MEMO *Published May 15, 2026 · 9 minute read*

What Are Frontier AI Models?

Mike Sexton

Takeaways

- Frontier AI models are the most advanced and capable AI models at a given time.
- Frontier models' emergent abilities are powerful and unpredictable, creating unprecedented opportunities and risks.
- While one metric—training compute—is a workable proxy to define the "frontier," definitions need to be dynamic and capability-focused to keep pace with future development.
- Getting AI regulation right means managing risks from the development and application of both frontier and non-frontier AI, preserving a competitive development edge against China, and putting increasingly powerful AI to use in ways that make Americans' lives better.

Artificial intelligence is rapidly becoming one of the most powerful forces shaping the global economy, national security, scientific advancement, and the future of work. At the center of this transformation are "frontier models," the most advanced AI models ever created, capable of reasoning through complex problems, writing software, and powering AI agents that use digital tools autonomously.

As lawmakers grapple with how to govern this technology, understanding what qualifies as a frontier model—and why it matters for national security, economic competitiveness, public safety, innovation, and scientific discovery—is an essential first step. This memo explains what frontier models are, how policymakers currently define them, and why developing safety systems to guide them will be central to responsible AI governance in the years ahead.

What Are Frontier Models?

"Frontier" is the label given to contemporary AI models that represent the state-of-the-art. Because AI developers constantly train new, more advanced models, the AIs considered frontier are constantly changing. Seven models from leading AI labs are considered frontier at the time of publication.

Since the state-of-the-art in AI is always moving forward, the capabilities that distinguish frontier models from prior generations are always changing. For example, Anthropic and OpenAI's latest models are so effective at cyber offense and defense that the labs are limiting access to select companies and organizations to expedite security patching. [1] Each new generation resets the questions we have to ask about how to advance, protect, and implement these systems.

- ChatGPT-5.5 [2] (from OpenAI)
- Claude Opus 4.7 [3] (from Anthropic)
- Gemini 3.1 Pro [4] (from Google)
- Muse Spark [5] (from Meta)
- Grok 4.3 [6] (from xAI)
- Mistral Large 3 [7] (from Mistral)
- DeepSeek V4 [8] (from DeepSeek)

The Frontier AI Ecosystem

Frontier AI models are only one piece of a larger and rapidly evolving technological, economic, and regulatory ecosystem. Their significance lies in their versatility: frontier models can be adapted to a wide range of high-impact uses, from software development and cybersecurity to scientific research, enterprise operations, education, and public services. One cutting edge AI use today connects models—sometimes frontier, sometimes not—to agentic platforms enabling them to control and automate computer functions. OpenAI Codex, Claude Code, and OpenClaw are examples of such platforms, often referred to as "harnesses." [9]

