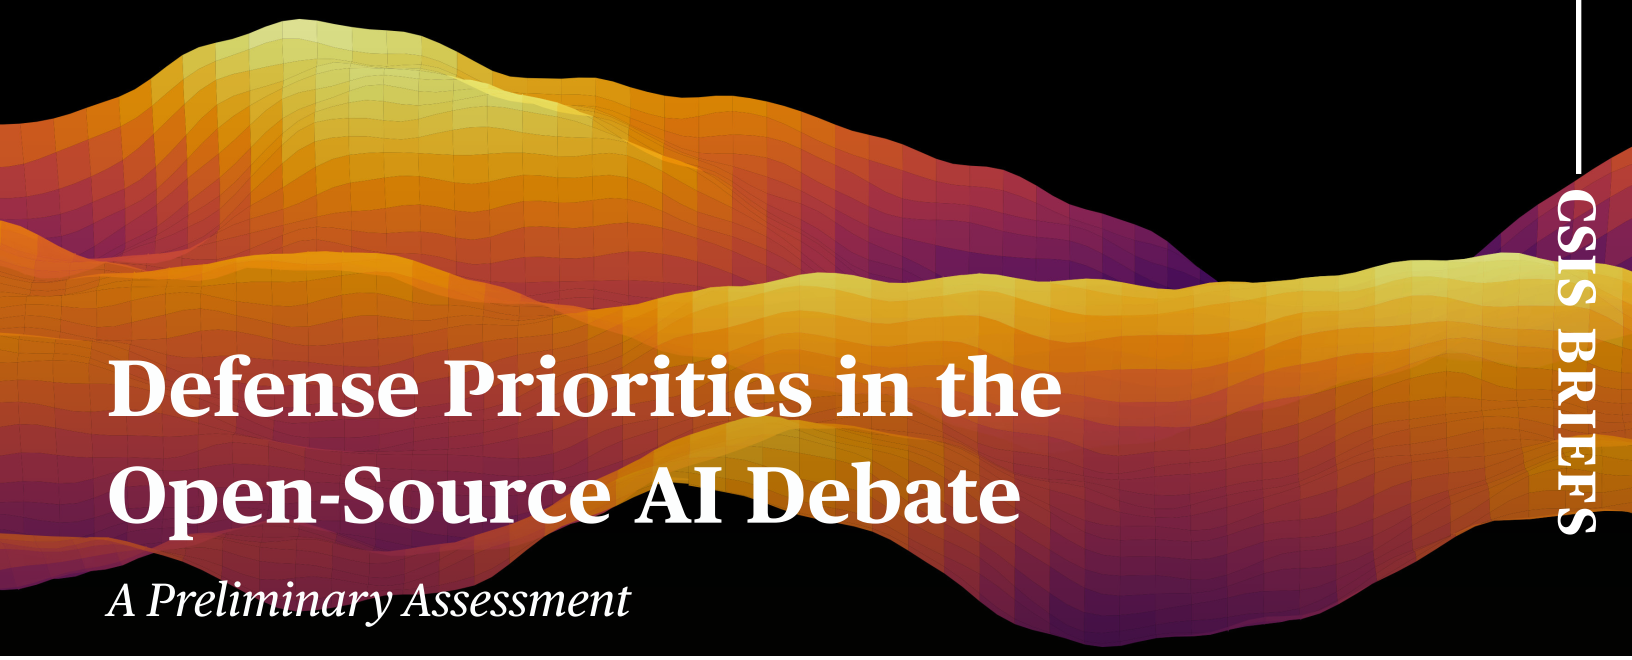

CSIS BRIEFS

Defense Priorities in the Open-Source AI Debate

A Preliminary Assessment

By Masao Dahlgren

AUGUST 2024

THE ISSUE

- **A spirited debate is taking place over the regulation of *open foundation models*—**artificial intelligence models whose underlying architectures and parameters are made public and can be inspected, modified, and run by end users.
- **Proposed limits on releasing open foundation models may have significant defense industrial impacts**. If model training is a form of defense production, these impacts deserve further scrutiny.
- **Preliminary evidence suggests that an open foundation model ecosystem could benefit the U.S. Department of Defense's supplier diversity, sustainment, cybersecurity, and innovation priorities.** Follow-on analyses should quantify impacts on acquisition cost and supply chain security.

Generative artificial intelligence (AI) has increasingly become a U.S. Department of Defense (DOD) priority. Software powered by generative AI *foundation models*—generalist systems that emulate human reasoning—might process reams of raw intelligence, automate Pentagon paperwork, or allow aircraft, trucks, and ships to navigate themselves.[1] Many advancements in this sector originate in commercial and academic research.[2] If generative AI sees wide adoption across the DOD, this base of commercial foundation model developers will become a critical part of the defense industrial base.[3] The Joint Force thus has a stake in the commercial foundation model ecosystem and how it evolves.

Indeed, DOD AI strategies hinge on continued commercial innovation in AI.[4] To that end, the Pentagon has assigned new funding to acquire AI-powered systems, such as for its Replicator drones and Joint All-Domain Command and Control battle network, and new organizations to manage them, empowering the Chief Digital and AI Office (CDAO), Task Force Lima, and others.[5]

Amid this institutional buildup, the Pentagon should appraise proposed commercial foundation model market regulations. As with spectrum auctions or shipbuilding, civil sector policymaking will shape the DOD's future choices.[6] Policies that create a competitive ecosystem of market players could improve the supply chain for future DOD programs. Conversely, policies that accelerate consolidation—like the Jones Act or the 1993 "Last Supper"—might threaten it.[7] The 2023 Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence has made unprecedented use of the Defense Production Act to regulate the commercial AI market.[8] It is now worth asking where regulation of commercial markets could affect defense production.

Figure 1: Example of 1990s-Era Defense Industrial Consolidation

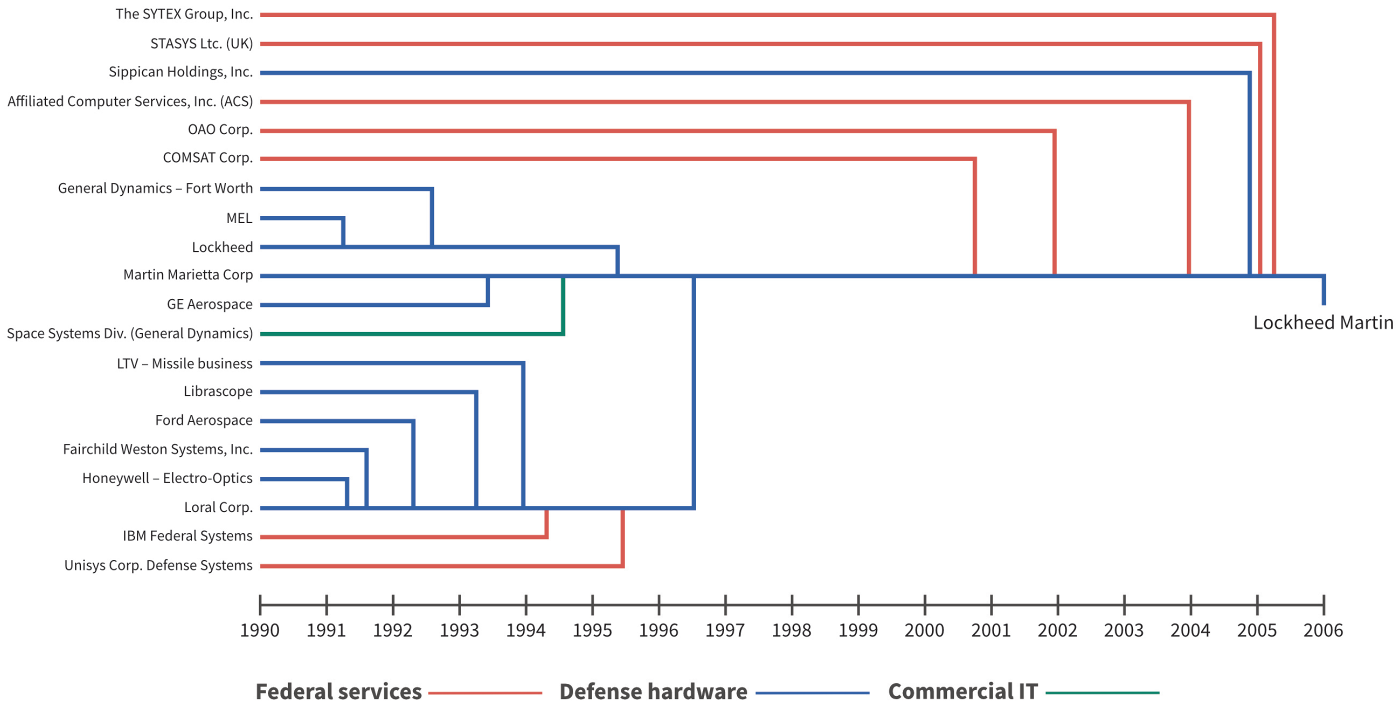

Source: CSIS.[9]

THE OPEN FOUNDATION MODEL DEBATE: A LITMUS TEST

Emerging civil society debates over AI safety—especially over open foundation models—merit particular attention. Unlike with closed models like GPT-4, developers of open foundation models like Llama, Mistral, or Qwen openly publish the models' underlying parameters ("weights"), allowing them to be inspected, modified, and operated by end users.[10] With the performance of open models approaching their closed counterparts (Figure 2), some have suggested that open model distribution could pose "extreme risks" for misuse.[11] Others, meanwhile, have highlighted open models' benefits for research, security, and national competitiveness.[12] Though outcomes remain uncertain, proposals to limit the distribution of open models—such as through California Senate Bill (SB) 1047—have recently gained legislative traction.[13]

How the open foundation model debate is resolved would have direct implications for the defense industrial base. As detailed in later sections, there are preliminary reasons to believe that a diverse open model ecosystem might benefit the DOD. The widespread availability of high-performance, open-source foundation models could improve the DOD's ability to (1) competitively source and sustain AI systems, (2) deploy AI securely, and (3) address novel use cases. Considering these impacts, the open model debate represents a test case for how civil society evaluates defense priorities in AI policy decisions.

Outlining these implications might also clarify, in White House Office of Science and Technology Policy director Arati Prabhakar's words, an often "garbled conversation about the implications, including safety implications, of AI technology."[14] Indeed, in its flagship report on the subject, the Biden administration suggested that "the government should not restrict the wide availability of model weights" but that "extrapolation based on current capabilities and limitations is too difficult to conclude whether open foundation models, overall, pose more marginal risks than benefits."[15] The administration has not endorsed open model restrictions nor foreclosed future regulation. An accounting of defense industrial benefits might therefore contribute to this ongoing conversation.

TERMS OF THE DEBATE

Open-source software and standards are already widespread in U.S. national security applications.[16] Army smartphones, Navy warships, and Space Force missile-warning satellites run on Linux-derived operating systems.[17] AI-powered F-16s run on open-source orchestration frameworks like Kubernetes, which is regularly updated, maintained, and tested by industry and the broader public.[18] Open-source software is ubiquitous, permeating over 96 percent of civil and military codebases, and will remain a core piece of defense infrastructure for years to come.[19]

What constitutes an "open" foundation model is less well defined. Developers can distribute foundation models at different levels of "openness"—from publishing white papers and basic technical information to releasing models entirely, including their underlying weights, training data, and the code used to run them.[20] By contrast, developers of closed models, including GPT-4 or Claude, release fewer details or data, only allowing user access through proprietary application programming interfaces.[21] In general, this brief defines "open" models as those with widely available weights, consistent with relevant categories in the 2023 AI executive order.[22] Many of the risks and benefits discussed here flow from these definitions.

