—including in government.
- Expand opportunity to more people and places so that the jobs and economic benefits of AI are not concentrated in highly educated coastal cities and a handful of tech hubs in the interior of the country.

- Encourage AI's adoption in medicine to reduce barriers to health care access and develop novel medical therapies like gene editing.
- Procure robots for construction, parks management, public sanitation and other underserved government functions.
- Accelerate autonomous vehicle and drone adoption with a focus on access and safety.
- Procure new military weaponry that capitalizes on AI capabilities and counteracts emerging technological threats from state and non-state adversaries.

Conclusion

Lawmakers stand at a defining crossroad. Government regulation is important. But political hostility toward innovators' success does nothing to help Americans benefit from AI and may, in fact, hinder progress. Now is the critical moment when Washington should steer AI development in ways that support Americans, not simply seek to punish AI companies for being profitable. Transformative innovation cannot be achieved through government alone, but government is a critical partner to ensure that innovation makes America smarter, safer, and more successful in the 21st century.

ENDNOTES

1. Future of Life Institute. "Pause Giant AI Experiments: An Open Letter," 22 March 2023. <https://futureoflife.org/open-letter/pause-giant-ai-experiments/>. Accessed 5 May 2025.
2. X (formerly Twitter). "AI Notkilleveryoneism Memes %u23F8%uFE0F (@AISafetyMemes) / X." <https://x.com/aisafetymemes>. Accessed 5 May 2025.
3. X (formerly Twitter). "Eliezer Yudkowsky %u23F9%uFE0F (@ESYudkowsky) / X." <https://x.com/esyudkowsky>. Accessed 5 May 2025.
4. Marantz, Andrew. "Among the A.I. Doomsayers." *The New Yorker*, 11 March 2024. <https://www.newyorker.com/magazine/2024/03/18/among-the-ai-doomsayers>. Accessed 5 May 2025.
5. Yudkowsky, Eliezer. "The Open Letter on AI Doesn't Go Far Enough." *TIME*, 29 March 2023. <https://time.com/6266923/ai-eliezer-yudkowsky-open-letter-not-enough/>. Accessed 5 May 2025.
6. Marantz.
7. Metz, Cade, and Gregory Schmidt. "Elon Musk and Others Call for Pause on A.I., Citing 'Profound Risks to Society.'" *The New York Times*, 29 March 2023, sec. Technology. <https://www.nytimes.com/2023/03/29/technology/ai-artificial-intelligence-musk-risks.html>. Accessed 29 May 2025.
8. Effective Altruism Forum. "Organizations and Projects in Effective Altruism - EA Forum." <https://forum.effectivealtruism.org/topics/organizations-and-projects-in-effective-altruism>. Accessed 29 May 2025.
9. OpenAI. "About." <https://openai.com/about/>. Accessed 5 May 2025.
10. Anthropic. "Home." <https://www.anthropic.com/>. Accessed 5 May 2025.

11. Sam Altman [@sama]. "It Is Possible at Some Point He Will Deserve the Nobel Peace Prize for This--I Continue to Think Short Timelines and Slow Takeoff Is Likely the Safest Quadrant of the Short/Long Timelines and Slow/Fast Takeoff Matrix." Tweet. *Twitter*, 23 February 2023.
<https://x.com/sama/status/1621621725791404032>. Accessed 5 May 2025.
12. Andreessen, Marc. "AI Will Save the World." *The Free Press*, 27 March 2025.
<https://www.thefp.com/p/why-ai-will-save-the-world>. Accessed 6 May 2025.
13. Andreessen, Marc. "The Techno-Optimist Manifesto." Andreessen Horowitz, 16 October 2023. <https://a16z.com/the-techno-optimist-manifesto/>. Accessed 6 May 2025.
14. Primack, Dan. "J.D. Vance's Short Career in Venture Capital." *Axios*, 16 July 2024.
<https://www.axios.com/2024/07/16/jd-vance-venture-capital-career>. Accessed 6 May 2025.
15. EU Artificial Intelligence Act. "The Act Texts."
<https://artificialintelligenceact.eu/the-act/>. Accessed 7 May 2025.
16. Weatherbed, Jess. "Meta Won't Release Its Multimodal Llama AI Model in the EU." *The Verge*, 18 July 2024. <https://www.theverge.com/2024/7/18/24201041/meta-multimodal-llama-ai-model-launch-eu-regulations>. Accessed 7 May 2025.
17. Vermeer, Michael J. D. "Could AI Really Kill Off Humans?" *RAND*, 9 May 2025.
<https://www.rand.org/pubs/commentary/2025/05/could-ai-really-kill-off-humans.html>. Accessed 4 June 2025.