

UNTANGLING AI OPENNESS

PARTH NOBEL,* ALAN Z. ROZENSHTEIN** & CHINMAYI SHARMA***

The debate over AI openness—whether to make components of an artificial intelligence system available for public inspection and modification—forces policymakers to balance innovation, democratized access, safety, and national security. By inviting startups and researchers into the fold, it enables independent oversight and inclusive collaboration. But technology giants can also use it to entrench their own power, while adversaries can use it to shortcut years and billions of dollars in building systems, like China’s DeepSeek, that rival our own. How we govern AI openness today will shape the future of AI and America’s role in it.

Policymakers and scholars grasp the stakes of AI openness, but the debate is trapped in a flawed premise: that AI is either “open” or “closed.” This dangerous oversimplification—inherited from the world of open source software—belies the complex calculus at the heart of AI openness. Unlike traditional software, AI is a composite technology built on a stack of discrete components—from compute to labor—controlled by different stakeholders with competing interests. Each component’s openness is neither a binary choice nor inherently desirable. Effective governance demands a nuanced understanding of how the relative openness of each component serves some goals while undermining others. Only then can we determine the trade-offs we are willing to make and how we hope to achieve them.

This Article aims to equip policymakers with the analytical toolkit to do just that. First, it introduces a novel taxonomy of “differential openness,” untangling AI into its constituent components and illustrating how each one has its own spectrum of openness. Second, it uses this taxonomy to systematically analyze how each component’s relative openness necessitates intricate trade-offs both within and between policy goals. Third, it

* Ph.D., Electrical Engineering, Stanford University. Nobel was supported in part by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1656518. Any opinions, findings, conclusions, or recommendations expressed in this Article are those of the author and do not necessarily reflect the views of the National Science Foundation.

** Associate Professor of Law, University of Minnesota Law School; Senior Editor and Research Director, Lawfare; Visiting Senior Fellow, Institute for Law & AI; Nonresident Senior Fellow, Brookings Institution. Rozenshtein consults on a range of technology law and policy issues.

*** Associate Professor of Law, Fordham Law School; Contributing Editor, Lawfare; Advisor, American Law Institute Principles of the Law, Civil Liability for Artificial Intelligence; Member, Microsoft Responsible AI Committee. For helpful comments, we thank Ryan Calo, James Grimmelman, Woody Hartzog, Asaf Lubin, John Speed Meyers, Paul Ohm, Cullen O’Keefe, Sana Pandey, Neal Parikh, Ashwin Ramaswami, Emily Royall, Elizabeth Seger, Andrew Selbst, Keith Winstein, Bianca Wylie, and participants at both We Robot 2025 and the Eighth Junior Faculty Forum for Law and STEM. For excellent research assistance, we thank Emma Haberman, Ben Evelev, and Audrey Kim.

operationalizes these insights by advancing a research agenda that shows how law can be analyzed and refined to support more precise configurations of component openness. AI openness is neither all or nothing nor inherently good or evil—it is a tool that must be wielded with precision if it has any hope of serving the public interest.

Introduction	171
I. A Taxonomy of AI Openness	178
A. Beyond the Open Source Software Analogy	179
1. Beyond Source Code	179
2. Beyond Altruism.....	183
3. Beyond the Developer	186
B. Disaggregating AI	190
1. Compute	191
2. Data	196
3. Source Code.....	198
4. Model Weights.....	199
5. System Prompts.....	201
6. Operational Control and Records	202
7. The Application Layer.....	204
8. The Human Layer	206
II. The Value of AI Openness.....	210
A. Safety	211
1. Benefits	212
2. Costs	214
B. Innovation and Economic Growth.....	216
1. Benefits	217
2. Costs	218
C. Democratic Access and Control	220
1. Benefits	221
2. Costs	222
D. National Security and Global Leadership	222
1. Benefits	223
2. Costs	224
E. Navigating Trade-offs in AI Openness	227
1. Trade-offs Within Policy Goals	227
2. Trade-offs Between Policy Goals	228
3. Deeper Structural Trade-offs	229
4. The Compounding Complexity of Interdependent Components	230
III. Calibrating Differential AI Openness	232
A. Liability	233
B. Competition.....	236
C. Intellectual Property	239
D. Trade	242
E. Government Support	244

Conclusion	247
Appendix: Openness of Select Frontier Models	249

INTRODUCTION

The question of “AI openness”—who controls artificial intelligence, who benefits from it, and who bears responsibility for its failures—has rapidly evolved from an obscure debate among scholars and programmers into a flashpoint in global policy, corporate strategy, and international affairs.¹ On the one hand, “open spectrum AI” (osAI)—a term this Article coins in lieu of the more common “open source AI” to more accurately capture the complexity of systems that are, to some degree, free and publicly available for inspection, use, and modification²—has shown itself capable of being a force for profound good. Google DeepMind’s AlphaFold predicts protein structures with revolutionary accuracy, accelerating drug discovery.³ On complex cases from a Boston teaching hospital, Meta’s Llama 3.1 performs on par with leading proprietary AI in generating differential diagnoses for complex cases, promising to enhance clinical decision support while protecting patient data.⁴ Meanwhile, conservationists deploy Wildbook, a system using computer vision to identify individual animals from photographs, creating a massive, crowdsourced database to track endangered species and inform conservation policy.⁵

Yet, in the shadows of this progress, the same powerful osAI technologies are weaponized. Cybercriminals unleash WormGPT, an AI built on the open GPT-J model, to craft malware and highly convincing

1. See, e.g., EXEC. OFF. OF THE PRESIDENT OF THE U.S., WINNING THE RACE: AMERICA’S AI ACTION PLAN 4–5 (2025), <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf> [<https://perma.cc/UB3X-6UZV>]; Iain Martin, *The EU Is Betting \$56 Million on Open Source AI*, FORBES (Feb. 6, 2025, at 10:27 ET), <https://www.forbes.com/sites/ianmartin/2025/02/02/the-eu-is-betting-56-million-on-open-source-ai/>; Troy Wolverton, *AI’s Openness Is Being Sharply Debated by Technologists, Policymakers*, S.F. EXAM’R (July 21, 2025), https://www.sfexaminer.com/news/technology/open-source-ai-debate-sharp-among-technologists-politicians/article_ab781b42-28e7-11ef-836b-9b118373b94c.html [<https://perma.cc/VZC8-SD5X>].

2. See *infra* Part I.

3. See Josh Abramson et al., *Accurate Structure Prediction of Biomolecular Interactions with AlphaFold 3*, 630 NATURE 493, 493, 496–97 (2024).

4. See Thomas A. Buckley, Byron Crowe, Raja-Elie E. Abdunour, Adam Rodman & Arjun K. Manrai, *Comparison of Frontier Open-Source and Proprietary Large Language Models for Complex Diagnoses*, JAMA HEALTH F., Mar. 2025, at 1, 2, <https://doi.org/10.1001/jamahealthforum.2025.0040>.

5. See Tanya Y. Berger-Wolf et al., *Wildbook: Crowdsourcing, Computer Vision, and Data Science for Conservation*, 2017 PROCS. DATA FOR GOOD EXCH. 1, 1, <https://doi.org/10.48550/arXiv.1710.08880>.

phishing emails with effortless precision.⁶ Meanwhile, pro-China influence operations deploy networks of fake social media accounts using StyleGAN-generated profile pictures—synthetic faces of nonexistent people—to amplify propaganda, discredit critics, and distort international discourse on human rights and global events.⁷ And a recent investigation into a widely used open training dataset uncovered thousands of images of child sexual abuse material, tainting the very foundation of popular image-generation models.⁸ This is the paradox of osAI: a single technological wellspring feeding both lifesaving innovation and sophisticated digital malice.

Meanwhile, openness is becoming a key driver of the AI market, which the United Nations projects to reach nearly \$5 trillion in less than a decade.⁹ Open models often lag behind their closed counterparts by only three months,¹⁰ a narrow gap that is fueling high-profile clashes and strategic moves across the industry. For example, Elon Musk sued OpenAI, accusing it of breaching a promise to put the public before profits and demanding it return to its open source roots.¹¹ In a widely publicized jab, he offered to drop the suit if the company simply renamed itself “ClosedAI.”¹² Subsequently, Musk’s own company, xAI, released its powerful Grok model under an open license.¹³ Meta has also made a strategic bet on openness, championing its Llama models as “open-

6. See Chuck Easttom, *Malicious Use of Artificial Intelligence*, 2025 IEEE 15TH ANN. COMPUTING & COMM’N WORKSHOP & CONF. (CCWC) 499, 500, <https://doi.org/10.1109/CCWC62904.2025.10903787>.

7. BENJAMIN STRICK, CTR. FOR INFO. RESILIENCE, ANALYSIS OF THE PRO-CHINA PROPAGANDA NETWORK TARGETING INTERNATIONAL NARRATIVES 4 (2024), https://www.info-res.org/app/uploads/2024/11/Analysis-of-the-Pro-China-Propaganda-Network-Targeting-International-Narratives_FINAL.pdf [<https://perma.cc/454N-NYQB>].

8. DAVID THIEL, STAN. INTERNET OBSERVATORY, IDENTIFYING AND ELIMINATING CSAM IN GENERATIVE ML TRAINING DATA AND MODELS 2, 13 (2023), <https://doi.org/10.25740/kh752sm9123>.

9. U.N. CONF. ON TRADE & DEV., TECHNOLOGY AND INNOVATION REPORT 2025: INCLUSIVE ARTIFICIAL INTELLIGENCE FOR DEVELOPMENT 6 (2025), https://unctad.org/system/files/official-document/tir2025_en.pdf.

10. Luke Emberson, *Open-Weight Models Lag State-of-the-Art by Around 3 Months on Average*, EPOCH AI (Oct. 30, 2025), <https://epoch.ai/data-insights/open-weights-vs-closed-weights-models> [<https://perma.cc/B7UV-KYF8>].

11. Complaint at 5–9, *Musk v. Altman*, No. CGC-24-612746 (Cal. Super. Ct. Feb. 29, 2024); Adam Satariano, Cade Metz & Tripp Mickle, *Elon Musk Sues OpenAI and Sam Altman for Violating the Company’s Principles*, N.Y. TIMES (Mar. 1, 2024), <https://www.nytimes.com/2024/03/01/technology/elon-musk-openai-sam-altman-lawsuit.html>.

12. Elon Musk (@elonmusk), X, *To ClosedAI and I will drop the lawsuit* (Mar. 6, 2024, at 10:10 CT), <https://x.com/elonmusk/status/1765409615070601417> [<https://perma.cc/K5HP-77X5>].

13. See *Open Release of Grok-1*, xAI (Mar. 17, 2024), <https://x.ai/news/grok-os>.

source” catalysts for innovation and security that can coexist with profitability.¹⁴ Even OpenAI’s CEO Sam Altman has acknowledged that the company’s closed approach may have placed it “on the wrong side of history”¹⁵—signaling a shift from one of the industry’s leading closed-model advocates—and it has since released a leading open spectrum AI model of its own.¹⁶ The debate has been further intensified by the rise of powerful osAI models from Chinese labs like DeepSeek, which now rival the performance of the proprietary Western systems upon which they were built¹⁷ and at significantly lower cost,¹⁸ adding a new layer of geopolitical urgency.¹⁹

The stakes of the debate over osAI could not be higher. In one view, openness is the key to unlocking unprecedented innovation, democratizing access to powerful technology, and enhancing safety by subjecting AI systems to broad, independent scrutiny. In the other view, osAI risks catastrophic misuse, threatens national security, and will inevitably be co-opted by dominant corporate interests, reinforcing the very power structures they claim to challenge.²⁰

But despite these high stakes, the discourse is dangerously oversimplified, flattening the concept of AI openness into an inaccurate “open-closed” binary.²¹ This misleading view has its roots in the history of open source software (OSS) but is ill-suited for governing open

14. See Mark Zuckerberg, *Open Source AI Is the Path Forward*, META (July 23, 2024), <https://about.fb.com/news/2024/07/open-source-ai-is-the-path-forward/> [https://perma.cc/ZA5J-VSJC]. Whether Meta’s commitment to AI openness continues remains to be seen. See Mark Zuckerberg, *Personal Superintelligence*, META (July 30, 2025), <https://www.meta.com/superintelligence/> [https://perma.cc/SX36-ZFNH] (“We’ll need to be rigorous about mitigating . . . risks and careful about what we choose to open source.”).

15. Kyle Wiggers, *Sam Altman: OpenAI Has Been on the ‘Wrong Side of History’ Concerning Open Source*, TECHCRUNCH (Jan. 31, 2025, at 15:34 PT), <https://techcrunch.com/2025/01/31/sam-altman-believes-openai-has-been-on-the-wrong-side-of-history-concerning-open-source/> [https://perma.cc/5DN7-UURV].

16. *Introducing gpt-oss*, OPENAI (Aug. 5, 2025), <https://openai.com/index/introducing-gpt-oss/>.

17. See Luis E. Romero, *ChatGPT, DeepSeek, or Llama? Meta’s LeCun Says Open-Source Is the Key*, FORBES (Jan. 28, 2025, at 18:07 ET), <https://www.forbes.com/sites/luisromero/2025/01/27/chatgpt-deepseek-or-llama-metas-lecun-says-open-source-is-the-key/>.

18. See Kevin Roose, *Why DeepSeek Could Change What Silicon Valley Believes About A.I.*, N.Y. TIMES (Jan. 28, 2025), <https://www.nytimes.com/2025/01/28/technology/china-deepseek-ai-silicon-valley.html>; Prithwiraj Choudhury, Natarajan Balasubramanian & Mingtao Xu, *Why DeepSeek Shouldn’t Have Been a Surprise*, HARV. BUS. REV. (Jan. 30, 2025), <https://hbr.org/2025/01/why-deepseek-shouldnt-have-been-a-surprise>.

19. See LAURIE HARRIS, CONG. RSCH. SERV., IF13051, DEEPSEEK AND THE RACE TO DEVELOP ARTIFICIAL INTELLIGENCE (2025).

20. See *infra* Part II.

21. See *infra* Section I.A.

spectrum AI. While openness in traditional software primarily meant access to source code, osAI systems are complex, layered technologies composed of multiple interdependent components: computational hardware that powers AI, training data that shapes capabilities, model weights that encode knowledge, source code that defines structure, operational records and controls that reveal performance characteristics, and the humans putting it all together.²² Each component exists on its own spectrum of openness and carries distinct implications for safety, innovation, democratic control, and national security.²³ Thus, in the context of AI ecosystems, openness refers to the degree to which these components are transparent in their operation, accessible to external scrutiny or use, and inclusive of diverse contributors throughout the development process.

By stripping the discourse around osAI of its necessary complexity, policymakers fail to address the nuanced trade-offs inherent in its governance and risk undermining the very goals they seek to achieve. With few exceptions,²⁴ policymakers tend to treat the openness of a model as a single, undifferentiated feature, without parsing the degree to which specific components are actually open. For example, the most thorough policy document on the topic, the National Telecommunication and Information Administration’s 2024 report, focused almost entirely on the risks of releasing open model weights.²⁵ The more recent Trump administration’s “AI Action Plan” does more or less the same, equating “open source” with “open weight.”²⁶ Abroad, the EU AI Act reflects a similarly oversimplified conception of AI openness.²⁷

Beyond distorting the debate, this oversimplified view dangerously casts osAI policy as an “all or nothing” decision, which makes outright prohibition a live option in a way it never was for OSS. In the OSS era, governments flirted with prohibiting narrow categories of code, such as

22. See *infra* Section I.B.

23. See *infra* Part II.

24. See, e.g., BIPARTISAN HOUSE TASK FORCE ON A.I., 118TH CONG., REPORT ON ARTIFICIAL INTELLIGENCE 155 (2024) (“Despite often being characterized as either open or closed, there is in fact a continuum of different forms of AI model availability and transparency. . . . [D]ifferent parts of a model can be made open while others remain closed.”).

25. NAT’L TELECOMMS. & INFO. ADMIN., DUAL-USE FOUNDATION MODELS WITH WIDELY AVAILABLE MODEL WEIGHTS 2–3 (2024), <https://www.ntia.gov/sites/default/files/publications/ntia-ai-open-model-report.pdf> [<https://perma.cc/5ZDY-3TP5>].

26. EXEC. OFF. OF THE PRESIDENT OF THE U.S., *supra* note 1, at 4–5.

27. See 2024 O.J. (L 1689) pmb. ¶ 102 (referencing “free and open-source licences . . . including the weights”).

strong encryption²⁸ or DVD-decryption tools,²⁹ and agencies briefly considered excluding GPL-licensed software from sensitive systems,³⁰ but OSS as such was never seriously on the brink of being outlawed. By contrast, today's osAI debate already includes proposals that would, in practice, foreclose open release of some frontier models altogether, such as licensing and certification schemes that osAI systems could not satisfy.³¹ Even where those efforts have been softened or rejected, the fact that policymakers are openly contemplating whether certain forms of osAI should be illegal rather than merely regulated underscores the urgency of introducing a more nuanced and precise framework that gives policymakers options beyond prohibiting osAI altogether.

AI openness cannot be treated as a simple binary; it must be assessed at the component level, with more scrutiny devoted to how the relative openness of each component impacts safety, innovation, democratic control, and national security. While academic work has begun to

28. See Juliette Garside, *Philip Zimmermann: King of Encryption Reveals His Fears for Privacy*, *GUARDIAN* (May 25, 2015, at 12:02 ET), <https://www.theguardian.com/technology/2015/may/25/philip-zimmermann-king-encryption-reveals-fears-privacy> [<https://perma.cc/R7ZK-A5YA>].

29. See *Unintended Consequences: Fifteen Years Under the DMCA*, ELEC. FRONTIER FOUND. (Mar. 2013), <https://www.eff.org/pages/unintended-consequences-fifteen-years-under-dmca> [<https://perma.cc/EUL2-FSFF>]; see also *Universal City Studios, Inc. v. Reimerdes*, 111 F. Supp. 2d 294, 342–46 (S.D.N.Y. 2000) (barring a developer from publishing an OSS project that defeated DVD-encryption as a Digital Millennium Copyright Act violation), *aff'd sub nom., Universal City Studios, Inc. v. Corley*, 273 F.3d 429 (2d Cir. 2001).

30. To resolve an internal debate, the Department of Defense commissioned a report from the MITRE Corporation to evaluate the costs and benefits of banning OSS use in the department. See generally MITRE CORP., *USE OF FREE AND OPEN-SOURCE SOFTWARE (FOSS) IN THE U.S. DEPARTMENT OF DEFENSE* (2003), https://dodcio.defense.gov/portals/0/documents/foss/dodfoss_pdf.pdf (advising that DOD not ban the use of open source software in its work in response to a request for outside recommendations regarding the proposal to ban open source software in DOD work).

31. California's S.B. 1047, vetoed in 2024, would have imposed safety certifications around monitoring and post-release obligations that open-weight frontier models cannot satisfy, because once weights are publicly released, developers cannot enforce downstream safeguards, update or patch proliferating copies, or prevent harmful uses. See Florence G'sell, Ashok Ayar & Zeke Gillman, *California's SB1047 vs EU AI Act: A Comparative Analysis of AI Regulation*, *SCIENCEPO* (Oct. 28, 2024), <https://www.sciencespo.fr/public/chaire-numerique/en/2024/10/28/californias-sb1047-vs-eu-ai-act-a-comparative-analysis-of-ai-regulation/> [<https://perma.cc/CG7U-KFWG>]; Ben Brooks, *California's AI Reforms Scare All Developers, Not Just Big Tech*, *TECH POL'Y PRESS* (Aug. 23, 2024), <https://www.techpolicy.press/californias-ai-reforms-scare-all-developers-not-just-big-tech/> [<https://perma.cc/ZH6U-JFED>]. Nothing comparable was ever proposed for OSS, where restrictions targeted narrow functions rather than prohibiting the publication of open tools. See LINUX FOUND., *UNDERSTANDING OPEN SOURCE TECHNOLOGY & US EXPORT CONTROLS 2* (2021), https://www.linuxfoundation.org/hubfs/UnderstandingOpenSourceTechnologyandUSExportControls_Whitepaper_071921.pdf [<https://perma.cc/JVD6-HB2R>].

recognize AI openness as a multifaceted issue,³² much of the literature remains focused on the accessibility of model weights,³³ and even more

32. See, e.g., ELIZABETH SEGER ET AL., CTR. FOR THE GOVERNANCE OF A.I., OPEN-SOURCING HIGHLY CAPABLE FOUNDATION MODELS 8–14, 26–29 (2023), https://cdn.governance.ai/Open-Sourcing_Highly_Capable_Foundation_Models_2023_GovAI.pdf [<https://perma.cc/NGY4-6GYB>]; Rishi Bommasani et al., *Considerations for Governing Open Foundation Models*, 386 SCI. 151, 153 (2024), <https://doi.org/10.1126/science.adp1848>; Irene Solaiman, *The Gradient of Generative AI Release: Methods and Considerations*, 6 PROCS. ACM CONF. ON FAIRNESS, ACCOUNTABILITY, & TRANSPARENCY (FACCT 2023) 111, 111, 113–14 (2023), <https://doi.org/10.1145/3593013.3593981>; David Gray Widder, Meredith Whittaker & Sarah Myers West, *Why ‘Open’ AI Systems Are Actually Closed, and Why This Matters*, 635 NATURE 827, 829–31 (2024) [hereinafter Widder, Whittaker & West, *Why ‘Open’ AI Systems Are Actually Closed*], <https://doi.org/10.1038/s41586-024-08141-1>; Tejas N. Narechania & Ganesh Sitaraman, *An Antimonopoly Approach to Governing Artificial Intelligence*, 43 YALE L. & POL’Y REV. 95, 120 (2024); Sayash Kapoor et al., *Position: On the Societal Impact of Open Foundation Models*, 235 PROCS. 41ST INT’L CONF. ON MACH. LEARNING 23082, 23082 (2024); Andreas Liesenfeld & Mark Dingemans, *Rethinking Open Source Generative AI: Open-Washing and the EU AI Act*, 2024 PROCS. ACM CONF. ON FAIRNESS, ACCOUNTABILITY, & TRANSPARENCY (FACCT ’24) 1774, 1777–78, <https://doi.org/10.1145/3630106.3659005>; Francisco Eiras et al., *Position: Near to Mid-Term Risks and Opportunities of Open-Source Generative AI*, 235 PROCS. 41ST INT’L CONF. ON MACH. LEARNING 12348, 12350–54 (2024); NIK MARDIA, JASMINE SUN & MARK SURMAN, PUBLIC AI: MAKING AI WORK FOR EVERYONE, BY EVERYONE 4–5 (2024), https://assets.mofoprod.net/network/documents/Public_AI_Mozilla.pdf [<https://perma.cc/3XQ7-RMD7>]; Matt White, Callean Osborne, Xiao-Yang Yanglet Liu, Keyi Wang & Sachin Mathew Varghese, *The Model Openness Framework: Promoting Completeness and Openness for Reproducibility, Transparency, and Usability in Artificial Intelligence*, in WORKSHOP ON REGULATABLE ML AT THE 39TH CONFERENCE ON NEURAL INFORMATION PROCESSING SYSTEMS 2–3 (2025); Tamara Paris, AJung Moon & Jin L.C. Guo, *Opening the Scope of Openness in AI*, 2025 PROCS. ACM CONF. ON FAIRNESS, ACCOUNTABILITY, & TRANSPARENCY 1293, 1296–1302, <https://doi.org/10.1145/3715275.3732087> (analyzing openness from an interdisciplinary perspective, focusing on interactivity, freedom, and inclusiveness); David Atkinson, *Open Shouldn’t Mean Exempt: Open-Source Exceptionalism and Generative AI 1*, 4 (July 23, 2025) (unpublished manuscript), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5355736 [<http://dx.doi.org/10.2139/ssrn.5355736>] (focusing on the open-closed source binary and the products open-source AI produces).

33. See, e.g., PREM M. TRIVEDI & NAT MEYSENBERG, OPEN TECH. INST., OPENNESS IN ARTIFICIAL INTELLIGENCE MODELS 10–12 (2024), <https://www.newamerica.org/oti/reports/openness-in-artificial-intelligence-models/> [<https://perma.cc/3A5C-VQFD>] (advocating for AI openness to address more than model weights and source code to include transparency in the development process, but stopping short of acknowledging other components in the AI stack); Matthew Leisten, *Open(?) AI 3–5* (Jan. 3, 2026) (unpublished manuscript), <https://mleisten.github.io/Research/index> (focusing exclusively on model weights and architecture); Kaige Gao, Youngjin Yoo & Aaron Schecter, *Open Source AI Community as “Trading Zone”: The Role of Open-Source Models in the Diffusion of Artificial Intelligence Innovation*, 45 INT’L CONF. ON INFO. SYS. 1, 2, 14 (2024), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5019689 [<http://dx.doi.org/10.2139/ssrn.5019689>] (focusing exclusively on models); Yujin Potter, Michael Potter & Dawn Song, “As an AI, I Believe AI Models Should Be Open Source” 6–7 (2024) (unpublished manuscript), https://rdi.berkeley.edu/research/uploads/LLM_open_vs_closed.pdf [<https://perma.cc/M2NV->

nuanced analyses often stop short of connecting component-level distinctions to concrete legal strategies.

This Article fills that gap by introducing a framework of “differential openness” that rejects oversimplified labels of “open AI” or “open source AI” for the more analytically precise “open spectrum AI.” It untangles osAI systems into their constituent components, mapping each along a gradient of openness and evaluating how specific configurations of openness advance or undermine public goals. This taxonomy provides policymakers with the analytical tools needed to navigate the complex trade-offs with precision and to craft targeted, calibrated interventions that maximize benefits while mitigating risks. This Article’s taxonomy and framework are designed primarily with frontier foundation models in mind. These models concentrate the most difficult policy trade-offs because of their scale, their general-purpose capabilities, and the number of actors shaping their development. Still, the core insight—that openness exists along a spectrum within each component and produces context-dependent trade-offs—extends to other forms of AI as well.

To develop this argument, this Article proceeds in three parts. Part I dismantles the flawed “open–closed” binary and introduces our taxonomy of what we call “differential openness” for osAI. It begins by challenging the open source software analogy, demonstrating that the governance frameworks built for it are technologically and culturally ill-suited for osAI systems, which are defined by many modular,

DBQW] (flattening AI openness into a binary); Mike Sexton, *Open-Source Is a National Security Imperative*, THIRD WAY (Jan. 30, 2025), <https://www.thirdway.org/report/open-source-ai-is-a-national-security-imperative> (flattening AI openness into a binary); Domen Vake, Bogdan Šinik, Jernej Vičič & Aleksandar Tošić, *Is Open Source the Future of AI? A Data Driven Approach*, 15 APPLIED SCI. 1, 16 (2025) (flattening AI openness into a binary); MASAO DAHLGREN, CTR. FOR STRATEGIC & INT’L STUD., DEFENSE PRIORITIES IN THE OPEN-SOURCE AI DEBATE 3 (2024), <https://www.csis.org/analysis/defense-priorities-open-source-ai-debate> [<https://perma.cc/FG26-SG5B>] (recognizing more components can be relevant but focusing exclusively on model weights); DIGIT. PUB. GOODS ALL. & UNICEF, CORE CONSIDERATIONS FOR EXPLORING AI SYSTEMS AS DIGITAL PUBLIC GOODS 2–5 (2023) <https://www.digitalpublicgoods.net/AI-CoP-Discussion-Paper.pdf> [<https://perma.cc/XN9R-6BU8>] (focusing exclusively on models and data); Tairu Zhang & Tianyi Feng, *Application and Technology of an Open Source AI Large Language Model in the Medical Field*, 2 RADIOLOGY SCI. 96, 100–01 (2023), <https://www.scienceopen.com/hosted-document?doi=10.15212/RADSCI-2023-0007> (focusing on a binary of openness and its effects on the medical field); Thibault Schrepel & Jason Potts, *Measuring the Openness of AI Foundation Models: Competition and Policy Implications*, 34 INFO. & COMM’N TECH. L. 279, 286–91 (2025), <https://doi.org/10.1080/13600834.2025.2461953> (advancing a broader multidimensional framework but ultimately flattening openness into a license-and-access-centric typology); Alex Engler, *How Open-Source Software Shapes AI Policy*, BROOKINGS INST. (Aug. 10, 2021), <https://www.brookings.edu/articles/how-open-source-software-shapes-ai-policy/> [<https://perma.cc/BF45-LJG6>] (flattening AI openness into a binary).

interconnected components and a complex ecosystem of corporate and state actors. It then untangles AI systems into their eight key components—compute, data, source code, model weights, system prompts, operational records and controls, applications, and labor—to establish a more precise vocabulary for analyzing how openness functions at each layer of the AI stack.

Part II uses this new taxonomy to systematically evaluate how different configurations of component-level openness advance or undermine core policy objectives: public safety, innovation and economic growth, democratic accountability, and national security. We analyze the complex, often contradictory, effects of differential component openness on each goal. This analysis reveals that while openness is often a powerful engine for progress, it is not an intrinsic good but instead an instrumental value whose desirability depends entirely on context, forcing policymakers to confront the difficult trade-offs inherent in osAI governance.

Finally, Part III moves beyond diagnosis to examine how existing legal and regulatory levers might be analyzed through the lens of our component-based framework. We assess how tools related to liability, competition policy, intellectual property, trade controls, and direct government support map onto different layers of the AI stack, revealing both their potential and their limitations. Rather than prescribing specific policy solutions, this analysis offers a research agenda that is informed by our taxonomy and applies our framework of differential openness to orient future work toward more targeted, nuanced interventions beyond blunt, system-level mandates.