Claims of extraordinary risk have motivated several recent proposals surrounding open-source AI. Analysts have expressed concern that malicious users might modify open foundation models to discover cybersecurity vulnerabilities or instruct users in the creation of chemical and biological weapons.[23] Others have argued that public distribution of model weights could aid adversaries in advancing their AI capabilities.[24] Given these apprehensions, some observers have proposed export controls, licensing rules, and liability regimes that would limit the distribution of open foundation models.[25]

A competing school of thought has emphasized the societal benefits of open foundation models.[26] Open distribution of weights, some argue, accelerates innovation and adoption: indeed, the key frameworks and innovations underpinning today's large language models (LLMs), like PyTorch and the transformer architecture itself, were distributed openly.[27] Others contend that the public scrutiny of model weights enables rapid discovery and repair of vulnerabilities, improves public transparency, and reduces the concentration of political and economic power as AI systems increase in importance.[28]

What is most clear, however, is that this risk-benefit assessment remains incomplete. The U.S. Department of Commerce's initial assessment is inconclusive, and AI safety literature has thus far lacked clear frameworks for identifying relative risk and benefit and whether they are unique to open models.[29] Despite concerns over AI models instructing untrained users in biological weapon development, for instance, recent red-teaming exercises concluded that LLM-equipped teams performed similarly to those without.[30] Similar concerns over AI-assisted cyber vulnerability discovery remain unclear, with some arguing that enhanced vulnerability detection may benefit cyber defenders over attackers, or that the balance of advantage would be case-dependent.[31] Malicious use, meanwhile, continues to take place with closed models.[32] In brief, more research remains necessary to unpack where the relative risks and benefits lie.[33] The purportedly catastrophic harms of tomorrow's foundation models have not yet come into clear view.[34]

Second, the pace of technical change has been so uncertain that evaluating future benefits, harms, and policy interventions can be challenging.[35] Whether a licensing regime is effective, for example, depends on how readily foundation model technologies will diffuse.[36] And whether export controls benefit national security hinges on which analogy becomes relevant: Is restricting open models like restricting nuclear weapons exports, or is it akin to Cold War bans (now repealed) on public-key cryptography, a technology which now underpins online banking, e-commerce, and a multi-trillion-dollar digital economy?[37] In the absence of a U.S. market presence, will Chinese open models take their place?[38]

Finally, questions remain on how to implement AI policy. Definitional challenges abound; early AI policy approaches, including the EU AI Act, AI executive order, and California SB 1047, apply thresholds for "systemic risk" to models exceeding a certain amount of computing power or cost used in their development.[39] However, such thresholds for triggering government review, such as the 10^{26} floating-point-operation threshold in the AI executive order, may incompletely capture the capabilities they aim to regulate.[40] How to balance resourcing for AI policy implementation against other cyber and biological threat mitigations, such as for material monitoring or new cyberdefense capabilities, remains another open question.[41]

Figure 2: Performance Comparison of Large Open-Weight and Closed Models as of July 2024

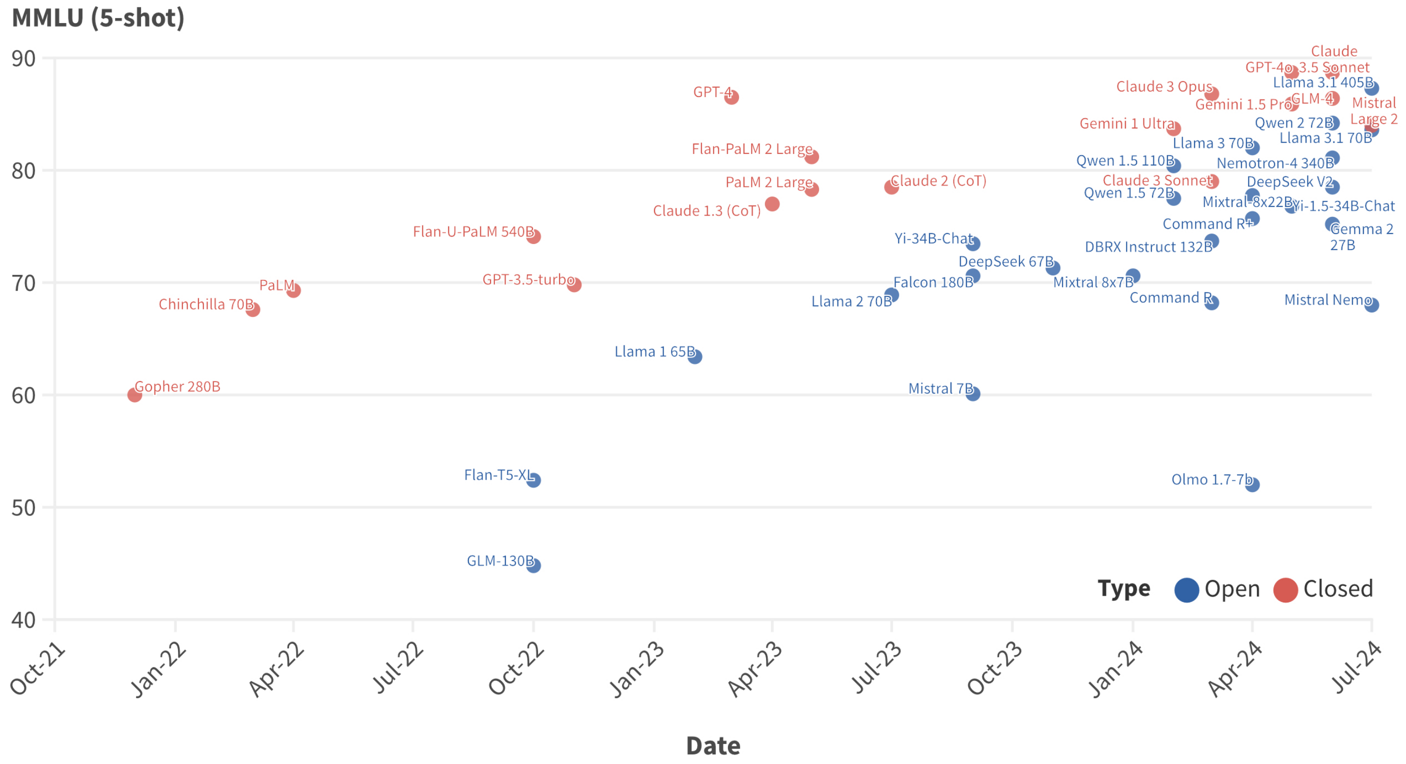

Source: CSIS (see appendix). Chart inspired by Maxime Labonne.[42]

OPEN SOURCE AND THE JOINT FORCE

A defense industrial assessment could thus contribute a valuable perspective to the AI risk debate. With AI industry trends favoring consolidation, the open foundation model ecosystem may become an increasingly important source of competition in the industrial base.[43] Because end users can modify and run open models directly, they have become increasingly relevant for developing local, secure applications and embedded systems–needed by military users demanding low power usage, security, and reliability. And because open models can be publicly inspected, red teamed, and verified, they may present defense-related cybersecurity advantages.[44]

To date, however, the DOD has largely focused on AI adoption.[45] In its flagship data, responsible AI, and adoption strategies, the DOD has focused on harnessing private sector innovations for national security end uses.[46] It has embedded chief data officers in combatant commands; tested AI use cases in major experimentation initiatives, like the Global Information Dominance Experiment, Project Convergence, and others; and developed the Responsible AI Framework, emphasizing the use of traceable, transparent AI systems. [47]

In August 2023, the DOD established Task Force Lima, an element within the CDAO tasked with "responsibly pursu[ing] the adoption" of generalist models.[48] Alongside the CDAO and Responsible AI Working Council, Lima was chartered to "accelerate" AI initiatives, "federate disparate developmental and research efforts," and engage across the interagency on the "responsible development and use of generative AI," with a final strategy due for release in early 2025.[49]

Clarifying potential AI use cases within the DOD is a valuable first step in "mak[ing] life easier for program offices who want to do AI or add AI."[50] A valuable second step would be to identify the broader trajectory of the AI industrial base. The DOD will often rely on industry expertise to develop and identify more generative AI use cases; a broad ecosystem of model and application developers will be critical for this process.[51]

In short, an assessment of defense industrial impacts is conspicuously missing from the broader debate on open foundation models. Arguments over model regulation are couched in national security language but should involve a broader swath of national security practitioners, including defense acquisition professionals.[52] Accordingly, DOD elements, including the CDAO, should independently assess the national capacity to develop AI-powered systems and the impact of open foundation models.

Arguments over model regulation are couched in national security language but should involve a broader swath of national security practitioners.

FROM ADOPTION STRATEGIES TO INDUSTRIAL STRATEGIES

A defense industrial accounting is needed because of preliminary evidence that open foundation models—and their supporting ecosystem—could be useful for the DOD. AI adoption remains a DOD priority, and open release has historically accelerated the rate of technology adoption.[53] Open-source development may have positive competitive implications for defense acquisition. And open-source communities are accelerating developments in on-premises deployment, fine-tuning for specialized applications, model reliability, and other desirable characteristics for defense end users.[54]

1. **Supplier diversity.** In its 2023 National Defense Industrial Strategy (NDIS), the DOD prioritized the diversification of national component supplier bases.[55] To improve competition for new entrants, for instance, it has developed new funding mechanisms, like Strategic Funding Increase contracts, and innovation organizations, like the Defense Innovation Unit.[56] A robust ecosystem of open foundation model developers might improve defense supplier diversity and prevent market consolidation.

   There are strong incentives for consolidation in the foundation model industry.[57] Unlike other software products, large foundation models are particularly capital-intensive to develop, rivaling the supply chain complexity of major defense hardware. Training a foundation model, open or closed, demands extremely large datasets, costly graphics processing units and accelerator chips, and talent, advantaging larger players.[58] Operating them—a process termed *inference*—also imposes high energy costs.[59] Due to these factors, industry leaders have spent tens of billions of dollars to develop and deploy so-called frontier foundation models.[60] This high cost of entry increases the risk of industry consolidation and, consequently, capacity pressure on future DOD acquisition programs.

   Would a robust open AI ecosystem relieve these pressures, creating competition for proprietary vendors and accelerating innovation?[61] Open foundation models have begun to show competitive performance with closed, proprietary alternatives.[62] To date, many competitive open foundation models have been developed by well-resourced actors.[63] But community modification and experimentation have sped the development of new architectures and reductions in training and inference costs.[64] Modified open models have also demonstrated high performance in highly specialized tasks; they might be similarly performant for those anticipated for defense applications, like flight control, business automation, or intelligence fusion.[65]

2. **Competitive sustainment.** More critically, the existence of open foundation models might mitigate dependence on single vendors when sustaining AI-powered defense systems.[66] Foundation models used in defense applications may need specialized retuning, including with government data, as the operating environment changes.[67] High-performance closed models have also shown performance drifts over time; access to model weights and other information would help vendors diagnose emergent problems.[68] In short, by building on open foundation models, the DOD could reduce barriers for vendors to compete on model integration, retuning, and sustainment.

   These imperatives resemble those motivating the Pentagon's Modular Open Systems Approach (MOSA) for major hardware purchases.[69] Lack of access to technical data has historically challenged the DOD's ability to reliably and affordably sustain military hardware.[70] Given these barriers, military services have increasingly demanded access to the data needed to service new helicopters, trucks, and ships. For its nearly $70 billion Future Long-Range Assault Aircraft program, for example, the U.S. Army rejected a $3.6 billion lower bid over incomplete MOSA requirements, opting for a vendor that provided full access to aircraft technical data packages (TDPs) required for maintenance.[71] Just as it is challenging to competitively source aircraft maintenance without access to TDPs, competitively sustaining AI-powered software may be difficult without access to model weights.[72]

It would be challenging, however, to adopt such a modular approach with closed-model-powered software. Closed model developers are unlikely to relinquish access to model weights for other vendors to modify while performing their sustainment contracts.[73] Moreover, to securely implement model-level changes to a defense application, national security customers would potentially need to certify both the application vendor and closed model vendor separately, adding additional friction to an already difficult acquisition process.[74] Evaluating appropriate data rights for closed models could demand a further institutional lift.[75]

A world with competitive open foundation models would allow the Pentagon to sidestep these sustainment challenges. While vendors would tune and adapt open models for proprietary applications, visibility into their underlying weights and architectures would make it easier for others to maintain and integrate them, potentially reducing the cost of sustainment.[76] The presence of open options, moreover, could improve taxpayer bargaining power in negotiations with closed providers on complex data-rights matters. The DOD has already invested heavily in competitively sustaining military hardware. Those same lessons apply when sustaining AI-powered software.

3. **Security and reliability.** Concerns over the confidentiality and reliability of foundation models remain key barriers to DOD generative AI adoption.[77] The potential to locally operate open foundation models on DOD infrastructure therefore becomes a key advantage. Open models present useful options for building AI-powered systems without needing to certify an external foundation model developer for sandboxed deployments.[78] Like with Linux, Kubernetes, or other open software libraries, these models might become a secure baseline for AI that vendors modify with classified or specialized information.[79]

   A robust open research community could also drive advancements in AI model reliability and interpretability, reducing the number of *hallucinations*—nonuniform responses that do not reflect the information a user needs.[80] Access to model weights is often crucial for diagnosing failure and evaluating model reliability in detail.[81] Insofar as hallucinations remain a "showstopper" for defense AI adoption, innovations in the open foundation model ecosystem are worth considering.[82]

   Finally, much like with open-source software, open access to weights might enable a greater possibility of detecting vulnerabilities.[83] Indeed, the AI executive order places special emphasis on red teaming future foundation models; more open approaches to model publication allow a wider peer review of model performance.[84] As Cybersecurity and Infrastructure Security Agency analysts have recently emphasized, there is "significant value in open foundation models to help strengthen cybersecurity, increase competition, and promote innovation."[85] Accordingly, future assessments might review the advantages and disadvantages open models present for hardened, defense-critical AI systems.

4. **Specialized use cases.** Lastly, a robust open foundation model ecosystem might enable AI use cases that receive less attention from closed-source providers. Because open models can be retrained, fine-tuned, and broadly customized, they can serve as a basis for national-security-specific applications.[86] Further, open-source initiatives can drive innovation in national-security-relevant topics, such as for document and data search with semantic retrieval.[87] These search algorithms leverage embedding models—a form of language model—to compare the semantic meaning of stored documents and return relevant results; open models largely dominate performance metrics for embedding generation.[88] Domains overlooked by major commercial players could benefit from the dynamism of open development.[89] Further investigation should assess use cases where open model performance meaningfully exceeds closed alternatives.

AN AI "LAST SUPPER"?