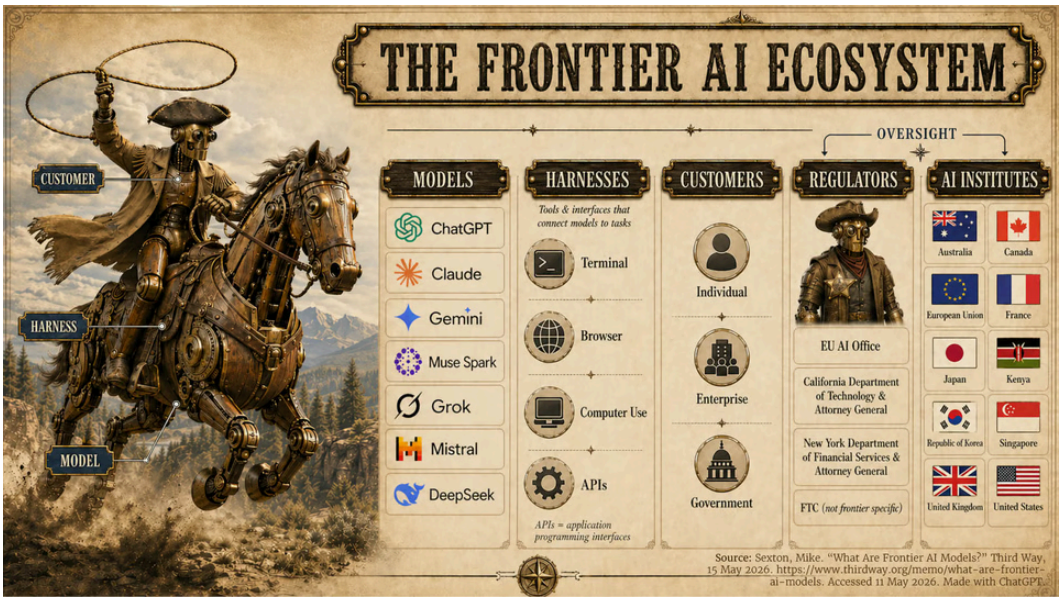

While the federal government does not currently regulate the development of frontier AI models, the European Union, California, and New York do. The EU AI Act requires frontier AI developers to publish copyright policies, training data summaries, and other technical documentation, and it places cybersecurity and reporting requirements on models deemed a "systemic risk." [10] California's SB 53 and New York's RAISE Act protect whistleblowers, mandate incident reporting, and require transparency reports that include model risk assessments. [11]

In addition to regulatory oversight, many governments have established AI institutes or tasked other agencies to monitor frontier AI development. In the United States, the Center for AI Standards and Innovation (CAISI) at the National Institute of Standards and Technology (NIST) tests and evaluates different models to inform policymakers about industry development, and it works with some AI developers on a voluntary basis for deeper analysis and safety checks. [12]

Frontier AI labs heavily interface with both the private sector and governments—the question is how to harmonize regulatory principles and optimize rules for risk management and scientific and economic progress.

Why Regulating Frontier Models Matters

Frontier AI development is increasingly framed—including by the major labs themselves—as a transition to superintelligence: AI systems that will exceed human capability. [13] How fast that transition unfolds, and what role policy plays in shaping it, are no longer abstract questions.

One reason is that AI systems are increasingly contributing to their own advancement. As of writing, two agentic AI systems—Claude Cowork and OpenAI Codex—were built with nearly 100% AI-written code. [14] Recursive self-improvement, in which AI systems help build the next generation of AI systems, is moving from theoretical possibility to observable trend. This feedback loop in the development cycle shortens iteration timelines and compresses the window in which policymakers can shape outcomes.

The stakes of this transition are difficult to overstate. The disruptions AI is already producing across the economy and national security will only intensify as systems grow more capable. That is why the choices made now—about how these systems are developed, deployed, and overseen—carry so much weight. Getting AI policy right is not a hedge against a hypothetical future; it is a precondition for navigating the one already arriving.

Nowhere is this clearer than with autonomous weapons, as the Department of War's dispute with Anthropic shows. A military operating superintelligent autonomous weapons without adequate safeguards could pose exactly the catastrophic dangers AI researchers fear most. And yet, without these capabilities, democracies risk falling behind better-equipped and less restrained adversaries—which is why AI labs remain willing to work with the US military despite the tension. [15]

Americans need reliable executive leadership and thoughtful federal AI regulation. That means strict, enforceable, and standardized guardrails for safety—and a pathway for responsible deployment in critical national security and economic functions. If oversight becomes a punitive dragnet, government will learn slower, industry will innovate elsewhere, and the United States will hand momentum to China. **The task for policymakers is to advance American leadership in frontier AI, protect society from its most dangerous risks, and implement the technology in ways that strengthen democratic power**.

From State Patchwork to Federal Framework

The most pressing question for frontier AI policy is not which definition of a 'frontier model' is technically correct—it is how to design a policy that captures the most concerning model capabilities without hampering frontier development.

The most common way laws define “frontier models” is by **how much compute**—i.e., how many computing operations—is used in the training process. At a basic level, modern digital computers perform operations as mathematical calculations in binary (zeros and ones). If we imagine these operations performed on an abacus, “compute” is roughly analogous to how many times the beads move to perform a given operation.