The “open–closed” binary and the assumption that it is an unmitigated good or evil is a siren song that has already led policy astray. Effective governance requires abandoning this simplistic lens and embracing a more sophisticated, differential openness framework for governing osAI—one that calibrates policy to the distinct risks and benefits of each component of the AI stack. Only by untangling AI in this way can we move beyond ideological debates and begin the difficult but essential work of crafting targeted rules for a technological future that best serves the public interest.

I. A TAXONOMY OF AI OPENNESS

The debate over open spectrum AI is distorted by its inheritance from the history of open source software. This legacy has generated two core misconceptions that skew policy: a false “open–closed” binary and the reflexive assumption that openness is an inherent good. This simplistic framing fails because it imports assumptions from a different technological and institutional era. The OSS world—driven largely by

individual developers and academics focused on opening source code—is fundamentally distinct from the modern osAI ecosystem, which involves a complex stack of interdependent components, each of which exists on its own spectrum of openness³⁴ and is controlled by one of a concentrated set of powerful corporate actors.³⁵

This Part dismantles the flawed OSS analogy and, in its place, introduces this Article’s core contribution: a taxonomy that untangles AI systems into their key technical and human components. By revealing osAI’s differential openness—the many dimensions along which openness actually varies—we demonstrate that the value of opening any single component is not innate but instrumental, capable of advancing some policy goals while simultaneously undermining others.

This taxonomic precision is essential for effective governance. It moves the analysis beyond asking *if* a system is open to asking more critical questions: What is open, how open is it, and to what end? Answering these questions is a prerequisite for crafting policies that can effectively balance the competing values at stake in AI development—safety, innovation, democratic control, and national security—the central task we undertake in Part II.

A. Beyond the Open Source Software Analogy

The debate over open spectrum AI inherits a great deal—some of it useful, much of it not—from the history of OSS. This impulse, however misguided, is understandable: Openness yielded considerable benefits, from innovation to decentralized governance—at least in the movement’s idealized self-understanding.³⁶ Although there are similarities in the core concepts of openness, the traditional software and AI are fundamentally distinct. To offer more rigor and nuance to the debate, this Section dismantles the flawed OSS analogy by establishing three distinctions that require us to look beyond source code, beyond individual developers, and beyond the AI labs themselves. Ultimately, it warns policymakers against assuming that osAI can be governed by the OSS playbook.

1. Beyond Source Code

The first reason that the OSS analogy fails as applied to osAI is that software is a comparatively simpler technology, and its model of openness—while revolutionary when it was developed—is insufficient to

34. Solaiman, *supra* note 32, at 111–14.

35. Widder, Whittaker & West, *Why ‘Open’ AI Systems Are Actually Closed*, *supra* note 32, at 827.

36. Kapoor et al., *supra* note 32, at 23082–83.

capture the complex, multi-layered reality of AI.³⁷ The success of OSS hinged on making openness legible, scalable, and enforceable by focusing on a single, critical component: source code.³⁸ The OSS community accomplished this with a technical mechanism for distributing open source code and a legal mechanism for enforcing its continued openness.

The technical dimension of software openness is straightforward. The value of an OSS project is unlocked almost entirely by making its source code free and publicly accessible. Code repositories like Microsoft's GitHub provide the infrastructure for this, creating a universal platform where the OSS source code lives, allowing developers, from the original authors to third-party contributors, to inspect, modify, and contribute to a project.³⁹ This can take the form of reporting bugs or flaws in the software, suggesting improvements (from safety enhancements to improved efficiency), or building off the software in innovative ways (from new use cases to new capabilities).⁴⁰

This technical access is made legally meaningful through a legal mechanism: a spectrum of OSS licenses. By default, copyright law grants exclusive rights to the creator.⁴¹ OSS licenses strategically override this default, creating a durable legal basis for permissionless use and collaboration.⁴² These licenses are not monolithic; they represent a range of trade-offs between ensuring downstream freedom, maximizing adoption, and retaining some proprietary control.

At one end, "copyleft" licenses, such as the GNU General Public License (GPL),⁴³ enforce downstream openness by imposing restrictive terms that require derivative works to impose the same viral license.⁴⁴ In doing so, they prevent users from locking up derivative works behind

37. See David Gray Widder, Meredith Whittaker & Sarah Myers West, *Open (For Business): Big Tech, Concentrated Power, and the Political Economy of Open AI 2* (Aug. 16, 2023) (unpublished manuscript) [hereinafter Widder, Whittaker & West, *Open (For Business)*], https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4543807 [<http://dx.doi.org/10.2139/ssrn.4543807>].

38. See Chinmayi Sharma, *Tragedy of the Digital Commons*, 101 N.C. L. REV. 1129, 1141–43 (2023) [hereinafter Sharma, *Tragedy of the Digital Commons*].

39. *Id.* at 1139.

40. *Id.* at 1142–43.

41. See *id.* at 1164–65; 17 U.S.C. § 106.

42. Sharma, *Tragedy of the Digital Commons*, *supra* note 38, at 1164–65.

43. RICHARD M. STALLMAN, *What Is Copyleft?*, in *FREE SOFTWARE, FREE SOCIETY: SELECTED ESSAYS OF RICHARD M. STALLMAN* 89, 89–90 (Joshua Gay ed., 1st ed. 2002); see also Michael J. Madison, *Reconstructing the Software License*, 35 LOY. U. CHI. L.J. 275, 283–84 (2003) (describing how open source licenses, particularly copyleft provisions, enforce ongoing source code disclosure across all participants in a software project).

44. See David McGowan, *Legal Implications of Open-Source Software*, 2001 U. ILL. L. REV. 241, 254.

closed systems; at times, this dampens OSS adoption. At the other end, “permissive” licenses like MIT, BSD, and Apache impose minimal restrictions, encouraging maximum OSS adoption and commercialization by requiring little more than attribution to the original OSS project.⁴⁵ While the former preserves downstream openness through brute force, the latter can foster openness by inviting more players to contribute to the OSS ecosystem without foregoing the possibility of financial gain.

Between these two poles is a growing class of source-available license configurations. Some make OSS source code visible and readable to users but impose significant restrictions such as prohibiting modification, redistribution, or commercial use without explicit permission from the original developer.⁴⁶ This achieves some of openness’s value—trust and oversight through transparency—but it limits generative collaborations. Others, such as the edtech company, Instructure, straddle both ends of the OSS spectrum a different way: By employing dual-licensing strategies, they can release a copyleft version of a project that contributes to the OSS community while maintaining a proprietary version with extended features to earn a profit.⁴⁷ These hybrid license configurations foreshadow similar strategies in osAI, reconciling the competing interests between realizing the benefits of openness and preserving profitability.⁴⁸

45. See Petr Pícha & Souhaila Serbout, *On the Adoption of Open Source Software Licensing - A Pattern Collection*, 29 EUROPLoP '24: PROCS. EUR. CONF. ON PATTERN LANGUAGES PROGRAMS, PEOPLE, & PRACS. 1, 5 (2024), <https://doi.org/10.1145/3698322.3698341> (describing the benefits of permissive licenses such as MIT, Apache 2.0, and BSD, including “[o]pen[ing] up possibilities for innovative and potentially profitable uses of the software”); Andre Morin, Jennifer Urban & Piotr Sliz, *A Quick Guide to Software Licensing for the Scientist-Programmer*, PLOS COMPUTATIONAL BIOLOGY, July 2012, at 1, 3 (2012), <https://doi.org/10.1371/journal.pcbi.1002598>; LAWRENCE ROSEN, OPEN SOURCE LICENSING: SOFTWARE FREEDOM AND INTELLECTUAL PROPERTY LAW 69–70 (2005).

46. See, e.g., MEGAsync, *Mega Limited Code License*, GITHUB, <https://github.com/meganz/MEGAsync/blob/master/LICENCE.md> [<https://perma.cc/L5XJ-8JCZ>] (last visited Jan. 26, 2026); *Adopting and Developing BSL Software*, MARIADB, <https://mariadb.com/bsl-faq-adopting/> [<https://perma.cc/5D4R-C5JV>] (last visited Jan. 26, 2026); Thomas Claburn, *Redis Has a License to Kill: Open-Source Database Maker Takes Some Code Proprietary*, REGISTER (Aug. 23, 2018, at 06:05 UTC), https://www.theregister.com/2018/08/23/redis_database_license_change/ [<https://perma.cc/5UG5-2NRW>]; *Licenses*, REDIS, <https://redis.io/legal/licenses/> [<https://perma.cc/F6YE-9UXY>] (last visited Jan. 26, 2026).

47. See *Our Open Source Strategy*, INSTRUMENT, <https://www.instructure.com/resources/blog/our-open-source-strategy> [<https://perma.cc/SZ6P-56AC>] (last visited Jan. 26, 2026).

48. See, e.g., Shirin Ghaffary, *Why Meta Is Giving Away Its Extremely Powerful AI Model*, VOX (July 28, 2023, at 05:00 CT), <https://www.vox.com/technology/2023/7/28/23809028/ai-artificial-intelligence-open-closed-meta-mark-zuckerberg-sam-altman-open-ai>; Bart de Witte, *Case Study: Meta’s Strategy for Open-Sourcing Llama: A Detailed Analysis*, HEALTHCARE INNOVATION LETTER (Aug. 5, 2024),

This dual focus on OSS source code—made available technically and legally—has created a powerful and legible narrative about openness. For example, the contrast between an open operating system like Linux and a closed one like Windows provides a clear case study in how these choices shaped market structures, pricing, and user control.⁴⁹ And while other factors shape the software ecosystem—including hardware lock-in, proprietary data formats, and anticompetitive behavior⁵⁰—the focal importance of source code and the clear differences between open and closed licenses have made the OSS model so influential.

Consequently, policymakers latched onto the elegantly simple single-component framework.⁵¹ Source code access became the proxy for procurement policies that sought transparency and security,⁵² antitrust analysis that examined whether closed systems created unfair market advantages,⁵³ and liability frameworks that treated OSS projects more

<https://blog.hippoi.org/metastategy-for-open-sourcing-Llama-a-detailed-analysis-hippogram-27/> [<https://perma.cc/RG4C-2QQC>].

49. *Compare Licensing*, MICROSOFT, <https://www.microsoft.com/en-us/licensing> (last visited Jan. 26, 2026) (describing Microsoft’s complex and restrictive commercial licensing framework, which governs access, use, and distribution through product and enterprise-specific terms), *with* Linux GPL-2.0 License, GITHUB, <https://github.com/torvalds/linux/blob/master/LICENSES/preferred/GPL-2.0> [<https://perma.cc/7AKM-CWQH>] (last visited Jan. 26, 2026) (providing a free and permissive license for the Linux kernel that guarantees users the rights to access, use, modify, and distribute the source code).

50. *See generally* TIM WU, *THE MASTER SWITCH: THE RISE AND FALL OF INFORMATION EMPIRES* (2010) (describing how the internet has moved from an open, generative space to a closed ecosystem of walled gardens); JONATHAN ZITTRAIN, *THE FUTURE OF THE INTERNET—AND HOW TO STOP IT* (2008) (same).

51. *See* Robert W. Hahn, *Government Policy Toward Open Source Software: An Overview*, in *GOVERNMENT POLICY TOWARD OPEN SOURCE SOFTWARE* 1, 4–7 (Robert W. Hahn ed., 2002); *see also* CONG. RSCH. SERV., RL32268, *INTELLECTUAL PROPERTY, COMPUTER SOFTWARE AND THE OPEN SOURCE MOVEMENT* 1–4, 11–13, 17–18 (2004) (providing an overview of open source code licensing, including how copyleft provisions enforce source code disclosure and the resulting tensions with intellectual property rights).

52. *See* Robert W. Gomulkiewicz, *Considering a Right to Repair Software*, 37 *BERKELEY TECH. L.J.* 943, 958–60 (2022); *see also* David S. Evans & Bernard J. Reddy, *Government Preferences for Promoting Open-Source Software: A Solution in Search of a Problem*, 9 *MICH. TELECOMM. TECH. L. REV.* 313, 315 (2003) (surveying government initiatives worldwide to promote open source software through procurement preferences, R&D subsidies, and standardization mandates).

53. *See, e.g.*, Press Release, U.S. Dep’t of Just., Justice Department Sues Google for Monopolizing Digital Advertising Technologies (Feb. 6, 2025), <https://www.justice.gov/archives/opa/pr/justice-department-sues-google-monopolizing-digital-advertising-technologies> [<https://perma.cc/TC4B-Z3YK>]; Press Release, U.S. Dep’t of Just., Justice Department Sues Apple for Monopolizing Smartphone Markets (Feb. 6, 2025), <https://www.justice.gov/archives/opa/pr/justice-department-sues-apple-monopolizing-smartphone-markets> [<https://perma.cc/33CT-CSVC>] (alleging that Apple maintained its smartphone monopoly by imposing contractual restrictions on developers and limiting interoperability with competing platforms); *United States v. Apple, Inc.*,

leniently when assigning responsibility for failures or vulnerabilities.⁵⁴ This entrenched the idea that governing technological openness was primarily a matter of governing source code.

The simplicity that offered such clarity for OSS, however, now becomes a liability. As we explain in detail in the next section,⁵⁵ osAI systems are not monolithic programs; they are layered systems composed of interdependent components—compute, data, source code, model weights, and more—whereas source code is just one piece of the puzzle, and often not the most important one. A myopic focus on source code obfuscates both the cascading effects of AI openness decisions and the policy levers available to governments.

2. Beyond Altruism

A second reason the OSS analogy fails when applied to osAI is because the primary actors driving openness in AI, and their motivations, are fundamentally different than in the software context.⁵⁶ Therefore, it is incumbent on policymakers to understand the power dynamics in the osAI ecosystem—why certain players may choose to open or close different AI components—to diagnose when openness decisions might serve policy goals and how policy levers can account for the incentives that drive them.⁵⁷

The ethical commitment to software openness evolved in a distinct institutional context, one defined by a decentralized community of academics, researchers, hobbyists, corporations, and government entities collaborating on an ecosystem of open protocols and widespread information sharing.⁵⁸ Many were driven in part by self-gain, but OSS is

No. 24-CV-4055, 2025 WL 1829127, at *15 (D.N.J. June 30, 2025) (“[T]he Amended Complaint alleges Apple maintains a market share of 65 percent in the smartphone market and 70 percent in the performance smartphone market, imposed several barriers to entry, and has engaged in anticompetitive conduct.”).

54. Compare Sharma, *Tragedy of the Digital Commons*, *supra* note 38, at 1134, 1137 (emphasizing that “[o]pen source is not the problem” and arguing that open source must be treated as critical public infrastructure that requires government-protected intervention), with Bryan H. Choi, *Tainted Source Code*, 39 HARV. J.L. & TECH. (forthcoming Fall 2025) (manuscript at 1), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5169060 [<http://dx.doi.org/10.2139/ssrn.5169060>] (arguing that OSS contributions should not be categorically exempt from liability and proposing a negligence-based framework for harms caused by defective open source code).

55. See *infra* Section I.B.

56. See Widder, Whittaker & West, *Why ‘Open’ AI Systems Are Actually Closed*, *supra* note 32, at 827.

57. See Widder, Whittaker & West, *Open (For Business)*, *supra* note 37, at 4.

58. See Sharma, *Tragedy of the Digital Commons*, *supra* note 38, at 1148, 1152 (describing the groundbreaking development of the Linux OS and Apache web server).

unique because, at its core, it believes openness serves the public interest.⁵⁹ Accordingly, the modern OSS movement responded to the trend of software commercialization⁶⁰ by codifying an ethical vision of software freedom: Users should have the right to run, study, modify, and share the code they use.⁶¹ By harnessing the engine of volunteer developers and corporate interests alike, the open source movement flourished, yielding technological advancements and improved accountability.⁶² Today, cornerstone OSS projects like Linux, Apache, and Python form the backbone of global computing infrastructure.⁶³

While many hope that openness in the AI ecosystem will catalyze similar collaborative innovation and build public trust, osAI is driven by a very different, much more centralized set of stakeholders.⁶⁴ These actors are motivated by different, often competing, equities—some prioritize economic growth, others national security or public accountability—and not all are driven by the altruistic motivations that gave rise to the OSS movement.⁶⁵

Meta's release of Llama's model weights, for instance, was not just a nod to open science but a calculated play to gain multiple strategic advantages.⁶⁶ By making its models free and widely available, Meta encourages a global community of developers to build tools and applications on its platform, effectively crowdsourcing innovation and making the Llama architecture a de facto industry standard.⁶⁷ This strategy seeks to commoditize the model layer of the AI stack, creating a competitive disadvantage for rivals like OpenAI who charge for access, while also serving as a powerful recruiting tool by allowing its researchers to be more public about their advances.

59. See *id.* at 1148–49.

60. See Martin Campbell-Kelly, *Not All Bad: An Historical Perspective on Software Patents*, 11 MICH. TELECOMM. TECH. L. REV. 191, 211–12 (2005).

61. See Richard Stallman, *Why Open Source Misses the Point of Free Software*, GNU OPERATING SYS., <https://www.gnu.org/philosophy/open-source-misses-the-point.en.html> [<https://perma.cc/T53T-GAEK>] (last visited Jan. 26, 2026).

62. Sharma, *Tragedy of the Digital Commons*, *supra* note 38, at 1148–52.

63. See Jesus M. Gonzalez-Barahona, *A Brief History of Free, Open Source Software and Its Communities*, IEEE COMPUT., Feb. 2021, at 75, 79; see also Paul Jansen, *TIOBE Index for January 2026*, TIOBE, <https://www.tiobe.com/tiobe-index/> [<https://perma.cc/H9K3-P7KB>] (last visited Jan. 26, 2026) (reporting Python as the most popular among all programming languages, open and closed).

64. See Widder, Whittaker & West, *Open (For Business)*, *supra* note 37, at 4.

65. See *id.* at 3.

66. See Ghaffary, *supra* note 48; de Witte, *supra* note 48. But see Eli Tan, *Meta's New Superintelligence Lab Is Discussing Major A.I. Strategy Changes*, N.Y. TIMES (July 14, 2025), <https://www.nytimes.com/2025/07/14/technology/meta-superintelligence-lab-ai.html>.

67. See Ghaffary, *supra* note 48.

Indeed, recent developments in the OSS movement’s own trajectory caution against assuming openness is always altruistic. What began as a grassroots, volunteer-driven movement is now largely powered by corporate developers—as of 2020, nearly half of OSS contributions today come from employees at firms like Google, Microsoft, and IBM.⁶⁸ Many of these companies have learned to strategically embrace openness, not as a value, but as a vehicle for shaping ecosystems and entrenching market position, often sacrificing true openness values at the altar of pursuing greater market power.⁶⁹ For example, Google openly released Android, its mobile operating system, which drove app innovation and therefore consumer adoption, without lowering the barriers to entry for innovating on the operating system itself—yielding more customers, no new competitors, and greater lock-in.⁷⁰ Other companies skirt the licenses enforcing downstream openness by capitalizing on what they learned from the OSS projects they used to build parallel systems they can profit from.⁷¹

These dynamics are increasingly visible not only in how AI components are released but also in how ostensibly “open” AI initiatives are governed. While open source software projects have long relied on foundation-based governance to steward community-driven development,⁷² emerging osAI initiatives often invert this sequence:

68. Sharma, *Tragedy of the Digital Commons*, *supra* note 38, at 1150–51; FRANK NAGLE, DAVID A. WHEELER, HILA LIFSHITZ-ASSAF, HAYLEE HAM & JENNIFER L. HOFFMAN, REPORT ON THE 2020 FOSS CONTRIBUTOR SURVEY 6 (2020), https://www.linuxfoundation.org/wp-content/uploads/2020FOSSContributorSurveyReport_121020.pdf [<https://perma.cc/U2C6-MA89>] (reporting that 48.7% of survey respondents “[were] paid by their employer to contribute to” open source projects).

69. See Sharma, *Tragedy of the Digital Commons*, *supra* note 38, at 1150 (“Companies consuming open source can also benefit from contributing to it by using the insight they gain from the projects to build complementary products or to influence project development in a way that supports their products.”); Salil Deshpande, *Lack of Leadership in Open Source Results in Source-Available Licenses*, TECHCRUNCH (May 30, 2019, at 14:00 PT), <https://techcrunch.com/2019/05/30/lack-of-leadership-in-open-source-results-in-source-available-licenses/> [<https://perma.cc/C9DK-R3NM>] (describing “the Amazon problem” as an example of companies copying open source projects and selling them as commercial services without contributing to the original project, forcefully taking over control of projects from their original creators, or building on top of popular open source projects to siphon customers away from the actual project to their own proprietary products).

70. See Widder, Whittaker & West, *Open (For Business)*, *supra* note 37, at 13.

71. See Stephen Shankland, *Google Gets Web Allies by Letting Outsiders Help Build Chrome’s Foundation*, CNET (Nov. 30, 2020, at 09:26 PT), <https://www.cnet.com/tech/mobile/google-gets-web-allies-by-letting-outsiders-help-build-chromes-foundation/> [<https://perma.cc/U2YZ-LADT>].

72. Many major open source software projects rely on foundation-based governance structures to manage legal, financial, and trademark concerns while leaving day-to-day technical governance to project-specific processes. See, e.g., *How the ASF Works*, APACHE SOFTWARE FOUND., <https://www.apache.org/foundation/how-it-works/>

Corporations design projects, define governance structures, and only later situate them within foundation frameworks to broaden adoption and legitimacy.⁷³ As we turn to osAI, these dynamics remind us that openness can be both a public good and a competitive tactic.

3. Beyond the Developer

Even a focus on the strategic motivations of developers is too narrow. The third failure of the OSS analogy is that it cannot account for the sprawling ecosystem of diverse stakeholders who control different, and often more critical, layers of the AI technology stack. OSS developers are the ones writing source code; they have the power to open

works.html [https://perma.cc/24ME-368M] (last visited Jan. 27, 2026) (explaining how the Apache Foundation Board manages corporate assets, including funds and intellectual property while Project Management Committees retain technical decision-making authority). Examples include the Apache Foundation (a 501(c)(3)) and the Linux Foundation (a 501(c)(6)). See *Public Records*, APACHE SOFTWARE FOUND., https://www.apache.org/foundation/records/ [https://perma.cc/M8T4-AV2C] (last visited Mar. 5, 2026); *About the Linux Foundation*, LINUX FOUND., https://www.linuxfoundation.org/about [https://perma.cc/S86Y-6J9H] (last visited Mar. 5, 2026). These foundations typically emerged in response to community needs to steward widely adopted, volunteer-driven projects, particularly to manage intellectual property and organizational continuity. DALIA TOPELSON RITVO, KIRA HESSEKIEL & CHRISTOPHER T. BAVITZ, ORGANIZATION & STRUCTURE OF OPEN SOURCE SOFTWARE DEVELOPMENT INITIATIVES: CHALLENGES & OPPORTUNITIES CONCERNING CORPORATE FORMATION, NONPROFIT STATUS, & GOVERNANCE FOR OPEN SOURCE PROJECTS 4 (2017), https://clinic.cyber.harvard.edu/wp-content/uploads/2017/03/2017-03_governance-FINAL.pdf [https://perma.cc/3NVL-JURR].

73. Two prominent examples are the PyTorch Foundation and the Agentic AI Foundation, both subsidiaries of the Linux Foundation. *Meta Transitions PyTorch to the Linux Foundation, Further Accelerating AI/ML Open Source Collaboration*, LINUX FOUND. (Sep. 12, 2022) [hereinafter *PyTorch Foundation*], https://www.linuxfoundation.org/press/press-release/meta-transitions-pytorch-to-the-linux-foundation [https://perma.cc/C6D3-NCFM]; *Linux Foundation Announces the Formation of the Agentic AI Foundation (AAIF)*, LINUX FOUND. (Dec. 9, 2025) [hereinafter *Agentic AI Foundation*], https://www.linuxfoundation.org/press/linux-foundation-announces-the-formation-of-the-agentic-ai-foundation [https://perma.cc/7XHV-J6JH]. Unlike many OSS foundations that formed around existing community-driven projects, these osAI foundations were established through the donation of corporate-developed initiatives. *PyTorch Foundation, supra*; *Agentic AI Foundation, supra*. Meta donated PyTorch, a project it had developed internally and released as open source over several years. *PyTorch Foundation, supra*. The Agentic AI Foundation was seeded through contributions including Anthropic's Model Context Protocol, OpenAI's AGENTS.md specification, and Block's Goose. *Agentic AI Foundation, supra*; see also AGENTS.MD, https://agents.md [https://perma.cc/3WEA-LQLW] (last visited Jan. 27, 2026) (specification page referenced as OpenAI's contribution); *What Is the Model Context Protocol (MCP)?*, MODEL CONTEXT PROTOCOL, https://modelcontextprotocol.io (last visited Jan. 27, 2026) (protocol documentation); *Goose*, BLOCK, https://block.github.io/goose/ [https://perma.cc/QB44-V4JC] (last visited Jan. 27, 2026) (project documentation).

source it. The entities that build osAI, on the other hand, are distributed across a wide range of powerful players, each building different components of the AI stack. Thus, many hands shape osAI's differential openness.

Some of the largest and most powerful players in the AI ecosystem are the companies that design, produce, and manage the specialized computational hardware AI requires.⁷⁴ Their core incentive is control over supply and demand.⁷⁵ Companies that design the specialized hardware, the often-overseas manufacturers that produce them, and the service providers that make them available to downstream consumers (leading AI labs, startups, and researchers alike) seek to maximize return on high-capital investments by controlling access to compute.⁷⁶ As hyperscale data centers proliferate, vertically integrated electric utilities and new generation companies are becoming part of this compute ecosystem as well, using control over grid interconnection timelines, power procurement, and tailored rate structures to determine which data center projects can secure the firm, large-scale electricity supply that frontier clusters require.⁷⁷ Essentially, these osAI ecosystem players dictate who gets to experiment with AI—and who gets priced out.⁷⁸ As actors with an economic interest in scarcity, they lack the incentive to change the status quo.

Beyond hardware, AI learns from data, and so it relies heavily on those who generate data (whether willingly or not) and those who transform raw data into usable datasets.⁷⁹ Because data quality, quantity, and type significantly impact model capabilities and biases, data providers who create, curate, and commodify datasets wield enormous

74. See Widder, Whittaker & West, *Open (For Business)*, *supra* note 37, at 7–8.

75. See *id.*

76. See Will Henshall, *Big Tech Companies Were Investors in Smaller AI Labs. Now They're Rivals*, *TIME* (May 13, 2024, at 09:29 CT), <https://time.com/6977424/ai-competition-openai-anthropic-microsoft-amazon/> [<https://perma.cc/F2HL-PULH>].

77. Jim McMahon, *Beyond Co-Location: The Emerging Opportunity for Vertically Integrated Utilities in the Data Center Boom*, *POWER* (July 29, 2025), <https://www.powermag.com/beyond-co-location-the-emerging-opportunity-for-vertically-integrated-utilities-in-the-data-center-boom/> [<https://perma.cc/3BFK-NF25>].

78. See Kevin M.K. Fodouop, *Promoting Access to Innovative AI*, 7 *J.L. & TECH. TEX.* 1, 9, 16, 29–30 (2023–24) (“The field has seen a continuous trend toward gigantic models, meaning that only the most resourced corporations can internally develop state-of-the-art innovative capabilities.”).

79. Kevin Roose, *The Data That Powers A.I. Is Disappearing Fast*, *N.Y. TIMES* (July 19, 2024), <http://nytimes.com/2024/07/19/technology/ai-data-restrictions.html>.

influence in the AI ecosystem.⁸⁰ These entities range from companies that happen to own vast archives of proprietary data they can sell to AI companies, such as news organizations like *The New York Times* or academic publishers like Wiley, to those who collect extensive user interaction data, such as social media platforms and e-commerce sites.⁸¹ Other companies specialize in collecting, refining, and labeling data—some even generate synthetic data specifically for AI development, including firms like Scale AI.⁸² Data providers can choose to filter out harmful content, avoid copyrighted material, and rectify biases—or not. They either opt to release datasets openly or restrict access through licensing or fees,⁸³ functionally determining, like hardware providers, who can play in the sandbox. They are driven by the desire to minimize legal risk while maximizing the ability to profit from a scarce resource.⁸⁴

The most visible members of the AI ecosystem are frontier model developers, including organizations such as OpenAI (ChatGPT), Google DeepMind (Gemini and Gemma), Meta (Llama), Anthropic (Claude), and xAI (Grok).⁸⁵ These companies are building the most powerful AI systems in the ecosystem today—the foundation models that everyone else builds upon—which is resource and expertise intensive.⁸⁶ Their incentive calculus blends public positioning with competitive strategy.⁸⁷ By releasing Llama, Meta sought to capture developer mindshare and

80. See *id.* (“[W]idespread data restrictions may pose a threat to A.I. companies, which need a steady supply of high-quality data to keep their models fresh and up-to-date.”).

81. See Diana Kwon, *Publishers Make Millions Selling Papers to Train AI*, 636 NATURE 529, 530 (2024); Annie Palmer, *Amazon AI Deal with New York Times Brings the Paper’s Content to Alexa*, CNBC (May 29, 2025, at 13:49 ET), <https://www.cnbc.com/2025/05/29/amazon-ai-new-york-times-alex.html> [https://perma.cc/VL2W-H42A]; *Meta Strikes Multiple AI Deals with News Publishers*, REUTERS (Dec. 6, 2025), <https://www.reuters.com/business/meta-strikes-multiple-ai-deals-with-news-publishers-axios-reports-2025-12-05/>.

