FOUNDATION FOR AMERICAN INNOVATION

September 17, 2025

Written Testimony of Samuel Hammond
Chief Economist, Foundation for the American Innovation (FAI)
Before the
U.S. House Oversight Subcommittee on Cybersecurity, Information Technology, and Government Innovation, September 17, 2025

"Shaping Tomorrow: The Future of Artificial Intelligence"

Chairwoman Mace, Ranking Member Brown and members of the Subcommittee, I thank you for the opportunity to testify today.

My name is Samuel Hammond, chief economist for the Foundation for American Innovation. FAI is a group of technologists and policy experts focused on developing technology, talent and ideas to support a freer and more abundant future.[1]

The capabilities of frontier AI systems are improving at a stunning rate. Five years ago, Large Language Models (LLMs) could barely generate coherent English text. Today, they can hold forth on any topic, reason through PhD-level math problems, and code entire applications from scratch.

Recent AI progress, including the rise of "reasoning models" and AI agents, has been largely driven by breakthroughs in reinforcement learning applied to LLMs. LLMs gain their raw intelligence by predicting sequences of text, but with reinforcement learning can be trained to follow instructions, use tools, and pursue complex goals. Scalable reinforcement learning for reasoning and goal-directed behavior in LLMs was only unveiled a year ago,[2] but is already driving rapid improvements in domains like math and programming. The scope of this breakthrough is still not fully appreciated. In principle, these techniques can be used to create superhuman AI agents in any domain where success can be objectively benchmarked. Math and software engineering are just the low-hanging fruit.

[1] See: www.thefai.org
[2] OpenAI. "Learning to Reason with LLMs." September 12, 2024. https://openai.com/index/learning-to-reason-with-llms/

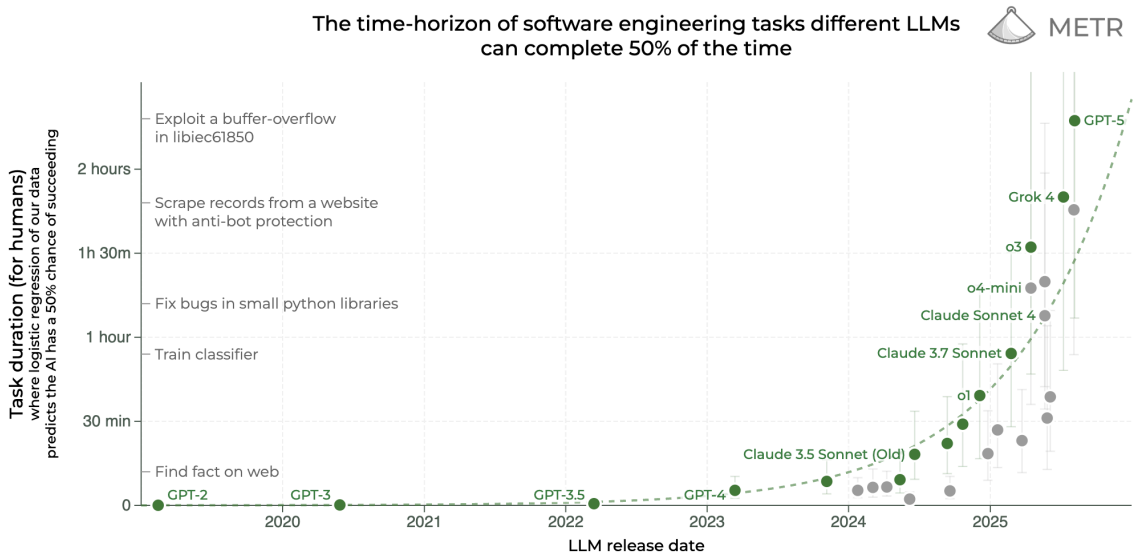

The AI research organization, METR, carefully measures progress in AI autonomy, and has found that the length of tasks AI agents can perform doubles every 4 to 7 months – a trend that has held for the last six years.[3] While the earliest chatbots could only perform tasks measured in seconds or minutes, OpenAI's latest model, GPT-5, can coherently execute tasks that take human engineers 2 hours and 17 minutes on average. If this trend continues, we are only two doublings (8 to 14 months) away from AI agents that can autonomously perform tasks that take humans a full eight hour work day.

Progress in non-verifiable and open-ended domains is also accelerating, recently leading to some of the first major examples of AIs making novel scientific discoveries.[4] It is now plausible that we will have the first superintelligent AI scientists and mathematicians by the year's end, portending a dramatic speed-up in the pace of R&D going forward. This includes AIs optimized for AI research itself, creating the glimmers of a self-improving feedback loop whereby AIs rapidly help build their own successors.

What happens when AIs get better at AI R&D than the best human researchers in the world? At a minimum, we should expect a discontinuous leap in the power and efficiency of frontier systems, but where this process tops-out is a matter of significant uncertainty.[5] It is possible that, even with fully automated AI R&D, progress will remain bottlenecked by

[3] METR. "Measuring AI Ability to Complete Long Tasks." March 19, 2025. https://metr.org/blog/2025-03-19-measuring-ai-ability-to-complete-long-tasks/
[4] See, for example: "AlphaEvolve: A Gemini-powered coding agent for designing advanced algorithms." Google DeepMind, May 14, 2025. https://deepmind.google/discover/blog/alphaevolve-a-gemini-powered-coding-agent-for-designing-advanced-algorithms/
[5] Jason Hausenloy. "Red Lines for Recursive Self-Improvement." The First Scattering, Jul 14, 2025. https://firstscattering.com/p/red-lines-for-recursive-self-improvement

the availability of compute, data, and energy. It is also possible that we are only one major breakthrough away from systems that can learn in a continuous and unbounded fashion. Regardless, the jump in capabilities unlocked by recursively self-improving AI is likely to be profound even within the bounds of existing infrastructure, and is coming sooner than many realize.