Preliminary evidence suggests that open foundation models might benefit the defense industrial base. What is now needed is a quantitative assessment of the open ecosystem's fiscal impacts. Beyond assessing pathways to adoption, defense policymakers should review the changing competitive landscape for foundation models and potential implications for the defense industrial base. Three recommendations follow:

1. **In its forthcoming review of DOD foundation model adoption, Task Force Lima and partners should compare open and closed models in generative AI use cases.** Different foundation model architectures have roles to play in defense applications. For document summarization tasks, for instance, developers might use open- and closed-source models in a variety of architectures, from retrieval-augmented generation to in-context learning. Although technical details on performance are beyond this paper's scope, a broad comparison of potential open- and closed-model use cases could provide useful context for acquisition professionals.
2. **NDIS implementation should include assessments of the foundation model industrial base, including data, infrastructure, human capital, competitive dynamics, and the impact of open model development.** The 2022 National Defense Strategy acknowledges that the DOD "will be a fast-follower where market forces are driving commercialization of . . . trusted artificial intelligence," while the NDIS asserts that it must "diversify [the] supplier base" and "foster competition within the defense market."[90] An independent assessment of those market forces, including competition in the AI market, would be critical to implementing the NDIS.
3. **The DOD should collect data on where open foundation models are used in its systems, the number of nontraditional performers who leverage open foundation models, and models' potential fiscal impacts.** As the DOD begins to adopt AI-powered systems, it should collect data on where open foundation models are being leveraged. Such data might inform future analyses on the health of the AI industrial base and the ability for new entrants to compete.

There is a considerable risk in disregarding defense industrial impacts in the debate over open foundation models. Major consolidation in other sectors, such as in shipbuilding and aerospace, has come alongside major declines in defense acquisition speed and capacity.[91] If open foundation models are indeed "dual-use"—and therefore critical to national security—the potential for consolidation deserves national security attention.

Civil sector policy decisions have created a shipbuilding base outproduced by China 230 to 1.[92] Other deliberate choices over the nation's industrial base have meant that demand now outpaces supply for missile defense systems, artillery shells, and AI talent.[93] If production is deterrence, these are the stakes of the open-source AI debate.[94]

Others will not wait for the United States if it falls back.[95] China and other states have made vast investments in stimulating their domestic AI industries, with top models "not far behind" Western counterparts.[96] The United States' competitors view open foundation models as a means of capturing global market share and advancing scientific and economic development.[97] While the development of AI is not an arms race, it is a broader economic and social competition—one where U.S. priorities on democracy, transparency, and security should define global standards.[98] The technology and value system to do so are already in place. Attention is all it needs.[99] ■

***Masao Dahlgren** is a fellow with the Missile Defense Project at the Center for Strategic and International Studies in Washington, D.C.*

The project team held two private roundtables, on July 10 and August 1, 2024, to inform the conclusions of this brief. Participants included subject matter experts from the U.S. Department of Defense, U.S. Department of Energy, AI and defense firms, investors, universities, nonprofits, and other stakeholder groups. In addition, the author interviewed multiple government and industry experts to support the research process.

Special thanks to Patrycja Bazylczyk for her assistance in roundtable facilitation, report review, and project performance. Further acknowledgments to Shaan Shaikh, Wes Rumbaugh, Tom Karako, Cynthia Cook, Greg Sanders, Charles Yang, Kevin Li, Michelle Fang, and the many external reviewers involved in refining this report.

This brief was made possible with support from the Omidyar Network and general support to the CSIS Missile Defense Project.

Cover Photo: CSIS

APPENDIX: SELECTED MODEL PERFORMANCE EVALUATIONS

The following table compiles self-reported Massive Multitask Language Understanding (MMLU) scores, one of many benchmarks used to evaluate foundation model performance. Typical scores reported are "5-shot," where five examples are provided before models are prompted with a question. "Open" models listed include those with publicly available model weights.

While MMLU is an incomplete representation of model performance, this benchmark was selected because it is among the oldest, allowing for consistent comparison over time. Writ large, the foundation model industry increasingly suffers from a benchmarking crisis, facing issues of model overfitting–internalizing existing benchmark results–and obsolescence.[100] Developing trusted benchmarks has thus become a major Department of Commerce priority.[101]

Date	Type	Country	Developer	Model	MMLU 5-shot[102]
May 2024	Open	China	01.AI	Yi-1.5-34B-Chat	76.8[103]
Sep. 2023	Open	China	01.AI	Yi-34B-Chat	73.46[104]
Feb. 2024	Open	China	Alibaba	Qwen 1.5 110B	80.4[105]
Feb. 2024	Open	China	Alibaba	Qwen 1.5 72B	77.5[106]
June 2024	Open	China	Alibaba	Qwen 2 72B (base)	84.2[107]
Apr. 2024	Open	United States	Allen Institute	OLMo 1.7-7B	52[108]
Apr. 2023	Closed	United States	Anthropic	Claude 1.3 (CoT)	77[109]
July 2023	Closed	United States	Anthropic	Claude 2 (CoT)	78.5[110]
Mar. 2024	Closed	United States	Anthropic	Claude 3 Opus	86.8[111]
Mar. 2024	Closed	United States	Anthropic	Claude 3 Sonnet	79[112]
June 2024	Closed	United States	Anthropic	Claude 3.5 Sonnet	88.7[113]
Mar. 2024	Open	Canada	Cohere	Command R	68.2[114]
Apr. 2024	Open	Canada	Cohere	Command R+	75.7[115]
Mar. 2024	Open	United States	Databricks	DBRX Instruct 132B	73.7[116]
Nov. 2023	Open	China	DeepSeek	DeepSeek 67B	71.3[117]
June 2024	Open	China	DeepSeek	DeepSeek V2	78.5[118]
Apr. 2023	Open	United States	EleutherAI	Pythia 12B	26.76[119]
Mar. 2022	Closed	United States	Google	Chinchilla 70B	67.6[120]
May 2023	Closed	United States	Google	Flan-PaLM 2 Large	81.2[121]
Oct. 2022	Open	United States	Google	Flan-T5-XL	52.4[122]
Oct. 2022	Closed	United States	Google	Flan-U-PaLM 540B	74.1[123]
May 2024	Closed	United States	Google	Gemini 1.5 Pro	85.9[124]
Feb. 2024	Closed	United States	Google	Gemini 1 Ultra	83.7[125]

Date	Type	Country	Developer	Model	MMLU 5-shot[102]
June 2024	Open	United States	Google	Gemma 2 27B	75.2[126]
Dec. 2021	Closed	United States	Google	Gopher 280B	60[127]
Apr. 2022	Closed	United States	Google	PaLM	69.3[128]
May 2023	Closed	United States	Google	PaLM 2 Large	78.3[129]
July 2023	Open	United States	Meta	Llama 2 70B	68.9[130]
Apr. 2024	Open	United States	Meta	Llama 3 70B (base)	79.5[131]
Apr. 2024	Open	United States	Meta	Llama 3 70B (instruct)	82[132]
July 2024	Open	United States	Meta	Llama 3.1 405B (base)	85.2[133]
July 2024	Open	United States	Meta	Llama 3.1 70B (base)	79.3[134]
July 2024	Open	United States	Meta	Llama 3.1 405B (instruct)	87.3[135]
July 2024	Open	United States	Meta	Llama 3.1 70B (instruct)	83.6[136]
Feb. 2023	Open	United States	Meta	Llama 1 65B	63.4[137]
Sept. 2023	Open	France	Mistral AI	Mistral 7B	60.1[138]
July 2024	Closed	France	Mistral AI	Mistral Large 2	84[139]
Jan. 2024	Open	France	Mistral AI	Mixtral 8x7B	70.6[140]
Apr. 2024	Open	France	Mistral AI	Mixtral-8x22B	77.75[141]
July 2024	Open	France, United States	Mistral AI, NVIDIA	Mistral Nemo	68[142]
June 2024	Open	United States	NVIDIA	Nemotron-4 340B	81.1[143]
May 2020	Closed	United States	OpenAI	GPT-3	43.9[144]
Nov. 2022	Closed	United States	OpenAI	GPT-3.5-turbo	69.8[145]
Mar. 2023	Closed	United States	OpenAI	GPT-4	86.5[146]
May 2024	Closed	United States	OpenAI	GPT-4o	88.7[147]
Sept. 2023	Open	United Arab Emirates	Technology Innovation Institute UAE	Falcon 180B	70.6[148]
Oct. 2022	Open	China	Zhipu AI	GLM-130B	44.8[149]
June 2024	Closed	China	Zhipu AI	GLM-4	86.4[150]

ENDNOTES

1 David Morgan, “Using Large Language Models in the DoD Context,” Defense Acquisition University, February 14, 2024, video, 35:57, https://media.dau.edu/media/t/1_p0m38z1q; Benjamin Jensen and Dan Tadross, “How Large-Language Models Can Revolutionize Military Planning,” War on the Rocks, April 12, 2023, https://warontherocks.com/2023/04/how-large-language-models-can-revolutionize-military-planning/; Chris Riotta, “DOD Wants AI to Help Automate Records Management,” Defense One, May 24, 2023, https://www.defenseone.com/technology/2023/05/dod-aims-leverage-ai-and-automation-records-management/386727/; Peter J. Schwartz et al., “AI-enabled wargaming in the military decision making process,” in ed. Tien Pham, Latasha Solomon, and Katie Rainey, *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications II*, Proceedings of SPIE 11413 (2020), doi:10.1117/12.2560494; and William N. Caballero and Phillip R. Jenkins, “On Large Language Models in National Security Applications,” arXiv, July 3, 2024, https://arxiv.org/pdf/2407.03453v1.

2 Nestor Maslej et al., *The AI Index 2024 Annual Report* (Stanford, CA: Stanford University Institute for Human-Centered AI, April 2024), https://aiindex.stanford.edu/wp-content/uploads/2024/04/HAI_2024_AI-Index-Report.pdf; Shervin Minaee et al., “Large Language Models: A Survey,” arXiv, February 9, 2024, doi:10.48550/arXiv.2402.06196; Tula Masterman et al., “The Landscape of Emerging AI Agent Architectures for Reasoning, Planning, and Tool Calling: A Survey,” arXiv, April 17, 2024, https://arxiv.org/pdf/2404.11584v1; Linxi Fan et al., “MineDojo: Building Open-Ended Embodied Agents with Internet-Scale Knowledge,*” Advances in Neural Information Processing Systems* 35 (December 6, 2022): 18343–62, https://proceedings.neurips.cc/paper_files/paper/2022/hash/74a67268c5cc5910f64938cac4526a90-Abstract-Datasets_and_Benchmarks.html; Tom Brown et al., “Language Models Are Few-Shot Learners,” *Advances in Neural Information Processing Systems* 33 (2020): 1877–1901, https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfcb4967418bfb8ac142f64a-Abstract.html; and Trieu Trinh and Thang Luong, “AlphaGeometry: An Olympiad-Level AI System for Geometry,” Google DeepMind, January 17, 2024, https://deepmind.google/discover/blog/alphageometry-an-olympiad-level-ai-system-for-geometry/.

3 Christine Michienzi, “Pentagon acquisition can no longer ignore the industrial base,” Defense News, May 17, 2024, https://www.defensenews.com/opinion/2024/05/17/pentagon-acquisition-can-no-longer-ignore-the-industrial-base/; and Luke A. Nicastro, *The U.S. Defense Industrial Base: Background and Issues for Congress* (Washington, DC: Congressional Research Service, October 12, 2023), https://crsreports.congress.gov/product/pdf/R/R47751.

4 U.S. Department of Defense, *Data, Analytics, and Artificial Intelligence Adoption Strategy: Accelerating Decision Advantage* (Washington, DC: U.S. Department of Defense, 2023), https://media.defense.gov/2023/Nov/02/2003333300/-1/-1/1/DOD_DATA_ANALYTICS_AI_ADOPTION_STRATEGY.PDF; U.S. Department of Defense, *Executive Summary: DoD Data Strategy* (Washington, DC: U.S. Department of Defense, 2020), https://media.defense.gov/2020/Oct/08/2002514180/-1/-1/0/DOD-DATA-STRATEGY.PDF; Eric Schmidt et al., *Final Report* (Washington, DC: National Security Commission on Artificial Intelligence, 2021), https://cybercemetery.unt.edu/nscai/20211005220330/https://www.nscai.gov/; Department of Defense Responsible AI Working Council, *U.S. Department of Defense Responsible Artificial Intelligence Strategy and Implementation Pathway* (Washington, DC: U.S. Department of Defense, June 2022), https://www.ai.mil/docs/RAI_Strategy_and_Implementation_Pathway_6-21-22.pdf; Tom Temin and David Brauchler, “Can agencies actually follow the White House AI order?,” April 18, 2024, in *the Federal Drive with Tom Temin*, produced by Tom Temin and Federal News Network, podcast, audio, 23:58, https://federalnewsnetwork.com/artificial-intelligence/2024/04/can-agencies-actually-follow-the-white-house-ai-order/; and Chief Information Officer, “DON Guidance on the Use of Generative Artificial Intelligence and Large Language Models,” U.S. Department of the Navy, September 6, 2023, https://www.doncio.navy.mil/ContentView.aspx?id=16442.