82. See, e.g., *Data Engine*, SCALE, <https://scale.com/data-engine> [https://perma.cc/8DYX-229A] (last visited Jan. 27, 2026).

83. Narechania & Sitaraman, *supra* note 32, at 122–23; see also Katie Paul & Anna Tong, *Inside Big Tech’s Underground Race to Buy AI Training Data*, REUTERS (Apr. 5, 2024, at 13:53 CT), <https://www.reuters.com/technology/inside-big-techs-underground-race-buy-ai-training-data-2024-04-05/> (reporting that major AI companies are reliant on open datasets but also spend between \$25 and \$50 million or more on licenses for proprietary datasets).

84. See Widder, Whittaker & West, Open (For Business), *supra* note 37, at 8–10.

85. Tim Lu, *Frontier Models Explained: What Defines the Cutting Edge of AI*, DATACAMP: BLOG (Jan. 13, 2026), <https://www.datacamp.com/blog/frontier-models>.

86. See MARKUS ANDERLJUNG ET AL., FRONTIER AI REGULATION: MANAGING EMERGING RISKS TO PUBLIC SAFETY 7–9 (2023), https://cdn.governance.ai/Frontier_AI_Regulation_Managing_Emerging_Risks.pdf [https://perma.cc/4V28-TNVC].

87. See Widder, Whittaker & West, Open (For Business), *supra* note 37, at 4.

ecosystem control.⁸⁸ In contrast, OpenAI’s closed approach may have helped protect its lead in fine-tuning and enterprise deployment—at least, so far.⁸⁹ Showcasing the market domination these behemoths seek, many frontier companies are proactively investing in owning—and therefore controlling—the hardware AI depends on, vertically integrating the stack and doubling their capacity to influence openness in the ecosystem.⁹⁰

Once a model is trained, downstream developers adapt foundational AI models to specific use cases for public consumption—often taking the form of applications.⁹¹ Their incentives revolve around defensibility, differentiation, and user trust.⁹² For instance, a healthcare technology firm might fine-tune an osAI model on proprietary medical data, creating a powerful diagnostic tool that they can share with others or withhold for themselves. Many downstream developers may make some parts of an osAI system open while restricting others to protect proprietary assets and minimize exposure.⁹³ This hybrid openness, reminiscent of similar arrangements in OSS licenses, is becoming more common as companies leverage the benefits of open source collaboration while keeping their own competitive edge.⁹⁴

Finally, governance stakeholders—including regulators, standards bodies, and civil society groups—are trying to balance different policy goals such as safety, innovation, democratic control, and national security.⁹⁵ These groups do not control the technology directly, but their decisions around data governance, operational transparency, ethical

88. Ghaffary, *supra* note 48; de Witte, *supra* note 48.

89. See Narechania & Sitaraman, *supra* note 32, at 120–21.

90. See, e.g., Amin Vahdat, *Introducing Google Axion Processors, Our New Arm-based CPUs*, GOOGLE CLOUD: BLOG (Apr. 9, 2024), <https://cloud.google.com/blog/products/compute/introducing-googles-new-arm-based-cpu> [https://perma.cc/C3CG-6KR5]; Jake Siegel, *With a Systems Approach to Chips, Microsoft Aims to Tailor Everything ‘from Silicon to Service’ to Meet AI Demand*, MICROSOFT (Nov. 15, 2023), <https://news.microsoft.com/source/features/ai/in-house-chips-silicon-to-service-to-meet-ai-demand> [https://perma.cc/NV5Z-FK45]; *AWS and NVIDIA*, AMAZON WEB SERVS., <https://aws.amazon.com/nvidia/> [https://perma.cc/VH3K-XUVU] (last visited Jan. 27, 2026); Katie Paul & Krystal Hu, *Exclusive: Meta Begins Testing Its First In-House AI Training Chip*, REUTERS (Mar. 11, 2025, at 08:37 CT), <https://www.reuters.com/technology/artificial-intelligence/meta-begins-testing-its-first-in-house-ai-training-chip-2025-03-11/>.

91. Narechania & Sitaraman, *supra* note 32, at 119 (“Downstream developers need to access the foundation models for fine-tuning and use in a particular application . . .”).

92. See Bommasani et al., *supra* note 32, at 151.

93. See *infra* Appendix; Bommasani et al., *supra* note 32, at 151.

94. See Ben Cottier, Josh You, Natalia Martemianova & David Owen, *How Far Behind Are Open Models?*, EPOCH AI (Nov. 4, 2024), <https://epoch.ai/blog/open-models-report#open-models-have-lagged-on-benchmarks-by-5-to-22-months> [https://perma.cc/8KAM-NX6A]

95. See *infra* Part II.

constraints, and public accountability shape the legal and ethical landscape in which AI operates.⁹⁶ How much transparency should be required? Should companies be allowed to release powerful models with no oversight? Should there be safeguards against AI monopolization? Their efficacy in this role, however, can be impacted by political strife, resource constraints, expertise scarcity, and the bludgeoning of corporate lobbyists.⁹⁷

Unsurprisingly, each of the different players in the AI ecosystem operates under different incentives, including as it relates to openness. The technological nature of software is not as reliant on hardware, data sources, and expertise as is AI, and policymakers can get away with a singular focus on developers who control source code openness.⁹⁸ AI stakeholders are diverse and diffuse, which means effective osAI policy must accommodate its technical and sociocultural distinctions from open source software.

B. Untangling AI

The simplicity that once made OSS governance tractable becomes a liability with osAI. To grapple with osAI's complexity, we introduce our core analytical innovation: the concept of untangling AI systems into constituent components and fitting them within a framework of differential openness. This Section dissects the technical layers of an AI system—compute, data, source code, model weights, system prompts, operational control and records, applications, and labor—identifying how openness manifests differently across them.

96. See, e.g., Press Release, Bureau of Indus. & Sec., Biden-Harris Administration Announces Regulatory Framework for the Responsible Diffusion of Advanced Artificial Intelligence Technology (Jan. 13, 2025), <https://www.bis.gov/press-release/biden-harris-administration-announces-regulatory-framework-responsible-diffusion-advanced-artificial> [<https://perma.cc/T57A-USTX>]; Matt O'Brien, *White House Says No Need to Restrict Open-Source AI, For Now*, PBS NEWS (July 30, 2024, at 14:02 ET), <https://www.pbs.org/newshour/nation/white-house-says-no-need-to-restrict-open-source-ai-for-now> [<https://perma.cc/Z8MH-3HST>]; THE WHITE HOUSE, NATIONAL CYBERSECURITY STRATEGY 1, 21 (2023); Brooks, *supra* note 31; Zuzanna Warso & Maximilian Gahntz, *How the EU AI Act Can Increase Transparency Around AI Training Data*, TECH POL'Y PRESS (Dec. 9, 2024), <https://www.techpolicy.press/how-the-eu-ai-act-can-increase-transparency-around-ai-training-data/> [<https://perma.cc/AXC9-P8MA>]; Pablo Chavez, *Sovereign AI in a Hybrid World: National Strategies and Policy Responses*, LAWFARE (Nov. 7, 2024, at 10:40 CT), <https://www.lawfaremedia.org/article/sovereign-ai-in-a-hybrid-world--national-strategies-and-policy-responses> [<https://perma.cc/U66E-ACQW>].

97. See Chinmayi Sharma, *AI's Hippocratic Oath*, 102 WASH. U. L. REV. 1101, 1137–41 (2025) [hereinafter Sharma, *AI's Hippocratic Oath*].

98. See Tejas N. Narechania, *Machine Learning as Natural Monopoly*, 107 IOWA L. REV. 1543, 1569–88 (2022).

These components are not static; they interact with each other both in development and post-deployment through feedback loops.⁹⁹ Information gathered by one component can inform the future development of another, and by controlling more than one component, a single entity can multiply its market advantage.¹⁰⁰ Through a regulatory lens, choices about openness at one layer can sharply condition what is possible at others. An open-weight model that can only be trained or served on a small number of proprietary graphical processing unit (GPU) clusters remains effectively constrained by the closed compute layer beneath it. An open source application that relies on a closed model-as-a-service inherits that model's limits on transparency and user control. Conversely, open hardware and interoperable cloud platforms can magnify the value of open models and data by making them realistically usable. In practice, the openness of an AI system is bounded not just by any single component, but by how those components interact across the stack.

By untangling AI into a legible taxonomy and explaining the concept of differential openness, we lay the foundation for demonstrating how policymakers can develop more targeted interventions: identifying which component is most relevant to their goals, who controls it, and how open that component should be. Although this taxonomy is indexed to the structure of frontier foundation models, which exhibit the most pronounced separation of components and institutional roles, its core logic is more general. Many non-frontier systems, including the smaller, task-specific models, map onto these components as well, even if their architectures are simpler and the stakes of openness differ. Figure 1 visualizes the relationships in the osAI stack.

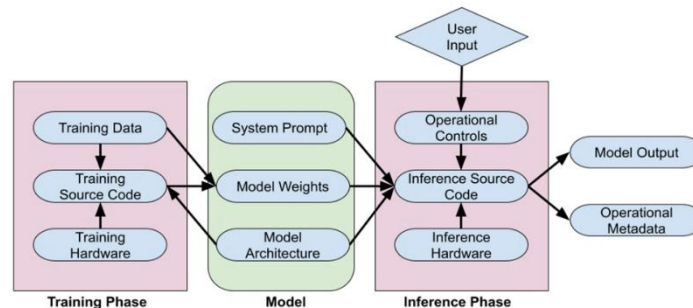


Figure 1. Illustration of an AI Model's Lifecycle

99. *AI Guide for Government: Understanding and Managing the AI Life Cycle*, U.S. GEN. SERVS. ADMIN.: CTRS. OF EXCELLENCE, <https://coe.gsa.gov/coe/ai-guide-for-government/understanding-managing-ai-lifecycle/> [https://perma.cc/95KC-HX68] (last visited Jan. 27, 2026) (explaining that model development continues after deployment based on information gathered).

100. See Narechania, *supra* note 98, at 1584.

This figure illustrates an AI model’s lifecycle, which comprises a *training phase* and an *inference phase*. In the training phase, *training source code* running on *training hardware* (normally a GPU or tensor processing unit (TPU)) processes *training data* to create *model weights*. These weights, which are numerical parameters, combine with the *model architecture* (the model’s underlying structure) and a text-based *system prompt* to form the complete *model*. During the subsequent inference phase, a *user input* is first managed by operational controls, such as safety filters, before being processed by *inference source code on inference hardware* (similar to training hardware). This code applies the multi-component model to the input to generate a *model output* as well as *operational metadata* like usage logs and audit trails. Note that safety features can also be part of the training process or be applied to the model’s final output.

1. Compute

At the foundation of AI systems lies compute—the physical hardware that powers both AI training and “inference,” the process by which a trained model produces output based on user input. Unlike traditional software, which can run on almost any computer, cutting-edge “frontier” AI models require massive, specialized, and often cost-prohibitive hardware like GPUs and TPUs.¹⁰¹ This reality has created a “hardware lottery,” where even the best theoretical advances rely on the happenstance of available computation to have practical impact.¹⁰² The hardware layer creates a significant bottleneck in the AI supply chain, with the market dominated by a few key firms: Nvidia and Google for chip design, Taiwan Semiconductor Manufacturing Company (TSMC) for manufacturing, and the Dutch company ASML for essential lithography equipment.¹⁰³ However, other parts of the AI stack are investing in their own compute to avoid reliance on these titans.¹⁰⁴ These emerging competitors, however, are few and no less powerful.

A complete account of compute must also include energy. Modern AI systems consume enormous amounts of electricity, making power availability a practical constraint on who can train or deploy frontier models and where. Analysts now treat electricity as a third core input to AI—alongside chips and data—because even abundant accelerators

101. Narechania, *supra* note 98, at 1569–88, 1575 n.137.

102. See Sara Hooker, *The Hardware Lottery*, COMM’NS ACM, Dec. 2021, at 58, 60–63.

103. Narechania & Sitaraman, *supra* note 32, at 112–13.

104. See *supra* note 90 and accompanying text.

cannot be fully utilized without sufficient power.¹⁰⁵ Data centers already draw hundreds of terawatt-hours per year globally, and AI-focused clusters can require gigawatt-scale loads, producing wait times of five to seven years for new grid interconnection in some regions.¹⁰⁶ This makes power availability a practical constraint on who can train or deploy advanced models and where they can do so. Indeed, a July executive order called for the acceleration of energy infrastructure projects, specifically to support development of more data centers.¹⁰⁷ Energy is such a critical choke point that, unless bottlenecks are addressed, AI companies may have to relocate AI infrastructure abroad, “potentially compromising the U.S. competitive advantage in compute and AI and increasing the risk of intellectual property theft.”¹⁰⁸ Compute’s openness is multifaceted. First there is the challenge of opening physical access. Chips and data centers are constrained by high costs, limited supply, exclusionary vendor relationships, and national export controls.¹⁰⁹ In some rare cases, only one company has access to core compute infrastructure: Google is the only entity that owns TPUs; others are forced to rent from it.¹¹⁰ For the many that cannot afford access to compute directly, they are reliant on renting compute—cloud services—from the handful of entities that own it.¹¹¹ Energy infrastructure is part of this physical access landscape. Some developers can only draw power from public grids at standard rates, while better-resourced actors can

105. CY MCGEADY, JOSEPH MAJKUT, BARATH HARITHAS & KARL SMITH, CTR. FOR STRATEGIC & INT’L STUD., *THE ELECTRICITY SUPPLY BOTTLENECK ON U.S. AI DOMINANCE 1–2* (2025), <https://www.csis.org/analysis/electricity-supply-bottleneck-us-ai-dominance> [<https://perma.cc/4FFA-EA2K>].

106. *Id.* at 2–3.

107. Exec. Order No. 14318, 90 Fed. Reg. 35385 (July 28, 2025).

108. KONSTANTIN F. PILZ, YUSUF MAHMOOD & LENNART HEIM, RAND, *AI’S POWER REQUIREMENTS UNDER EXPONENTIAL GROWTH: EXTRAPOLATING AI DATA CENTER POWER DEMAND AND ASSESSING ITS POTENTIAL IMPACT ON U.S. COMPETITIVENESS*, at v (2025).

109. See Dashveenjit Kaur, *2025’s AI Chip Wars: What Enterprise Leaders Learned About Supply Chain Reality*, AINews (Jan. 6, 2026), <https://www.artificialintelligence-news.com/news/ai-chip-shortage-enterprise-ctos-2025/> [<https://perma.cc/7DEB-7LEM>].

110. See Kenrick Cai & Krystal Hu, *Exclusive: OpenAI Taps Google in Unprecedented Cloud Deal Despite AI Rivalry, Sources Say*, REUTERS (June 10, 2025, at 12:22 CT), <https://www.reuters.com/business/retail-consumer/openai-taps-google-unprecedented-cloud-deal-despite-ai-rivalry-sources-say-2025-06-10/>.

111. See Cade Metz, Karen Weise & Mike Isaac, *Nvidia’s Big Tech Rivals Put Their Own A.I. Chips on the Table*, N.Y. TIMES (Jan. 29, 2024), <https://www.nytimes.com/2024/01/29/technology/ai-chips-nvidia-amazon-google-microsoft-meta.html> (identifying the tech companies large enough to invest in their own compute hardware); Erin Griffith, *The Desperate Hunt for the A.I. Boom’s Most Indispensable Prize*, N.Y. TIMES (Aug. 16, 2023), <https://www.nytimes.com/2023/08/16/technology/ai-gpu-chips-shortage.html>. (explaining that most companies rent compute power from cloud services to avoid building their own data centers).

negotiate bespoke utility agreements or build dedicated power generation that effectively excludes smaller competitors from reaching frontier-scale training capacity.¹¹²

Second, hardware architecture is usually closed: Designs for specialized chips like those from Google and Nvidia are proprietary, preventing independent replication or modification.¹¹³ Third, software stacks create lock-in: To use Nvidia’s market-leading hardware, developers are functionally required to use its proprietary CUDA programming interface, making it difficult to move to another platform when their systems are engineered around one.¹¹⁴ The RISC-V movement, which is trying to enable independent chip design by opening chip blueprints,¹¹⁵ can simultaneously challenge the duopoly that controls this hardware layer while building increased reliance on the centralized entities controlling compute’s software layer.¹¹⁶

Fourth, energy itself sits on an openness spectrum. At one end are transparent, public-grid-based systems with uniform rates and open interconnection rules; at the other are highly closed arrangements, including private power purchase agreements, grid-bypassing generation, and confidential utility deals that give a small number of actors privileged access to the electrical capacity required for frontier training. Including energy in the compute layer of this taxonomy makes clear that AI capacity depends not only on specialized hardware but also on the

112. Matt Perault, *Speed-to-Power: An Energy Policy Agenda for a Thriving AI Market*, ANDREESSEN HOROWITZ (Nov. 17, 2025), <https://a16z.com/speed-to-power-an-energy-policy-agenda-for-a-thriving-ai-market/> [<https://perma.cc/MA7E-RCZD>]. *But see* Tim Fernholz, *The White House Wants AI Companies to Cover Rate Hikes. Most Have Already Said They Would.*, TECHCRUNCH (Feb. 25, 2026, at 12:42 PT), <https://techcrunch.com/2026/02/25/the-white-house-wants-ai-companies-to-cover-rate-hikes-most-have-already-said-they-would/> [<https://perma.cc/K6RJ-99UQ>] (reporting that the White House is pushing major AI companies to honor their pledges to cover the power grid cost hikes arising from their energy demand to avoid spreading the cost to the wider consumer base).

113. *See* Widder, Whittaker & West, *Open (For Business)*, *supra* note 37, at 7.

114. *See* NVIDIA CORP., *CUDA COMPATIBILITY 3* (2025), https://docs.nvidia.com/deploy/pdf/CUDA_Compatibility.pdf [<https://perma.cc/H6FY-XRH4>] (explaining that use of CUDA requires an Nvidia driver). *But see* Emre Çitak, *Nvidia to Bring CUDA Platform Support to the RISC-V*, DATAANOMY MEDIA (July 21, 2025), <https://dataanomy.com/2025/07/21/nvidias-cuda-platform-now-officially-supports-risc-v-cpus/> [<https://perma.cc/J9JG-LB2Q>].

115. *See* RISC-V, *License*, GITHUB, <https://github.com/riscv/riscv-isa-manual/blob/main/LICENSE> [<https://perma.cc/D5JU-ZLBW>] (last visited Jan. 28, 2026) (released under a “Creative Commons Attribution 4.0 International License”); Che Pan & Brenda Goh, *Exclusive: China to Publish Policy to Boost RISC-V Chip Use Nationwide, Sources Say*, REUTERS (Mar. 4, 2025), <https://www.reuters.com/technology/china-publish-policy-boost-risc-v-chip-use-nationwide-sources-2025-03-04/> (reporting on China’s support for the RISC-V movement, which is attempting to enable independent chip design by opening chip blueprints).

116. *See* Çitak, *supra* note 114.

accessibility, reliability, and distribution of the electricity system that enables large-scale model training and deployment.

Fully open compute infrastructure would entail something akin to a public option: universally accessible hardware components with low to no barriers to entry.¹¹⁷ However, this reality is unlikely to manifest. And while some initiatives—like decentralized compute networks¹¹⁸ or research cloud credits¹¹⁹—aim to expand access, these are generally partial, selective, and funnel market power back to incumbents. So, in most cases, only those with significant capital can afford the high-end GPUs or cloud services necessary to meaningfully experiment with large-scale AI today.¹²⁰

In practice, compute arrangements span a wide range between these endpoints. Access to high-end Nvidia GPUs, for example, is neither fully open nor fully closed: While demand far outpaces supply and allocation is tightly controlled,¹²¹ much of this hardware is made available through cloud or “neocloud” providers¹²² that compete to resell compute capacity

117. See Eleanor Shearer, Matt Davies & Mathew Lawrence, *The Role of Public Compute*, ADA LOVELACE INST. (Apr. 24, 2024), <https://www.adalovelaceinstitute.org/blog/the-role-of-public-compute/> [<https://perma.cc/QS9Q-MRWR>].

118. See Will Knight, *These Startups Are Building Advanced AI Models Without Data Centers*, WIRED (Apr. 30, 2025, at 12:00 CT), <https://www.wired.com/story/these-startups-are-building-advanced-ai-models-over-the-internet-with-untapped-data/> [<https://perma.cc/7JTW-FXP8>] (explaining that while startups are exploring decentralized frontier model development, most AI companies require “huge quantities of compute concentrated inside data centers stuffed with advanced GPUs”).

119. See *Apply for Google Cloud Research Credits*, GOOGLE CLOUD, https://edu.google.com/intl/ALL_us/programs/credits/research/ [<https://perma.cc/7FN7-USMV>] (last visited Jan. 28, 2026); *AWS Cloud Credit for Research*, AMAZON WEB SERVS., <https://aws.amazon.com/government-education/research-and-technical-computing/cloud-credit-for-research/> [<https://perma.cc/K8YP-T2TV>] (last visited Jan. 28, 2026).

120. See JAI VIPRA & SARAH MYERS WEST, AI NOW INST., *COMPUTATIONAL POWER AND AI 8* (2023), https://ainowinstitute.org/wp-content/uploads/2023/09/AI-Now_Computational-Power-an-AI.pdf [<https://perma.cc/MYY2-C8LL>] (explaining that “[c]ompute is scarce” and therefore a bottleneck for AI development, which amplifies the market power of the few companies providing it).

121. See Stephen Nellis & Aditya Soni, *Nvidia’s Supply Snags Limit Deliveries Even as Demand Booms*, REUTERS (Nov. 21, 2024), <https://www.reuters.com/technology/nvidias-supply-snags-hurting-deliveries-mask-booming-demand-2024-11-21/>.

122. See, e.g., Channy Yun (윤석찬), *New – Amazon EC2 P5 Instances Powered by NVIDIA H100 Tensor Core GPUs for Accelerating Generative AI and HPC Applications*, AMAZON WEB SERVS.: NEWS BLOG (July 26, 2023), <https://aws.amazon.com/blogs/aws/new-amazon-ec2-p5-instances-powered-by-nvidia-h100-tensor-core-gpus-for-accelerating-generative-ai-and-hpc-applications/> [<https://perma.cc/N3EZ-UATG>]; Ishan Sharma & Tanvi Srivastava, *Announcing Smaller Machine Types for A3 High VMs*, GOOGLE CLOUD: BLOG (Jan. 24, 2025), <https://cloud.google.com/blog/products/compute/announcing-smaller-machine-types-for-a3-high-vm> [<https://perma.cc/SV8Y-F6UB>]; Nidhi Chappell & Eric Boyd, *Scale Generative AI with New Azure AI*

to downstream users.¹²³ This model separates ownership from use: A small number of actors control the hardware itself, while a broader set of developers can access compute indirectly subject to pricing, prioritization, and contractual limits set by intermediaries. These intermediate arrangements underscore that compute openness is often partial, mediated, and contingent—varying across actors, uses, and stages of the AI lifecycle.

2. Data

Data is the fuel that powers AI and is among the most contested components in the ecosystem. The capabilities of traditional software are defined by its source code; AI, however, “learns” from extensive datasets¹²⁴ that range from raw, unstructured data (*e.g.*, social media posts or biotech sensor readings) to carefully curated training data (*e.g.*, labeled images) to highly specialized datasets for refining (or fine-tuning) model performance in particular domains or for specific tasks (*e.g.*, medical diagnoses or legal document analysis).¹²⁵

Privacy concerns, copyright laws, and competitive advantages all determine what data is shared, who can use it, and under what terms. In this way, they also dictate how open the models built on them can truly be.¹²⁶ Consequently, AI training, fine-tuning, and testing data is far more contested and controlled than traditional source code, adding yet another

Infrastructure Advancements and Availability, MICROSOFT AZURE: BLOG (Aug. 7, 2023), <https://azure.microsoft.com/en-us/blog/scale-generative-ai-with-new-azure-ai-infrastructure-advancements-and-availability/> [<https://perma.cc/PY7F-JFPG>].

123. Krystal Hu & Kenrick Cai, *CoreWeave to Offer Compute Capacity in Google's New Cloud Deal with OpenAI, Sources Say*, REUTERS (June 11, 2025), <https://www.reuters.com/business/coreweave-offer-compute-capacity-googles-new-cloud-deal-with-openai-sources-say-2025-06-11/>; Yawen Chen, *Neoclouds' Fine Print Is a Silver Lining of Sorts*, REUTERS (Oct. 22, 2025, at 00:00 CT), <https://www.reuters.com/commentary/breakingviews/neoclouds-fine-print-is-silver-lining-sorts-2025-10-22/>.

124. See Christopher S. Yoo, *Beyond Algorithmic Disclosure For AI*, 25 COLUM. SCI. & TECH. L. REV. 314, 318–21 (2024).

125. See Jenny Quang, *Does Training AI Violate Copyright Law?*, 36 BERKELEY TECH. L.J. 1407, 1411 (2021).

126. See Yoo, *supra* note 124, at 321–24; see also Katie Knibbs, *Meta Secretly Trained Its AI on a Notorious Piracy Database, Newly Unredacted Court Docs Reveal*, WIRED (Jan. 9, 2025, at 17:33 CT), <https://www.wired.com/story/new-documents-unredacted-meta-copyright-ai-lawsuit/> [<https://perma.cc/EC64-6AYK>] (reporting that unredacted court documents revealed Meta secretly used a pirated book database to train its AI models, despite employees acknowledging that the material was copyrighted); Tifani Sadek et al., *Artificial Intelligence Impacts on Privacy Law*, RAND (Aug. 8, 2024), https://www.rand.org/pubs/research_reports/RRA3243-2.html [<https://perma.cc/MCV7-VBAV>] (examining how AI's reliance on large-scale data collection conflicts with existing privacy law frameworks, including notice-and-consent requirements and data minimization principles).

layer of complexity to AI's differential openness.¹²⁷ For example, OpenAI is the poster child for the liability exposure that emerges when datasets are visible, facing a slew of lawsuits accusing it of training models on hordes of copyrighted material.¹²⁸ Even models known for their openness, such as Mistral 7B, still withhold access to, or even information about, datasets, citing competitive pressures.¹²⁹

Some datasets, such as Common Crawl—a massive, publicly archived crawl of the web—are fully open, allowing unrestricted access to raw training materials.¹³⁰ Others are partially open, meaning they are available under certain conditions, such as being licensed for research use, but restricted for commercial applications.¹³¹ Many datasets, however, remain completely closed, either through proprietary licenses or strict contractual agreements to protect privacy, competitive advantage, or intellectual property rights.¹³² Proprietary medical, financial, or corporate datasets, for example, are often off-limits to all but the companies that own them.¹³³ Use of this data, no matter how open or closed, imports the risks embedded in them.

Importantly, true data openness is about more than just access to a dataset; it is also determined by the transparency of its curation.¹³⁴ Meaningful access is often dependent on the upstream models, such as smaller neural networks, that are used for filtering, classification, or

127. See Mehtab Khan & Alex Hanna, *The Subjects and Stages of AI Dataset Development: A Framework For Dataset Accountability*, 19 OHIO ST. TECH. L.J. 171, 179 (2023) (“The vast majority of the training data from [frontier] models are private.”); Roose, *supra* note 79.

128. See Kyle Jahner, *OpenAI Sued by New Set of Authors Over Training Data Copyrights*, BLOOMBERG L. (July 2, 2025, at 13:52 CT), <https://news.bloomberglaw.com/ip-law/openai-sued-by-new-set-of-authors-over-training-data-copyrights>.

129. See Arthur Mensch (@arthurmensch), HUGGING FACE (Oct. 12, 2023), <https://huggingface.co/mistralai/Mistral-7B-v0.1/discussions/8> [<https://perma.cc/X2E5-D56D>] (“Unfortunately we’re unable to share details about the training and the datasets . . . due to the highly competitive nature of the field.”).

130. See *Our Mission*, COMMON CRAWL, <https://commoncrawl.org/mission> [<https://perma.cc/6MN2-8SR8>] (last visited Jan. 28, 2026) (“Small startups or even individuals can now access high quality crawl data that was previously only available to large search engine corporations.”).

131. For example, ImageNet, a large database of labeled images, is freely available for non-commercial use. IMAGENET (Mar. 11, 2021), <https://www.image-net.org/> [<https://perma.cc/6Q76-BVUX>].

132. See Fodouop, *supra* note 78, at 15; Sydney Rouser, *Unfair Competition in the Creative Industries: The Impact of AI Scraping*, 16 TENN. J.L. & POL’Y 134, 144 (2024).

133. See, e.g., Isabelle Rose I. Alberto et al., *The Impact of Commercial Health Datasets on Medical Research and Health-Care Algorithms*, 5 LANCET DIGIT. HEALTH e288, e288 (2023).

134. See Widder, Whittaker & West, *Open (For Business)*, *supra* note 37, at 9.

ranking—models that are themselves often neither open nor auditable.¹³⁵ This upstream opacity means that even a notionally “open” dataset may be shaped by hidden selection biases, creating a black box at the very start of the AI pipeline.¹³⁶

3. Source Code

While not the sole determinant of functionality, source code remains a critical component of the AI stack and a factor in osAI’s differential openness. First, “inference code” shapes the potential capabilities and functions the system can perform by defining its architecture, or the structure of how the model processes input data into predictions.¹³⁷ Second, “training code” details how the model learns from its training data.¹³⁸ Openness in inference code facilitates visibility and allows stakeholders to understand and assess a model’s theoretical capabilities, while openness in training code enables replication, verification, and potential improvement of the original results.¹³⁹ Together, inference and training code define the model’s architecture, shaping how training data becomes model weights and how those weights produce outputs.