It is worth emphasizing that creating superintelligent AIs capable of outperforming humans in every domain is the explicit goal of every leading U.S. AI company, including OpenAI, Anthropic, GoogleDeepMind, Meta and Xai. While some dismiss this as science fiction or marketing hype, I assure you the leaders of these companies are deadly serious. As for timing, Anthropic cofounder Jack Clark testified earlier this year that he expects transformative AI to arrive as soon as the end of 2026.[6]

Even if these forecasts are on trend, AI capabilities will remain uneven for at least several more years. For a brief moment, we will have superintelligent AIs that can prove new math theorems but still struggle at many things that humans find trivial. This is especially true in areas like robotics where, despite remarkable progress, we are still years away from machines outperforming humans in every physical domain given the paucity of high quality training data.

As we run headlong into this new world, I see four major takeaways for national policymakers:

First, monitoring frontier AI capabilities in real time is a national security imperative of the U.S. government. Early and differential access to developments at the frontier can provide policymakers and national security advisors with foresight into the sorts of capabilities that will eventually proliferate as costs come down, giving us time to prepare and adapt.[7]

Second, as AI systems become human-level and beyond, geopolitical power will be increasingly proxied by the global distribution of computing resources.[8] America's existing

[6] Jack Clark. "On "Algorithms and Authoritarians: Why U.S. AI Must Lead." Testimony before the United States House of Representatives Select Committee on the Chinese Communist Party (June 25, 2025). https://docs.house.gov/meetings/ZS/ZS00/20250625/118428/HHRG-119-ZS00-Wstate-ClarkJ-20250625.pdf

[7] Shaun Eu. "Asymmetry by Design: Boosting Cyber Defenders with Differential Access to AI." Institute for AI Policy and Strategy, May 23, 2025. https://www.iaps.ai/research/differential-access

[8] Samuel Hammond. "The Scramble for AI Computing Power." American Affairs, Summer 2024 / Volume VIII, Number 2. https://americanaffairsjournal.org/2024/05/the-scramble-for-ai-computing-power/

lead is downstream of our advantages in AI hardware and data centers, but this is tenuous at best.[9] With China out-building us on new energy generation, we must double-down on chip and semiconductor export controls or be leapfrogged.[10]

This should include either reversing or minimizing the damage associated with the decision to license Nvidia's H20 chip to Chinese customers. The H20 is a powerful chip for AI inference – precisely what China needs to compete on reasoning models and agents. If all 1.3 million H20s that Nvidia previously reported on their orderbook go through, China's inference compute capacity will roughly double. Recent reports that Nvidia may also be approved to sell China their new B30A chip are even more troubling, as the B30A is around 12x more performant than the H20.[11] At a minimum, Nvidia should be required to first offer these chips to U.S. buyers, as proposed in the GAIN AI Act from Senator Jim Banks. Going forward, action must also be taken to crack down on China's chip smuggling efforts, such as by requiring exported chips to adopt basic location verification mechanisms.[12]

Third, we need to quickly advance the frontier of AI control and interpretability research; review our laws and regulations for their compatibility with powerful AI; and invest in much more robust cyber- and infrastructure security – all priorities of President Trump's AI Action Plan.[13]

Fourth and finally, we must open our minds to radically new ways of structuring our institutions. From the printing press to the Industrial Revolution, every major technological transition has driven equally transformative changes to our system of government.[14] The AI revolution will be no different, but raises unique challenges given AI's uses for surveillance and censorship, as seen in China's model of digital authoritarianism. Reconciling the advent of powerful AI systems with America's traditions of individual liberty and limited government is thus the challenge of our time.

Thank you, and I look forward to your questions.

---

[9] Konstantin Pilz, Robi Rahman, James Sanders and Lennart Heim, "GPU Clusters." EpochAI, July 18, 2024. https://epoch.ai/data/gpu-clusters
[10] Samuel Hammond. "U.S. Companies Are Helping China Win the AI Race." CityJournal, May 17, 2025. https://www.city-journal.org/article/artificial-intelligence-china-deepseek-nvidia-broadcom-openai
[11] Liam Mo and Fanny Potkin. "Nvidia working on new AI chip for China that outperforms the H20, sources say." Reuters, August 19, 2025. https://www.reuters.com/world/china/nvidia-working-new-ai-chip-china-that-outperforms-h20-sources-say-2025-08-19/
[12] Samuel Hammond. "The Chip Security Act: A Bipartisan Solution to Chip Smuggling." Foundation for American Innovation, July 18, 2025. https://www.thefai.org/posts/the-chip-security-act-a-bipartisan-solution-to-chip-smuggling
[13] Samuel Hammond. "Reacting to the AI Action Plan." Foundation for American Innovation, July 24, 2025. https://www.thefai.org/posts/reacting-to-the-ai-action-plan
[14] Samuel Hammond. "AI and Leviathan: Part I." SecondBest, August 23, 2023. https://www.secondbest.ca/p/ai-and-leviathan-part-i