5 Matt “Nomad” Strohmeyer, “Accelerating CJADC2 Through Experimentation,” Defense Data and AI Symposium, February 21, 2024, https://www.dvidshub.net/publication/issues/69981; Sam LaGrone, “Pentagon Will Spent $1B on First Round of Replicator Drones,” USNI News, March 11, 2024, https://news.usni.org/2024/03/11/pentagon-will-spend-1b-on-first-round-of-replicator-drones; U.S. Department of Defense, “DOD Announces Establishment of Generative AI Task Force,” press release, August 10, 2023, https://www.defense.gov/News/Releases/Release/Article/3489803/dod-announces-establishment-of-generative-ai-task-force/; Jaspreet Gill, “Say goodbye to JAIC and DDS, as offices cease to exist as independent bodies June 1,” Breaking Defense, May 24, 2022, https://breakingdefense.com/2022/05/say-goodbye-to-jaic-and-dds-as-offices-cease-to-exist-as-independent-bodies-june-1/; Kathleen Hicks, “Establishment of the Chief Digital and Artificial Intelligence Officer,” U.S. Department of Defense, memorandum, December 8, 2021, https://media.defense.gov/2021/Dec/08/2002906075/-1/-1/1/MEMORANDUM-ON-ESTABLISHMENT-OF-THE-CHIEF-DIGITAL-AND-ARTIFICIAL-INTELLIGENCE-OFFICER.PDF; Colin Demarest, “Pentagon achieves ‘minimum viable’ version of CJADC2, Hicks says,” C4ISRNet, February 21, 2024, https://www.c4isrnet.com/battlefield-tech/c2-comms/2024/02/21/pentagon-achieves-minimum-viable-version-of-cjadc2-hicks-says/; and Tim Grayson, “Air and Space Force Perspective on CJADC2,” (speech, 2023 JADC2: All Domain Warfare Symposium, Alexandria, VA, July 19, 2023), https://www.ndia.org/events/2023/7/18/384d---jadc2.

6 Mikayla Easley, “Pentagon, telecom industry’s battle over spectrum symptomatic of a troubled system,” DefenseScoop, July 31, 2023, https://defensescoop.com/2023/07/31/pentagon-spectrum-access-5g-commercial/; and James Wallar, “U.S. Coast-Wise Shipping: No Place for Old Jones in New Industrial Policy,” CSIS, November 30, 2022, https://www.csis.org/analysis/us-coast-wise-shipping-no-place-old-jones-new-industrial-policy.

7 William LaPlante and Seth G. Jones, “Strengthening the U.S.

Industrial Base with Hon. Dr. William A. LaPlante," (discussion, CSIS, Washington, DC, September 26, 2023), https://www.csis.org/analysis/strengthening-us-industrial-base-hon-dr-william-laplante; Heidi Shyu et al., "A Call to Arms: Mobilizing Industry and Unlocking Innovation," (discussion, CSIS, Washington, DC, April 24, 2024), video, 1:14:21, https://www.youtube.com/watch?v=L8IBwse5sgQ; and Sanders et al., "Is Bigger Better? Concentration, Competition, and Defense Contracting Outcomes," (presentation, CSIS, Washington, DC, February 11, 2019), video, 1:30:28, https://www.csis.org/events/bigger-better-concentration-competition-and-defense-contracting-outcomes.

[8] Michael F. Doyle et al., "President Biden Issues Wide-Ranging Executive Order on Artificial Intelligence," K&L Gates Hub (blog), K&L Gates LLP, November 3, 2023, https://www.klgates.com/President-Biden-Issues-Wide-Ranging-Executive-Order-on-Artificial-Intelligence-11-3-2023; Nicol Turner Lee et al., "Will the White House AI Executive Order deliver on its promises?" TechTank (blog), Brookings Institution, November 2, 2023, https://www.brookings.edu/articles/will-the-white-house-ai-executive-order-deliver-on-its-promises/; *White House Overreach on AI: Hearing before the Subcommittee on Cybersecurity, Information Technology, and Government Innovation of the Committee on Oversight and Accountability, U.S. House of Representatives,* 118th Cong., 2nd sess. (March 21, 2024), https://www.govinfo.gov/content/pkg/CHRG-118hhrg55220/html/CHRG-118hhrg55220.htm; Isaiah Poritz, "Biden's AI Order Taps Cold War Crisis Powers for Tech Oversight," Bloomberg Law, November 7, 2023, https://news.bloomberglaw.com/tech-and-telecom-law/bidens-ai-order-taps-cold-war-crisis-powers-for-tech-oversight; and James Andrew Lewis, Emily Benson, and Michael Frank, "The Biden Administration's Executive Order on Artificial Intelligence," CSIS, *Commentary,* October 31, 2023, https://www.csis.org/analysis/biden-administrations-executive-order-artificial-intelligence.

[9] Pierre A. Chao et al., *The Structure and Dynamics of the U.S. Federal Professional Services Industrial Base, 1995-2005* (Washington, DC: CSIS, May 2007), https://csis-website-prod.s3.amazonaws.com/s3fs-public/legacy_files/files/media/csis/pubs/070501_psc.pdf.

[10] National Telecommunications and Information Administration, *Dual-Use Foundation Models with Widely Available Model Weights* (Washington, DC: National Telecommunications and Information Administration, July 2024), https://www.ntia.gov/sites/default/files/publications/ntia-ai-open-model-report.pdf; and Kyle Miller, "Open Foundation Models: Implications of Contemporary Artificial Intelligence," Center for Security and Emerging Technology, March 12, 2024, https://cset.georgetown.edu/article/open-foundation-models-implications-of-contemporary-artificial-intelligence.

[11] Elizabeth Seger et al., "Open-sourcing highly capable foundation models: An evaluation of risks, benefits, and alternative methods for pursuing open-source objectives," LawAI Working Paper Series 2 (October 2023), Center for the Governance of AI, doi:10.2139/ssrn.4596436; National Telecommunications and Information Administration, *Dual-Use Foundation Model*s 12–33; Markus Anderljung et al., "Frontier AI Regulation: Managing Emerging Risks to Public Safety," arXiv, July 6, 2023, doi:10.48550/arXiv.2307.03718; Ryan Heath, "AI's next battle: open or closed," Axios, June 26, 2023, https://www.axios.com/2023/06/26/ais-next-battle-open-closed-chatgpt; and Hamza Tariq Chaudhry, "Request for Comment (NTIA–2023–0009) on Dual-Use Foundation Artificial Intelligence Models with Widely Available Model Weights," Future of Life Institute, March 27, 2024, https://www.regulations.gov/comment/NTIA-2023-0009-0251.

[12] Rishi Bommasani et al., *Considerations for Governing Open Foundation Models* (Palo Alto, CA: Stanford University, 2023), https://hai.stanford.edu/issue-brief-considerations-governing-open-foundation-models; National Telecommunications and Information Administration*, Dual-Use Foundation Models* 12–33; Garry Tan to National Telecommunications and Information Administration, March 25, 2024, https://www.regulations.gov/comment/NTIA-2023-0009-0144; and Camille François et al., "Joint Statement on AI Safety and Openness," Mozilla Foundation, October 31, 2023, https://open.mozilla.org/letter/.

[13] Scott Weiner, "Senator Wiener's Landmark AI Safety And Innovation Bill Passes Assembly Privacy Committee," press release, June 19, 2024, https://sd11.senate.ca.gov/news/senator-wieners-landmark-ai-safety-and-innovation-bill-passes-assembly-privacy-committee; Benj Edwards and Kyle Orland, "From sci-fi to state law: California's plan to prevent AI catastrophe," Ars Technica, July 29, 2024, https://arstechnica.com/information-technology/2024/07/from-sci-fi-to-state-law-californias-plan-to-prevent-ai-catastrophe/; Timothy B. Lee, "A California bill could seriously limit open-weight models," *Understanding AI* (blog), July 26, 2024, https://www.understandingai.org/p/a-california-bill-could-seriously; Zvi Mowshowitz, "Why is Everyone Suddenly Furious About AI Regulation?" Asterisk, May 2024, https://asteriskmag.com/issues/06/why-is-everyone-suddenly-furious-about-ai-regulation; Nathan Benaich and Alex Chalmers, "California's AI bill was an avoidable disaster," *Air Street Press* (blog), June 6, 2024, https://press.airstreet.com/p/californias-ai-bill-was-an-avoidable; and Jake Denton, "California's Gift to Big Tech," Compact, June 10, 2024, https://www.compactmag.com/article/californias-gift-to-big-tech/.

[14] Nilay Patel, "Biden's top tech adviser says AI is a 'today problem,'" *The Verge*, July 15, 2024, https://www.theverge.com/24197237/arati-prabhakar-ostp-director-tech-policy-science-ai-regulation-decoder-podcast.

[15] National Telecommunications and Information Administration, *Dual-Use Foundation Models,* 34–36.

[16] Red Hat, "Lockheed Martin Taps Red Hat to Accelerate F-22 Raptor Upgrades," press release, May 6, 2019, https://www.redhat.com/en/about/press-releases/lockheed-martin-taps-red-hat-accelerate-f-22-raptor-upgrades; and Ritwik Gupta, "The Department of Defense is Prioritizing Open Source Software. Here's How Open Source Projects Can Benefit," *Gradients* (blog), February 7, 2022, https://ritwikgupta.me/dod-opportunities-for-open-source-software/.

[17] Sean Gallagher, "The Navy's newest warship is powered by Linux," Ars Technica, October 18, 2013, https://arstechnica.com/information-technology/2013/10/the-navys-newest-warship-is-powered-by-linux/; Space Systems Command, "Space Systems Command's Space-Based Infrared System Baseline Release Con-

trol Authority Transferred to Space Operations Command," press release, April 18, 2024, https://www.ssc.spaceforce.mil/Portals/3/Documents/PRESS%20RELEASES/Space%20Systems%20Commands%20Space-Based%20Infrared%20System%20Baseline%20Release%20Control%20Authority%20Transferred%20to%20Space%20Operations%20Command.pdf?ver=I1tZvCqIx8d6nqdxU9vumQ%3D%3D; Sam Skove, "In a first, Army uses Slack-style battlefield software to coordinate field exercises," Defense One, January 24, 2024, https://www.defenseone.com/technology/2024/01/first-army-uses-slack-style-battlefield-software-field-exercises/393585/; Denise Chow, "Troops Call for Military Airstrike? There's an App for That," NBC News, October 21, 2013, https://www.nbcnews.com/id/wbna53336049; and Mackenzie Eaglen, "Tech-Challenged Pentagon Searches for a Silicon Ally," *Wall Street Journal*, January 31, 2016, https://www.wsj.com/articles/tech-challenged-pentagon-searches-for-a-silicon-ally-1454282727.

[18] Gareth Corfield, "K8s on a plane! US Air Force slaps Googly container tech on yet another war machine to 'run advanced ML algorithms,'" The Register, October 8, 2020, https://www.theregister.com/2020/10/08/usaf_u2_spyplane_kubernetes_britten_norman/; and Tom Krazit, "How the U.S. Air Force Deployed Kubernetes and Istio on an F-16 in 45 days," The New Stack, December 24, 2019, https://thenewstack.io/how-the-u-s-air-force-deployed-kubernetes-and-istio-on-an-f-16-in-45-days/.

[19] Synposys Inc., *2024 Open Source Security and Risk Analysis Report* (Sunnyvale, CA: Synopsys Inc., 2024), https://www.synopsys.com/software-integrity/resources/analyst-reports/open-source-security-risk-analysis.html; Knut Blind et al., *The impact of open source software and hardware on technological independence, competitiveness and innovation in the EU economy* (Brussels: European Commission, 2021), https://digital-strategy.ec.europa.eu/en/library/study-about-impact-open-source-software-and-hardware-technological-independence-competitiveness-and; Manuel Hoffmann, Frank Nagle, and Yanuo Zhou, "The Value of Open Source Software," Working Paper No. 24-038, Harvard Business School Strategy Unit, 2024, doi:10.2139/ssrn.4693148; and Ben FitzGerald, Jacqueline Parziale, and Peter L. Levin, *Open Source Software and the Department of Defense* (Washington, DC: Center for a New American Security, 2016), https://www.cnas.org/publications/reports/open-source-software-and-the-department-of-defense.

[20] Irene Solaiman, "The Gradient of Generative AI Release: Methods and Considerations," in *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency,* FAccT '23 (New York, NY: Association for Computing Machinery, 2023): 111–22, https://doi.org/10.1145/3593013.3593981; Nathan Lambert, "The koan of an open-source LLM," *Interconnects* (blog), March 6, 2024, https://www.interconnects.ai/p/an-open-source-llm?publication_id=48206; Megan Morrone, "'Open' software needs an AI rethink," Axios, February 15, 2024, https://www.axios.com/2024/02/15/open-source-ai-definition-openai-meta; and Edd Gent, "The tech industry can't agree on what open-source AI means. That's a problem." *MIT Technology Review*, March 25, 2024, https://www.technologyreview.com/2024/03/25/1090111/tech-industry-open-source-ai-definition-problem/.