Openness of AI source code generally aligns with the established spectrum of OSS licenses employing permissive licenses, such as MIT, for maximal accessibility,¹⁴⁰ or copyleft licenses, like GPL, to ensure ongoing openness of derivative works.¹⁴¹ But many leading systems—such as OpenAI’s GPT series, Anthropic’s Claude, and Google’s Gemini—keep source code entirely proprietary, preventing external

135. See, e.g., Catherine Arnett, Eliot Jones, Ivan P. Yamshchikov & Pierre-Carl Langlais, *Toxicity of the Commons: Curating Open-Source Pre-Training Data* (Nov. 18, 2024) (unpublished manuscript), <https://arxiv.org/pdf/2410.22587> [<https://perma.cc/5RWK-M897>] (example of open classifier for toxicity filtering).

136. See STEFAN BAACK ET AL., *TOWARDS BEST PRACTICES FOR OPEN DATASETS FOR LLM TRAINING* 8–9 (2025), <https://arxiv.org/pdf/2501.08365> [<https://perma.cc/VG3J-WGH6>].

137. See *Version 1.0*, OPEN SOURCE INITIATIVE, <https://opensource.org/ai/open-source-ai-definition> [<https://perma.cc/9S5Z-XKHD>] (last visited Jan. 28, 2026) (distinguishing code used to guide training from code used for model architecture).

138. See *id.*

139. See Andrew D. Mitchell, Dominic Let & Lingxi Tang, *AI Regulation and the Protection of Source Code*, 31 INT’L J.L. & INFO. TECH. 283, 286–87, 291–92 (2023).

140. See, e.g., DeepSeek, *MIT License*, HUGGING FACE, <https://huggingface.co/deepseek-ai/DeepSeek-R1/blob/main/LICENSE> [<https://perma.cc/5J8R-FARL>] (last visited Jan. 28, 2026).

141. See, e.g., *GNU Affero General Public License v3.0*, HUGGING FACE (Nov. 19, 2007), <https://huggingface.co/datasets/choosealicense/licenses/blob/main/markdown/agpl-3.0.md> [<https://perma.cc/22X3-QM2H>].

developers and researchers from inspecting or experimenting on the model architecture's inner workings.¹⁴²

4. Model Weights

But even complete openness in source code—covering both architecture and training scripts—does not fully define or predict an AI system's behavior. That largely depends on the next crucial piece: model weights. An AI system emerges from training with a static “model” that determines *how* the model makes predictions based on inputs. Model weights—billions (and soon-to-be trillions) of numerical parameters—store this learning as compressed knowledge.¹⁴³ They shape everything from a model's writing style to its ability to recognize images. While AI source code establishes model architecture—like a building's blueprint—it is the model's weights that dictate what the model knows and how well it performs, much like furniture dictates how a building is actually used.

Model weight openness exists along a spectrum of licenses, similar to source code, that either enables or restricts transparency, experimentation, and reuse.¹⁴⁴ At one extreme are fully closed models, like Anthropic's Claude and OpenAI's ChatGPT.¹⁴⁵ These models do not release model weights and only allow access to chat interfaces or application programming interfaces (APIs)—restrictive protocols for machine-to-machine communication. Desires to preserve a competitive edge or curb misuse risks end up preventing independent scrutiny or external modification.¹⁴⁶ Some companies, like OpenAI, make their model weights accessible to select research institutes, such as the UK AI

142. Parth Nobel, Alan Z. Rozenshtein & Chinmayi Sharma, *Open-Access AI: Lessons from Open-Source Software*, LAWFARE (Oct. 25, 2024, at 10:00 CT), <https://www.lawfaremedia.org/article/open-access-ai--lessons-from-open-source-software> [https://perma.cc/9NL4-GZ2E].

143. Giorgio Franceschelli, Claudia Cevenini & Mirco Musolesi, *Training Foundation Models as Data Compression: On Information, Model Weights and Copyright Law 1* (Mar. 12, 2025) (unpublished manuscript), <https://arxiv.org/pdf/2407.13493> [https://perma.cc/2WE5-SA8V].

144. See Eiras et al., *supra* note 32, at 2–3, 18; NAT'L TELECOMMS. & INFO. ADMIN., *supra* note 25, at 8–9.

145. Sharon Goldman, *Why Anthropic and OpenAI Are Obsessed with Securing LLM Model Weights*, VENTUREBEAT (Dec. 15, 2023), <https://venturebeat.com/ai/why-anthropic-and-openai-are-obsessed-with-securing-llm-model-weights/> [https://perma.cc/VN75-3CXX].

146. See, e.g., SEGER et al., *supra* note 32, at 5; *Reasoning Models*, OPENAI, <https://platform.openai.com/docs/guides/reasoning> (last visited Aug. 1, 2025) (“[W]e don't expose the raw reasoning tokens emitted by the model”). OpenAI also generally lacks information on accessing weights.

Safety Institute, but even this accessibility is strictly controlled.¹⁴⁷ On the other end are models, like DeepSeek’s R1, that release weights under copyleft or permissive licenses.¹⁴⁸ In between are models, such as Meta’s Llama, that release weights but limit their use.¹⁴⁹

The release of model weights under permissive licenses is a necessary condition for a model to be truly open, but not a sufficient one. Many of the largest models labeled “open source” are merely open-weight, and often minimally so. The most notable example is Meta’s Llama, the popular “open source” American model¹⁵⁰—though it would be more accurate to call it “non-proprietary” for two reasons: (1) Many of its components, like training data and training code, are not public; and (2) openly released components, like model weights, come under a highly restrictive license—neither copyleft nor permissive.¹⁵¹ In particular, Meta restricts use of Llama by requiring companies above a defined scale threshold—those with more than 700 million monthly users—to obtain a bespoke license, effectively limiting who can deploy the model at scale.¹⁵² Meta further limits downstream deployment through an acceptable use policy that enumerates prohibited use cases and fields, constraining what users can do with the released weights in practice.¹⁵³ These constraints sharply limit the use, redistribution, and innovation that genuine openness is meant to support.

147. See *Pre-Deployment Evaluation of OpenAI’s o1 Model*, AI SEC. INST. (Dec. 18, 2024), <https://www.aisi.gov.uk/work/pre-deployment-evaluation-of-openais-o1-model> [<https://perma.cc/SXD2-VJQ8>]; see also OPENAI, GPT-4O SYSTEM CARD 3–4 (2024), <https://cdn.openai.com/gpt-4o-system-card.pdf> [<https://perma.cc/TG4V-MDCB>] (discussing that “OpenAI worked with more than 100 external red teamers, speaking a total of 45 different languages, and representing geographic backgrounds of 29 different countries” (footnote omitted)).

148. DeepSeek, *DeepSeek-R1*, HUGGING FACE, <https://huggingface.co/deepseek-ai/DeepSeek-R1> [<https://perma.cc/8Z26-2FUV>] (last visited Jan. 28, 2026) (“DeepSeek-R1 series support commercial use, allow for any modifications and derivative works, including, but not limited to, distillation for training other LLMs.”).

149. See Widder, Whittaker & West, *Why ‘Open’ AI Systems Are Actually Closed*, *supra* note 32, at 828–29.

150. See Kyle Wiggers, *Meta Says Its Llama AI Models Have Been Downloaded 1.2B Times*, TECHCRUNCH (Apr. 29, 2025, at 10:41 PT), <https://techcrunch.com/2025/04/29/meta-says-its-llama-ai-models-have-been-downloaded-1-2b-times/> [<https://perma.cc/D8J7-9MU3>].

151. See Michael Nolan, *Llama and ChatGPT Are Not Open-Source*, IEEE SPECTRUM (July 27, 2023), <https://spectrum.ieee.org/open-source-llm-not-open> [<https://perma.cc/L59W-SCM3>]; Nobel, Rozenshtein & Sharma, *supra* note 142.

152. Meta Llama, *Llama 3.1 Community License Agreement*, HUGGING FACE (July 23, 2024), <https://huggingface.co/meta-llama/Llama-3.1-8B/blob/main/LICENSE> [<https://perma.cc/H4VD-TQXD>].

153. Meta Llama, *Llama 3.1 Acceptable Use Policy*, <https://huggingface.co/meta-llama/Llama-3.1-8B/blob/main/README.md> [<https://perma.cc/3ZPG-CUHL>] (last visited Jan. 28, 2026); Stefano Maffulli, *Meta’s LLaMa License Is*

This exploitation of differential openness has been termed “openwashing,”¹⁵⁴ which risks misleading the public and policymakers into believing these systems are more open than they are, while allowing companies to benefit from a reputational boost and, in some cases, regulatory benefits of choosing “openness.” For example, the EU AI Act privileges certain models that make model weights open under vaguely defined open licenses without requiring disclosure of critical elements like training data or fine-tuning methods.¹⁵⁵

5. System Prompts

Another increasingly important layer in deployed AI systems is system prompts: the foundational instructions or configuration strings that set the behavior of the model at runtime. Separate from a model’s learned knowledge, and unbeknownst to users, system prompts govern tone, style, boundaries, and behavioral defaults, significantly shaping outputs without changing the underlying model weights or architecture.¹⁵⁶ It accomplishes this by appending content to the user’s own input before submitting the query to the model—sometimes, it reframes the query and other times, it overrides the user’s explicit request.¹⁵⁷ For instance, a system prompt might instruct a model: “You are a professional research assistant. Your tone must be neutral and objective. You must refuse any request for personal opinions or political commentary.”

Because these prompts are often crafted through iterative experimentation and internal alignment processes, they can lead to unintended, sometimes downright abhorrent, outcomes. Google’s misguided system prompts in an early version of Gemini led to the widely derided generation of troubling, historically inaccurate images—such as multiracial Nazis.¹⁵⁸ And when Grok’s system prompts were updated to

Not Open Source, OPEN SOURCE INITIATIVE (July 20, 2023), <https://opensource.org/blog/metals-llama-2-license-is-not-open-source> [<https://perma.cc/2X2K-45ER>].

154. See Sarah Kessler, *Openwashing*, N.Y. TIMES (May 17, 2024), <https://www.nytimes.com/2024/05/17/business/what-is-openwashing-ai.html>; see also Liesenfeld & Dingemans, *supra* note 32, at 1774, 1778 (“Our survey yields 40 text generators that are described as ‘open source’ or ‘open.’ . . . We also find a large number of systems (roughly the bottom third) that make only model weights available but share little to no detail about other parts of their system.”).

155. See Liesenfeld & Dingemans, *supra* note 32, at 1774; *supra* note 28 and accompanying text.

156. See Anna Neumann, Elisabeth Kirsten, Muhammed Bilal Zafar & Jatinder Singh, *Position Is Power: System Prompts as a Mechanism of Bias in Large Language Models (LLMs)*, 2025 PROCS. ACM CONF. ON FAIRNESS, ACCOUNTABILITY, & TRANSPARENCY 573, 573–75, <https://doi.org/10.1145/3715275.3732038>.

157. *Id.* at 573–74.

158. See James Grimmelmann, Blake E. Reid & Alan Z. Rozenshtein, *Generative Baseline Hell and the Regulation of Machine-Learning Foundation Models*,

“not shy away from making claims which are politically incorrect, as long as they are well substantiated,” it dubbed itself “MechaHitler” within days.¹⁵⁹

Openness of system prompts also exists on a spectrum. Some models, such as Anthropic’s Claude and xAI’s Grok, publish their system prompts, allowing users and the public to understand, copy, and modify the behavioral guardrails guiding the model.¹⁶⁰ Others, including many commercial offerings like OpenAI’s GPT models and Google’s Gemini, treat system prompts as proprietary—hiding them from public view to retain control, limit gaming, or obscure alignment choices.¹⁶¹ While seemingly minor, these hidden prompts play an outsized role in shaping downstream applications and safety, making their disclosure an increasingly relevant axis of differential openness.

6. Operational Control and Records

Beyond the model itself, an AI system’s behavior is also influenced by operational controls, which are layers on top of the model that further enhance how a model behaves in the real world, and operational data, which documents everything from the system’s development process to its real world interactions.

Operational controls facilitate model behavior as they transition from development to real-world deployment.¹⁶² They include content filters, moderation tools, and risk management layers.¹⁶³ Their

LAWFARE (May 8, 2024, at 11:09 CT), <https://www.lawfaremedia.org/article/generative-baseline-hell-and-the-regulation-of-machine-learning-foundation-models> [<https://perma.cc/ZD4U-5WXL>].

159. Lisa Hagen, Huo Jingnan & Audrey Nguyen, *Elon Musk’s AI Chatbot, Grok, Started Calling Itself ‘MechaHitler,’* NPR (July 9, 2025, at 15:12 ET), <https://www.npr.org/2025/07/09/nx-s1-5462609/grok-elon-musk-antisemitic-racist-content> [<https://perma.cc/3B8U-X5KL>].

160. See, e.g., *System Prompts*, CLAUDE API DOCS, <https://platform.claude.com/docs/en/release-notes/system-prompts> [<https://perma.cc/7UCR-N29V>] (last visited Jan. 29, 2026); xai-org, *Grok Prompts*, GITHUB, <https://github.com/xai-org/grok-prompts> [<https://perma.cc/YCY9-G9SW>] (last visited Jan. 29, 2026).

161. See Nobel, Rozenshtein & Sharma, *supra* note 142.

162. See Rosario Cammarota et al., *Trustworthy AI Inference Systems: An Industry Research View* 1–2, 8 (Feb. 10, 2023) (unpublished manuscript), <https://arxiv.org/pdf/2008.04449> [<https://perma.cc/RF8Z-YKRK>].

163. Daniel Maggen, *Predict and Suspect: The Emergence of Artificial Legal Meaning*, 23 N.C. J.L. & TECH. 67, 98 (2021) (discussing how various algorithms in AI systems take on “triage responsibilities” such as processing, classifying, and filtering information); Pranav Gade, Simon Lermen, Charlie Rogers-Smith & Jeffrey Ladish, *Cheaply Removing Safety Fine-Tuning from Llama 2-Chat 13B*, at 1 (May 28, 2024) (unpublished manuscript), <https://arxiv.org/pdf/2311.00117> [<https://perma.cc/6BV4-ZNYT>]; *Our Approach to AI Safety*, OPENAI (Apr. 5, 2023), <https://openai.com/index/our-approach-to-ai-safety/>.

importance and specific design are highly dependent on the AI model's intended use. For instance, a chatbot designed for medical diagnosis will incorporate strict safety protocols to prevent faulty information, whereas one for casual companionship may employ less rigorous safeguards.

Complementing these active controls are operational records—model cards, data cards, technical reports, and system design papers—that provide behind-the-scenes static documentation of how an AI system is built, trained, tested, and deployed.¹⁶⁴ These records do not directly interact with AI models but are instead references for internal and external stakeholders to guide model iteration. When openly available, they facilitate third party experimentation, since AI systems are too complex to understand by reading source code alone. Even with traditional OSS, the accessibility of a system for inspection, experimentation, and reuse relies on effective documentation.¹⁶⁵ Similarly, operational records also allow external researchers, policymakers, and users to better understand and evaluate the AI system's intentions, limitations, and risks.¹⁶⁶ Conversely, vague or missing documentation impairs external oversight and makes it harder to diagnose or prevent harms such as bias or flaws.

Beyond such static documentation, operational records include dynamic metadata generated with every system interaction. This data includes monitoring logs, audit trails, and performance metrics—elements that reveal how an AI model behaves in the wild.¹⁶⁷ They are essential for accountability, risk detection, and safety improvement based on insight from real world interactions; pre-deployment testing can only go so far. Yet, despite their value, this type of metadata is rarely, if ever,

164. KASIA CHMIELINSKI ET AL., THE CLEAR DOCUMENTATION FRAMEWORK FOR AI TRANSPARENCY: RECOMMENDATIONS FOR PRACTITIONERS AND CONTEXT FOR POLICYMAKERS 4, 8–9, 25–26 (2024), https://shorensteincenter.org/wp-content/uploads/2024/05/CleAR_KChmielinski_FINAL.pdf [<https://perma.cc/2ANJ-FYAE>].

165. See TOPELSON RITVO, HESSEKIEL & BAVITZ, *supra* note 72, at 22 (“Though it is sometimes overlooked, the history of the open source movement shows us that the projects that defined their corporate structure and governance practices early and concretely set themselves up for success.”).

166. See, e.g., CHMIELINSKI ET AL., *supra* note 164, at 2 (“Documentation of [datasets, models, and AI systems] is crucial and serves several purposes, including: (1) Supporting responsible development and use, as well as mitigation of downstream harms, by providing transparency into the design, attributes, intended use, and shortcomings of datasets, models, and AI systems; (2) Motivating dataset, model, or AI system creators and curators to reflect on the choices they make; and (3) Facilitating dataset, model, and AI system evaluation and auditing.”).

167. See Dominik Kreuzberger, Niklas Kühl & Sebastian Hirschl, *Machine Learning Operations (MLOps): Overview, Definition, and Architecture*, 11 IEEE ACCESS 31866, 31870–71 (2023), <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=10081336>.

included in traditional software openness frameworks.¹⁶⁸ Without transparency in this area, even an osAI system that is fully open in terms of code and weights can still be a black box when it comes to real-world deployment and actual user experience.¹⁶⁹

7. The Application Layer

If model weights are the engine, the application layer is the vehicle that lets users drive.¹⁷⁰ This is the layer where humans actually encounter AI systems through chat interfaces, mobile apps, voice assistants, or other tools.¹⁷¹ It includes user-facing and orchestration software, which together handle the mechanics of interaction and control: transforming raw model calls into coherent conversations, managing context across engagements, and connecting the model to external data sources and tools.¹⁷²

In practice, the application layer includes several distinct components. User interfaces and interaction managers govern how prompts are collected, how outputs are displayed, and which controls (such as “retry,” “stop,” or “feedback” buttons) users have at their disposal.¹⁷³ Retrieval and memory systems supply the model with

168. The Open Source definition makes no mention of usage records. *See* OPEN SOURCE INITIATIVE, *supra* note 137. While some OSS that collects telemetry publishes it, it is generally “anonymized and aggregated to ensure user privacy” and not part of the software release. *Firefox Public Data Report*, FIREFOX, <https://data.firefox.com/> [<https://perma.cc/NSK5-E52W>] (last visited Jan. 29, 2026).

169. *See* Lisa Lee, *What Is Metadata in AI?*, SALESFORCE, <https://www.salesforce.com/blog/what-is-metadata/> [<https://perma.cc/PQU8-N6U3>] (last visited Jan. 29, 2026); Everton Gomedé, *The Importance of Metrics and Metadata in Data Observability*, MEDIUM (Mar. 6, 2024), <https://pub.aimind.so/the-importance-of-metrics-and-metadata-in-data-observability-f6d571fd2269>.

170. *See* Narechania & Sitaraman, *supra* note 32, at 126–27.

171. Mohammad (Matt) Namvarpour & Afsaneh Razi, *The Art of Talking Machines: A Comprehensive Literature Review of Conversational User Interfaces*, 7 CUI '25: PROCS. ACM CONF. ON CONVERSATIONAL USER INTERFACES 1, 1–2, <https://doi.org/10.1145/3719160.3736621>.

172. *See* INFO. TECH. INDUS. COUNCIL, *THE AI TECHNOLOGY STACK AND WHY IT MATTERS FOR AI POLICY AND GOVERNANCE 7–8* (2025), https://www.itic.org/documents/artificial-intelligence/ITI_AITechnologyStack.pdf [<https://perma.cc/7D3Q-J6F4>]; Vanna Winland & Joshua Noble, *What Is LLM Orchestration?*, IBM, <https://www.ibm.com/think/topics/llm-orchestration> [<https://perma.cc/ZKF5-NLDN>] (last visited Jan. 29, 2026); ELLIOT JONES, MAHI HARDALUPAS & WILLIAM AGNEW, ADA LOVELACE INST., *UNDER THE RADAR?: EXAMINING THE EVALUATION OF FOUNDATION MODELS* 16 fig. 1 (2024), <https://www.adalovelaceinstitute.org/wp-content/uploads/2024/09/Ada-Lovelace-Institute-Under-the-radar-230924.pdf> [<https://perma.cc/PD8Y-2SQR>].

173. *See* Sunil Ramlochan, *Beyond the Bot—Why ChatGPT’s Interface Was the Real Innovation*, PROMPT ENG’G (Sep. 3, 2023), <https://promptengineering.org/beyond-the-bot-why-chatgpts-interface-was-the-real-innovation/> [<https://perma.cc/XL2F-D694>].

external context, such as documents the user uploads, proprietary databases, or up-to-date web information, and can store conversational history or user profiles for later use.¹⁷⁴ Tool frameworks and plugin systems allow the model to call external services, execute code, or perform actions in the world,¹⁷⁵ while identity and authorization layers determine who can access which capabilities.¹⁷⁶ Finally, safety middleware and telemetry monitor and filter inputs and outputs, log interactions, and support debugging and oversight.¹⁷⁷ Two systems that share identical weights can behave very differently in the world simply because they are wrapped in different application layers.

Openness at the application layer exists on a spectrum and often diverges from openness at the model level. At one end are tightly controlled, vertically integrated products in which the provider controls the interface, plugin ecosystem, safety middleware, and telemetry, and exposes only a narrow chat user interface (UI) or API.¹⁷⁸ In the middle are hybrid approaches: for example, proprietary models that allow third-party open source clients, open standards for plugins on otherwise closed platforms, or transparent but non-modifiable safety pipelines.¹⁷⁹ At the most open end are application stacks whose code, configuration, and orchestration logic are available for inspection, modification, and self-hosting, so that different actors can build their own interfaces, safety

174. *What Is Retrieval Augmented Generation (RAG)?*, IBM, <https://www.ibm.com/think/topics/retrieval-augmented-generation> [https://perma.cc/YUE2-GK5U] (last visited Jan. 29, 2026).

175. Isobel Moure, Tim O'Reilly & Ilan Strauss, *Open Protocols Can Prevent AI Monopolies*, AI FRONTIERS (July 30, 2025), <https://ai-frontiers.org/articles/open-protocols-prevent-ai-monopolies> [https://perma.cc/535X-V8G3].

176. Wenxin Du & Jessica Chan, *How to Build User Authentication into Your Gen AI App—Accessing a Database*, GOOGLE CLOUD: BLOG (July 19, 2024), <https://cloud.google.com/blog/products/ai-machine-learning/build-user-authentication-into-your-genai-app-accessing-database> [https://perma.cc/5BWP-JUGF].

177. Muqsit Azeem et al., *Monitizer: Automating Design and Evaluation of Neural Network Monitors*, 36 INT'L CONF. ON COMPUT. AIDED VERIFICATION, PROCS., PART II, 265, 265–66 (2024), <https://link.springer.com/book/10.1007/978-3-031-65630-9>; *LLM Security & Guardrails*, LANGFUSE, <https://langfuse.com/docs/security-and-guardrails> [https://perma.cc/B9RR-BMNC] (last visited Jan. 30, 2026).

178. *See, e.g.*, Andreas Liesenfeld, Alianda Lopez & Mark Dingemans, *Opening up ChatGPT: Tracking Openness, Transparency, and Accountability in Instruction-Tuned Text Generators*, 5 PROCS. INT'L CONF. ON CONVERSATIONAL USER INTERFACES (CUI '23) (2023), <https://arxiv.org/pdf/2307.05532> [https://perma.cc/RF26-HJRH].

179. *See, e.g.*, Moure, O'Reilly & Strauss, *supra* note 175 (advocating for open protocols for AI models); Sam Adler, *Interoperable Agentic AI: Unlocking the Full Potential of AI Specialization*, TECH POL'Y PRESS (Dec. 3, 2024), <https://www.techpolicy.press/interoperable-agentic-ai-unlocking-the-full-potential-of-ai-specialization/> [https://perma.cc/E9HW-N9FT] (same).

wrappers, and integrations on top of shared models.¹⁸⁰ Where a system falls along this spectrum determines who can see how it works, who can adapt it to new domains, and who can constrain or extend its behavior.

Because it controls how capabilities are packaged, the relative openness of the application layer has significant impact. From a safety perspective, it is often cheaper and more effective to add guardrails, usage constraints, and contextual warnings in the application than to retrain or redesign the underlying model, but those same controls can be used to silently suppress information exposure and permissible speech. From a competition perspective, open application-layer code and plugin ecosystems can lower entry barriers and allow third parties to differentiate on user experience, even when models and compute remain concentrated, while highly closed interfaces can turn model providers into gatekeepers that decide which uses and business models are allowed.¹⁸¹ In terms of accessibility and accountability, the application layer effectively determines how easy it is to deploy a system at scale, which populations it reaches, and how observable its behavior is to regulators and civil society.

8. The Human Layer

AI systems are not just technical artifacts—they are built by human hands and minds. The talent, expertise, and institutions that train and organize AI professionals are fundamental components of the AI stack and determinants of osAI’s differential openness.¹⁸² Like source code or training data, this human layer can be more or less open, and its degree of openness shapes innovation, concentration, and accountability across the ecosystem.

The state of the human layer is not merely a workforce issue—it is a governance mechanism. The answer to the question of who gets to contribute, switch jobs, start labs, or critique dominant approaches determines whose values are embedded in AI. In OSS, labor is relatively open: Contributors from anywhere, with any degree of formal training,

180. See, e.g., *Open WebUI*, GITHUB, <https://github.com/open-webui/open-webui> [<https://perma.cc/YC2S-TSFZ>] (last visited Jan. 30, 2026) (GitHub repository); JAN, <https://www.jan.ai/> [<https://perma.cc/TA88-QJ5Y>] (last visited Jan. 30, 2026); *WebLLM Chat*, GITHUB, <https://github.com/mlc-ai/web-llm-chat> [<https://perma.cc/QM8J-8UKP>] (last visited Jan. 30, 2026) (GitHub repository).

181. Thibault Schrepele & Alex ‘Sandy’ Pentland, *Competition Between AI Foundation Models: Dynamics and Policy Recommendations*, 34 INDUS. & CORP. CHANGE 1085, 1087, 1090, 1092–93 (2025).

182. Gordon Hanson, *Immigration and Regional Specialization in AI*, in ROBOTICS AND AI 180, 180–81 (Lili Yan Ing & Gene M. Grossman eds., 2022).

can submit code and build reputational capital.¹⁸³ But the AI ecosystem—where physical presence (especially in San Francisco and Silicon Valley) is still crucial¹⁸⁴—is characterized by three key forms of closure: (1) restrictive pipelines that limit access to a diverse talent pool, (2) institutional constraints that inhibit professional mobility, and (3) corporate controls that suppress the diffusion of critical knowledge.

First, access to the field is constrained at both the domestic and international levels. Domestically, STEM education and talent pipelines fail to produce enough qualified individuals to meet demand, and the domestic talent that exists is unevenly distributed by race, gender, geography, and institutional prestige.¹⁸⁵ This lack of diversity means that communities most impacted by AI often have no voice in its design. The ecosystem’s reliance on foreign talent faces even steeper barriers. As Nvidia CEO Jensen Huang has noted, half of the world’s top AI researchers are Chinese, highlighting the global nature of expertise.¹⁸⁶ However, unlike in OSS, where anyone can contribute via platforms like GitHub, meaningful participation in frontier AI development typically requires being hired by a dominant firm and moving to America. Consequently, as demand outpaces domestic supply, visa bottlenecks and restrictive immigration policies have become critical choke points,¹⁸⁷

183. See Stephanie Susnjara & Ian Smalley, *What Is Open Source Software?*, IBM, <https://www.ibm.com/think/topics/open-source> [<https://perma.cc/KCV4-FMKP>] (last visited Jan. 30, 2026).

184. See Kristina McElheran et al., *AI Adoption in America: Who, What, and Where*, 33 J. ECON. & MGMT. STRATEGY 375, 376 (2024).

185. See Darrell M. West, *Improving Workforce Development and STEM Education to Preserve America’s Innovation Edge*, BROOKINGS INST. (July 26, 2023), <https://www.brookings.edu/articles/improving-workforce-development-and-stem-education-to-preserve-americas-innovation-edge/> [<https://perma.cc/J7AM-V999>] (“[A]ccording to a Deloitte study, there are fewer than 100,000 U.S. graduates with electrical engineering and computer science degrees each year, which is below the number that will be required in the coming decade.”); NAT’L SCI. BD., *THE STEM LABOR FORCE: SCIENTISTS, ENGINEERS, AND SKILLED TECHNICAL WORKERS 16–19* (2024), <https://nces.nsf.gov/pubs/nsb20245/representation-of-demographic-groups-in-stem> [<https://perma.cc/7K7S-63RP>] (“[T]he proportion of men in STEM occupations remained higher than that of women in 2011, 2016, and 2021 STEM workers were disproportionately Asian and White”).

186. Tanya Rawat, *Nvidia CEO Jensen Huang Sounds the Alarm as 50% of AI Researchers are Chinese, Urges America to Reskill Amid ‘Infinite Game,’* YAHOO FIN. (May 1, 2025), <https://finance.yahoo.com/news/nvidia-ceo-jensen-huang-sounds-035916833.html>.