The EU AI Act sets the compute threshold for frontier models subject to regulatory scrutiny at 10^{25} floating point operations (or FLOP). [16] The European Commission and EU AI Office also have discretion to designate lower-threshold models as systemically risky based on their capabilities on a case-by-case basis. New York and California's AI safety laws—the RAISE Act and SB 53—set the threshold 10 times higher, at 10^{26} FLOP, without the EU's explicit carveout to designate models based on qualitative assessments. [17]

However, as the frontier pushes forward, the case for capability-based flexibility will likely grow stronger. According to the nonprofit Epoch AI, at least 30 AI models crossed the EU's 10^25 FLOP threshold by June 2025 while only a handful are believed to have surpassed New York and California's threshold of 10^26 FLOP. [18] By 2030, there are projected to be 239 models above 10^26 FLOP and 357 above 10^25 FLOP. [19]

For this reason, any statutory definition of the frontier should be dynamic and, ideally, capability-based, allowing policymakers and technology experts to modify the threshold as more training-intensive and capable models emerge. The EU AI Act incorporates this principle, empowering the EU AI Office to adjust training thresholds and define risky models based on capabilities. [20]

In the United States, only the federal government can guarantee compliance with one unified, up-to-date standard across state and national borders—which is how frontier models are deployed in reality. Capability-defining authority is therefore best vested in the federal government, where it can also draw on guidance from CAISI and the International Network of AI Safety Institutes. [21] The challenge is that multiple states have adopted frameworks before the federal government, complicating development of a universal standard.

Until Congress acts, harmonizing state-level requirements is an immediate priority. That will keep safety risks from falling through gaps in a regulatory patchwork and prevent an uneven legal landscape for frontier model development. But the ultimate goal for policymakers should be federal codification: a bipartisan federal AI safety framework that builds on that state convergence, allows for capability-based designations in coordination with CAISI, and locks in a single national standard.

Done well, this approach also creates room to address the broader AI regulatory landscape, including the substantial risks posed by non-frontier systems.

Why Regulating Frontier Models Isn't Enough

Regulating AI development, especially of the highest-risk, highest-capability models, is an urgent policy priority. However, frontier-level regulations are not sufficient to mitigate the breadth of AI risks.

AI development has branched: OpenAI Codex, Claude Cowork, and the open-source OpenClaw are not frontier AI models, but agentic AI systems. These systems won't be covered by model-level regulations and will pose different risks and challenges than the models themselves. As more agentic systems are released, policymakers will also have to grapple with questions of accountability, oversight, and transparency. As part of that, policymakers should make clear that humans ultimately bear responsibility for their use of AI systems, agentic or not.

This principle also applies to misuse of non-frontier AI, such as AI-enabled discrimination, deepfakes, non-consensual intimate imagery, scams, and more. Legislation like the Take It Down Act, which federally criminalizes the distribution of non-consensual intimate imagery, is demonstrative: the law addresses real risks from novel technology by holding perpetrators accountable without imposing technically infeasible or impractical restrictions on the underlying tool used by the perpetrator.

Conclusion

A workable federal approach must do three things simultaneously: keep the frontier governable, keep the United States competitive, and help more people in more places benefit from the technology. This requires meaningful engagement with stakeholders across the AI ecosystem—developers, deployers, sector regulators, AI institutes, and the communities that AI impacts—whose combined knowledge is what turns policy into practice. Building superintelligence and useful AI applications safely is only a third of the battle. Regulating AI and policing misuse are also vital priorities. And ultimately, a whole-of-society movement is needed to put AI to use in service of human flourishing.

ENDNOTES

1. Ostrovsky, Nikita. "'Too Dangerous to Release' Is Becoming AI's New Normal." Business. *Time*, 24 April 2026. https://time.com/article/2026/04/24/claude-mythos-chatgpt-rosalind-release-dangerous/. Accessed 27 April 2026.

2. Capoot, Ashley. "OpenAI Announces GPT-5.5, Its Latest Artificial Intelligence Model." CNBC, 23 April 2026. https://www.cnbc.com/2026/04/23/openai-announces-latest-artificial-intelligence-model.html. Accessed 24 April 2026.

3. Anthropic. "Introducing Claude Opus 4.7." 16 April 2026. https://www.anthropic.com/news/claude-opus-4-7. Accessed 24 April 2026.

4. Google. "Gemini 3.1 Pro: A Smarter Model for Your Most Complex Tasks." 19 February 2026. https://blog.google/innovation-and-ai/models-and-research/gemini-models/gemini-3-1-pro/. Accessed 24 April 2026.

5. Vanian, Jonathan. "Meta Debuts New AI Model, Attempting to Catch Google, OpenAI after Spending Billions." CNBC, April 8, 2026. https://www.cnbc.com/2026/04/08/meta-debuts-first-major-ai-model-since-14-billion-deal-to-bring-in-alexandr-wang.html.