[21] Tiernan Ray, "With GPT-4, OpenAI opts for secrecy versus disclosure," ZDNet, March 16, 2023, https://www.zdnet.com/article/with-gpt-4-openai-opts-for-secrecy-versus-disclosure/; OpenAI, *GPT-4 System Card,* (San Francisco, CA: OpenAI, March 23, 2023), https://cdn.openai.com/papers/gpt-4-system-card.pdf; and Anthropic PBC, *The Claude 3 Model Family: Opus, Sonnet, Haiku* (San Francisco, CA: Anthropic PBC, March 2024), https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf.

[22] Exec. Order No. 14110, 88 F.R. 75191 (2023), https://www.federalregister.gov/d/2023-24283.

[23] Miller, "Open Foundation Models"; Seger et al., "Open-sourcing highly capable foundation models"; Alan Chan et al., "Hazards from Increasingly Accessible Fine-Tuning of Downloadable Foundation Models," arXiv, December 22, 2023, https://doi.org/10.48550/arXiv.2312.14751; Anjali Gopal et al., "Will releasing the weights of future large language models grant widespread access to pandemic agents?" arXiv, October 25, 2023, doi:10.48550/arXiv.2310.18233; Yisroel Mirsky et al., "The Threat of Offensive AI to Organizations," arXiv, June 30, 2021, doi:10.48550/arXiv.2106.15764; Richard Moulange et al., "Towards Responsible Governance of Biological Design Tools," arXiv, November 27, 2023, doi:10.48550/arXiv.2311.15936; Jonas B. Sandbrink, "Artificial intelligence and biological misuse: Differentiating risks of language models and biological design tools," arXiv, June 24, 2023, doi:10.48550/arXiv.2306.13952; Dmitrii Volkov, "Badllama 3: removing safety finetuning from Llama 3 in minutes," arXiv, July 1, 2024, doi:10.48550/arXiv.2407.01376; Pablo Chavez, "An AI Challenge: Balancing Open and Closed Systems," *Bandwidth* (blog), Center for European Policy Analysis, May 30, 2023, https://cepa.org/article/an-ai-challenge-balancing-open-and-closed-systems/; and Zoë Brammer, *How Does Access Impact Risk? Assessing AI Foundation Model Risk Along a Gradient of Access* (Oakland, CA: Institute for Security and Technology, 2023), https://securityandtechnology.org/wp-content/uploads/2023/12/How-Does-Access-Impact-Risk-Assessing-AI-Foundation-Model-Risk-Along-A-Gradient-of-Access-Dec-2023.pdf.

[24] Meaghan Tobin and Cade Metz, "China Is Closing the A.I. Gap With the United States," *New York Times*, July 25, 2024, https://www.nytimes.com/2024/07/25/technology/china-open-source-ai.html; John Luttig, "The future of foundation models is closed-source," *Luttig's Learnings* (blog), May 21, 2024, https://blog.johnluttig.com/p/the-future-of-foundation-models-is; Vinod Khosla, Todd Young, and Jacob Helberg, "Perspectives on the US-China Techno-Economic War and the National Security Implications of AI," (panel discussion, Hill and Valley Forum, Washington, DC, May 1, 2024), video, 32:20–51:35, https://www.youtube.com/watch?v=RqxE3ub7wWA; and Dwarkesh Patel and Leopold Aschenbrenner, "China/US Super Intelligence Race, 2027 AGI, & The Return of History," June 4, 2024, in the *Dwarkesh Podcast*, produced by Dwarkesh Patel, podcast, audio, 4:31:17, https://www.dwarkeshpatel.com/p/leopold-aschenbrenner.

[25] Alexandra Alper, "US eyes curbs on China's access to AI software behind apps like ChatGPT," Reuters, May 8, 2024, https://www.reuters.com/technology/us-eyes-curbs-chinas-access-ai-software-behind-apps-like-chatgpt-2024-05-08/; Nathan Lambert, "SB 1047, AI regulation, and unlikely allies for open models,"

Interconnects (blog), July 17, 2024, https://www.interconnects.ai/p/sb-1047-and-open-weights; Owen J. Daniels, "California AI bill becomes a lightning rod–for safety advocates and developers alike," *Bulletin of the Atomic Scientists*, June 17, 2024, https://thebulletin.org/2024/06/california-ai-bill-becomes-a-lightning-rod-for-safety-advocates-and-developers-alike/; Lee, "A California bill could seriously limit"; Benaich and Chalmers, "California's AI bill"; Denton, "California's Gift"; James Pethokoukis, "AI, Regulation, and Economic Thinking," *AEIdeas* (blog), American Enterprise Institute, April 29, 2024, https://www.aei.org/economics/ai-regulation-and-economic-thinking/; Lara Korte and Dustin Gardiner, "'Little Tech' brings a big flex to Sacramento," *Politico*, June 21, 2024, https://www.politico.com/newsletters/california-playbook/2024/06/21/little-tech-brings-a-big-flex-to-sacramento-00164369; and Dean W. Ball, "California's Effort to Strangle AI," *Hyperdimensional* (blog), February 9, 2024, https://www.hyperdimensional.co/p/californias-effort-to-strangle-ai.

[26] Stella Biderman, "Principles," (presentation, Workshop on Responsible and Open Foundation Models, online, September 21, 2023), https://sites.google.com/view/open-foundation-models; Mark Surman and Camille François, "A Third Way on AI," *Distilled* (blog), Mozilla Foundation, November 7, 2023, https://blog.mozilla.org/en/mozilla/a-third-way-on-ai/; Ben Thompson, "Attenuating Innovation (AI)," *Stratechery* (blog), November 1, 2023, https://stratechery.com/2023/attenuating-innovation-ai/; and Andrew Ng, "The AI Application Ecosystem," (presentation, Promoting Competition on Artificial Intelligence Workshop, Stanford, CA, May 30, 2024), https://www.gsb.stanford.edu/events/promoting-competition-ai; and National Telecommunications and Information Administration, *Dual-Use Foundation Models*,12–31.

[27] Bommasani et al., Considerations for Governing; Alex Engler, "The EU's attempt to regulate open-source AI is counterproductive," *TechTank* (blog), Brookings Institution, August 24, 2022, https://www.brookings.edu/articles/the-eus-attempt-to-regulate-open-source-ai-is-counterproductive/; "pytorch/LICENSE," Github, accessed March 4, 2022, https://github.com/pytorch/pytorch/blob/main/LICENSE; and Ali Farhadi et al., "Android Moment of Open Source AI," (panel discussion, Cerebral Valley New York, New York, NY, July 1, 2024), video, 19:55, https://www.youtube.com/watch?v=UVjeqjijb6s&t=352s; 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/; Fei-Fei Li, "'The Godmother of AI' says California's well-intended AI bill will harm the U.S. ecosystem," *Fortune,* August 6, 2024, https://fortune.com/2024/08/06/godmother-of-ai-says-californias-ai-bill-will-harm-us-ecosystem-tech-politics/; Ng, "The AI Application Ecosystem"; Biderman, "Principles"; and Thompson, "Attenuating Innovation."

[28] "An Open Letter to the European Parliament: Protecting Open-Source AI for a Safe, Secure, and Sovereign Digital Future," LAION e.V., n.d., https://laion.ai/documents/open-letter-to-eu-parliament.pdf; Dean W. Ball, "Open-Source AI Has Overwhelming Support," *Hyperdimensional* (blog), April 7, 2024, https://hyperdimensional.substack.com/p/open-source-ai-has-overwhelming-support; Martin Casado and Katherine Boyle, "AI Talks Leave 'Little Tech' Out," *Wall Street Journal,* May 15, 2024, https://www.wsj.com/articles/ai-talks-leave-little-tech-out-homeland-security-adversaries-open-source-board-46e3232d; Fox Business, "Big Tech wants to make a buck and they don't care how our kids pay for it: Kara Frederick," video, 4:11, June 17, 2024, https://www.foxbusiness.com/video/6355184067112; Josh Wolfe, Theresa Carlson, and Zoe Weinberg, "The Geopolitical Stakes for LLMs," (panel discussion, Cerebral Valley New York, NY, July 2, 2024), video, 21:38, https://www.youtube.com/watch?v=strioMkVv_Q; Nick Clegg, "Openness on AI is the way forward for tech," *Financial Times*, July 11, 2023, https://www.ft.com/content/ac3b585a-ce50-43d1-b71d-14dfe6dce999; and Derek Slater and Betsy Masiello, "Will Open Source AI Shift Power from 'Big Tech'? It Depends." Tech Policy Press, June 16, 2023, https://www.techpolicy.press/will-open-source-ai-shift-power-from-big-tech-it-depends/.

[29] Sayash Kapoor et al., "On the Societal Impact of Open Foundation Models," arXiv, February 27, 2024, http://arxiv.org/abs/2403.07918; "Open-source AI and the Defense Industrial Base," workshop, CSIS, Washington, DC, July 10, 2024; and National Telecommunications and Information Administration, *Dual-Use Foundation Models,* 33–50.

[30] Christopher A. Mouton, Caleb Lucas, and Ella Guest*, The Operational Risks of AI in Large-Scale Biological Attacks: Results of a Red-Team Study* (Santa Monica, CA: RAND Corporation, January 25, 2024), https://www.rand.org/pubs/research_reports/RRA2977-2.html; and Dean W. Ball, "AI Biorisk: A Dose of Reality," *Hyperdimensional* (blog), January 31, 2024, https://hyperdimensional.substack.com/p/ai-biorisk-a-dose-of-reality.

[31] Jenny Jun, "How will AI Change Cyber Operations?" War on the Rocks, April 30 2024, https://warontherocks.com/2024/04/how-will-ai-change-cyber-operations/; Ben Buchanan*, A National Security Research Agenda for Cybersecurity and Artificial Intelligence* (Washington, DC: Center for Security and Emerging Technology, May 2020), https://cset.georgetown.edu/publication/a-national-security-research-agenda-for-cybersecurity-and-artificial-intelligence/; "Disrupting malicious uses of AI by state-affiliated threat actors," OpenAI, February 14, 2024, https://openai.com/index/disrupting-malicious-uses-of-ai-by-state-affiliated-threat-actors/; and Toby Shevlane and Allan Dafoe, "The Offense-Defense Balance of Scientific Knowledge: Does Publishing AI Research Reduce Misuse?" in *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society,* AIES '20 (New York, NY: Association for Computing Machinery, 2020), 173–79, doi:10.1145/3375627.3375815.

[32] James Vincent, "'As an AI language model': the phrase that shows how AI is polluting the web," *The Verge*, April 25, 2023, https://www.theverge.com/2023/4/25/23697218/ai-generated-spam-fake-user-reviews-as-an-ai-language-model; Yi Liu et al., "Jailbreaking ChatGPT via Prompt Engineering: An Empirical Study," arXiv, May 23, 2023, https://arxiv.org/abs/2305.13860; and Will Oremus, "The clever trick that turns ChatGPT into its evil twin," *Washington Post*, February 14, 2023, https://www.washingtonpost.com/technology/2023/02/14/chatgpt-dan-jailbreak/.

[33] Tianyu Cui et al., "Risk Taxonomy, Mitigation, and Assessment Benchmarks of Large Language Model Systems," arXiv, January 11, 2024, doi:10.48550/arXiv.2401.05778; Kapoor et al., "On the Societal Impact," 8–9; and National Telecommunications and

Information Administration, *Dual-Use Foundation Models,* 43–45.

[34] James Andrew Lewis, "A Real Risk for Artificial Intelligence," CSIS, *Commentary,* June 11, 2024, https://www.csis.org/analysis/real-risk-artificial-intelligence; Brammer, *How Does Access Impact Risk?*; and Dean W. Ball, "Llama 3.1 and the 'Path Forward,'" *Hyperdimensional* (blog), June 20, 2024, https://www.hyperdimensional.co/p/llama-31-and-the-path-forward?publication_id=2244049&utm_campaign=email-post-title&r=1qn0ni&utm_medium=email;

[35] Lewis, "A Real Risk"; and National Telecommunications and Information Administration, *Dual-Use Foundation Models*, 33–50.

[36] Mike Watson, "IAEA for AI? That Model Has Already Failed," Wall Street Journal, June 1, 2023, https://www.wsj.com/articles/iaea-for-ai-that-model-has-already-failed-chaptgpt-technology-nuclear-proliferation-4339543b; Sayash Kapoor and Arvind Narayanan, "Licensing is neither feasible nor effective for addressing AI risks," AI Snake Oil (blog), June 10, 2023, https://www.aisnakeoil.com/p/licensing-is-neither-feasible-nor; and Michael C. Horowitz, "Artificial Intelligence, International Competition, and the Balance of Power," *Texas National Security Review* 1, no. 3 (May 2018), doi:10.15781/T2639KP49.

[37] Henry Corrigan-Gibbs, "Keeping Secrets," *Stanford Magazine*, November/December 2014, https://stanfordmag.org/contents/keeping-secrets; Jordan Schneider and Lily Ottinger, "Scale's Alex Wang on the US-China AI Race," *ChinaTalk* (blog), June 25, 2024, https://www.chinatalk.media/p/scales-alex-wang-on-the-us-china; and Tina Highfill and Christopher Surfield, *New and Revised Statistics of the U.S. Digital Economy, 2005–2021* (Washington, DC: Bureau of Economic Analysis, November 2022), https://www.bea.gov/system/files/2022-11/new-and-revised-statistics-of-the-us-digital-economy-2005-2021.pdf.