187. See Remco Zwetsloot et al., *Skilled and Mobile: Survey Evidence of AI Researchers’ Immigration Preferences*, 2021 AIES ‘21: PROCS. AAAI/ACM CONF. ON AI, ETHICS & SOC’Y 1050, 1052, <https://dl.acm.org/doi/epdf/10.1145/3461702.3462617> (“AI PhDs who chose to leave the U.S. were likely to cite immigration-related concerns (23%) and the U.S. immigration system (33%) as highly relevant”); *Sharifmoghammad v. Blinken*, No. 23-CV-1472, 2024 WL 939991, at *6 (D.D.C. Mar. 5, 2024) (acknowledging that “delays in visa processing [had] interfer[ed] with

failing to retain international students and attract foreign researchers, thus locking out the very talent America needs to lead.¹⁸⁸

Second, for those who do gain access, institutional constraints then limit mobility. Noncompete agreements have long blocked researchers from switching companies or launching startups,¹⁸⁹ locking expertise inside a handful of dominant firms and slowing the diffusion of knowledge.¹⁹⁰ This is not theory; it has been proven. Many credit Silicon Valley's meteoric growth to California's aggressive stance against enforcing noncompete clauses.¹⁹¹ By contrast, an open labor market empowers experts to join competitors or found new ventures that may better align with their ethical or scientific values.

Third, knowledge itself is trapped in institutional silos, weakening the broader diffusion of expertise. Restrictions on academic moonlighting,¹⁹² opaque clearance processes for publications,¹⁹³ and overreliance on proprietary data further trap knowledge within

[plaintiff's] career progression and ability to contribute to artificial intelligence research in the United States" but still denying visa application).

188. See *Technology, Trade, and Military-Civil Fusion: China's Pursuit of Artificial Intelligence, New Materials, and New Energy: Hearing Before the U.S.-China Economic & Security Review Commission*, 116th Cong. 34 (2019) [hereinafter Toner] (statement of Helen Toner, Ctr. for Sec. & Emerging Tech., Geo. Univ.) ("More than 85% of Chinese and Indian students in U.S. computer science and engineering PhD programs state that they intend to stay after graduation [but are often unable to.]").

189. Caitlin Harrington, *Innovation-Killing Noncompete Agreements Are Finally Dying*, WIRED (Dec. 4, 2023, at 07:00 CT), <https://www.wired.com/story/innovation-killing-noncompete-agreements-finally-dying/> ("35 percent of people working in computer- and math-related vocations . . . work under noncompetes, the highest share of workers in all industries . . .").

190. Aminu Abdullahi, *Some AI Talent 'In Despair' Reportedly About Google DeepMind's Noncompete Rules*, TECH REPUBLIC (Apr. 14, 2025), <https://www.techrepublic.com/article/news-deepmind-noncompete-clauses-ai-talent-wars> ("DeepMind . . . is using aggressive noncompete clauses Some senior-level researchers are subject to a full year of paid 'garden leave' . . .").

191. Mike McPhate, *California Today: Silicon Valley's Secret Sauce*, N.Y. TIMES (May 19, 2017), <https://www.nytimes.com/2017/05/19/us/california-today-silicon-valley.html>.

192. See e.g., *Research Policy Handbook 4.3: Consulting and Other Outside Professional Activities by Members of the Academic Council and University Medical Line Faculty*, STAN. UNIV.: DORESEARCH, <https://doresearch.stanford.edu/policies/research-policy-handbook/conflicts-commitment-and-interest/consulting-and-other-outside-professional-activities-members-academic-council-and-medical-center-line-faculty> [https://perma.cc/K9HY-PS3P] (last visited Jan. 30, 2026) ("The maximum number of Consulting days permissible for a member of the Academic Council or the University Medical Line Faculty on a full-time appointment is 13 days per academic quarter.").

193. Thomas Klebel et al., *Peer Review and Preprint Policies Are Unclear at Most Major Journals*, PLOS ONE, Oct. 21, 2020, at 1, 3, <https://journals.plos.org/plosone/article/file?id=10.1371/journal.pone.0239518&type=printable>.

institutional silos.¹⁹⁴ Experts within companies are on a short leash when it comes to academic contributions and outside academics are simultaneously locked out of the systems they seek to research. Adherence to open science principles—transparent peer review, public dissemination of results, and reproducibility requirements—is a direct countermeasure, enhancing labor openness by reducing barriers to entry for contributors outside the monoculture of elite and profit-driven institutions.¹⁹⁵

These layers of closure are starkly illustrated by the labor-intensive process of Reinforcement Learning from Human Feedback (RLHF). This work, essential for model alignment and safety, is performed by armies of human annotators—often low-paid contractors in non-English-speaking countries like Kenya, India, or the Philippines—who rate and label millions of model outputs, making the answer keys from which AI learns.¹⁹⁶ These workers are invisible in governance frameworks: Their contributions are essential, yet they have no voice in system design or deployment. Openness in this context would demand transparency about who performs this difficult work, ensure fair compensation, and mandate their inclusion in feedback and oversight processes. Treating these workers as integral contributors, not interchangeable cogs, strengthens both accountability and system integrity.

Ultimately, just as open data or public weights can democratize technology, so can open labor practices diffuse expertise, reduce capture, and foster accountability. Conversely, just as closed infrastructure or proprietary models entrench power, labor constraints can centralize control over the direction and pace of AI innovation. Differential openness for osAI must therefore treat labor not as a background condition but as a core, governable component of the AI stack.

Part I has provided the essential analytical toolkit of differential openness. By rejecting the flawed “open–closed” binary and untangling AI into its eight core components, we have seen why OSS governance models fall short: Every osAI system represents a unique configuration of component-level choices, each existing on its own spectrum of

194. Anat Lior, *Private and Academic AI Collaboration: Opportunities and Challenges to Open Science in the US*, J. OPEN ACCESS TO L., 2023, at 1, 3, <https://ojs.law.cornell.edu/index.php/joal/article/view/144/128> [<https://perma.cc/WEH4-3BW9>].

195. See Erin C. McKiernan et al., *Point of View: How Open Science Helps Researchers Succeed*, eLIFE SCIENCES, July 2016, at 1, 3–10, <https://doi.org/10.7554/eLife.16800>.

196. Sarah, *Reinforcement Learning from Human Feedback (RLHF): A Simple Explainer*, BLUE DOT IMPACT (May 15, 2025), <https://bluedot.org/blog/rlhf-explainer> [<https://perma.cc/FY5U-9EEU>]; Rina Chandran, Adam Smith & Mariejo Ramos, *AI Boom Is Dream and Nightmare for Workers in Global South*, CONTEXT (Mar. 14, 2023), <https://www.context.news/ai/ai-boom-is-dream-and-nightmare-for-workers-in-global-south> [<https://perma.cc/6HG3-LFL5>].

openness. With this granular understanding in place, Part II will apply this framework to the central challenge of osAI policy: promoting and navigating the trade-offs between safety, innovation, democratic control, and national security.

II. THE VALUE OF AI OPENNESS

The concept of openness, as inherited from traditional software development, carries unhelpful baggage into the discourse on artificial intelligence. Beyond the flawed analogy of applying a software-centric model to a complex AI stack, the most distorting piece of this legacy is the assumption that openness is, by definition, an intrinsic good.¹⁹⁷

This Part challenges that premise. We argue that AI openness is more accurately characterized as an *instrumental* good. It is not inherently good or bad; it is a powerful policy tool whose desirability is entirely contingent on which components are opened, to what degree, and to what end.¹⁹⁸ Consequently, there is no universal answer to how open or closed any particular AI system should be. Instead, the differential openness of osAI must be calibrated with precision to achieve specific policy objectives.

To build this case, this Part systematically evaluates how differential openness affects four key policy objectives. Sections A through D analyze each objective in turn—public safety, innovation and economic growth, democratic control, and national security—demonstrating how, for each goal, openness functions as a double-edged sword, creating both profound benefits and acute risks. Finally, Section E synthesizes this analysis, moving from the tensions *within* each goal to the unavoidable trade-offs *between* them. This reveals the complex balancing act policymakers face, where every decision to open or close a component of the AI stack necessarily prioritizes one value over another.

197. See Richard Stallman, *Free Software Is Even More Important Now*, GNU OPERATING SYS., <https://www.gnu.org/philosophy/free-software-even-more-important.html> [https://perma.cc/YH5Z-5ZUK] (last visited Jan. 30, 2026).

198. See JON BATEMAN ET AL., CARNEGIE ENDOWMENT FOR INT’L PEACE, BEYOND OPEN VS. CLOSED: EMERGING CONSENSUS AND KEY QUESTIONS FOR FOUNDATION AI MODEL GOVERNANCE 4 (2024), https://assets.production.carnegie.fusionary.io/static/files/Bateman%20et%20al_Foundation%20AI%20Models_final.pdf [https://perma.cc/592Q-XXLP].

A. Safety

AI already shapes critical, sometimes life-or-death, decisions. It detects cancer in medical scans,¹⁹⁹ approves or denies mortgages,²⁰⁰ flags security threats,²⁰¹ and determines what billions of people see online.²⁰² Its failures do not unfold in the abstract—they manifest in hospitals,²⁰³ courtrooms,²⁰⁴ financial markets,²⁰⁵ and battlefields.²⁰⁶ When AI goes wrong, people lose jobs, homes, access to critical services, and sometimes even their lives.

Governing AI safety requires addressing three distinct challenges: (1) ensuring models are accurate (producing correct and unbiased outputs) and reliable (doing so consistently), (2) maintaining alignment so that their outputs do not cause harm through misuse or unintended behavior, and (3) enabling auditability through sufficient transparency and control to diagnose and remedy failures.

Differential openness is the primary tool for meeting these challenges, but it is a double-edged sword. While the OSS ethos that

199. Rebecca C. Fitzgerald et al., *The Future of Early Cancer Detection*, 28 NATURE MED. 666, 666–67 (2022).

200. See, e.g., Elijah Clark, *Rocket Mortgage's AI Technology: The Future of Mortgage Lending*, FORBES (Apr. 15, 2024, at 13:54 ET), <https://www.forbes.com/sites/elijahclark/2024/04/15/rocket-mortgages-ai-technology-the-future-of-mortgage-lending/>; Kori Hale, *A.I. Bias Caused 80% of Black Mortgage Applicants to Be Denied*, FORBES (Sep. 3, 2021, at 09:35 ET), <https://www.forbes.com/sites/korihale/2021/09/02/ai-bias-caused-80-of-black-mortgage-applicants-to-be-denied/>.

201. Aya H. Salem, Safaa M. Azzam, O.E. Emam & Amr A. Abohany, *Advancing Cybersecurity: A Comprehensive Review of AI-Driven Detection Techniques*, J. BIG DATA, Dec. 2024, at 1, 16, <https://link.springer.com/article/10.1186/s40537-024-00957-y> [<https://perma.cc/2MAF-6EZA>].

202. See Sang Ah Kim, *Social Media Algorithms: Why You See What You See*, 2 GEO. L. TECH. REV. 147, 150–51 (2017).

203. Laure Wynants et al., *Prediction Models for Diagnosis and Prognosis of Covid-19: Systematic Review and Critical Appraisal*, BMJ, Apr. 2020, at 1, 1, <https://doi.org/10.1136/bmj.m1328> (reviewing 731 AI systems purporting to diagnose or predict prognosis for COVID patients and finding five prognostic tools “showed adequate predictive performance in studies at low risk of bias”).

204. Jess Weatherbed, *Errors Found in US Judge's Withdrawn Decision Stink of AI*, VERGE (July 25, 2025, at 05:30 CT), <https://www.theverge.com/news/713653/judge-withdraws-cormedix-case-ai-citation-errors> [<https://perma.cc/M4YV-8VSZ>].

205. See Paolo Giudici & Emanuela Raffinetti, *SAFE Artificial Intelligence in Finance*, FIN. RSCH. LETTERS, Sep. 2023, at 1, 1–3, 12, <https://doi.org/10.1016/j.frl.2023.104088>.

206. Michael Biesecker, Sam Mednick & Garance Burke, *As Israel Uses US-Made AI Models in War, Concerns Arise About Tech's Role in Who Lives and Who Dies*, A.P. NEWS (Feb. 18, 2025, at 06:06 CT), <https://apnews.com/article/israel-palestinians-ai-technology-737bc17af7b03e98c29cec4e15d0f108> [<https://perma.cc/4FJF-6WFY>].

“given enough eyeballs, all bugs are shallow”²⁰⁷ applies with equal force to osAI, the same transparency that enables public auditing can also be exploited by malicious actors. The central question for safety, then, is not whether to make AI open or closed wholesale, but where targeted openness can meaningfully reduce harm without creating unacceptable risks.²⁰⁸

1. Benefits

To see how this works in practice, consider an autonomous vehicle that strikes a Black woman in a wheelchair crossing a street at night. To understand what went wrong—and to prevent it from happening again—investigators need visibility into the entire AI stack. Was the model trained on diverse and representative data? Did it struggle in low-light conditions? Was the decision logic flawed, the hardware malfunctioning, or the system manipulated?

Access to the training data would allow external researchers the capacity to assess whether the dataset the autonomous vehicle system was trained on included enough nighttime scenarios or sufficient representation of people with different skin tones or disabilities. Without that access, it is impossible to know whether the model was ever given the chance to learn how to recognize someone like the victim.²⁰⁹ Separately, data openness often uncovers systemic flaws, such as routine scraping of illegal content or the lack of responsible filtering that internal teams might miss or ignore.²¹⁰ For example, researchers have found ample child sexual abuse material in image datasets.²¹¹

207. ERIC S. RAYMOND, *THE CATHEDRAL AND THE BAZAAR* 30 (Tim O’Reilly ed., rev. ed. 2001).

208. Alondra Nelson et al., Comment Letter to Dep’t of Com. on Request for Comment on Dual Use Foundation Artificial Intelligence Models with Widely Available Model Weights 2–3 (Mar. 27, 2024), <https://hai-production.s3.amazonaws.com/files/2024-03/Response-NTIA-RFC-Open-Foundation-Models.pdf> [<https://perma.cc/UC3P-9JFU>].

209. See Jack Cable & Aeva Black, *With Open Source Artificial Intelligence, Don’t Forget the Lessons of Open Source Software*, CYBERSECURITY & INFRASTRUCTURE SEC. AGENCY: BLOG (July 29, 2024), <https://www.cisa.gov/news-events/news/open-source-artificial-intelligence-dont-forget-lessons-open-source-software> [<https://perma.cc/NVA2-ASKE>].

210. See KEVIN KLYMAN ET AL., *HUMAN-CENTERED A.I. AT STAN. UNIV., SAFEGUARDING THIRD-PARTY RESEARCH 2* (2025), <https://hai.stanford.edu/assets/files/hai-policy-brief-safeguarding-third-party-ai-research.pdf> [<https://perma.cc/SM5N-NGGD>].

211. David Thiel, *Investigation Finds AI Image Generation Models Trained on Child Abuse*, STAN. UNIV. CYBER POL’Y CTR. (Dec. 20, 2023), <https://cyber.fsi.stanford.edu/news/investigation-finds-ai-image-generation-models-trained-child-abuse>.

Moving deeper into the model's architecture, transparent model weights and system prompts would allow independent experts to investigate the AI's decision-making processes directly. They could identify if the system's logic deprioritizes pedestrians with certain demographic characteristics or in certain conditions—for example, wearing certain types of clothing, walking a pet, or standing in poor lighting. Looking under the hood at the model's decision-making processes can also enable crucial research into broader problems like AI hallucinations, where large language models generate false but highly plausible-seeming information.²¹² Insight into safety prompts can expose misguided or malicious instructions appended to use inputs that can lead to unsafe outputs.²¹³

Once a system is deployed, operational records become equally critical.²¹⁴ Transparent metadata—including detailed logs and decision trails—surface critical failures, particularly important in high-stakes applications like medical diagnostics, hiring systems, or autonomous vehicles.²¹⁵ Logs might pinpoint exactly when and why the system failed—whether it detected the car accident victim at all, whether it misclassified her as a shadow or background object, or whether it delayed braking. This data is also essential for detecting adversarial interference such as malicious tampering.

Finally, access to operational controls and compute enables proactive, not merely reactive, safety work.²¹⁶ Open access to a model's bias-detection algorithms or adversarial testing frameworks allows independent researchers to perform “red-teaming”—simulating a range of edge cases—such as pedestrians with different skin tones, clothing styles, or body types, to surface blind spots *before* they cause harm.²¹⁷ This requires computational resources to be available beyond companies and a handful of well-funded safety labs,²¹⁸ as opening safety tools

212. See, e.g., Sebastian Farquhar et al., *Detecting Hallucinations in Large Language Models Using Semantic Entropy*, 630 NATURE 625, 625 (2024).

213. See NAT'L TELECOMMS. & INFO. ADMIN., *supra* note 25, at 17–18.

214. NAT'L INST. OF STANDARDS & TECH., DEP'T OF COM., ARTIFICIAL INTELLIGENCE RISK MANAGEMENT FRAMEWORK (AIRMF 1.0), at 15–16, 35 (2023).

215. U.S. DEP'T OF HOMELAND SEC. & A.I. SAFETY & SEC. BOARD, ROLES AND RESPONSIBILITIES FRAMEWORK FOR ARTIFICIAL INTELLIGENCE IN CRITICAL INFRASTRUCTURE 19 (2024) (highlighting the importance of maintaining operational records for critical infrastructure uses of AI).

216. See Vinita Fordham, Allie Diehl & David Caswell, *Securing Government Against Adversarial AI*, DELOITTE (Apr. 11, 2023), <https://www2.deloitte.com/us/en/insights/industry/public-sector/adversarial-ai.html> [<https://perma.cc/S5ZJ-V9RG>].

217. U.S. DEP'T OF HOMELAND SEC. & A.I. SAFETY & SEC. BOARD, *supra* note 215, at 19 & n.24.

218. Manish Parashar, *Enabling Responsible Artificial Intelligence Research and Development Through the Democratization of Advanced Cyberinfrastructure*, HARV.

without providing the infrastructure on which to run them is an empty gesture.

At the root of all of this is the human component. When human workers—from data labelers in RLHF to internal engineers—are properly trained, protected, empowered, and embedded in transparent workflows, they can serve as an early warning system, flagging unsafe outputs, flawed incentives, or rushed deployments at their source. These individuals are best positioned to uncover issues, especially when they stem from secretive development processes.²¹⁹

2. Costs

But for all the ways that openness can strengthen safety, it also introduces serious and often irrevocable risks. The same transparency that enables oversight can be exploited by malicious actors, and the diffusion of powerful tools containing hidden flaws, biases, or security gaps can amplify harm, even when users are well-meaning. This trade-off is not abstract; it manifests at each layer of the AI stack.

The most acute risk lies in the dissemination of model weights.²²⁰ Once released, an AI model's core knowledge cannot be recalled.²²¹ It remains indefinitely available to be repurposed, modified, and potentially weaponized.²²² Any embedded flaw—a bias, unsafe instruction, or alignment gap—can be replicated into perpetuity. There is no practical mechanism to compel bad actors to cease its use or alert all well-meaning users of emergent harms. Consequently, open-weight models are routinely stripped of safeguards to generate extremist propaganda, nonconsensual deepfakes, and automated social engineering scams.²²³

Open data creates parallel dangers. The release of training sets containing private health records, intimate photos, or copyrighted material—even in the name of transparency—can constitute a massive

DATA SCI. REV., Apr. 2, 2024, at 1, 2–3, <https://hdrs.mitpress.mit.edu/pub/fysjbutp/release/2> [<https://perma.cc/2EYA-ETZB>].

219. Sharma, *AI's Hippocratic Oath*, *supra* note 97, at 1157–58.

220. P'SHIP ON AI, THE PARTNERSHIP ON AI RESPONSE TO THE NTIA REQUEST FOR COMMENT (RFC) ON DUAL USE FOUNDATION ARTIFICIAL INTELLIGENCE MODELS WITH WIDELY AVAILABLE MODEL WEIGHTS 6 (2024), <https://partnershiponai.org/wp-content/uploads/2024/04/PAI-Response-to-NTIA-RFC-Open-Foundation-Models.pdf> [<https://perma.cc/5BGY-6D7L>].

221. Kapoor et al., *supra* note 32, at 23084 (“Once the weights for a foundation model are made widely available, little recourse exists for the foundation model developer to rescind access.”).

222. Edd Gent, *Protesters Decry Meta's “Irreversible Proliferation” of AI*, IEEE SPECTRUM (Oct. 6, 2023), <https://spectrum.ieee.org/meta-ai> [<https://perma.cc/Q4BT-RVUF>].

223. See NAT'L TELECOMMS. & INFO. ADMIN., *supra* note 25, at 24–26.

violation of privacy and property rights.²²⁴ The use of scraped social media content and personal images has already led to real-world harms, from non-consensual deepfake pornography²²⁵ to government targeting and surveillance of dissidents.²²⁶

Finally, operational controls and records as well as application layer logic, while essential to oversight, can be weaponized. Transparent safety benchmarks and bias detection tools can be reverse engineered by adversaries to learn how to evade them. For example, deepfake creators can use open detection models as a training tool, fine-tuning their outputs until they beat the very systems designed to stop them.²²⁷ Exposing application-layer logic and interfaces follows the same logic: It helps adversaries identify how to bypass client-side filters, abuse plugin interfaces, or script high-volume misuse through automated front-ends.²²⁸ Likewise, public operational data—from error logs to performance metrics—can provide a detailed roadmap to an AI system’s blind spots, allowing cyberattackers to design exploits that target known

224. See, e.g., Jahner, *supra* note 128.

225. See Kate Tenbarge, *Found Through Google, Bought with Visa and Mastercard: Inside the Deepfake Porn Economy*, NBC NEWS (Mar. 27, 2023, at 10:56 CT), <https://www.nbcnews.com/tech/internet/deepfake-porn-ai-mr-deep-fake-economy-google-visa-mastercard-download-rcna75071> [<https://perma.cc/TG3E-AX2U>]; David Evan Harris, *Open-Source AI Is Uniquely Dangerous: But the Regulations that Could Rein It in Would Benefit All of AI*, IEEE SPECTRUM (Jan. 12, 2024), <https://spectrum.ieee.org/open-source-ai-2666932122> [<https://perma.cc/FJ55-FYZ5>].

226. Jay Stanley, *Machine Surveillance Is Being Super-Charged by Large AI Models*, ACLU (Mar. 21, 2025), <https://www.aclu.org/news/privacy-technology/machine-surveillance-is-being-super-charged-by-large-ai-models> [<https://perma.cc/NX8Q-2S93>]; Darrell M. West, *How AI Can Enable Public Surveillance*, BROOKINGS INST. (Apr. 15, 2025), <https://www.brookings.edu/articles/how-ai-can-enable-public-surveillance/> [<https://perma.cc/3JNL-AU55>].

227. In the OSS setting, spam detection is often not implemented as open source. See, e.g., Jim O’Leary, *Improving First Impressions on Signal*, SIGNAL (Nov. 1, 2021), <https://signal.org/blog/keeping-spam-off-signal/> [<https://perma.cc/48LZ-DEEB>] (“If we put this code on the Internet alongside everything else, spammers would just read it and adjust their tactics to gain an advantage in the cat-and-mouse game of keeping spam off the network.”).

228. See Yigitcan Kaya et al., *When AI Meets the Web: Prompt Injection Risks in Third-Party AI Chatbot Plugins*, 2026 IEEE SYMPOSIUM ON SEC. & PRIVACY (forthcoming 2026), <https://arxiv.org/html/2511.05797v1> [<https://perma.cc/M7PN-8XBT>] (explaining that knowledge of the design of application-layer features such as client-side input and output filters as well as plugin interfaces makes it easier for adversaries to design attacks to bypass or exploit vulnerabilities or gaps); Matteo Lupinacci et al., *The Dark Side of LLMs: Agent-Based Attacks for Complete Computer Takeover* (July 11, 2025) (unpublished manuscript), <https://arxiv.org/pdf/2507.06850v3> [<https://arxiv.org/pdf/2507.06850v3>] (explaining that AI agents are able to automate and weaponize at scale exploitations of application layer vulnerabilities, for example by using prompt injections to facilitate complete system takeovers).

weaknesses.²²⁹ Even operational records stored in the name of safety can, if opened without sufficient anonymization, create so-called “radioactive” piles²³⁰ of sensitive personal data.²³¹ Openness here accelerates the learning curve for attackers to overtake safety advances.

B. Innovation and Economic Growth

Artificial intelligence is both a powerful driver of innovation and a critical battleground for economic growth. Unlike traditional software, which can often be replicated with minimal resources, frontier AI requires immense computational power, vast datasets, and optimized model architectures.²³² Control over these core components is currently highly centralized, with a few dominant firms creating significant risks of monopolistic behavior, including discriminatory pricing, vendor lock-in, and “kill zones” that stifle new entrants. As scholars like Tejas Narechania and Ganesh Sitaraman have explored, this concentration threatens to stagnate the very innovation that AI promises to deliver.²³³

In this high-stakes environment, well-calibrated differential openness for osAI can serve as a potent anti-concentration tool, countering monopolistic tendencies by democratizing access to cutting-edge technology. However, its effectiveness is not guaranteed. While opening certain components of the AI stack can accelerate progress and broaden participation, strategic, partial openness can also be subverted into a tool for incumbent entrenchment—a form of “openwashing” that creates an illusion of accessibility while keeping true market power consolidated.²³⁴

229. Chuan Guo, Jacob R. Gardner, Yurong You, Andrew Gordon Wilson & Kilian Q. Weinberger, *Simple Black-Box Adversarial Attacks*, 97 PROCS. 36TH INT’L CONF. ON MACH. LEARNING 2484, 2484 (2019) (demonstrating that attackers can use seemingly innocuous information, such as confidence scores associated with an AI system output, to infer vulnerabilities to exploit).

230. Trey Herr, *Protecting Society from Radioactive Data*, TECH POL’Y PRESS (July 21, 2025), <https://www.techpolicy.press/protecting-society-from-radioactive-data/> [https://perma.cc/K6M4-N2FK].

231. Kevin Bankston, *In ChatGPT Case, Order to Retain All Chats Threatens User Privacy*, CTR. FOR DEMOCRACY & TECH. (June 25, 2025), <https://cdt.org/insights/in-chatgpt-case-order-to-retain-all-chats-threatens-user-privacy/> [https://perma.cc/9PGR-JCPB].

232. Tim Hwang, *Computational Power and the Social Impact of Artificial Intelligence* 4–5, 9–10 (Mar. 23, 2018) (unpublished manuscript), <https://arxiv.org/pdf/1803.08971> [https://perma.cc/Q2RM-QUQW].

233. Narechania & Sitaraman, *supra* note 32, at 128–37.

234. *See supra* note 154 and accompanying text.

1. Benefits

Optimal configurations of osAI's differential openness can dramatically lower barriers to entry and accelerate the pace of technological progress. By reducing redundancy and fostering collaboration, they enable new entrants to build upon existing advances without the prohibitive cost of developing foundational models from scratch.

This pro-competitive effect has been demonstrated by the impact of open-weight models like Meta's Llama, Stability AI's Stable Diffusion, Alibaba's Qwen3-Coder, and EleutherAI's GPT-Neo. These releases have empowered a global community of researchers, startups, and independent developers to create domain-specific applications in fields such as biomedical AI, climate modeling, and materials science—all without requiring billions of dollars in training costs.²³⁵ This decentralization is further supported by osAI development frameworks like TensorFlow and PyTorch²³⁶ and publicly available datasets such as Common Crawl²³⁷ and the Pile,²³⁸ which provide a shared foundation for innovation that is distributed widely rather than siloed within a few

235. See generally Niklas Muennighoff et al., s1: Simple Test-Time Scaling (Mar. 1, 2025) (unpublished manuscript), <https://arxiv.org/pdf/2501.19393> [<https://perma.cc/G7XB-ZYSP>] (mathematical reasoning experiments dependent on open weight in Qwen and open traces in Gemini); Shrey Pandit et al., MedHallu: A Comprehensive Benchmark for Detecting Medical Hallucinations in Large Language Models (Feb. 20, 2025) (unpublished manuscript), <https://arxiv.org/pdf/2502.14302> [<https://perma.cc/R26Q-WQ2E>] (study of medical hallucinations on open and closed models); Jimeng Shi et al., Deep Learning and Foundation Models for Weather Prediction: A Survey (Jan. 12, 2025) (unpublished manuscript), <https://arxiv.org/pdf/2501.06907> [<https://perma.cc/B5R3-EMB4>] (see section 3 for a list of weather modeling examples); Yingheng Tang et al., MatterChat: A Multi-Modal LLM for Material Science (Apr. 26, 2025) (unpublished manuscript), <https://arxiv.org/pdf/2502.13107> [<https://perma.cc/JV62-VVEG>] (material science LLM based on pretrained models).

236. See Cade Metz, *Google Just Open Sourced TensorFlow, Its Artificial Intelligence Engine*, WIRED (Nov. 9, 2015, at 09:00 CT), <https://www.wired.com/2015/11/google-open-sources-its-artificial-intelligence-engine/> [<https://perma.cc/WL2F-KKFX>]; Adam Paszke et al., *PyTorch: An Imperative Style, High-Performance Deep Learning Library*, 32 PROCS. OF THE 33RD INT'L CONF. ON NEURAL INFO. PROCESSING SYS. (NEURIPS 2019) 1, 2 (2019), https://proceedings.neurips.cc/paper_files/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf [<https://perma.cc/DG99-FNLK>]; Akshay Agrawal et al., *TensorFlow Eager: A Multi-Stage, Python-Embedded DSL for Machine Learning*, 1 PROCS. 2ND MACH. LEARNING & SYS. CONF. (MLSYS 2019) 1, 1 (2019), <https://arxiv.org/pdf/1903.01855> [<https://perma.cc/CL7V-6EFE>].