6. Grok. "Grok Release Notes." 17 April 2026. https://grok.com/release-notes. Accessed 24 April 2026.

7. Mistral AI. "Introducing Mistral 3." https://mistral.ai/news/mistral-3. Accessed 10 February 2026.

8. DeepSeek. "DeepSeek V4 Preview Release | DeepSeek API Docs." https://api-docs.deepseek.com/news/news260424. Accessed 27 April 2026.

9. Ethan Mollick, "A Guide to Which AI to Use in the Agentic Era," *One Useful Thing*, 17 February 2026. https://www.oneusefulthing.org/p/a-guide-to-which-ai-to-use-in-the. Accessed 27 April 2026.

10. "High-Level Summary of the AI Act | EU Artificial Intelligence Act." *The EU Artificial Intelligence Act*, 27 February 2024. https://artificialintelligenceact.eu/high-level-summary/. Accessed 1 May 2026.

11. Gluck, Justine. "The RAISE Act vs. SB 53: A Tale of Two Frontier AI Laws." Blog. *Https://Fpf.Org/*, 8 January 2026. https://fpf.org/blog/the-raise-act-vs-sb-53-a-tale-of-two-frontier-ai-laws/. Accessed 1 May 2026.

12. "International Network for Advanced AI Measurement, Evaluation, and Science Publishes Consensus Areas on Practices for Automated Evaluations." *NIST*, 13 February 2026. https://www.nist.gov/news-events/news/2026/02/international-network-advanced-ai-measurement-evaluation-and-science. Accessed 27 April 2026.

13. OpenAI. "Industrial Policy for the Intelligence Age." 6 April 2026. https://openai.com/index/industrial-policy-for-the-intelligence-age/. Accessed 27 April 2026.

14. Morrone, Megan. "Anthropic's Viral New Work Tool Wrote Itself." Axios, 13 January 2026. https://www.axios.com/2026/01/13/anthropic-claude-code-cowork-vibe-coding. Accessed 13 February 2026.

15. Amodei, Dario. "Statement from Dario Amodei on Our Discussions with the Department of War." Anthropic, 26 February 2026. https://www.anthropic.com/news/statement-department-of-war. Accessed 4 March 2026.

16. High-Level Summary of the AI Act | EU Artificial Intelligence Act. 27 February 2024. https://artificialintelligenceact.eu/high-level-summary/. Accessed 10 February 2026.

17. Kearns, Kyle. “SB 53: What California’s New AI Safety Law Means for Developers.” *Wharton AI & Analytics Initiative*, 14 November 2025. https://ai-analytics.wharton.upenn.edu/wharton-accountable-ai-lab/sb-53-what-californias-new-ai-safety-law-means-for-developers/. Accessed 10 February 2026.

Loring, M. “New York’s RAISE Act: What Frontier Model Developers Need to Know.” Jones Walker LLP - New York’s RAISE Act: What Frontier Model Developers Need to Know, 2 January 2026. https://www.joneswalker.com/en/insights/blogs/ai-law-blog/new-yorks-raise-act-what-frontier-model-developers-need-to-know.html. Accessed 10 February 2026.

18. Rahman, Robi. “Over 30 AI Models Have Been Trained at the Scale of GPT-4.” Epoch AI, 30 January 2025. https://epoch.ai/data-insights/models-over-1e25-flop. Accessed 25 February 2026.

Cottier, Ben. “How Many AI Models Will Exceed Compute Thresholds?” Epoch AI, 30 May 2025. https://epoch.ai/blog/model-counts-compute-thresholds. Accessed 10 February 2026.

19. Ibid.

20. European Commission. “General-Purpose AI Models in the AI Act – Questions & Answers | Shaping Europe’s Digital Future.” 9 September 2025. https://digital-strategy.ec.europa.eu/en/faqs/general-purpose-ai-models-ai-act-questions-answers. Accessed 12 February 2026.

21. U.S. Department of Commerce. “FACT SHEET: U.S. Department of Commerce & U.S. Department of State Launch the International Network of AI Safety Institutes at Inaugural Convening in San Francisco.” 20 November 2024. https://www.commerce.gov/news/fact-sheets/2024/11/fact-sheet-us-department-commerce-us-department-state-launch-international. Accessed 28 April 2026.