[38] Rebecca Arcesati and Caroline Meinhardt, *China bets on open-source technologies to boost domestic innovation* (Berlin: Mercator Institute for China Studies, May 19, 2021), https://merics.org/en/report/china-bets-open-source-technologies-boost-domestic-innovation; Keegan McBride, "Open Source AI: The Overlooked National Security Imperative," Just Security, June 6, 2024, https://www.justsecurity.org/96422/open-source-ai-the-overlooked-national-security-imperative/#:~:text=The%20AI-driven%20world%20of,on%20our%20shared%20digital%20future; Nathan Benaich and Alex Chalmers, "The State of Chinese AI," *Air Street Press* (blog), July 18, 2024, https://press.airstreet.com/p/the-state-of-chinese-ai; Luttig, "The future of foundation models"; Irene Zhang, "Putting China's Top LLMs to the Test," ChinaTalk, December 5, 2023, https://www.chinatalk.media/p/putting-chinas-top-llms-to-the-test; and Yiwen Lu, "OpenAI Pulls the Plug on China," ChinaTalk, July 11, 2024, https://www.chinatalk.media/p/openai-pulls-the-plug-on-china.

[39] "The AI Act Explorer," EU Artificial Intelligence Act, European Commission, 2024, https://artificialintelligenceact.eu/ai-act-explorer/; and David A. Simon et al., "Latest Text of EU AI Act Proposes Expanding Obligations for High-Risk and General AI Systems and Banning a Third Category," *AI Insights* (blog), Skadden, Arps, Slate, Meagher & Flom LLP and Affiliates, February 5, 2024, https://www.skadden.com/insights/publications/2024/02/latest-text-of-eu-ai-act-proposes-expanding-obligation.

[40] Lambert, "SB 1047, AI Regulation, and Unlikely Allies."

[41] "Open-source AI Roundtable: Assessing Defense-Industrial Impacts," workshop, CSIS, online, August 1, 2024; U.S. Department of Defense, 2023 Biodefense Posture Review (Washington, DC: U.S. Department of Defense, 2023), https://media.defense.gov/2023/Aug/17/2003282337/-1/-1/1/2023_BIODEFENSE_POSTURE_REVIEW.PDF; Arielle D'Souza*, Why the Department of Defense Should Invest in Biosurveillance* (Washington, DC: Institute for Progress, 2022), https://ifp.org/investing-in-biosurveillance/; Bipartisan Commission on Biodefense, *The National Blueprint for Biodefense* (Washington, DC: Bipartisan Commission on Biodefense, April 2024), https://biodefensecommission.org/wp-content/uploads/2024/05/National-Blueprint-for-Biodefense-2024_final_digital.pdf; John Sakellariadis, "CISA avoids major cuts in fiscal 2024 funding bill," *Politico*, March 21, 2024, https://subscriber.politicopro.com/article/2024/03/cisa-avoids-major-cuts-in-fiscal-2024-funding-bill-00148260; and Tim Starks, "Easterly appeals to Congress on CISA funding, citing Chinese threats to critical infrastructure," CyberScoop, April 30, 2024, https://cyberscoop.com/jen-easterly-cisa-funding-congress-critical-infrastructure-china/.

[42] Maxime Labonne (@@maximelabonne), "I made the closed-source vs. open-weight models figure for this moment," X, July 24, 2024, 3:12 a.m., https://twitter.com/maximelabonne/status/1816008591934922915.

[43] Manuel Hoffmann, Frank Nagle, and Yanuo Zhou, "The Value of Open Source Software," Working Paper No. 24-038, Harvard Business School Strategy Unit, 2024, doi:10.2139/ssrn.4693148; Mark Surman et al.*, Accelerating Progress Toward Trustworthy AI* (San Francisco, CA: Mozilla Foundation, February 22, 2024), https://foundation.mozilla.org/en/research/library/accelerating-progress-toward-trustworthy-ai/whitepaper/; Dylan Patel and Afzal Ahmad, "Google 'We Have No Moat, And Neither Does OpenAI': Leaked Internal Google Document Claims Open Source AI Will Outcompete Google and OpenAI," Semianalysis, May 4, 2023, https://www.semianalysis.com/p/google-we-have-no-moat-and-neither; and Dwarkesh Patel and Mark Zuckerberg, "Llama 3, Open Sourcing $10b Models, & Caesar Augustus," April 18, 2024, in *the Dwarkesh Podcast*, produced by Dwarkesh Patel, podcast, audio, 1:17:54, https://www.dwarkeshpatel.com/p/mark-zuckerberg.

[44] Andy Zou et al. "Universal and Transferable Adversarial Attacks on Aligned Language Models," arXiv, July 27, 2023, doi:10.48550/arXiv.2307.15043; and National Telecommunications and Information Administration, *Dual-Use Foundation Models,* 17–19.

[45] William Streilein, "Accelerating Adoption of AI for Decision Advantage: The Digital Ecosystem," (presentation, Advantage DoD 24: Defense Data & AI Symposium, February 2024).

[46] U.S. Department of Defense*, Department of Defense Data, Analytics, and Artificial Intelligence Adoption Strategy;* U.S. Department of Defense, "Executive Summary: DoD Data Strategy," in *DoD Data Strategy* (Washington, DC: U.S. Department of Defense, 2020), https://media.defense.gov/2020/Oct/08/2002514180/-1/-1/0/

DOD-DATA-STRATEGY.PDF; Schmidt et al., *Final Report;* and DoD Responsible AI Working Council, *U.S. Department of Defense Responsible Artificial Intelligence Strategy*.

[47] Kathleen Hicks, "Implementing Responsible Artificial Intelligence in the Department of Defense," memorandum, U.S. Department of Defense, May 26, 2021, https://media.defense.gov/2021/May/27/2002730593/-1/-1/0/IMPLEMENTING-RESPONSIBLE-ARTIFICIAL-INTELLIGENCE-IN-THE-DEPARTMENT-OF-DEFENSE.PDF; DoD Responsible AI Working Council, *U.S. Department of Defense Responsible Artificial Intelligence Strategy;* Bill Streilein, "AI/ML Scaffolding and AI Assurance," presentation, CDAO Industry Day, May 8, 2023, https://www.ai.mil/docs/IndDay/6_AI_Scaffolding.pdf; Gregory Allen, Xavier Lugo, and Matthew Strohmeyer, "Scaling AI-enabled Capabilities at the DOD: Government and Industry Perspectives," (presentation, CSIS, March 28, 2024), https://www.csis.org/analysis/scaling-ai-enabled-capabilities-dod-government-and-industry-perspectives; Glen VanHerck and Tom Karako, "Rethinking Homeland Defense: Global Integration, Domain Awareness, Information Dominance and Decision Superiority," (discussion, CSIS, August 17, 2021), https://www.csis.org/analysis/rethinking-homeland-defense-global-integration-domain-awareness-information-dominance-and; and Maggie Gray, "CDAO and Accelerating DoD AI Adoption," *Gray Matters* (blog), March 25, 2024, https://maggiegray.substack.com/p/cdao-and-accelerating-dod-ai-adoption.

[48] Mikayla Easley, "Defense Department stands up generative AI task force," DefenseScoop, August 10, 2023, https://defensescoop.com/2023/08/10/defense-department-generative-ai-task-force/.

[49] Ibid.; and Kathleen Hicks, "Establishment of Chief Digital and Artificial Intelligence Officer Generative Artificial Intelligence and Large Language Models Task Force, Task Force LimaU.S. Department of Defense," memorandum, August 10, 2023, https://media.defense.gov/2023/Aug/10/2003279040/-1/-1/1/ESTABLISHMENT_OF_CDAO_GENERATIVE_AI_AND_LARGE_LANGUAGE_MODELS_TASK_FORCE_TASK_FORCE_LIMA_OSD006491-23_RES_FINAL.PDF.

[50] Allen, Lugo, and Strohmeyer, "Scaling AI-enabled Capabilities at the DOD"; and Edward Graham, "DOD's generative AI task force looks for 'blind spots,'" NextGov, February 23, 2024, https://www.nextgov.com/artificial-intelligence/2024/02/dods-generative-ai-task-force-looks-blind-spots/394404/.

[51] Brandi Vincent, "Inside Task Force Lima's exploration of 180-plus generative AI use cases for DOD," DefenseScoop, November 6, 2023, https://defensescoop.com/2023/11/06/inside-task-force-limas-exploration-of-180-plus-generative-ai-use-cases-for-dod/.

[52] Andrew Imbrie, Owen J. Daniels, and Helen Toner, *Decoding Intentions: Artificial Intelligence and Costly Signals* (Washington, DC: Center for Security and Emerging Technology, October 2023), https://cset.georgetown.edu/publication/decoding-intentions/; and Macarena Bazan, "AI in the Spotlight: Breaking Down President Biden's Executive Order," *University of Miami Business Law Review Insights* (blog), University of Miami School of Law, November 28, 2023, https://business-law-review.law.miami.edu/ai-in-the-spotlight-breaking-down-president-bidens-executive-order/.

[53] National Telecommunications and Information Administration, *Dual-Use Foundation Models*, 33–50; and Tony DeMartino, "In AI we trust: how DoD's Task Force Lima can safeguard generative AI for warfighters," Breaking Defense, October 31, 2023, https://breakingdefense.com/2023/10/in-ai-we-trust-how-dods-task-force-lima-can-safeguard-generative-ai-for-warfighters/.

[54] Tim Dettmers et al., "QLoRA: Efficient Finetuning of Quantized LLMs," arXiv, May 23, 2023, doi:10.48550/arXiv.2305.14314; Nathan Benaich, *State of AI Report 2023* (London, UK: Air Street Capital, 2023), https://www.stateof.ai/; Shuming Ma et al., "The Era of 1-Bit LLMs: All Large Language Models Are in 1.58 Bits," arXiv, February 27, 2024, doi:10.48550/arXiv.2402.17764; Maxime Labonne, "I was curious to see how the latest batch of open-weight models changes the dynamics with closed-source LLMs," X, April 13, 2024, https://twitter.com/maximelabonne/status/1779123021480865807; and Ryan Heath, "AI hits trust hurdles with U.S. military," Axios, May 1, 2024, https://www.axios.com/2024/05/01/pentagon-military-ai-trust-issues.

[55] U.S. Department of Defense, *National Defense Industrial Strategy* (Washington, DC: U.S. Department of Defense, 2023), https://www.businessdefense.gov/docs/ndis/2023-NDIS.pdf.

[56] Courtney Albon, "Defense Innovation Unit prepares to execute $800 million funding boost," C4ISRNet, April 23, 2024, https://www.c4isrnet.com/unmanned/2024/04/23/defense-innovation-unit-prepares-to-execute-800-million-funding-boost/; and Tim Tresslar, "AFVentures launches 2024 STRATFI/TACFI Notice of Opportunity," Air Force Research Laboratory, August 11, 2023, https://www.afrl.af.mil/News/Article-Display/Article/3491081/afventures-launches-2024-stratfitacfi-notice-of-opportunity/.

[57] Jai Vipra and Anton Korinek, *Market concentration implications of foundation models: The Invisible Hand of ChatGPT* (Washington, DC: Brookings Institution, September 7, 2023), https://www.brookings.edu/articles/market-concentration-implications-of-foundation-models-the-invisible-hand-of-chatgpt/; Anton Korinek and Jai Vipra, "Concentrating Intelligence: Scaling Laws and Market Structure in Generative AI," University of Virginia and Centre for the Governance of AI, February 28, 2024; Leah Nylen and Lizette Chapman, "FTC's Khan Backs Open AI Models in Bid to Avoid Monopolies," Bloomberg, July 25, 2024, https://www.bloomberg.com/news/articles/2024-07-25/ftc-s-khan-backs-open-ai-models-in-bid-to-avoid-monopolies; and Federal Trade Commission, "Generative AI Raises Competition Concerns," *Technology Blog*, Federal Trade Commission, June 29, 2023, https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2023/06/generative-ai-raises-competition-concerns.

[58] Naomi Nix, Cat Zakrzewski, and Gerrit De Vynck, "Silicon Valley is pricing academics out of AI research," *Washington Post*, March 10, 2024, https://www.washingtonpost.com/technology/2024/03/10/big-tech-companies-ai-research/; Jaime Sevilla and Edu Roldán, *Training Compute of Frontier AI Models Grows by 4-5x per Year* (San Jose, CA: Epoch AI, May 28, 2024), https://epochai.org/blog/training-compute-of-frontier-ai-models-grows-by-4-5x-per-year; and Guido Appenzeller, Matt Bornstein, and Martin Casado, "Navigating the High Cost of AI Compute," Andreessen Horowitz, April 27, 2023, https://a16z.com/navigating-the-high-cost-of-ai-compute/.