237. Common Crawl is available under a limited license. *Terms of Use*, COMMON CRAWL, <https://commoncrawl.org/terms-of-use> [<https://perma.cc/AL2Q-35FE>] (last updated Jan. 31, 2026).

238. Leo Gao et al., *The Pile: An 800GB Dataset of Diverse Text for Language Modeling 1–2* (Dec. 31, 2020) (unpublished manuscript), <https://arxiv.org/pdf/2101.00027> [<https://perma.cc/C9W8-6DM4>].

dominant firms. Open components such as client libraries, plugin standards, and orchestration frameworks can play a similar role, reducing fixed deployment costs. This allows small teams to assemble production-ready services from shared building blocks, competing on differentiated services at the application layer rather than recreating interfaces, safety wrappers, and integration code from scratch.

Crucially, this form of openness also spreads the *use* of AI, which is itself a powerful driver of innovation and economic growth. Because open-weight models can be downloaded and run on local hardware, users bypass the pay-per-use fees of controlled API endpoints.²³⁹ This empowers a much broader base of individuals and researchers to experiment with and integrate powerful AI into novel applications, fostering a more dynamic, ground-up form of economic growth that cannot be achieved through controlled platforms alone.²⁴⁰

2. Costs

However, as Narechania and Sitaraman wisely caution, we must reject the “[f]alse [p]romise” that openness “will completely address the problems with an unregulated AI oligopoly.”²⁴¹ While opening model weights or source code can spur experimentation, this alone does not ensure broad competition if other critical parts of the ecosystem remain closed. As we have emphasized, AI development is not just about access to code; it is also about who controls the surrounding infrastructure that makes AI usable and scalable.²⁴² If control over compute, proprietary data, deployment pathways, application layers, and expert labor remain tightly controlled, then models that only open weights or source code merely offer the illusion of an open marketplace.

This is where openness can be subverted into a strategy for incumbent entrenchment—a form of “openwashing” where companies claim the reputational benefits of openness while withholding key components.²⁴³ Meta’s high-profile Llama release is a prime example: It made a self-interested business decision while positioning itself as a

239. *Cf. Pricing*, OPENAI, <https://platform.openai.com/docs/pricing> (last visited Jan. 31, 2026) (showing how OpenAI charges per API use).

240. *See* Robert Wolfe et al., *Laboratory-Scale AI: Open-Weight Models Are Competitive with ChatGPT Even in Low-Resource Settings*, 2024 PROCS. ACM CONF. ON FAIRNESS, ACCOUNTABILITY, & TRANSPARENCY (FACCT ’24) 1199, 1200–01, <https://dl.acm.org/doi/epdf/10.1145/3630106.3658966>.

241. Narechania & Sitaraman, *supra* note 32, at 153–54.

242. *See supra* Section I.B; Nobel, Rozenshtein & Sharma, *supra* note 142.

243. Liesenfeld & Dingemans, *supra* note 32, at 1776.

champion of openness.²⁴⁴ While Meta’s motivation to crowd out competitor proprietary systems may seem pro-competitive, its ultimate goal is not to dismantle the oligopolistic nature of the industry, but rather to reinforce its own position within it.

The strategic value of a splashy “open source” release is often found not in what is shared, but in the critical components that are held back. Control over four core areas—(1) compute infrastructure, (2) deployment access, (3) proprietary training data, and (4) labor—allows incumbents to retain real market power even as they gesture toward openness.

(1) *Compute*. Deploying advanced AI, open or not, requires specialized GPUs and cloud-scale infrastructure that are functionally inaccessible to startups or researchers without vast financial resources. A legal-tech or biomedical startup may fine-tune an open-weight model, but it cannot compete if it cannot afford the cloud resources to deploy it at scale—resources often controlled by the very firms developing the models. This forces smaller players into dependency on dominant providers like Amazon Web Services (AWS), Microsoft Azure, and Google Cloud, reinforcing the incumbents’ market position.²⁴⁵

(2) *Deployment access*. Making an AI model open is not the same as making it usable. Deploying an AI system in a real-world product—for example, in a legal-research app or a medical-imaging tool—requires cloud hosting and interface tools, which are often controlled by the same few companies. If developers are forced to rely on these gatekeepers, they risk getting locked into a particular company’s ecosystem—subject to its pricing, policies, and terms of use. While open standards like Anthropic’s widely adopted Model Context Protocol (MCP) promise interoperability that would make it easier to swap one model for another,²⁴⁶ their effectiveness depends on continued adoption by

244. See Mike Isaac, *How A.I. Made Mark Zuckerberg Popular Again in Silicon Valley*, N.Y. TIMES (May 31, 2024), <https://www.nytimes.com/2024/05/29/technology/mark-zuckerberg-meta-ai.html>.

245. See Paul Sandle, *Dominance of Amazon and Microsoft in Cloud Harming Competition, UK Says*, REUTERS (July 31, 2025), <https://www.reuters.com/legal/litigation/dominance-amazon-microsoft-cloud-harming-competition-uk-says-2025-07-31/>.

246. See *Introducing the Model Context Protocol*, ANTHROPIC (Nov. 25, 2024), <https://www.anthropic.com/news/model-context-protocol> [https://perma.cc/P94X-945Y]; see also Benj Edwards, *MCP: The New “USB-C for AI” That’s Bringing Fierce Rivals Together*, ARS TECHNICA (Apr. 1, 2025, at 06:30 CT), <https://arstechnica.com/information-technology/2025/04/mcp-the-new-usb-c-for-ai-thats-bringing-fierce-rivals-together/> (describing the Model Context Protocol); Kyle Wiggers, *Google to Embrace Anthropic’s Standard for Connecting AI Models to Data*, TECHCRUNCH (Apr. 9, 2025, at 16:18 PT), <https://techcrunch.com/2025/04/09/google-says-itll-embrace-anthropics-standard-for-connecting-ai-models-to-data/> [https://perma.cc/EY6S-UK9E] (describing Google’s adoption of the protocol); Kyle Wiggers, *OpenAI Adopts Rival Anthropic’s Standard for Connecting AI Models to Data*, TECHCRUNCH (Mar. 26, 2025, at 11:18

dominant firms who have historically used such standards to drive adoption of their services before reverting to restrictive, proprietary interfaces once their market position is secure.²⁴⁷

(3) *Data*. Training data remains one of the most valuable and least open parts of the AI stack. While models like Llama release weights, they do not release the massive proprietary datasets used to train them.²⁴⁸ Since a model's performance depends heavily on what it is trained on, incumbents who hoard their data can preserve a decisive competitive edge, regardless of who has access to their model weights. Open-weight models trained on publicly available datasets like Common Crawl rarely match the performance of those trained on curated, private corpora.²⁴⁹

(4) *Labor*. A thriving AI ecosystem requires open opportunities for people to shape and build them. However, when expertise is locked within a few corporate or academic labs through restrictive employment practices and a failure to invest in broad talent development, the human capacity needed to realize the potential of openness is stifled. This creates a facade of access while concentrating the most critical resource—human talent—in the hands of a few.²⁵⁰

C. Democratic Access and Control

Democratizing AI is a twofold goal. It requires liberalizing access to ensure powerful tools are available to the many, not just the few. It also demands the establishment of democratic societal control, so the public has a say in how these technologies are developed and deployed. True democratization, therefore, means both expanding the number of people who can build with and benefit from AI and ensuring that its evolution reflects public values rather than the narrow interests of powerful corporations or states.

PT), <https://techcrunch.com/2025/03/26/openai-adopts-rival-anthropics-standard-for-connecting-ai-models-to-data/> [<https://perma.cc/P4TL-8FTQ>] (describing OpenAI's adoption of the protocol).

247. See Chinmayi Sharma, *Concentrated Digital Markets, Restrictive APIs, and the Fight for Internet Interoperability*, 50 U. MEM. L. REV. 441, 455–61 (2019).

248. Training data is a “mix of publicly available, licensed data and information from Meta's products and services.” Meta Llama, *Model Information*, GITHUB, https://github.com/meta-llama/llama-models/blob/main/models/llama4/MODEL_CARD.md [<https://perma.cc/T4QU-DZ9L>] (last visited Jan. 31, 2026).

249. See *Leaderboard Overview*, ARENA, <https://lmarena.ai/leaderboard> (last visited Jan. 31, 2026) (no model on this AI leaderboard has been trained on Common Crawl or another similarly open dataset).

250. Toner, *supra* note 188, at 5.

1. Benefits

Openness is a powerful force for liberalizing access to technology. By dramatically lowering the immense cost of entry, open-source AI frameworks like PyTorch and open-weight models like Meta's Llama have fundamentally altered the AI landscape. They allow startups, academic labs, and independent researchers to build upon state-of-the-art foundations, moving AI development beyond a handful of elite corporate labs and fostering a more vibrant and competitive ecosystem. In a global context, this distribution of technical capability helps mitigate the risk that AI prowess remains confined to the U.S.–China duopoly. Teams in Nairobi or São Paulo can fine-tune frontier models for local agriculture, public health, or language-revitalization projects, creating opportunities for local adaptation where centralized systems have historically failed.

Beyond broadening inclusive participation, openness is a critical tool for enabling democratic control. Openness in documentation, system prompts, safety benchmarks, and bias detection tools empowers civil society, allowing journalists, advocates, and independent auditors to put corporate and government claims to the test, to challenge discriminatory outcomes, and to hold powerful institutions accountable. For example, automated fairness audits have revealed racial disparities in credit-scoring algorithms, enabling regulatory enforcement and consumer protection.²⁵¹ Open compliance tools allow external experts to stress-test AI for bias, fraud, and security vulnerabilities, preventing companies from self-policing in ways that prioritize corporate interests over public welfare.²⁵² Openness at the application layer can play a similar role: Open source clients, safety pipelines, and plugin orchestration code allow external researchers and deployers to test alternative guardrail configurations, compare different interaction designs, and build independent safety layers on top of shared models.

Finally, enhancing the diversity of the human labor behind AI—through inclusive hiring, equitable training pathways, and protections for whistleblowers and annotators—strengthens democratic control over systems. Too often the communities most affected by AI are excluded from its development, resulting in systems that are less likely to serve their interests.²⁵³ Expanding the range of backgrounds and institutions shaping these tools is therefore essential for ensuring that societal control is not just a theoretical ideal but a practical reality.

251. Mark McCarthy, *Fairness in Algorithmic Decision-Making*, BROOKINGS INST. (Dec. 9, 2019), <https://www.brookings.edu/articles/fairness-in-algorithmic-decision-making/> [https://perma.cc/789H-XRRM].

252. See Jose-Miguel Bello y Villarino & Simon Bronitt, *AI-Driven Corporate Governance: A Regulatory Perspective*, 33 GRIFFITH L. REV. 355, 362 (2024).

253. See Solaiman, *supra* note 32, at 112.

2. Costs

Unfortunately, openness does not automatically lead to a more equitable or controllable AI ecosystem. If not carefully structured, it can paradoxically undermine the very democratic goals it purports to serve by creating an illusion of control while consolidating power, or by fragmenting authority so completely that collective governance becomes impossible.

First, openness can create a facade of democratic control that masks a deep consolidation of corporate power.²⁵⁴ Even if a model's weights are released for free, the power to deploy it at scale remains centralized in the hands of the few corporations that control the underlying compute cloud infrastructure and the model itself.²⁵⁵ When those same firms control the dominant application layers, including the interfaces and plugin platforms through which most users encounter AI, they can translate technical openness elsewhere in the stack into a tightly curated menu of permissible uses. These firms can become de facto private regulators, exercising unfettered control over the terms of AI access, development, and use. When a handful of unaccountable companies can decide which political speech, scientific research, or social tools are allowed to run on their platforms, the power to shape society shifts from democratic institutions to corporate boardrooms.

Second, openness can undermine society's ability to exert collective control by leading to uncontrollable fragmentation. As previously discussed, once an open-weight model is released, the ability to enforce terms of use, responsible practices, and even legal compliance is severely hampered.²⁵⁶ The decentralized nature of the ecosystem also curtails the ability to identify a single actor to hold accountable for harm. This allows companies to distance themselves from the consequences of downstream misuse, shifting the externalities of their products onto the public. While this form of openness empowers the individual user, it critically weakens the power of society as a whole to set and enforce rules.

D. National Security and Global Leadership

AI is not just an economic and technological asset—it is a strategic resource that will shape military capabilities, intelligence dominance, and geopolitical influence for decades to come.²⁵⁷ The race to control AI is already a defining factor in global power struggles, determining which

254. See Widder, Whittaker & West, *Why 'Open' AI Systems Are Actually Closed*, *supra* note 32, at 831.

255. See *id.*

256. See *supra* notes 220–23 and accompanying text.

257. See TRIVEDI & MEYSENBURG, *supra* note 33, at 20–22.

nations lead in technological advancement, economic strength, and security.²⁵⁸ AI models underpin cybersecurity, intelligence gathering, autonomous military systems, and economic stability, making their regulation a matter of national security as much as technological governance.²⁵⁹

1. Benefits

Strategic openness in AI can enhance national power and secure global leadership. By championing osAI, the U.S. can establish its technology as the global standard, shoring up its economic power while preventing adversaries from spreading their technological influence, as China has tried to do by, for example, integrating its models into Saudi Arabia's national oil company.²⁶⁰ Much like how Google's open source Android operating system secured American influence over the mobile technology landscape,²⁶¹ osAI creates a powerful gravitational pull, drawing international users into an ecosystem that naturally favors U.S. cloud infrastructure and hardware. Control over widely adopted application layers—including chat interfaces, agent frameworks, and plugin ecosystems—is part of that pull, because the defaults they embed about security, privacy, and lawful use can quietly become global expectations even when underlying models circulate more freely.

Far from symbolic, this leadership is critical for shaping the future of digital governance. When the most widely adopted AI models, architectures, and regulatory frameworks come from open, democratic sources, global norms may have a better chance of favoring civil liberties over surveillance, pluralism over censorship, and accountability over opacity. Conversely, if authoritarian states dominate the global AI

258. See Barry Pavel et al., *AI and Geopolitics*, RAND (Nov. 3, 2023), <https://www.rand.org/pubs/perspectives/PEA3034-1.html> [<https://perma.cc/LW7R-YLR8>].

259. See *id.*

260. See Malcolm Moore, *Saudi Aramco Chief Says DeepSeek AI Makes 'Big Difference' to Operations*, FIN. TIMES (Mar. 4, 2025, at 04:43 CT), <https://www.ft.com/content/0d24dcf4-b53b-48e5-b49c-99606958a96d>; see also Mohammed Soliman, *Realigning US-Saudi Relations for the AI Era*, MIDDLE E. INST. (May 5, 2025), <https://www.mei.edu/publications/realigning-us-saudi-relations-ai-era> [<https://perma.cc/4HBM-TJ5Q>] (emphasizing that the U.S. needs to keep ahead of “competing tech corridors being built by China”).

261. See Dieter Bohn, *Android at 10: The World's Most Dominant Technology*, VERGE (Sep. 26, 2018, at 10:00 CT), <https://www.theverge.com/2018/9/26/17903788/google-android-history-dominance-marketshare-apple>.

ecosystem, they will embed values of control and repression into core infrastructures.²⁶²

Furthermore, American technological dominance has historically rested on fostering a deep talent pool of both domestic and foreign experts. Sustaining this advantage requires proactively maintaining pathways into the national talent pipeline through robust and inclusive STEM education and modernized immigration systems. Training international students in American universities, for example, remains one of the country's most effective forms of soft power, seeding global influence while attracting the world's brightest minds—whether the talent it cultivates remains in the country or not. Restrictive immigration policies, such as those currently being pursued by the present administration, undermine our global leadership and create opportunities for adversaries to snatch up the experts we myopically reject.²⁶³

Finally, openness strengthens military alliances by enabling deep technological collaboration.²⁶⁴ When the U.S. and its partners build upon shared, open AI frameworks, they dismantle the technical barriers that have historically complicated joint operations. This common foundation enables seamless data fusion, shared intelligence pictures, and integrated command-and-control systems, allowing allied forces to act with a speed and cohesion that closed, proprietary systems cannot match. A collaborative approach, fostered by osAI, allows allies to pool resources and talent to out-innovate adversaries, bolstering collective defense.

2. Costs

Despite these benefits, an unrestricted approach carries severe and direct risks to national security. When advanced models release weights publicly, they can provide a powerful accelerant to adversary states, allowing them to bypass years and billions of dollars of research and development and to adapt these models for military and intelligence applications. While some researchers doubt whether today's osAI actually increases adversary capabilities, the pace at which the

262. Zeyi Yang, *Here's How DeepSeek Censorship Actually Works—and How to Get Around It*, WIRED (Jan. 31, 2025, at 14:33 CT), <https://www.wired.com/story/deepseek-censorship/> [<https://perma.cc/53ER-QJJ5>].

263. Ben Greenho, Silva Mathema & Rosa Barrientos-Ferrer, *The Trump Administration's Hostility to Legal Immigration Harms America's Global Leadership in Innovation*, CTR. FOR AM. PROGRESS (Nov. 13, 2025), <https://www.americanprogress.org/article/the-trump-administrations-hostility-to-legal-immigration-harms-americas-global-leadership-in-innovation/>.

264. BEN FITZGERALD, PETER L. LEVIN & JACQUELINE PARZIALE, CTR. FOR A NEW AM. SEC., OPEN SOURCE SOFTWARE & THE DEPARTMENT OF DEFENSE 9 (2016), <https://s3.us-east-1.amazonaws.com/files.cnas.org/documents/CNASReport-OpenSourceSoftware.pdf> [<https://perma.cc/KD8L-YMSD>].

technology is advancing makes it a possibility that cannot be ignored.²⁶⁵ Open or lightly regulated application-layer tooling can similarly accelerate adversaries' ability to deploy frontier capabilities at scale by providing turnkey chatbots, coding assistants, and agent platforms that sit on top of whatever models and hardware they can acquire.

Open model and application risk is dangerously compounded by openness in the hardware supply chain. The widespread commercial availability of high-performance computing chips gives America's adversaries a powerful toolkit to close the technological gap. American-designed chips, particularly those from Nvidia, have powered much of China's AI development.²⁶⁶ China is also leveraging other open compute components, such as chip design and the software that connects it to applications' chip power, to build out its own compute infrastructure, allowing it to move away from western compute providers and undermining the efficacy of export controls.²⁶⁷

The threat also extends beyond rival nations to non-state actors. The same open models that empower startups can be weaponized by terrorist cells, transnational criminal organizations, or hacktivist collectives. Low barriers to entry mean that groups with limited resources can suddenly access previously unattainable capabilities, dramatically increasing the risk of sophisticated, AI-powered disinformation campaigns, automated cyberattacks against critical infrastructure, and even the design of biological or chemical weapons.²⁶⁸

The emergence of DeepSeek, a Chinese AI company whose models now rival the best proprietary systems in the West, serves as a powerful case study. Built in record time on open Western architectures²⁶⁹ and

265. See CHRISTOPHER A. MOUTON, CALEB LUCAS & ELLA GUEST, RAND, THE OPERATIONAL RISKS OF AI IN LARGE-SCALE BIOLOGICAL ATTACKS 1 (2024), https://www.rand.org/pubs/research_reports/RRA2977-2.html [<https://perma.cc/XF4W-YMWJ>].

266. Zijing Wu & Eleanor Olcott, *Nvidia AI Chips Worth \$1bn Smuggled to China After Trump Export Controls*, FIN. TIMES (July 25, 2025, at 05:52 CT), <https://www.ft.com/content/6f806f6e-61c1-4b8d-9694-90d7328a7b54>; Che Pan & Casey Hall, *Nvidia AI Chips: Repair Demand Booms in China for Banned Products*, REUTERS (July 24, 2025), <https://www.reuters.com/world/china/china-repair-demand-banned-nvidia-ai-chipsets-booms-2025-07-24/>.

267. Che Pan & Brenda Goh, *Exclusive: China to Publish Policy to Boost RISC-V Chip Use Nationwide, Sources Say*, REUTERS (Mar. 4, 2025), <https://www.reuters.com/technology/china-publish-policy-boost-risc-v-chip-use-nationwide-sources-2025-03-04/> (“China plans to issue guidance to encourage the use of open-source RISC-V chips nationwide for the first time . . .”).

268. See Shlomit Wagman & Sarah Hubbard, *Weaponized AI: A New Era of Threats and How We Can Counter It*, HARV. KENNEDY SCH: ASH CTR. FOR DEMOCRATIC GOVERNANCE & INNOVATION (Apr. 8, 2025), <https://ash.harvard.edu/articles/weaponized-ai-a-new-era-of-threats/> [<https://perma.cc/2APK-V5K9>].

269. Specifically, DeepSeek distilled Llama to create some of its models, which in turn relied on the transformer architecture invented by Google. DEEPSEEK, DEEPSEEK-

likely powered by high-end American chips available before export restrictions took full effect,²⁷⁰ DeepSeek's success illustrates the peril of osAI: Openness catalyzes rapid innovation, but not always in ways aligned with U.S. strategic interests.²⁷¹ The fact that DeepSeek was built by a hedge fund rather than created by tech behemoths such as Baidu or Tencent²⁷² highlights a broader risk: osAI is not just enabling China's largest firms but fostering a diverse ecosystem of smaller competitors that are more resilient to Western restrictions.²⁷³

The long-term implications of this proliferation are complex. While some hope that adversaries adopting open Western components might also adopt more democratic technical norms, this outcome is far from certain; such hopes were shared when China joined the World Trade Organization, but economic integration did not produce political liberalization.²⁷⁴ Moreover, an honest assessment requires looking inward. The U.S. is not immune to the problematic uses of AI, such as deploying AI-powered surveillance against immigrants, raising concerns about whether American leadership always aligns with democratic ideals.²⁷⁵

R1: INCENTIVIZING REASONING CAPABILITY IN LLMs VIA REINFORCEMENT LEARNING 11 (2025), https://github.com/deepseek-ai/DeepSeek-R1/blob/main/DeepSeek_R1.pdf [<https://perma.cc/C8UC-PWQ9>].

270. See Nathan Lambert, *The American DeepSeek Project*, INTERCONNECTS (July 4, 2025), <https://www.interconnects.ai/p/the-american-deepseek-project> [<https://perma.cc/83R2-AY86>].

271. See Paul Mozur, John Liu & Cade Metz, *China's Rush to Dominate A.I. Comes with a Twist: It Depends on U.S. Technology*, N.Y. TIMES (Feb. 21, 2024), <https://www.nytimes.com/2024/02/21/technology/china-united-states-artificial-intelligence.html>.

272. Cade Metz, *What to Know About DeepSeek and How It Is Upending A.I.*, N.Y. TIMES (Jan. 27, 2025), <https://www.nytimes.com/2025/01/27/technology/what-is-deepseek-china-ai.html>.

273. See Stefan Stein, *CrowdStrike Research: Security Flaws in DeepSeek-Generated Code Linked to Political Triggers*, CROWDSTRIKE (Nov. 20, 2025), <https://www.crowdstrike.com/en-us/blog/crowdstrike-researchers-identify-hidden-vulnerabilities-ai-coded-software/> [<https://perma.cc/U4R3-W6B7>] (showing that DeepSeek-R1, an open source Chinese LLM, introduces severe security vulnerabilities in response to prompts containing politically sensitive terms such as "Tibet," "Uyghurs," or "Falun Gong," a pattern CrowdStrike warns could expose U.S. organizations to compromise because up to 90 percent of developers rely on AI coding assistants and DeepSeek's biases appear embedded in the model weights themselves).

274. See *What Happened When China Joined the WTO?*, COUNCIL ON FOREIGN RELS. (Feb. 6, 2025), <https://education.cfr.org/learn/reading/what-happened-when-china-joined-wto> [<https://perma.cc/5NDB-BZRR>].

275. See Steven Hubbard, *Invisible Gatekeepers: DHS' Growing Use of AI in Immigration Decisions*, AM. IMMIGR. COUNCIL (May 9, 2025), <https://www.americanimmigrationcouncil.org/blog/invisible-gatekeepers-dhs-growing-use-of-ai-in-immigration-decisions/> [<https://perma.cc/4LCD-2TB2>].

Ultimately, the rapid pace of AI progress suggests that a permanent technological lead may be impossible for any single nation. If osAI empowers competitors like China, it may paradoxically improve global security by shifting the strategic calculus from zero-sum competition to mutual risk management.²⁷⁶ Recognizing a shared interest in preventing a future that no one can unilaterally control could create powerful new incentives for cooperation on global safety standards, a dynamic that echoes the logic of nuclear arms control.²⁷⁷

E. Navigating Trade-offs in AI Openness

Regulating osAI is an exercise in strategic prioritization. The central challenge is not whether AI should be “open” or “closed,” but how differential openness at the component level creates trade-offs along two axes. These tensions exist *within* single policy goals and *between* competing ones, all of which are layered upon deeper *structural* conflicts. Strategic governance, therefore, is not about ideology but calibration: weighing the costs and benefits of opening each component of an AI system. AI is simultaneously an economic asset and a security risk, a public good and a proprietary investment—contradictions that are built into its very design.

1. Trade-offs Within Policy Goals

The decision to open any specific AI component is often a double-edged sword, capable of both advancing and undermining the same policy objective. For public safety, which depends on oversight and control, openness is critical. Making components like system prompts, operational metadata, and control layers transparent allows regulators, researchers, and companies to monitor failures and intervene early. Yet, the same transparency that enables oversight also invites exploitation. Security benchmarks designed to build trust in a fraud detection system can double as roadmaps for adversaries seeking to evade it. Similarly, while greater labor inclusivity can improve error detection by bringing in diverse perspectives, expanding the talent pool too quickly without shared standards can lead to inconsistent and risk-prone practices, particularly in high-stakes domains.

276. See Steven Adler, *Are We Ready for a ‘DeepSeek for Bioweapons’?*, LAWFARE (May 29, 2025, at 12:00 CT), <https://www.lawfaremedia.org/article/are-we-ready-for-a--deepseek-for-bioweapons> [https://perma.cc/7JH7-R2WR].

277. See Simon Goldstein & Peter N. Salib, *DeepSeek Points Toward U.S.-China Cooperation, Not a Race*, LAWFARE (Mar. 5, 2025, at 12:33 CT), <https://www.lawfaremedia.org/article/deepseek-points-toward-u.s.-china-cooperation--not-a-race> [https://perma.cc/DA62-KGD2].

Innovation, by contrast, thrives on experimentation, flexibility, and rapid iteration, which are fueled by open-weight models, transparent architectures, available compute, and accessible training data. These components maximize experimentation across a broader ecosystem of players. However, openness does not inherently create competition; it can also entrench dominance. For instance, companies that release model weights openly while keeping components like cloud services, fine-tuning expertise, or proprietary hosting platforms locked behind paywalls can create dependencies that stifle the very flexibility on which innovation relies.

Accountability, in turn, relies on transparency, traceability, and interpretability. Opening components like operational data and system prompts make it possible to audit, challenge, and correct AI decisions. But more openness does not guarantee more accountability. Open training data can expose bias but also diffuse responsibility; if a model trained on public datasets produces discriminatory outcomes, it becomes difficult to assign blame. Furthermore, transparency can invite regulatory arbitrage. Full visibility into evaluation benchmarks may lead developers to optimize for test performance rather than real-world fairness or robustness, turning a push for oversight into a playbook for compliance theater.

Finally, national security depends on maintaining a strategic advantage through controlled access to powerful technology. Restricting model weights, proprietary training data, and advanced systems helps ensure that critical AI tools stay in trusted hands. But excessive secrecy can backfire. If U.S. systems remain too closed, global users may turn to alternative ecosystems, eroding American influence and control.

2. Trade-offs Between Policy Goals

These trade-offs become even more acute when different policy goals, each with its own logic for openness or closure, come into direct conflict. The most persistent tension exists between the need for secrecy in the name of public safety and the demand for transparency to foster innovation. This conflict plays out across multiple components. For example, transparency into the human alignment pipeline—the layers of people that curate data, design prompts, or flag edge cases—and operational records enhances safety by allowing external scrutiny. However, it can also chill innovation if researchers fear legal or professional retaliation for pursuing controversial ideas. Conversely, the openness intended to promote innovation, such as releasing model weights or training datasets, can complicate efforts to enforce safety and accountability. Open weights allow anyone to strip safeguards, while

open datasets built on sensitive sources introduce risks of bias, privacy, and security.

A similar clash occurs between innovation and accountability. The very components that ensure that AI decisions can be audited and corrected—such as transparent system prompts and operational data—often introduce regulatory friction. Requiring their openness imposes compliance costs, delays deployment, and can raise the bar for entry, particularly for small players. A company that develops an AI system capable of dramatically improving cancer detection may be forced to delay or redesign it if laws demand full interpretability, a standard that many cutting-edge deep learning models, by their very nature, cannot meet.²⁷⁸

A parallel tension emerges between competition and national security. Openness in datasets, architectures, and evaluation benchmarks is essential for breaking up AI monopolies by lowering entry barriers for startup academic labs. But the same openness that fosters domestic competition can erode the technological asymmetry that national security depends upon by accelerating adversarial capabilities. If frontier AI models or military-grade training data were made fully open, adversarial states would gain immediate access to capabilities once held exclusively by a few AI leaders. At the same time, excessive secrecy in the name of national security risks stifling domestic competition just as much as it hinders foreign rivals.

