

**Written Testimony of Jack Clark,  
Co-Founder and Head of Policy, Anthropic**

**For a hearing on “Federal Science Agencies and the Promise of AI in  
Driving Scientific Discoveries”**

**Before the Committee on Science, Space, and Technology  
United States House of Representatives  
February 6<sup>th</sup>, 2024**

Chair Lucas, Ranking Member Lofgren, and members of the committee, thank you for the opportunity to speak with you today about how the American research community can retain its leadership in the field of artificial intelligence (AI).

I am Jack Clark, Anthropic’s Co-Founder and Head of Policy. We founded Anthropic in 2021 as an AI safety and research company to build reliable, interpretable, and steerable AI systems. At Anthropic, we believe the impact of AI might be comparable to that of the industrial and scientific revolutions. We expect this level of impact could start to arrive soon—perhaps in the coming decade—and the benefits of AI will be profound. In the next few years, AI could greatly accelerate treatments for diseases, lower the cost of energy, revolutionize education, improve efficiency throughout government, and much more. Even as we see enormous potential benefits from the technology, we are concerned about a spectrum of risks it may also bring, from immediate and near-term concerns, to potential more catastrophic future risks. While we do not believe that the systems available today pose catastrophic risks, we also believe that we need to lay the groundwork now to ensure future, more powerful systems are safe.

For decades, federal science agencies have played a central role in America’s ability to lead the world in emerging technologies. As generative AI technologies continue to rapidly advance, it is critical that the United States proactively stakes a leadership position in steering responsible AI development, including by democratizing access to generative AI models. Federal initiatives and funding that promote equitable access to the critical inputs to generative models (data and compute), while researching techniques to strengthen oversight and mitigate harm, will help establish America as a global leader in advancing AI safely and for the common good.

At Anthropic, we are proud to have endorsed the CREATE AI Act and to be one of the 25 non-governmental partners supporting the recently launched National AI Research Resource (NAIRR) pilot at the National Science Foundation. We believe these are important steps towards setting the U.S. on a course to set the pace in building and developing safe rules of the road for today’s transformative technology.

In my testimony today, I will share insights about the costs of developing frontier AI systems today, how those costs are likely to grow in the coming years, and why these costs make it critical that the United States invests ambitiously in a National AI Research Resource. I will also share how Anthropic supports public research and close with some ideas for how a NAIRR can be successful.

## Costs of Frontier AI Systems

To understand why we need the NAIRR, it's helpful to understand the costs of developing frontier AI systems. Before 2020, developing a large-scale AI system would cost between \$10,000 to \$100,000—a significant amount of money, but one that could be gained through normal funding mechanisms, and which could easily be done at the level of individual universities. However, since then, the costs of training frontier AI systems has grown significantly. This is because a type of technical analysis called “Scaling Laws” has made it easier and more predictable to develop AI systems. Scaling Laws research has made it easier for people to work out how much performance they can expect in relation to resources (mostly computational power) they spend in training the system. As a consequence of Scaling Laws, the increased performance capabilities and return on investment in expanding the frontier is clearer, and the costs of developing frontier systems has increased significantly.

Based on data from our own work at Anthropic, as well as analysis based on reading research papers from across the AI industry, we believe the costs of staying on the frontier of AI development have grown significantly. Specifically, we believe that it cost between \$1-\$5 million to train a frontier AI system in 2020; between \$10-20 million in 2021 and 2022; and in 2022 and 2023 the costs grew to between \$50 million and \$100 million. At Anthropic, we have trained AI systems that fall within all these ranges during these time periods specified—and we don't believe we're alone.

To get a sense of just how strategic compute is seen by those developing AI systems, we'd note that in January a major technology company announced it would have 350,000 H100 GPUs from NVIDIA by the end of 2024<sup>1</sup> — at an estimated retail cost of around ~\$30,000 per GPU, that represents an investment of \$10.5 billion dollars.

Assuming that “Scaling Laws” and other evidence we observe today holds, we think it's not inconceivable that systems which cost *hundreds of millions* are trained in 2024 and 2025, and it's not beyond the realm of possibility that systems which cost \$1 billion are trained towards the end of 2025 and beyond.

These numbers illustrate the magnitude of the challenge before us and the importance of the federal government's support towards ensuring that generative AI research and development is

---

<sup>1</sup> Paul, Nellis & Cherney, Meta to deploy in-house custom chips this year to power AI drive - memo, Reuters (Feb. 1, 2024), available at: <https://www.reuters.com/technology/meta-deploy-in-house-custom-chips-this-year-power-ai-drive-memo-2024-02-01/> (accessed Feb. 2, 2024).

a shared enterprise at a societal level. A key way the federal government can support greater research and development in AI across academia and civil society is to fund its own experimental infrastructure to develop AI that is on par with (or near to) industry capabilities. Doing so will help diversify the actors that can help shape the future of AI.

Many experts agree that governments should be investing in greater computational capabilities to unlock more AI research possibilities. In Canada, Turing award winner Yoshua Bengio has advocated that the Canadian government spend \$1 billion (\$750 million USD) on developing its own large-scale AI system<sup>2</sup>, and in late 2023 the UK government committed £300 million (\$380 million) to support its own national AI research resource<sup>3</sup>.