59 Jonathan Vanian and Kif Leswing, "ChatGPT and generative AI are booming, but the costs can be extraordinary," CNBC, March 13, 2023, https://www.cnbc.com/2023/03/13/chatgpt-and-generative-ai-are-booming-but-at-a-very-expensive-price.html; Hugo Huang, "What CEOs Need to Know About the Costs of Adopting GenAI," *Harvard Business Review*, November 15, 2023, https://hbr.org/2023/11/what-ceos-need-to-know-about-the-costs-of-adopting-genai; Sevilla and Roldán*, Training Compute of Frontier AI Models;* Appenzeller, Bornstein, and Casado, "Navigating the High Cost of AI Compute"; and Nikhil Sardana and Jonathan Frankle, "Beyond Chinchilla-Optimal: Accounting for Inference in Language Model Scaling Laws," arXiv, December 31, 2023, https://arxiv.org/pdf/2401.00448v1.

60 Maslej et al., *AI Index 2024 Annual Report;* Ben Cottier et al., *How Much Does It Cost to Train Frontier AI Models?* (San Jose, CA: Epoch AI, June 3, 2024), https://epochai.org/blog/how-much-does-it-cost-to-train-frontier-ai-models; and Matthew Fox, "How Nvidia is dominating an AI-obsessed earnings season without even reporting yet," *Business Insider*, May 5, 2024, https://www.businessinsider.com/nvidia-impact-q1-earnings-season-ai-nvda-stock-price-capex-2024-5.

61 Federal Trade Commission, "On Open-Weights Foundation Models," *Technology Blog*, Federal Trade Commission, July 10, 2024, https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2024/07/open-weights-foundation-models; Patel and Ahmad, "Google 'We Have No Moat, And Neither Does OpenAI'"; and Lina M. Khan et al., "Joint Statement on Competition in Generative AI Foundation Models and AI Products," Federal Trade Commission, July 23, 2024, https://www.ftc.gov/legal-library/browse/joint-statement-competition-generative-ai-foundation-models-ai-products.

62 Maslej et al., *AI Index 2024 Annual Report*; and Nathan Lambert, "Model commoditization and product moats," *Interconnects* (blog), March 13, 2024, https://www.interconnects.ai/p/gpt4-commoditization-and-moats?publication_id=48206.

63 Gabriel Nicholas and Paul Friedl, *Regulating Large Language Models: A Roundtable Report* (Washington, DC: Center for Democracy and Technology, 2023), http://arxiv.org/abs/2403.15397.

64 Maslej et al., *Artificial Intelligence Index Report;* Nathan Lambert, "Artifacts Log 2: Gemma 2, more Chinese LLMs, high quality datasets, and domain-specific training," *Interconnects* (blog), July 10, 2024, https://www.interconnects.ai/p/artifacts-log-2; and Anson Ho et al., "Algorithmic progress in language models," arXiv, March 9, 2024, https://arxiv.org/pdf/2403.05812.

65 Benaich, *State of AI Report 2023*; and Suchin Gururangan et al., "Don't Stop Pretraining: Adapt Language Models to Domains and Tasks," arXiv, April 23, 2020, doi:10.48550/arXiv.2004.10964; Celso de Melo, "Research Challenges for Large Pre-Trained Models," (presentation, CDAO Advantage DoD 24: Defense Data & AI Symposium, Washington, DC, February 20–22, 2024), https://www.ai.mil/docs/ADOD24/ResearchChallengesforLPTMsCDAOAdvantagedeMeloEtAl_v4.pdf; and David Wallace-Wells, "How Long Will A.I.'s 'Slop' Era Last?" *New York Times*, July 24, 2024, https://www.nytimes.com/2024/07/24/opinion/ai-annoying-future.html.

66 Virginia L. Wydler, *Gaining Leverage over Vendor Lock to Improve Acquisition Performance and Cost Efficiencies* (McLean, VA: MITRE Corporation, April 1, 2014), https://www.mitre.org/sites/default/files/publications/gaining-leverage-over-vendor-lock-14-1262.pdf.

67 CSIS, "Open-source AI Roundtable"; Iain Cruikshank and Shane Kohtz, "Planning for AI Sustainment: A Methodology for Maintenance and Cost Management," (presentation, NPS Acquisition Research Symposium, Monterey, CA, May 9, 2024), video, 1:11:57, https://www.youtube.com/watch?v=A_yOE1CfmXk&t=3541s; and Sydney J. Freedberg Jr., "'Fight tonight' vs 'Getting it right': CDAO's Martell wants to build AI to last," Breaking Defense, February 21, 2024, https://breakingdefense.sites.breakingmedia.com/2024/02/fight-tonight-vs-getting-it-right-cdaos-martell-wants-to-build-ai-to-last/.

68 Arvind Narayanan and Sayash Kapoor, "Is GPT-4 getting worse over time?" AI Snake Oil (blog), July 19, 2023, https://www.aisnakeoil.com/p/is-gpt-4-getting-worse-over-time; Lingjiao Chen, Matei Zaharia, and James Zou, "How is ChatGPT's Behavior Changing Over Time?" *Harvard Data Science Review* 6, no. 2 (Spring 2024), doi:10.1162/99608f92.5317da47; Lauren Leffer, "Yes, AI Models Can Get Worse over Time," *Scientific American*, August 2, 2023, https://www.scientificamerican.com/article/yes-ai-models-can-get-worse-over-time/; and Anthony Sepci et al., "Artificial Intelligence and Model Risk Management," KPMG LLP, 2023, https://kpmg.com/kpmg-us/content/dam/kpmg/pdf/2024/1a-ai-and-model-risk-slipsheet.pdf.

69 Defense Acquisition University, "Modular Open Systems Approach," n.d., https://www.dau.edu/acquipedia-article/modular-open-systems-approach-mosa.

70 Frank Camm et al., *Data Rights Relevant to Weapon Systems in Air Force Special Operations Command* (Santa Monica, CA: Rand Corporation, 2021), https://www.rand.org/content/dam/rand/pubs/research_reports/RR4200/RR4298/RAND_RR4298.pdf; and Wydler, *Gaining Leverage*, 5–7.

71 Ashley Roque, "Sikorsky-Boeing's FLRAA bid was much cheaper, but couldn't offset 'unacceptable' design metric: GAO," Breaking Defense, April 14, 2023, https://breakingdefense.com/2023/04/sikorsky-boeings-flraa-bid-was-much-cheaper-but-couldnt-offset-unacceptable-design-metric-gao/.

72 Stephen Casper et al., "Black-Box Access Is Insufficient for Rigorous AI Audits," arXiv, January 25, 2024, doi:10.48550/arXiv.2401.14446; Andrew P. Hunter and Alexis Lasselle Ross, "Implementing Innovation: The Army's IP Strategy," (presentation, CSIS, Washington, DC: June 18, 2019), video, 54:49, https://www.youtube.com/watch?v=fvJmsRhJyJs; and Kshitij Gupta et al., "Continual Pre-Training of Large Language Models: How to (re)warm your model?" arXiv, August 8, 2023, doi:10.48550/arXiv.2308.04014.

73 Herbert C. Kemp, Lawrence Stutzriem, and Heather Penney, "Data Requirements and Rights: Time for a Reassessment," Mitchell Institute Policy Papers 14 (June 2018), https://mitchellaerospacepower.org/wp-content/uploads/2021/02/a2dd91_37e7f66f7bda4ca8b55a27604694094b.pdf; and Heidi M. Peters*, Intellectual Property and Technical Data in DOD Acquisitions* (Washington, DC: Congressional Research Service, April 22, 2022), https://crsreports.congress.gov/product/pdf/IF/IF12083.

[74] Maggie Gray and Max Dauber, “Simplifying AI Deployment for Defense,” *Gray Matters* (blog), June 18, 2024, https://maggiegray.substack.com/p/simplifying-ai-deployment-for-defense.

[75] Kemp, Stutzriem and Penney, “Data Requirements and Rights”; Peters, *Intellectual Property and Technical Data*; and Hunter and Ross, “Implementing Innovation.”

[76] CSIS, “Open-source AI and the Defense Industrial Base”; Colin Raffel, “Building Machine Learning Models Like Open Source Software,” *Communications of the ACM* 66, no. 2 (January 20, 2023): 38–40, doi:10.1145/3545111.

[77] Jon Harper, “Army set to issue new policy guidance on use of large language models,” DefenseScoop, May 9, 2024, https://defensescoop.com/2024/05/09/army-policy-guidance-use-large-language-models-llm/.

[78] CSIS, “Open-source AI and the Defense Industrial Base”; Max Lamparth and Jacquelyn Schneider, “Why the Military Can’t Trust AI,” *Foreign Affairs*, April 29, 2024, https://www.foreignaffairs.com/united-states/why-military-cant-trust-ai; and Celso de Melo, “Research Challenges for Large Pre-Trained Models,” (presentation, CDAO Advantage DoD 24: Defense Data & AI Symposium, February 20–22, 2024), https://www.ai.mil/docs/ADOD24/ResearchChallengesforLPTMsCDAOAdvantagedeMeloEtAl_v4.pdf.

[79] CSIS, “Open-source AI and the Defense Industrial Base.”

[80] Stella Biderman et al., “*Pythia*: A Suite for Analyzing Large Language Models Across Training and Scaling,” arXiv, May 31, 2023, https://arxiv.org/pdf/2304.01373; Miles Brundage et al., “Toward Trustworthy AI Development: Mechanisms for Supporting Verifiable Claims,” arXiv, April 15, 2020, doi:10.48550/arXiv.2004.07213.; Lei Huang et al., “A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions,” arXiv, November 9, 2023, https://arxiv.org/abs/2311.05232; Michael Hanna, Ollie Liu, and Alexandre Variengien, “How does GPT-2 compute greater-than?: Interpreting mathematical abilities in a pre-trained language model,” arXiv, April 30, 2023, https://arxiv.org/abs/2305.00586; Neel Nanda et al., “Progress measures for grokking via mechanistic interpretability,” arXiv, January 12, 2023, https://arxiv.org/abs/2301.05217; Nora Belrose et al., “Neural Networks Learn Statistics of Increasing Complexity,” arXiv, February 6, 2024, https://arxiv.org/abs/2402.04362; Eva Dou, Nitasha Tiku, and Gerrit De Vynck, “Pentagon explores military uses of large language models,” *Washington Post,* February 20, 2024, https://www.washingtonpost.com/technology/2024/02/20/pentagon-ai-llm-conference/; and Beren Millidge, “Open source AI has been vital for alignment,” *Beren’s Blog*, November 5, 2023, http://www.beren.io/2023-11-05-Open-source-AI-has-been-vital-for-alignment/.

[81] Casper et al., “Black-Box Access.”

[82] Nanda et al., “Progress measures for grokking”; Sydney J. Freedberg Jr., “Pentagon should experiment with AIs like ChatGPT–but don’t trust them yet: DoD’s ex-AI chiefs,” Breaking Defense, April 6, 2023, https://breakingdefense.com/2023/04/dods-ex-ai-chiefs-pentagon-should-experiment-with-ais-like-chatgpt-but-dont-trust-them-yet/; and Jon Harper, “Kendall: Generative AI tech like ChatGPT currently has limited military utility,” DefenseScoop, June 22, 2023, https://defensescoop.com/2023/06/22/kendall-generative-ai-tech-like-chatgpt-currently-has-limited-military-utility/.

[83] David M. Wennergren, “Clarifying Guidance Regarding Open Source Software (OSS),” memorandum, U.S. Department of Defense, October 16, 2009, https://dodcio.defense.gov/Portals/0/Documents/OSSFAQ/2009OSS.pdf.

[84] Shayne Longpre et al., “A Safe Harbor for AI Evaluation and Red Teaming,” arXiv, March 7, 2024, doi:10.48550/arXiv.2403.04893; and White House 2023, Exec. Order No. 14110.

[85] Jack Cable and Aeva Black, “With Open Source Artificial Intelligence, Don’t Forget the Lessons of Open Source Software,” Cybersecurity and Infrastructure Security Agency, July 29, 2024, https://www.cisa.gov/news-events/news/open-source-artificial-intelligence-dont-forget-lessons-open-source-software.

[86] CSIS, “Open-source AI and the Defense Industrial Base”; and Alison Smith and Drew Farris, “The case for open-source generative AI in government,” Booz Allen Hamilton, n.d., https://www.boozallen.com/insights/ai/the-case-for-open-source-generative-ai-in-government.html.

[87] Allen, Lugo, and Strohmeyer, “Scaling AI-enabled Capabilities at the DOD.”

[88] Niklas Muennighoff et al., “MTEB: Massive Text Embedding Benchmark,” arXiv, October 13, 2022, doi:10.48550/arXiv.2210.07316; “Massive Text Embedding Benchmark (MTEB) Leaderboard,” Huggingface, accessed August 2, 2024, https://huggingface.co/spaces/mteb/leaderboard/; and Matthew Lynley, “Let’s check back in on the vector databases,” *Supervised* (blog), April 4, 2024, https://www.supervised.news/p/lets-check-back-in-on-the-vector?publication_id=1459978.

[89] Nur Ahmed and Neil C. Thompson, “What should be done about the growing influence of industry in AI research?” Brookings Institution, December 5, 2023, https://www.brookings.edu/articles/what-should-be-done-about-the-growing-influence-of-industry-in-ai-research/; and Kaiyuan Gao et al., “Examining User-Friendly and Open-Sourced Large GPT Models: A Survey on Language, Multimodal, and Scientific GPT Models,” arXiv, August 27, 2023, doi:10.48550/arXiv.2308.14149.