3. Deeper Structural Trade-offs

These policy trade-offs are layered onto a deeper structural tension between centralization and decentralization. Beyond who gets access to AI lies the question of who builds and governs it. From a safety, accountability, and national security perspective, centralizing development within a few “national champion” firms simplifies top-down enforcement, making it easier for regulators to secure sensitive capabilities, enforce safeguards, maintain professional standards, and oversee compliance.²⁷⁹ But while a centralized, security-first model may make oversight easier, it can create single points of failure in the ecosystem with cascading effects.²⁸⁰ It also risks starving the ecosystem

278. See Emrullah ŞAHİN, Naciye Nur Arslan & Durmuş Özdemir, *Unlocking the Black Box: An In-Depth Review on Interpretability, Explainability, and Reliability in Deep Learning*, 37 NEURAL COMPUTING & APPLICATIONS 859, 868 (2025).

279. See Jack M. Balkin, *Old-School/New-School Speech Regulation*, 127 HARV. L. REV. 2296, 2298–99 (2014).

280. Press Release, U.S. Dep’t of the Treasury, *New Treasury Report Assesses Opportunities, Challenges Facing Financial Sector Cloud-Based Technology Adoption* (Feb. 8, 2023), <https://home.treasury.gov/news/press-releases/jy1252> [<https://perma.cc/5ESY-D4K4>].

of the flexibility and diversity of thought needed for long-term innovation and global leadership. Decentralization, by contrast, is foundational to democratization, as it redistributes power from a few powerful actors to the many communities impacted by the technology.

Beneath this lies an even more fundamental conflict: the imperative for control versus the ideal of freedom. Control is essential for national security, safety, and even certain forms of market-driven innovation, as it involves limiting access, embedding safeguards, and maintaining oversight. Yet, this control—whether over model behavior through alignment filters or over components as proprietary assets—inevitably constrains the freedom to iterate, experiment, and build a decentralized osAI ecosystem. A model marketed as reducing harmful speech, for example, can become a form of soft censorship if its operational controls are not transparent. Likewise, controlling the labor pipeline to ensure high standards may protect national assets from economic espionage but may also erode the ethos of widespread public participation.

All these tensions are intensified by geopolitical urgency. In a global race for AI leadership, especially between the U.S. and China, speed is often treated as a proxy for strength, creating immense pressure to prioritize rapid deployment over rigorous oversight. In this climate, responsible approval cycles, fairness audits, and compliance thresholds could easily be cast as obstacles in a zero-sum contest for dominance where even prudent caution can feel intolerable.

None of these dilemmas can be “solved” by picking openness or closedness. Rather, they are enduring trade-offs of differential openness that must be strategically managed. Every decision is an act of prioritization, requiring a sophisticated, context-aware governance approach that calibrates differential openness at the component level to strike a deliberate, evolving balance between competing values.

4. The Compounding Complexity of Interdependent Components

Even once policymakers identify the right balance of trade-offs for a single component, the task is complicated by the fact that no layer of the AI stack stands alone. Each component sits on a spectrum of openness, and moving it toward greater or lesser openness reshapes the landscape for every other component.

These interactions create both “cascading openness” and “cascading closedness” by shifting incentives or capabilities across the stack in nonlinear or counterintuitive ways, which add a second order of complexity to any trade-off analysis. Opening one component can widen the range of viable openness configurations elsewhere—interoperable cloud interfaces or transparent training pipelines can magnify the benefits

of selectively releasing model weights or partially open datasets.²⁸¹ But the reverse is just as common: A single closed bottleneck in compute, training infrastructure, or deployment environments can sharply limit what openness at other layers can achieve, even when the balance of trade-offs at those layers is independently well chosen.²⁸² A decision to adopt a moderately open application layer, for instance, cannot deliver meaningful user autonomy if the model beneath it is accessible only through a closed API that enforces opaque safety or logging requirements.²⁸³

The interplay across these spectrums produces a large and dynamic design space—what is effectively a permutation problem. Some downstream or upstream effects can be anticipated with reasonable confidence: Tightening export controls on compute will narrow which models can be trained²⁸⁴ and will likely push developers toward closed deployment architectures.²⁸⁵ But others cannot be reliably predicted until they emerge in practice, especially when different actors control different parts of the stack and adjust their behavior in response to each other.²⁸⁶

281. See Narechania & Sitaraman, *supra* note 32, at 105 (explaining that the true value of openness cannot be unlocked without openness at the compute layer); see, e.g., *Together*, HUGGING FACE, <https://huggingface.co/togethercomputer> [<https://perma.cc/Y3LU-VLA3>] (last visited Jan. 31, 2026) (“Together is building the first decentralized cloud dedicated to efficiently working with large foundation models. Together will enable researchers, developers and companies to leverage and improve artificial intelligence with an intuitive platform combining data, models and computation.”).

282. Matthew Solnik, *Open R1 vs DeepSeek: Security Implications of Open vs. Closed Data*, WITNESSAI (Mar. 7, 2025), <https://witness.ai/blog-open-r1-vs-deepseek-security-implications-of-open-vs-closed-data/> [<https://perma.cc/2N37-69S5>] (explaining that, although DeepSeek is touted as one of the most “open” AI systems available, its closed training dataset and pipeline limit inspection, testing, and reproducibility).

283. See Dale Wesdorp, *Why You Should Not Build Your Application on Top of OpenAI’s APIs*, MIYAGAMI (Jan. 15, 2025), <https://www.miyagami.com/insights/why-not-build-application-on-openai> [<https://perma.cc/CA37-G7D4>].

284. See generally Janet Egan & Lennart Heim, Oversight for Frontier AI Through a Know-Your-Customer Scheme for Compute Providers (Oct. 20, 2023) (unpublished manuscript), <https://arxiv.org/pdf/2310.13625> [<https://perma.cc/K3MA-S4TA>] (describing how export controls on compute are driving developers to pivot towards smaller models that require less compute power).

285. See Ritwik Gupta, Leah Walker & Andrew W. Reddie, *Whack-a-Chip: The Futility of Hardware-Centric Export Controls*, UC BERKELEY RISK & SEC. LAB (Nov. 21, 2024), <https://arxiv.org/pdf/2411.14425> [<https://perma.cc/EWK7-ZKZA>] (explaining how compliance with regulations like export controls requires a great deal of oversight and control over downstream uses of a product or service, which is easier to employ in closed systems).

286. For example, Apple optimized its phone hardware to run Stable Diffusion, a popular open AI image-generation model, counterintuitively promoting an osAI system rather than developing its own proprietary version, but also advancing its closed hardware ecosystem because users can only benefit from the Stable Diffusion optimization by using the model on Apple devices. Atila Orhon, Michael Siracusa & Aseem Wadhwa, *Stable*

These realities introduce a fundamental uncertainty into navigating trade-offs. Policymakers are not simply choosing the “right” openness position for each component; they are choosing those positions in a system where each component’s value depends on the evolving openness of the others. A differential openness framework does not eliminate this uncertainty, but it makes the interactions legible enough that policymakers can revisit and recalibrate earlier choices as consequences surface. In this sense, interdependence is not a reason to avoid making trade-offs, but a reason to approach them with humility, to expect both intended and unintended cross-layer effects, and to design policies with the flexibility required for revision over time.

III. CALIBRATING DIFFERENTIAL AI OPENNESS

The default posture toward osAI today is largely reactive, shaped more by commercial strategy and technical happenstance than by deliberate governance. As Part II demonstrated, the openness of specific components—from model weights to datasets—has direct, often conflicting implications for safety, innovation, and accountability.

This Part sketches a research agenda for calibrating these trade-offs. We do not attempt a comprehensive analysis of the complex legal regimes governing AI; such a task is beyond the scope of this Article. Instead, we identify five key policy levers—liability, competition, intellectual property, trade, and government support—where a component-specific analysis is most urgent. In each area, we show how current frameworks fail to account for the nuance of the AI stack and suggest how the differential openness framework can help researchers and policymakers move beyond blunt, system-level mandates toward more precise, component-level governance.

Diffusion with Core ML on Apple Silicon, APPLE MACH. LEARNING RSCH. (Dec. 2022), <https://machinelearning.apple.com/research/stable-diffusion-coreml-apple-silicon> [https://perma.cc/ZP4K-MXER]. Separately, both Google and Microsoft have responded to the exploding growth of the osAI ecosystem by announcing partnerships with Hugging Face that entail redesigning aspects of their compute products and cloud services to be more compatible with open models. Google will reconfigure its cloud services to improve speeds for Hugging Face models and will offer native support for TPUs for all Hugging Face open models, allowing osAI developers to run their models on Google TPUs and not just Nvidia GPUs. Ryan J. Salva, *Expanding Support for AI Developers on Hugging Face*, GOOGLE CLOUD: BLOG (Nov. 13, 2025), <https://cloud.google.com/blog/products/ai-machine-learning/expanding-support-for-ai-developers-on-hugging-face> [https://perma.cc/96WX-Z2BV]. Microsoft similarly announced it has reconfigured its systems to facilitate running Hugging Face models on Azure cloud services. Jeff Boudier, Simon Pagezy & Alvaro Bartolome, *Microsoft and Hugging Face Expand Collaboration to Make Open Models Easy to Use on Azure*, HUGGING FACE (May 19, 2025), <https://huggingface.co/blog/azure-ai-foundry> [https://perma.cc/L3FY-W3KV].

A. Liability

Liability frameworks are among the most powerful tools for shaping behavior in technological ecosystems, yet their current application to osAI creates perverse incentives that misalign with public safety and innovation. These tensions arise because liability doctrines rarely distinguish between components, even though the taxonomy shows that transparency at some layers—like documentation or system prompts—improves oversight, while openness at others—like raw model weights or datasets—eliminates meaningful control and amplifies risk. A critical research question is how to recalibrate liability rules to encourage responsible transparency while discouraging reckless proliferation, in ways that reflect both the spectrum of openness within each component and the interdependence among them.

Currently, the legal baseline often incentivizes the wrong kind of openness. Developers who release powerful model weights or un-curated datasets without safeguards may find shelter in established common law defenses. Tort law’s economic loss doctrine generally blocks recovery for purely financial harm,²⁸⁷ while open source licenses almost universally disclaim warranties and liability.²⁸⁸ Furthermore, by releasing components “as is”²⁸⁹ and forfeiting control over downstream use, developers can argue they lack the proximate cause required for negligence or that, having forfeited control over downstream uses, they cannot be held responsible for failing to do the impossible: include non-removable safety features in a system designed to be modified and stripped of safeguards.²⁹⁰ This effect is amplified by the decentralized nature of osAI development: Once a component spreads into the ecosystem, harms often cannot be traced to a single actor with control over its deployment, making negligence and causation doctrines difficult

287. Michael D. Scott, *Tort Liability for Vendors of Insecure Software: Has the Time Finally Come?*, 67 MD. L. REV. 425, 470 (2008).

288. See, e.g., *GNU General Public License*, GNU OPERATING SYS., <https://www.gnu.org/licenses/gpl-3.0.en.html> [<https://perma.cc/6WZ6-CS8M>] (last visited Jan. 31, 2026) (providing software “‘AS IS’ WITHOUT WARRANTY OF ANY KIND”); *The MIT License*, OPEN SOURCE INITIATIVE, <https://opensource.org/license/mit> [<https://perma.cc/5N6Y-SVTQ>] (last visited Jan. 31, 2026) (similar); *Apache License, Version 2.0*, APACHE SOFTWARE FOUND., <https://www.apache.org/licenses/LICENSE-2.0> [<https://perma.cc/JL37-5QX8>] (last visited Jan. 31, 2026) (similar).

289. See Choi, *supra* note 54 (manuscript at 6–7); see also Michael L. Rustad & Thomas H. Koenig, *The Tort of Negligent Enablement of Cybercrime*, 20 BERKELEY TECH. L.J. 1553, 1565 (2005) (noting courts’ position “in favor of broad enforceability of mass market license agreements”).

290. See KETAN RAMAKRISHNAN, GREGORY SMITH & CONOR DOWNEY, RAND, U.S. TORT LIABILITY FOR LARGE-SCALE ARTIFICIAL INTELLIGENCE DAMAGES: A PRIMER FOR DEVELOPERS AND POLICYMAKERS 29 (2024), https://www.rand.org/pubs/research_reports/RRA3084-1.html [<https://perma.cc/N8XL-BFRL>].

to satisfy.²⁹¹ Paradoxically, this means that the most irreversible form of openness—releasing the untangled components of a dangerous system—may carry the least legal risk.

Conversely, liability rules can chill the forms of openness most vital for accountability. The transparency inherent in open system prompts, operational logs, or internal safety evaluations creates a “paper trail” that plaintiffs can use to build a case for negligence. Developers may thus be incentivized to keep these components closed, creating “black box” systems where failures are harder to trace. For commercial users, the picture is equally complicated: Greater transparency in certain components can signal an expectation that deployers will review available documentation, vulnerability disclosures, or safety evaluations before relying on an open component, but courts have not clarified when such transparency actually creates a downstream duty of care. And even where such duties are recognized, users may be unable to satisfy them if other layers—such as training datasets, alignment pipelines, or client-side application code—remain closed, underscoring the need for liability rules that align inspection obligations with the taxonomy of the osAI stack and scale them to the partial or uneven access that different openness configurations allow.

The challenge for legal scholars and policymakers is to design a liability regime that targets components with precision. Ideally, such a regime would penalize the reckless distribution of unsecured, high-risk components—perhaps by narrowing defenses for emotional²⁹² or physical harm²⁹³—while creating safe harbors for the transparency needed to audit and improve them. Any such regime will also need to differentiate among intermediate configurations of openness—for example, audit-only access to datasets, limited-weight releases that reduce extractability, or sandboxed application-layer interfaces—rather than assuming that fully open and fully closed components exhaust the relevant legal categories. Yet, this calibration is fraught with risk: Any expansion of liability threatens to crush the community-driven projects that drive open innovation. Unlike large corporations, individual and small-team osAI developers lack the resources to weather litigation, meaning that a blunt liability regime could effectively deter new entrants and consolidate power in the hands of incumbent firms. Recent legislative interventions,

291. *See id.* at 24.

292. *See, e.g., Garcia v. Character Techs., Inc.*, 785 F. Supp. 3d 1157, 1181 (M.D. Fla. 2025) (permitting a lawsuit against AI chatbot company alleged to have caused a user’s suicide).

293. *See* Scott, *supra* note 287, at 471 (“Arguments can be made, however, that some claims arising from the failure of security software should be recoverable despite the economic loss rule. For example, a company’s reputation is an interest protected by tort law. Additionally, the data contained in the computer is property separate and apart from the software itself.” (cleaned up)).

such as the accountability frameworks debated in California’s Senate Bill 1047²⁹⁴ or the EU AI Act,²⁹⁵ illustrate this difficulty; they attempt to impose forward-looking duties but often struggle to distinguish between pre-commercial developers and the “economic operators” who actually commercialize risks.²⁹⁶ This challenge is heightened by the resource disparity within the osAI ecosystem: Community-driven projects typically lack the financial reserves or insurance coverage needed to absorb litigation risk, meaning poorly calibrated rules could suppress precisely the forms of open, non-commercial innovation that the taxonomy identifies as most socially valuable.

A more surgical approach might shift focus to other actors interacting with different components in the stack. For example, courts could extend liability to infrastructure providers who host open models, finding proximate cause where downstream misuse was foreseeable and preventable by the platform, though this risks chilling the provision of essential services.²⁹⁷ Alternatively, policy could calibrate the liability of osAI *users* by establishing a scalable “duty to inquire,” requiring them to take precautionary measures—such as consulting model cards or vulnerability disclosures—commensurate with their deployment risk, though overly burdensome requirements could chill adoption. Reforms could also target the labor component, bolstering whistleblower protections to encourage the “opening” of internal safety concerns even when technical components remain closed, while requiring careful safeguards against abuse of such mechanisms.²⁹⁸

Even a liability regime that successfully navigates these practical hurdles may ultimately collide with constitutional limits. The First Amendment may constrain efforts to regulate the publication of certain components—such as source code, datasets, or model weights—if courts

294. S.B. 1047, 2024 Leg., Reg. Sess. (Ca. 2024).

295. 2024 O.J. (L 1689) art. 53.

296. See 2024 O.J. (L 2853) pmb. ¶¶ 2–3 (focusing liability provisions exclusively on “economic operators,” which is a capaciously defined term that does not distinguish between pre-commercial component developers and commercial manufacturers).

297. See, e.g., David Evan Harris, LINKEDIN, *Update: Hugging Face hosts child porn AI, still!* (Dec. 18, 2024, at 15:37 ET), <https://www.linkedin.com/posts/david-evanharris-mitigating-the-risk-of-generative-ai-models-activity-7275247850702217216-OK7I/> [https://perma.cc/TL25-XFNU] (reporting that Hugging Face refuses to take down projects known to be built on datasets containing high quantities of child sexual abuse material).

298. See Charlie Bullock & Mackenzie Arnold, *Protecting AI Whistleblowers*, LAWFARE (June 25, 2025, at 05:00 CT), <https://www.lawfaremedia.org/article/protecting-ai-whistleblowers> [https://perma.cc/3GKG-ABTE].

view them as protected speech under the *Bernstein*²⁹⁹ line of cases. While the contours of constitutional protection for machine-readable components remain unclear,³⁰⁰ policymakers must be prepared to navigate the murky line between speech rights and the need for osAI governance.

B. Competition

The AI ecosystem is highly concentrated. A handful of firms dominate every layer of the stack—compute infrastructure, training data, foundation models, deployment platforms, and the expert talent needed to build systems.³⁰¹ This concentration has persisted despite recent signals of more aggressive enforcement, such as the 2024 DOJ-FTC investigations into Nvidia, Microsoft, and OpenAI,³⁰² and it remains unclear whether such enforcement will continue under the new administration’s deregulatory approach.³⁰³

Whether this concentration is problematic is contested. Defenders argue it may be strategically necessary: “National champions” with massive scale can compete against state-backed rivals like China, and a smaller number of actors is easier for regulators to monitor and hold accountable.³⁰⁴ Critics counter that consolidation strips firms of

299. *Bernstein v. U.S. Dep’t of State*, 922 F. Supp. 1426, 1435 (N.D. Cal. 1996), *aff’d sub nom.*, *Bernstein v. U.S. Dep’t of Justice*, 176 F.3d 1132, 1135 (9th Cir. 1999), *reh’g granted, op. withdrawn*, 192 F.3d 1308 (mem.) (9th Cir. 1999).

300. Compare Alan Z. Rozenshtein, *There Is No General First Amendment Right to Distribute Machine-Learning Model Weights*, LAWFARE (Apr. 4, 2024, at 08:02 CT), <https://www.lawfaremedia.org/article/there-is-no-general-first-amendment-right-to-distribute-machine-learning-model-weights> [https://perma.cc/3QGU-FDTD] (arguing that model weights are not covered by the First Amendment), and Doni Bloomfield, U.S. Expert Controls of AI Models 25–32 (June 3, 2024) (unpublished manuscript), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4741033 (arguing the First Amendment permits some regulation of models), with Michael Paradis, *Regulations Targeting Large Language Models Warrant Strict Scrutiny Under the First Amendment*, LAWFARE (July 26, 2024, at 12:00 CT), <https://www.lawfaremedia.org/article/regulations-targeting-large-language-models-warrant-strict-scrutiny-under-the-first-amendment> [https://perma.cc/BS23-6WW2] (arguing that the First Amendment does not permit most regulations of models).

301. See Narechania & Sitaraman, *supra* note 32, at 99–100.

302. See David McCabe, *U.S. Clears Way for Antitrust Inquiries of Nvidia, Microsoft and OpenAI*, N.Y. TIMES (June 5, 2024), <https://www.nytimes.com/2024/06/05/technology/nvidia-microsoft-openai-antitrust-doj-ftc.html>.

303. See David McCabe, *What to Know About Trump’s Antitrust Efforts Against Tech Giants*, N.Y. TIMES (Apr. 21, 2025), <https://www.nytimes.com/2025/04/13/technology/trump-tech-antitrust-cases.html>.

304. See, e.g., Dakota Foster, *Antitrust Investigations Have Deep Implications for AI and National Security*, BROOKINGS INST. (June 2, 2020), <https://www.brookings.edu/articles/antitrust-investigations-have-deep-implications-for-ai-and-national-security/> [https://perma.cc/WQC3-PCUG]; see also SATYA MARAR, MERCATUS

competitive pressure to invest in safety or to prioritize democratic accountability over market control,³⁰⁵ while concentrating critical national capabilities in a few high-value targets vulnerable to disruption.³⁰⁶

For policymakers seeking to reduce concentration, the challenge is applying component-level analysis. Effective competition policy must recognize how interdependence across the stack constrains what any single intervention can accomplish, because choices at one layer shape competitive conditions at others—opening model weights achieves little if compute remains controlled by a handful of firms,³⁰⁷ while even closed models may support meaningful competition if application-layer interfaces and plugin ecosystems are open. True decentralization requires addressing multiple layers: the ability to access and build upon core technical components and the mobility of human talent that fuels innovation.

Competition policy offers several potential interventions, each with significant trade-offs. Ex ante structural rules could mandate separation, preventing firms from operating across multiple stack layers to avoid self-preferencing, raising prices for rivals, or using visibility into downstream usage to replicate successful applications³⁰⁸—though this sacrifices integration efficiencies and may fragment oversight.³⁰⁹ Alternatively, nondiscrimination requirements could compel dominant providers to offer fair terms to competitors,³¹⁰ or interoperability mandates could standardize technical connections to reduce vendor lock-in,³¹¹ though both demand intensive regulatory monitoring.³¹² This is particularly salient in compute, where proprietary software stacks like CUDA effectively lock developers into a single hardware ecosystem and give incumbents structural leverage over osAI competitors.

Ex post enforcement presents different opportunities. Stricter merger controls could block both direct acquisitions and the increasingly

CTR., ARTIFICIAL INTELLIGENCE AND ANTITRUST LAW: A PRIMER 13 (2024), <https://www.mercatus.org/media/document/4815mararaiantitrustlawssv2pdf> [<https://perma.cc/5NKW-QK5Z>] (warning that competition enforcement can drive innovation overseas).

305. Narechania & Sitaraman, *supra* note 32, at 140–43.

306. *Id.* at 139; Press Release, U.S. Dep’t of the Treasury, *supra* note 280.

307. *See supra* Section II.B.2.

308. Narechsnia & Sitaraman, *supra* note 32, at 159–60.

309. *See* Richard J. Gilbert, *Separation: A Cure for Abuse of Platform Dominance?*, INFO. ECON. & POL’Y, Mar. 2021, at 1, 11 (“Some types of innovations require coordination between complementary businesses, which is impeded if the businesses are confined to separate companies.”).

310. Narechania & Sitaraman, *supra* note 32, at 160–62.

311. *See, e.g.,* *Introducing the Model Context Protocol*, *supra* note 246.

312. Narechania & Sitaraman, *supra* note 32, at 161.

common “quasi-acquisitions” through major investments and licensing deals,³¹³ while heightened scrutiny could target predatory pricing or the practice of offering initially open APIs before exploiting third-party dependence.³¹⁴ Regulators might also treat “openwashing”—claiming openness benefits while withholding critical components—as deceptive trade practices, using clear definitional benchmarks.³¹⁵ Yet, enforcement must remain calibrated: Blocking mergers between major data providers may counter oligopoly, while blocking all startup acquisitions could eliminate the “exit” incentive that attracts founders to the ecosystem.

Finally, competition policy intersects with the labor component through employment restrictions like noncompete agreements. The debate here mirrors the broader concentration question. Critics argue that noncompetes restrict the mobility of expertise and slow knowledge diffusion, pointing to Silicon Valley’s innovation success under California’s ban.³¹⁶ The FTC has endorsed this view, emphasizing that innovation requires “talented individuals with innovative ideas be permitted to move freely.”³¹⁷ Yet, defenders of noncompetes counter that such agreements protect firms’ investments in training and proprietary knowledge, incentivizing employers to develop talent that they might otherwise hoard or underinvest in.³¹⁸ The optimal approach to labor mobility in the AI stack—like concentration more broadly—remains an open question for policymakers.

313. See, e.g., Press Release, Fed. Trade Comm’n, *FTC Sues to Block \$40 Billion Semiconductor Chip Merger* (Dec. 2, 2021), <https://www.ftc.gov/news-events/news/press-releases/2021/12/ftc-sues-block-40-billion-semiconductor-chip-merger> [<https://perma.cc/3LA7-65SE>] (detailing FTC lawsuit to block Nvidia’s purchase of UK chip designer Arm); Mike Isaac, *Cognition AI Buys Windsurf as A.I. Frenzy Escalates*, N.Y. TIMES (July 14, 2025), <https://www.nytimes.com/2025/07/14/technology/cognition-ai-windsurf.html> (describing Google’s takeover of AI coding startup Windsurf by poaching its executives and top talent as well as licensing its technology, leaving what was left to either wither or, as they ultimately chose to, agree to an outright acquisition by AI coding competitor Cognition).

314. See generally Chinmayi Sharma, *Concentrated Digital Markets, Restrictive APIs, and the Fight for Internet Interoperability*, 50 U. MEM. L. REV. 441 (2019) (discussing this trend in other software environments).

315. See 15 U.S.C. § 45(a).

316. See Ronald J. Gilson, *The Legal Infrastructure of High Technology Industrial Districts: Silicon Valley, Route 128, and Covenants Not to Compete*, 74 N.Y.U. L. REV. 575, 589–91 (1999).

317. *Generative AI Raises Competition Concerns*, FTC: OFF. TECH. BLOG (June 29, 2023), <https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2023/06/generative-ai-raises-competition-concerns> [<https://perma.cc/ZX7Z-V9CQ>].

318. See generally Paul H. Rubin & Peter Shedd, *Human Capital and Covenants Not to Compete*, 10 J. LEGAL STUD. 93 (1981) (offering support for covenants not to compete).

C. Intellectual Property

Copyright law creates a perilous asymmetry for osAI development. While proprietary developers shield their training data and pipelines as trade secrets, transparent osAI projects must document methods, data sources, and design choices—transparency critical for reproducibility and accountability that simultaneously creates a roadmap for infringement litigation.³¹⁹ At the same time, while large corporations can absorb lawsuits as routine costs and indemnify customers,³²⁰ the same litigation can extinguish smaller community-driven projects. This dynamic incentivizes developers to withhold valuable components to avoid liability, chilling the openness essential for safety research and innovation.

The primary defense, fair use, offers uncertain protection. Fair use is an affirmative defense to copyright infringement that courts assess using four statutory factors,³²¹ but in practice the analysis often centers on two competing considerations: whether the use transforms copyrighted material into something new rather than merely substituting for it (favoring uses like research and education),³²² and whether the use harms the market for the original work (protecting creators' economic interests).³²³ This creates conflicting pressures for osAI: Transparency and collaborative goals support transformative use claims,³²⁴ yet open datasets containing copyrighted works risk market substitution,³²⁵ and jailbroken model weights can extract near-copies of training data, directly

319. See Mark A. Lemley & Bryan Casey, *Fair Learning*, 99 TEX. L. REV. 743, 746 (2021). There have been some recent attempts to train AI systems entirely on public domain content. See, e.g., Nikhil Kandpal et al., *The Common Pile v0.1: An 8TB Dataset of Public Domain and Openly Licensed Text 2* (June 5, 2025) (unpublished manuscript), <https://arxiv.org/pdf/2506.05209> [<https://perma.cc/H3FZ-FM4G>]. It is unclear, however, whether such projects will be competitive with the most advanced AI models. See Nitasha Tiku, *AI Firms Say They Can't Respect Copyright. These Researchers Tried.*, WASH. POST (June 5, 2025), <https://www.washingtonpost.com/politics/2025/06/05/tech-brief-ai-copyright-report/> [<https://perma.cc/7EFW-3NSJ>].

320. See, e.g., Ron Miller, *Adobe Indemnity Clause Designed to Ease Enterprise Fears About AI-Generated Art*, TECHCRUNCH (June 26, 2023, at 03:13 PT), <https://techcrunch.com/2023/06/26/adobe-indemnity-clause-designed-to-ease-enterprise-fears-about-ai-generated-art/> [<https://perma.cc/U2SC-JY92>].

321. 17 U.S.C. § 107.

322. § 107(1).

323. § 107(4).

324. See *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569, 579 (1994).

325. See, e.g., *Harper & Row, Publishers, Inc. v. Nation Enters.*, 471 U.S. 539, 566 (1985) (describing the market effect factor as “undoubtedly the single most important element of fair use”); see also Suneal Bedi & Mike Schuster, *Measuring Fair Use's Market Effect*, 2022 WIS. L. REV. 1467, 1469 (“Although courts often employ all four factors, one has received substantial attention in fair use determinations—the so called ‘market effects’ . . . factor.” (footnote omitted)).

competing with originals.³²⁶ With legal scholarship divided³²⁷ and caselaw nascent,³²⁸ fair use functions more as a litigation gamble than a safe harbor.