Given how transformative and powerful those of us developing AI technologies expect it to be, we think it's important for government research and investment to have a role in shaping the future of AI.

### **Public and Private Research in the Age of Large-Scale Frontier AI Systems**

A logical question this invites is “well, if this is so expensive, why don’t we find a way to harness the innovations within industry for the public benefit?” I would view this as a necessary but not sufficient response to the challenges in front of us, and I’ll explain why.

At Anthropic, we today provide subsidized or free access to our systems for third-party researchers and we mostly do this by prioritizing researchers interested in improving the safety and reliability of the kinds of systems we develop. As mentioned, one of our goals is to advance the safety and reliability of AI systems, so by providing access we’re able to do this. However, providing such access comes with its own challenges, which I’d put into three categories: money, training, and selection.

- **Money:** Using our systems costs money—after all, we have a business where we charge people money to tap into our systems, and we pay a cost for each query run through our systems. Therefore, providing researcher access costs money - that’s something we choose to do as a public benefit organization, but not all companies may have the same incentives. At the end of the day, researcher access is a line item that has to be accounted for.
- **Training:** To get the most out of AI systems, you need to know how to interface with them. At Anthropic, we find that making our third-party researchers successful typically requires us to spend some time educating them on our systems - how to use them to get the most performance out of them, and how to debug any problems they run into. This costs another kind of resource - staff time at Anthropic. This, given the fact we’re a small startup with a lot of work to do and not many people to do it, is also costly.

---

<sup>2</sup> <https://www.cbc.ca/news/canada/montreal/bengio-asks-canada-to-build-ai-supercomputer-1.7094858>

<sup>3</sup> <https://www.ukri.org/news/300-million-to-launch-first-phase-of-new-ai-research-resource/>

- **Selection:** Finally, as I mentioned, we choose to primarily prioritize the safety of our systems because it's something we're familiar with. And, with the just-announced NSF NAIRR pilot, we've also provided research access for work based on using AI to better understand issues relating to a changing global climate. But are these things the "best" things researchers could be doing with our systems? We don't know. There isn't an easy way to evaluate all the research that could be done on our systems and make educated judgment calls.

## **The Importance of the NAIRR**

Anthropic believes that the NAIRR will encourage public-private partnerships amongst industry, academia, and the federal government that will support the development of AI systems for the public interest. To that end, we hope the NAIRR will:

### **1. Enable Academic Research Using Industry Models**

The NAIRR is an opportunity for the United States to shape how the private sector can work most effectively with academic researchers. The NAIRR can do valuable work in bringing scientists together, identifying areas of enquiry that would be improved via access to cutting-edge systems, then using that information to have an informed conversation about model access with the private sector. This would also enable the NAIRR to develop training resources that could help academic researchers use frontier models and it would help the scientific community prioritize projects.

### **2. Broaden the Federal Government's Capacity to Develop its Own frontier AI Systems**

It is my earnest belief that the best way to make progress on grand technological problems is to ensure that a broad and diverse set of people are able to work on them. Therefore, I believe the NAIRR should seek to create the infrastructure necessary to help U.S. academia work at the same scale as industry. Some of the reasons for doing this include:

- Flooding resources towards developing beneficial applications of AI that may not be prioritized by the private sector.
- Creating large amounts of computational resources available for studying the safety properties of existing large-scale systems.
- Building the necessary infrastructure so that the NAIRR can benefit from open source and openly accessible AI systems, giving it the ability to host and run experiments on openly disseminated models. This will have a variety of benefits, ranging from providing training for how to use these systems, speeding up the science of developing measurement and evaluation systems for frontier models, and creating a legitimate third-party for working on difficult questions relating to assessing the safety and potential misuses of such models.

### **3. Train the Next Generation of American Scientists**

It's an unfortunate fact today that the clearest path to learn how to train and evaluate extremely large-scale AI systems is in the private sector. This means that academia and the government often lose out on talent, and in the case of universities, there are fewer people to train the next generation. The NAIRR is one solution to this emerging challenge. Instead of having to work solely in the private sector to gain access to cutting-edge systems and develop certain skills, the NAIRR would promote a robust job market and talent pipeline by giving the next generation of scientists a meaningful choice between industry, academia, and government. The NAIRR talent pipeline could take the form of offering scientists longer stays in academia working directly on the NAIRR, or federal government opportunities in roles connected to the NAIRR. Anthropic believes that this would improve the talent available to the United States.

### **Conclusion**

The NAIRR will ensure that American scientists keep our nation at the forefront of frontier compute technology — creating important research and workforce opportunities and new avenues for economic growth. It is essential that the U.S. government and academic researchers have access to the tools needed to conduct essential research in areas that should not be left solely to private companies. Anthropic strongly supports targeted investments in the National Science Foundation, the Department of Energy, the National Institute of Standards and Technology, and other agencies that will help ensure the U.S. retains its lead in the fast-evolving AI sector.

Thank you for the opportunity to be here today and I look forward to answering your questions.