The challenge for policymakers is designing interventions that account for component-level differences. Legislative clarification of fair use application—following jurisdictions like the EU,³²⁹ Japan,³³⁰ and Singapore³³¹—could direct courts to weigh public benefits differently across components: favoring transparency in system architecture and operational records while maintaining protection against market harm from raw training data distribution. But as the taxonomy in Part I suggests, components do not fall into simple categories of “open” or “closed.”³³² Between fully proprietary corpora and fully open datasets or model weights lie a range of intermediate configurations—licensed access under non-redistribution terms, filtered or synthetic corpora designed to reduce exposure to copyrighted works, or limited-weight releases that

326. See, e.g., Nicholas Carlini et al., *Extracting Training Data from Diffusion Models*, 32 PROCS. USENIX SEC. SYMP. 5253, 5253–54 (2023), <https://www.usenix.org/system/files/usenixsecurity23-carlini.pdf> [<https://perma.cc/2QA4-3K9D>].

327. Compare, e.g., Lemley & Casey, *supra* note 319, at 748–50 (arguing that copying for the purpose of training machine learning models should generally be protected by fair use), and James Grimmelmann, *Copyright for Literate Robots*, 101 IOWA L. REV. 657, 664 (2016) (“Verbatim copying of a complete work will be protected as fair use if the copy is used solely as input to a process that does not itself use the works expressively. Or, to put it a little more provocatively, nonexpressive uses do not count as reading.”), with Robert Brauneis, *Copyright and the Training of Human Authors and Generative Machines*, 48 COLUM. J.L. & ARTS 1, 58–59 (2024) (“The current case that generative AI training is a fair use is weak.”). See also U.S. COPYRIGHT OFF., COPYRIGHT AND ARTIFICIAL INTELLIGENCE, PART 3: GENERATIVE AI TRAINING 74 (pre-publication ed. 2025), <https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-3-Generative-AI-Training-Report-Pre-Publication-Version.pdf> [<https://perma.cc/LDF7-7RGE>] (“[T]he copying of expressive works from pirate sources in order to generate unrestricted content that competes in the marketplace, when licensing is reasonably available, is unlikely to qualify as fair use.”).

328. See generally *Thomson Reuters Enter. Ctr. GMBH v. Ross Intel. Inc.*, 765 F. Supp. 3d 382 (D. Del. 2025) (holding that fair-use defense did not apply to a second AI legal research platform trained on data compilation built from first AI legal research platform), *motion to certify appeal granted*, No. 20-CV-613, 2025 WL 1488015 (D. Del. May 23, 2025); *Bartz v. Anthropic PBC*, 787 F. Supp. 3d 1007 (N.D. Cal. 2025) (finding transformative fair use for AI training and digitizing purchased books, but not for retaining pirated copies); *Kadrey v. Meta Platforms, Inc.*, 787 F. Supp. 3d 1026 (N.D. Cal. 2025) (holding the use of copyrighted books for AI training to be a transformative fair use).

329. Council Directive 2019/790, arts. 3–4, 2019 O.J. (L 130) 92.

330. Chosakuken Hō [Copyright Act], Law No. 48 of 1970, art. 30-4, *translated in* (Japanese Law Translation [JLT DS]), https://www.japaneselawtranslation.go.jp/en/laws/view/4207#je_ch2sc3sb5at4 [<https://perma.cc/54DQ-U47M>].

331. Copyright Act 2021 (No. 22 of 2021) § 244.

332. See *supra* Part I.

prevent extraction of training data.³³³ These middle-ground arrangements change both the substitution risk and the public-interest value of openness, and a credible fair-use reform must therefore explicitly account for these gradations rather than assuming that all openness carries identical legal consequences. These reforms must also distinguish between strategically open releases by commercial actors seeking market entrenchment³³⁴ and genuine non-commercial academic projects.

Alternative approaches present different trade-offs. A compulsory licensing regime, analogous to digital music streaming,³³⁵ would require rights-holders to make works available for AI training under standardized terms and could establish tiered fees favoring academic projects.³³⁶ This, however, requires complex administration and may inadequately compensate creators. Other interventions include safe harbors for reproducibility-supporting components,³³⁷ text-to-data-mining exceptions for specific users and purposes,³³⁸ or procedural reforms like fee-shifting provisions deterring frivolous suits,³³⁹ each of which could protect smaller developers yet risks creating loopholes that undermine legitimate copyright enforcement. Here too, the policy analysis must track the taxonomy: Interventions will operate differently when applied to raw datasets, to partially synthetic corpora, or to model-weight releases with varying degrees of extractability. Calibrating copyright rules to these intermediary forms of openness is essential to avoid pushing developers toward either total secrecy or reckless full public release, preserving a spectrum of responsible disclosure that supports safety, reproducibility, and competition. The optimal intellectual property regime remains contested, requiring calibration that neither sacrifices creator rights nor

333. See *supra* Part I.

334. See *supra* Section I.A.1.

335. See 17 U.S.C. § 115; see also Mariana L. Orbay, *Songwriters v. Spotify: Is Spotify the Problem or a Symptom of the Problem?*, 48 PEPP. L. REV. 785, 796–804 (2021) (describing statutory music copyright regimes).

336. See, e.g., *Artificial Intelligence and Copyright*, 88 Fed. Reg. 59942, 59947 (Aug. 30, 2023).

337. See *generally* ELEONORA ROSATI, EUR. PARL., PE 604.942, THE EXCEPTION FOR TEXT AND DATA MINING (TDM) IN THE PROPOSED DIRECTIVE ON COPYRIGHT IN THE DIGITAL SINGLE MARKET—TECHNICAL ASPECTS (2018), [https://www.europarl.europa.eu/RegData/etudes/BRIE/2018/604942/IPOL_BRI\(2018\)604942_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/BRIE/2018/604942/IPOL_BRI(2018)604942_EN.pdf) (describing how text and data mining (TDM) exceptions to copyright laws can promote research and innovation).

338. See Matthew Sag & Peter K. Yu, *The Globalization of Copyright Exceptions for AI Training*, 74 EMORY L.J. 1163, 1184–91 (2025).

339. See 17 U.S.C. § 505 (permitting fee-shifting in copyright cases); see *generally* David E. Shipley, *Discouraging Frivolous Copyright Infringement Claims: Fee Shifting Under Rule 11 or 28 U.S.C. § 1927 as an Alternative to Awarding Attorney’s Fees Under Section 505 of the Copyright Act*, 24 J. INTELL. PROP. L. 33 (2016) (discussing several ways that attorney’s fees can be awarded in copyright infringement suits).

stifles the component-specific transparency essential for safety and competition.

D. Trade

Current U.S. export policy exempts open-weight AI models from the restrictions imposed on advanced closed-weight models,³⁴⁰ reflecting a strategic calculation that American competitive advantage lies in its open research ecosystem.³⁴¹ This approach prioritizes domestic innovation over proliferation concerns.³⁴² Policymakers recognize that effective restrictions on foreign dissemination would require restricting internal U.S. development as well—given the global nature of open development on the internet—thereby stifling domestic progress.³⁴³ They also note the lack of evidence that current open models meaningfully increase adversary capabilities.³⁴⁴ However, the rapid rise of powerful Chinese open models—DeepSeek-R1 and Qwen3³⁴⁵—suggests that the capability gap between open and closed systems may be closing, forcing a reconsideration of this balance.

Export controls on osAI face severe obstacles beyond just domestic innovation. Geopolitically, making U.S. open innovation inaccessible would drive smaller nations—unable to adapt frontier technology for local needs—toward alternative, potentially adversarial ecosystems, weakening democratic technology partnerships. Constitutionally, the *Bernstein* cases suggest that source code can be protected speech and that

340. Framework for Artificial Intelligence Diffusion, 90 Fed. Reg. 4544, 4544, 4547 (Jan. 15, 2025).

341. See John Villasenor, *The Tension Between AI Export Control and U.S. AI Innovation*, BROOKINGS INST. (Sep. 24, 2024), <https://www.brookings.edu/articles/the-tension-between-ai-export-control-and-u-s-ai-innovation/> [https://perma.cc/9BBY-Y7PU].

342. Crucially, any attempt to restrict open-weight publication would trigger the “deemed export” rules of the Export Administration Regulations, which treat the release of controlled technology to foreign persons as an export. 15 C.F.R. § 734.13(b) (2026) (“Any release in the United States of ‘technology’ or source code to a foreign person is a deemed export to the foreign person’s most recent country of citizenship or permanent residency.”). Because osAI development is conducted on globally accessible platforms like GitHub and Hugging Face, such restrictions would functionally prohibit open publication altogether, collapsing domestic open research as surely as restricting exports abroad.

343. ANGELA LUNA, AM. ACTION F., *AI EXPORT CONTROLS: BALANCING NATIONAL SECURITY AND AI INNOVATION 3* (2024), <https://www.americanactionforum.org/insight/ai-export-controls-balancing-national-security-and-ai-innovation/> [https://perma.cc/T6Y3-N422].

344. MOUTON, LUCAS & GUEST, *supra* note 265, at 1.

345. *Leaderboard Overview*, *supra* note 249.

export restrictions may constitute unconstitutional prior restraints,³⁴⁶ raising serious questions—even if the extension to non-human-readable model weights remains uncertain—about treating their publication as regulatable conduct rather than protected expression.

Alternative regulatory strategies present their own challenges. Compute hardware offers a more controllable choke point—the Biden administration’s chip controls³⁴⁷ aimed to limit adversary access to critical training infrastructure—but this strategy proved politically fragile, with the Trump administration rescinding key restrictions³⁴⁸ and approving substantial investments in Saudi Arabia and the UAE.³⁴⁹

Because compute is the most physically controllable component of the osAI stack,³⁵⁰ export controls imposed at this layer reshape the effective openness of components above it. Even when model weights or datasets are formally open, their functional openness depends on access to the compute required to train, fine-tune, or meaningfully run them.³⁵¹ By restricting high-end accelerators, export controls can turn nominally open components into de facto closed ones for foreign actors, curtailing proliferation risks of open components higher in the stack than compute, such as data or weights, while simultaneously increasing dependence on U.S.-based cloud providers who remain subject to domestic monitoring and licensing regimes.³⁵² Conversely, loosening compute controls—

346. *Bernstein v. U.S. Dep’t of State*, 922 F. Supp. 1426, 1435, 1437–38 (N.D. Cal. 1996), *aff’d sub nom.*, *Bernstein v. U.S. Dep’t of Justice*, 176 F.3d 1132, 1135 (9th Cir. 1999), *reh’g granted, op. withdrawn*, 192 F.3d 1308 (mem.) (9th Cir. 1999).

347. Implementation of Additional Export Controls: Certain Advanced Computing Items; Supercomputer and Semiconductor End Use; Updates and Corrections, 88 Fed. Reg. 73458 (Oct. 25, 2023) (codified at 15 C.F.R. §§ 732, 734, 736, 740, 742, 744, 746, 748, 758, 770, 772, 774 (2026)).

348. Press Release, Bureau of Indus. & Sec., U.S. Dep’t of Com., Department of Commerce Announces Recission of Biden-Era Artificial Intelligence Diffusion Rule, Strengthens Chip-Related Export Controls (May 13, 2025), <https://www.bis.gov/press-release/department-commerce-announces-recission-biden-era-artificial-intelligence-diffusion-rule-strengthens-chip> [<https://perma.cc/TW7E-2Q6N>]; *see also* David Sacks (@davidsacks47), X, *TO WIN THE AI RACE, THE AI DIFFUSION RULE MUST GO* (May 8, 2025, at 13:15 CT), <https://x.com/davidsacks47/status/1920543150449569992> [<https://perma.cc/R3XD-SY6X>] (articulating the political resistance the original Biden regulations received).

349. *See* Tripp Mickle & Ana Swanson, *Outsourcer in Chief: Is Trump Trading Away America’s Tech Future?*, N.Y. TIMES (May 15, 2025), <https://www.nytimes.com/2025/05/15/business/economy/trump-chips-ai-uae.html>.

350. *See* Lennart Heim, Markus Anderljung, Emma Bluemke & Robert Trager, *Computing Power and the Governance of AI*, GOVAI (Feb. 14, 2024), <https://www.governance.ai/analysis/computing-power-and-the-governance-of-ai> [<https://perma.cc/US25-SYST>].

351. *See* Narechania & Sitaraman, *supra* note 32, at 120.

352. *See* Lennart Heim, *China’s AI Models Are Closing the Gap—but America’s Real Advantage Lies Elsewhere*, RAND (May 2, 2025), <https://www.rand.org/perspectives/2025/05/02/china-ai-models-are-closing-the-gap-but-america-s-real-advantage-lies-elsewhere>.

whether through exemptions, foreign investment approvals, or expanded licensing—can dramatically increase the reach and usability of downstream osAI components. In this way, compute policy operates as a “force multiplier”: Trade restrictions at the hardware layer cascade upward, altering the practical openness, diffusion speed, and strategic significance of every component built on top of it.³⁵³ Import restrictions face even steeper obstacles: Proposals like Senator Hawley’s “DeepSeek ban”³⁵⁴ confront practical infeasibility, as controlling model weights once disseminated online requires invasive surveillance that raises First Amendment concerns.³⁵⁵

The challenge for policymakers is achieving component-level precision. Because openness at one layer shapes the practical openness of others, export controls must be calibrated with an explicit understanding of how restrictions cascade through the stack. A sustainable trade regime must therefore treat openness not as a binary but as an interdependent system, balancing national security and technological leadership by shaping the *configuration* of openness across components rather than bluntly constraining any single one. The optimal calibration of this trade-off between national security and domestic technological leadership remains a central question in osAI governance.

E. Government Support

The economics of AI development inherently favor large corporations that can afford the immense costs of compute and data

rand.org/pubs/commentary/2025/05/chinas-ai-models-are-closing-the-gap-but-americas-real.html [https://perma.cc/2QW9-5JCM].

353. Barath Harithas & Andreas Schumacher, *Where the Chips Fall: U.S. Export Controls Under the Biden Administration from 2022 to 2024*, CTR. FOR STRATEGIC & INT’L STUD. (Dec. 12, 2024), <https://www.csis.org/analysis/where-chips-fall-us-export-controls-under-biden-administration-2022-2024> [https://perma.cc/EK9E-VZE5] (“In September 2022, National Security Advisor Jake Sullivan declared that securing ‘as large a lead as possible’ in force multiplier technologies like AI was a national security imperative.”).

354. Decoupling America’s Artificial Intelligence Capabilities from China Act of 2025, S. 321, 119th Cong. (2025); *see also* Press Release, Sen. Josh Hawley, Hawley Introduces Legislation to Decouple American AI Development from Communist China (Jan. 29, 2025), <https://www.hawley.senate.gov/hawley-introduces-legislation-to-decouple-american-ai-development-from-communist-china> [https://perma.cc/X98Q-GK2Z] (announcing legislation to ban AI “import[s] from [and] export[s] to China,” “prohibit American companies from conducting AI research in China,” and bar U.S. investment in Chinese AI development).

355. *See* Alan Z. Rozenshtein, *There Is No General First Amendment Right to Distribute Machine-Learning Model Weights*, LAWFARE (Apr. 4, 2024, at 08:02 CT), <https://www.lawfaremedia.org/article/there-is-no-general-first-amendment-right-to-distribute-machine-learning-model-weights> [https://perma.cc/9DNY-62D9].

acquisition.³⁵⁶ Strategic public investment represents a potential lever to address these resource disparities across the compute, data, and labor components of the AI stack. While policy documents like the White House’s AI Action Plan call for direct government investment in fostering AI openness,³⁵⁷ actual implementation remains limited and the optimal approach contested.³⁵⁸ Public investment can support very different openness configurations—from fully open releases to controlled-access resources or shared evaluation tools—raising questions about which points along that spectrum best serve safety, innovation, democratic accountability, and national interests.

The central challenge for policymakers is whether to strengthen the existing commercial market or create genuinely public alternatives. Market-strengthening approaches—such as subsidizing semiconductor fabrication³⁵⁹ or improving the financial infrastructure for compute access³⁶⁰—leverage existing capacity but tend to flow disproportionately to firms already positioned to capture them, reinforcing rather than redistributing market power. True public infrastructure would require different governance models, raising questions about administrative capacity and efficiency. Either way, these choices play out differently across the compute, data, and labor components of the AI stack.

One prominent proposal, the National AI Research Resource (NAIRR), would provide researchers with computing infrastructure otherwise accessible only to large firms.³⁶¹ The critical design question is whether such programs function as government-run infrastructure or as subsidy mechanisms channeling funds through incumbent cloud providers.³⁶² Government-run infrastructure could enable independent innovation and, if governed through transparent structures with civil

356. See Sharma, *AI’s Hippocratic Oath*, *supra* note 97, at 1140–42.

357. EXEC. OFF. OF THE PRESIDENT OF THE U.S., *supra* note 1, at 4–5.

358. Sharma, *AI’s Hippocratic Oath*, *supra* note 97, at 1135–42.

359. CHIPS and Science Act, Pub. L. No. 117-167, 136 Stat. 1366, 1372–79 (2022).

360. EXEC. OFF. OF THE PRESIDENT OF THE U.S., *supra* note 1, at 4–5.

361. *Id.* at 5; see NAT’L A.I. RSCH. RES. PILOT, <https://nairrpilot.org/> [<https://perma.cc/M4T8-HG3M>] (last visited Feb. 1, 2026); Madison Adler, *National Science Foundation Rolls Out NAIRR Pilot with Industry, Agency Support*, FEDSCOOP (Jan. 24, 2024), <https://fedscoop.com/nsf-launches-nairr-pilot/> [<https://perma.cc/UE35-5TCE>] (“The pilot for the resource, referred to as the NAIRR, is composed of contributions from 11 federal agencies and 25 private sector partners, including Microsoft, Amazon Web Services, NVIDIA, Intel, and IBM. Those contributions range from use of the Department of Energy’s Summit supercomputer to datasets from NASA and the National Oceanic and Atmospheric Administration to access for models from OpenAI, Anthropic, and Meta.”).

362. See Narechania & Sitaraman, *supra* note 32, at 165–66.

society representation,³⁶³ align resource allocation with democratic values—though such governance mechanisms remain largely theoretical and demand substantial technical capacity. Alternatively, subsidy programs leverage existing infrastructure but may deepen vendor dependence on platforms whose incentives diverge from the public interest. Equally important is how such programs structure openness across components—for example, whether they attach conditions to data sharing, require open documentation of experiments, or support shared application-layer safety scaffolds rather than only raw compute time.

Public resources also offer potential for the data component. Federal agencies possess high-quality datasets—scientific archives, public health records—that could serve as curated, machine-learning-ready, portable, and openly accessible alternatives to the unvetted web-scraped corpora often used to train models.³⁶⁴ A sustainable ecosystem of public-option data may also benefit from grant-supported research contributing datasets back into shared corpora in standardized, machine-learning-ready formats, rather than allowing publicly funded resources to remain siloed within individual projects. Preferential access for noncommercial developers could prevent simply subsidizing incumbents. A public data repository of this magnitude requires governance safeguards against misuse.³⁶⁵ Government procurement policy offers a complementary approach: Incorporating requirements for transparency and interoperability in public-sector AI contracts could create sustainable revenue streams for open components while setting ecosystem-wide standards, as demonstrated in open source software markets.³⁶⁶ Because procurement can specify different openness requirements for different layers—such as mandating open documentation and interoperable application interfaces even where underlying models remain

363. See Shearer, Davies & Lawrence, *supra* note 117; see, e.g., Mark Coeckelbergh, *Artificial Intelligence, the Common Good, and the Democratic Deficit in AI Governance*, 5 AI ETHICS 1491, 1494–96 (2025).

364. See *AI-Ready Open Data*, BIPARTISAN POL’Y CTR. (Feb. 17, 2023), <https://bipartisanpolicy.org/explainer/ai-ready-open-data/> (“Government’s vast amount of open data can fill this gap: McKinsey estimates that open data can help unlock \$3 trillion to \$5 trillion in economic value annually . . .”).

365. See, e.g., Chinmayi Sharma, Thomas E. Kadri & Sam Adler, *Brokering Safety*, 114 CALIF. L. REV. (forthcoming 2026) (manuscript at 38), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5143114 [<https://dx.doi.org/10.2139/ssrn.5143114>]; Janet Freilich, *Government Misinformation Platforms*, 172 U. PA. L. REV. 1537, 1540–45 (2024).

366. See, e.g., Iain G. Mitchell, *Public Sector and Open Source*, in OPEN SOURCE LAW, POLICY AND PRACTICE 429, 453–65 (Amanda Brock ed., 2d ed. 2022); Eunice Mercado-Lara, Shannon Dosemagen & Alison Parker, *Unlocking Innovation: Why Federal Procurement Should Embrace Open Source*, TECH POL’Y PRESS (May 23, 2025), <https://www.techpolicy.press/unlocking-innovation-why-federal-procurement-should-embrace-open-source/> [<https://perma.cc/RR98-DXM9>].

proprietary—it is a particularly flexible tool for shaping the overall configuration of differential openness across the stack.

Beyond technical components, government support could address talent concentration. Direct funding for academic institutions and nonprofits enables non-commercial AI research outside elite firms,³⁶⁷ many of which also try to lock in expertise. The government can foster collaboration by funding cross-border initiatives and mandating adherence to open-science principles in the research it supports. Cross-border initiatives like ELIAS in Europe illustrate how formal partnerships, joint training programs, and open-access publishing can foster shared expertise and prevent intellectual concentration.³⁶⁸

Immigration reform—particularly stable visa pathways for students and researchers working on public-benefit AI projects—could expand the contributor base,³⁶⁹ though such policies face significant political obstacles. Whether public investment can effectively diversify the expertise pool that shapes AI development priorities remains an open question for policymakers.

CONCLUSION

The discourse around AI openness has profound implications for the future of technology, governance, and global power dynamics. Yet, as this Article has demonstrated, the conventional framing of AI openness as a natural extension of open source software invites a misleading binary characterization of AI as “open” or “closed” and an assumption that openness is an inherent good. This misunderstanding leads to regulatory approaches that are ill-equipped to govern AI, potentially stifling innovation, undermining accountability, or creating new security risks.

Because effective governance of osAI demands the utmost precision, this Article has proposed a more sophisticated approach by untangling AI into its constituent components—compute, data, source code, model weights, operational controls, applications, and labor—and mapping each across its own spectrum of openness. In doing so, we introduce the

367. See Nathan Lambert, *The White House’s Plan for Open Models & AI Research in the U.S.*, INTERCONNECTS (July 23, 2025), <https://www.interconnects.ai/p/the-white-houses-plan-for-open-models> [<https://perma.cc/8YZ2-8FW8>].

368. See *About*, ELIAS, <https://elias-ai.eu/about/> [<https://perma.cc/4LRA-6PKK>] (last visited Feb. 1, 2026) (“ELIAS is a consortium of 34 top European institutions from 17 countries. . . . committed to advancing fundamental research in AI . . .”).

369. See ZACHARY ARNOLD, ROXANNE HESTON, REMCO ZWETSLOOT & TINA HUANG, CTR. FOR SEC. & EMERGING TECH. AT GEO. UNIV., IMMIGRATION POLICY AND THE U.S. AI SECTOR 13–14, 16 (2019); Judy Wang & Nicol Turner Lee, *Trump’s Immigration Policies May Threaten American AI Leadership*, BROOKINGS INST. (July 21, 2025), <https://www.brookings.edu/articles/trumps-immigration-policies-may-threaten-american-ai-leadership> [<https://perma.cc/BH49-8B36>].

concept of differential openness to capture the matrix of many possible permutations of component openness, each creating distinct risk profiles and governance challenges.

The different policy goals that motivate osAI regulation—such as safety, innovation, democratic access, and national security—often pull in contradictory directions, creating inevitable trade-offs. Openness is often a double-edged sword. For instance, releasing model weights can democratize access and spur innovation, but it also lowers the barrier for malicious actors and makes it impossible to recall a dangerous model once it has proliferated. The framework of disentangled AI enables policymakers to make these trade-offs explicit and to calibrate regulatory approaches—whether through liability, competition policy, intellectual property, trade, or government support—to specific components rather than applying blunt, one-size-fits-all mandates.

AI is here to stay, and with it, openness. Policymakers cannot remain blithely unaware of osAI's complexity if we are to have any hope of shaping a future that is in the public's best interest.

APPENDIX: OPENNESS OF SELECT FRONTIER MODELS³⁷⁰

This survey examines existing leading models and the openness of their components. All models with publicly available weights also disclose their architectures. “Disclosed training hardware” means the type of hardware used for training is known, even if access to that hardware is restricted—for example, by Nvidia. “Public inference hardware” indicates that the model can be run on commercially available systems. While “operational metadata” is a broad category, this classification limits it to the availability of ongoing audit logs and performance metrics. The row labeled “OSAI” refers not to a specific model, but to the Open Source Initiative’s Open Source AI definition.

Model	Training Data	Training Code	Training Hardware	Inference Code	Inference Hardware	Model Weights	Operational Metadata
OSAI ³⁷¹	Permissive License	Permissive License	Agnostic	Permissive License	Agnostic	Permissive License	Agnostic
Alibaba Qwen3 ³⁷²	Private	Private	Private	Public, Permissive License	Public	Public, Permissive License	Private
Anthropic Claude 4 ³⁷³	Private	Private	Private	Private	Private	Private	Private
DeepSeek R1 ³⁷⁴	Private	Private	Disclosed	Public, Permissive License	Public	Public, Permissive License	Private
Google Gemini 2.5 ³⁷⁵	Private	Private	Disclosed	Private	Private	Private	Private

370. We found the data in the table by looking through the myriad published sources that accompany a model’s release and operations. We looked at the website for the model, its license, the GitHub account for the organization publishing the model, and its account on Hugging Face. Most developers publish a “model card” in some form that tends to summarize the available information about their models. We conclude information is privately held when, after searching the venues where information is customarily published, we found no relevant records. Further, as practitioners in the AI field, we have high levels of familiarity with what is and isn’t being published by the frontier labs. The results of our searches are consistent with our expectations.

371. See OPEN SOURCE INITIATIVE, *supra* note 137.

372. See Qwen, *Qwen3*, HUGGING FACE, <https://huggingface.co/collections/Qwen/qwen3-67dd247413f0e2e4f653967f> (last visited Feb. 1, 2026).

373. See *Introducing Claude 4*, ANTHROPIC (May 22, 2025), <https://www.anthropic.com/news/claude-4> [<https://perma.cc/R3J7-HQD5>].

374. See DeepSeek, *DeepSeek-R1*, HUGGING FACE, <https://huggingface.co/deepseek-ai/DeepSeek-R1/> [<https://perma.cc/YE8N-4GUF>] (last visited Feb. 1, 2026).

375. See *Gemini 2.5 Pro*, GOOGLE CLOUD, <https://cloud.google.com/vertex-ai/generative-ai/docs/models/gemini/2-5-pro> [<https://perma.cc/L6QX-HQTP>] (last visited Feb. 1, 2026). See generally GHEORGHE COMANICI ET AL., GEMINI 2.5: PUSHING THE FRONTIER WITH ADVANCED REASONING, MULTIMODALITY, LONG CONTEXT, AND NEXT GENERATION AGENTIC CAPABILITIES (2025), <https://storage.googleapis.com/deepmind->

Model	Training Data	Training Code	Training Hardware	Inference Code	Inference Hardware	Model Weights	Operational Metadata
Google Gemma 3 ³⁷⁶	Private	Private	Disclosed	Public, Permissive License	Public	Public, Restricted License	Private
Meta Llama 4 ³⁷⁷	Private	Private	Disclosed	Public, Permissive License	Public	Public, Restricted License	Private
Nvidia Nemotron 4-340b ³⁷⁸	Partially Open	Permissive License	Disclosed	Public, Permissive License	Public	Public, Permissive License	Private
OpenAI o4-mini ³⁷⁹	Private	Private	Private	Private	Private	Private	Private
xAI Grok 1 ³⁸⁰	Private	Private	Private	Public, Permissive License	Public	Public, Permissive License	Private
xAI Grok 4.1 ³⁸¹	Private	Private	Private	Private	Private	Private	Private

media/gemini/gemini_v2_5_report.pdf [https://perma.cc/Q52A-RXRZ] (providing more technical detail about the model, its development, and its capabilities).

376. See Google, *Gemma 3 Release*, HUGGING FACE, <https://huggingface.co/collections/google/gemma-3-release-67c6c6f89c4f76621268bb6d> [https://perma.cc/DNF9-9HNU] (last visited Feb. 1, 2026). See generally AISHWARYA KAMATH ET AL., *GEMMA 3 TECHNICAL REPORT* (2025), <https://storage.googleapis.com/deepmind-media/gemma/Gemma3Report.pdf> [https://perma.cc/9YSX-SBDD] (providing more technical detail about the model, its development, and its capabilities).

377. See Meta Llama, *Llama 4*, HUGGING FACE, <https://huggingface.co/collections/meta-llama/llama-4-67f0c30d9fe03840bc9d0164> [https://perma.cc/SK3P-CPUJ] (last visited Feb. 1, 2026).

378. See Nvidia, *Nemotron 4 340B*, HUGGING FACE, <https://huggingface.co/collections/nvidia/nemotron-4-340b> [https://perma.cc/PG3A-P2CN] (last visited Feb. 1, 2026).

379. See *Introducing OpenAI o3 & o4-mini*, OPENAI (Apr. 16, 2025), <https://openai.com/index/introducing-o3-and-o4-mini/>.

380. See *Open Release of Grok-1*, xAI (Mar. 17, 2024), <https://x.ai/news/grok-os>.

381. xAI, *GROK 4.1 MODEL CARD* (2025), <https://data.x.ai/2025-11-17-grok-4-1-model-card.pdf> [https://perma.cc/6NUR-7YXH].