[90] U.S. Department of Defense*, 2022 National Defense Strategy of the United States of America* (Washington, DC: U.S. Department of Defense, 2022), https://media.defense.gov/2022/Oct/27/2003103845/-1/-1/1/2022-NATIONAL-DEFENSE-STRATEGY-NPR-MDR.PDF; and U.S. Department of Defense, *National Defense Industrial Strategy*, 2023, 19–20, https://www.businessdefense.gov/docs/ndis/2023-NDIS.pdf.

[91] Office of the Under Secretary of Defense for Acquisition and Sustainment, *State of Competition within the Defense Industrial Base* (Washington, DC: U.S. Department of Defense, February 2022), https://media.defense.gov/2022/Feb/15/2002939087/-1/-1/1/STATE-OF-COMPETITION-WITHIN-THE-DEFENSE-INDUSTRIAL-BASE.PDF; Doug Cameron, “Lagging Arms Production Makes Pentagon Wary of Further Industry Consolidation,” *Wall Street Journal*, January 3, 2023, https://www.wsj.com/articles/lagging-arms-production-makes-pentagon-wary-of-further-industry-con-

solidation-11672754128?mod=article_inline; Bill Greenwalt, “FTC activism and Ukraine signal a new era for the US defense industrial base,” Breaking Defense, August 10, 2022, https://breakingdefense.com/2022/08/ftc-activism-and-ukraine-signal-a-new-era-for-the-us-defense-industrial-base/; and Sanders et al., “Is Bigger Better?”

92 Alexander Palmer, Henry H. Carroll, and Nicholas Velazquez, “Unpacking China’s Naval Buildup,” CSIS, June 5, 2024, https://www.csis.org/analysis/unpacking-chinas-naval-buildup; Michael Lee, “Chinese shipbuilding capacity over 200 times greater than US, Navy intelligence says,” Fox News, September 14, 2023, https://www.foxnews.com/world/chinese-shipbuilding-capacity-over-200-times-greater-than-us-navy-intelligence-says; Matthew P. Funaiole, “The Threat of China’s Shipbuilding Empire,” CSIS, *Commentary,* May 10, 2024, https://www.csis.org/analysis/threat-chinas-shipbuilding-empire; Ronald O’Rourke, *Navy Force Structure and Shipbuilding Plans: Background and Issues for Congress* (Washington, DC: Congressional Research Service, March 2024), https://crsreports.congress.gov/product/pdf/RL/RL32665; Lori Ann LaRocco, “Biden promise to rival China on shipbuilding faces a big economic problem,” CNBC, April 25, 2024, https://www.cnbc.com/2024/04/25/bidens-plan-to-rival-china-shipbuilders-has-a-big-economic-problem.html; Jack Reed and Stacie Pettyjohn, “Fireside Chat with Senator Jack Reed,” (presentation, Center for a New American Security, Washington DC, April 24, 2023), https://www.cnas.org/publications/transcript/fireside-chat-with-senator-jack-reed; and Wallar, “U.S. Coast-Wise Shipping.”

93 Noah Robertson, “The Pentagon wants industry to transform again to meet demand. Can it?” *Defense News,* February 20, 2024, https://www.defensenews.com/industry/2024/02/20/the-pentagon-wants-industry-to-transform-again-to-meet-demand-can-it/; and Jen Judson, “How companies plan to ramp up production of Patriot missiles,” *Defense News*, April 9, 2024, https://www.defensenews.com/land/2024/04/09/how-companies-plan-to-ramp-up-production-of-patriot-missiles/.

94 Jim Garamone, “Hicks Again Makes Case for Strengthening Industrial Base, Eliminating Continuing Resolutions,” *DOD News*, U.S. Department of Defense, March 20, 2024, https://www.defense.gov/News/News-Stories/Article/Article/3713186/hicks-again-makes-case-for-strengthening-industrial-base-eliminating-continuing/.

95 Seth G. Jones and Alexander Palmer, *Rebuilding the Arsenal of Democracy: The U.S. and Chinese Defense Industrial Bases in an Era of Great Power Competition* (Washington, DC: CSIS, March 2024), https://csis-website-prod.s3.amazonaws.com/s3fs-public/2024-03/240306_Jones_Rebuilding_Democracy_0.pdf?VersionId=sCRmR1UN.8dKYYu4hyh8PMB.U32ww58D.

96 Jeffrey Ding and Jenny Xiao, *Recent Trends in China’s Large Language Model Landscape* (Oxford, UK: Centre for the Governance of AI, 2023), https://www.governance.ai/research-paper/recent-trends-chinas-llm-landscape; Michael C. Horowitz et al., *Strategic Competition in an Era of Artificial Intelligence* (Washington, DC: Center for a New American Security, July 25, 2018), https://www.cnas.org/publications/reports/strategic-competition-in-an-era-of-artificial-intelligence; Jing Cheng and Jinghan Zeng, “Shaping AI’s Future? China in Global AI Governance,” *Journal of Contemporary China* 32, no. 143 (September 3, 2023): 794–810, doi:10.1080/10670564.2022.2107391; Daniel Araya, “Who will lead in the age of artificial intelligence?” Brookings Institution, February 26, 2019, https://www.brookings.edu/articles/who-will-lead-in-the-age-of-artificial-intelligence/; Kai-Fu Lee, *AI Superpowers: China, Silicon Valley, and the New World Order* (New York, NY: Harper Business, 2018); Will Knight, “This Chinese Startup Is Winning the Open Source AI Race,” *Wired*, January 23, 2024, https://www.wired.com/story/chinese-startup-01-ai-is-winning-the-open-source-ai-race/; and Rita Liao, “Valued at $1B, Kai-Fu Lee’s LLM startup unveils open source model,” TechCrunch, November 5, 2023, https://techcrunch.com/2023/11/05/valued-at-1b-kai-fu-lees-llm-startup-unveils-open-source-model/.

97 Zeyi Yang, “Why Chinese companies are betting on open-source AI,” *MIT Technology Review,* July 24, 2024, https://www.technologyreview.com/2024/07/24/1095239/chinese-companies-open-source-ai/; Arcesati and Meinhardt, *China bets on open-source technologies*; Klon Kitchen, “The UAE’s AI Dreams,” *The Kitchen Sync*, American Enterprise Institute, March 25, 2024, https://www.aei.org/op-eds/the-uaes-ai-dreams/; Lisa Barrington, “Abu Dhabi makes its Falcon 40B AI model open source,” Reuters, May 25, 2023, https://www.reuters.com/technology/abu-dhabi-makes-its-falcon-40b-ai-model-open-source-2023-05-25/; and Islam Al Khatib, “Beyond Techwashing: The UAE’s AI Industrial Policy as a Security Regime,” in *AI Nationalism(s): Global Industrial Policy Approaches to AI* (New York, NY: AI Now Institute, March 12, 2024), https://ainowinstitute.org/wp-content/uploads/2024/03/AI-Nationalisms-Chapter-7.pdf.

98 Andrew Imbrie and Elsa B. Kania, *AI Safety, Security, and Stability Among Great Powers: Options, Challenges, and Lessons Learned for Pragmatic Engagement* (Washington, DC: Center for Security and Emerging Technology, December 2019), https://cset.georgetown.edu/wp-content/uploads/CSET-AI-Safety-Security-and-Stability-Among-the-Great-Powers.pdf; and Khosla, Young, and Helberg, “Perspectives on the US-China Techno-Economic War.”

99 Ashish Vaswani et al., “Attention Is All You Need,” in *Advances in Neural Information Processing Systems* 30 (2017), https://papers.nips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.

100 Kevin Roose, “A.I. Has a Measurement Problem,” *New York Times,* April 15, 2024, https://www.nytimes.com/2024/04/15/technology/ai-models-measurement.html.

101 “Evaluating Generative AI Technologies,” National Institute of Standards and Technology, accessed July 29, 2024, https://ai-challenges.nist.gov/genai.

102 Dan Hendrycks et al., “Measuring Massive Multitask Language Understanding,” arXiv, September 7, 2020, doi:10.48550/arXiv.2009.03300. 5-shot, non-CoT results selected where specified.

103 “Yi-1.5-34B-Chat,” Huggingface, accessed July 25, 2024, https://huggingface.co/01-ai/Yi-1.5-34B-Chat.

104 Ibid.

105 “Hello Qwen2,” QwenLM, June 7, 2024, https://qwenlm.github.io/blog/qwen2/.

106 “Introducing Qwen1.5,” QwenLM, February 4, 2024, https://qwenlm.github.io/blog/qwen1.5/.

107 “Hello Qwen2.”

108 “OLMo 1.7-7B: A 24 point improvement on MMLU,” *AI2 Blog*, April 17, 2024, https://blog.allenai.org/olmo-1-7-7b-a-24-point-improvement-on-mmlu-92b43f7d269d.

109 “Model Card and Evaluations for Claude Models,” Anthropic, accessed July 25, 2024, https://www-cdn.anthropic.com/5c49cc247484cecf107c699baf29250302e5da70/claude-2-model-card.pdf.

110 Ibid.

111 Anthropic, *The Claude 3 Model Family*.

112 Ibid.

113 Anthropic, “Claude 3.5 Sonnet Model Card Addendum,” Anthropic, n.d., https://www-cdn.anthropic.com/fed9cc193a14b84131812372d8d5857f8f304c52/Model_Card_Claude_3_Addendum.pdf.

114 “c4ai-command-r-plus,” Huggingface, accessed July 25, 2024, https://huggingface.co/CohereForAI/c4ai-command-r-plus.

115 Ibid.

116 Mosaic AI Research, “Introducing DBRX: A New State-of-the-art Open LLM,” Databricks, March 27, 2024, https://www.databricks.com/blog/introducing-dbrx-new-state-art-open-llm.

117 “DeepSeek-LLM,” Github, accessed July 25, 2024, https://github.com/deepseek-ai/DeepSeek-LLM.

118 DeepSeek-AI et al., “DeepSeek-V2: A Strong, Economical, and Efficient Mixture-of-Experts Language Model,” arXiv, May 7, 2024, http://arxiv.org/abs/2405.04434.

119 “pythia-12b,” Huggingface, accessed July 25, 2024, https://huggingface.co/EleutherAI/pythia-12b/blob/refs%2Fpr%2F2/README.md.

120 Rohan Anil et al., “PaLM 2 Technical Report,” arXiv, September 13, 2023, doi:10.48550/arXiv.2305.10403.

121 Ibid.

122 Harry Mellor, “Flan-T5: Sweet Results with the Smaller, More Efficient LLM,” Graphcore, May 30, 2023, https://www.graphcore.ai/posts/flan-t5-sweet-results-with-the-smaller-more-efficient-llm.

123 Anil et al., “PaLM 2 Technical Report.”

124 “Gemini Models,” Google DeepMind, accessed July 25, 2024, https://deepmind.google/technologies/gemini/.

125 Ibid.

126 “gemma-2-27b,” Huggingface, accessed July 25, 2024, https://huggingface.co/google/gemma-2-27b.

127 Anil et al., “PaLM 2 Technical Report.”

128 Aakanksha Chowdhery et al., “PaLM: Scaling Language Modeling with Pathways,” arXiv, April 5, 2022, http://arxiv.org/abs/2204.02311.

129 Anil et al., “PaLM 2 Technical Report.”

130 “Llama / Model Card,” Github, February 24, 2023, accessed July 25, 2024, https://github.com/meta-llama/llama/blob/main/MODEL_CARD.md.

131 “Meta-Llama-3.1-70B,” Huggingface, accessed July 31, 2024, https://huggingface.co/meta-llama/Meta-Llama-3.1-70B.

132 “Llama3 / Model Card,” Github, April 11, 2024, https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md.

133 “Meta-Llama-3.1-70B.”

134 Ibid.

135 Ibid.

136 Ibid.

137 “Llama / Model Card.”

138 “Mistral 7B,” Mistral AI, September 27, 2023, https://mistral.ai/news/announcing-mistral-7b/.

139 “Large Enough,” Mistral AI, July 24, 2024, https://mistral.ai/news/mistral-large-2407/.

140 “Mixtral of experts,” Mistral AI, December 11, 2023, https://mistral.ai/news/mixtral-of-experts/.

141 “Cheaper, Better, Faster, Stronger,” Mistral AI, April 17, 2024, https://mistral.ai/news/mixtral-8x22b/.

142 “mistral-nemo-12b-instruct,” Nvidia, accessed July 25, 2024, https://build.nvidia.com/nv-mistralai/mistral-nemo-12b-instruct/modelcard.

143 Nvidia et al., “Nemotron-4 340B Technical Report,” arXiv, June 17, 2024, http://arxiv.org/abs/2406.11704.

144 Mellor, “Flan-T5: Sweet Results.”

145 “GPT-4o mini: advancing cost-efficient intelligence,” OpenAI, July 18, 2024, https://openai.com/index/gpt-4o-mini-advancing-cost-efficient-intelligence/.

146 “Hello GPT-4o,” OpenAI, May 13, 2024, https://openai.com/index/hello-gpt-4o/.

147 Ibid.

148 Ebtesam Almazrouei et al., “The Falcon Series of Open Language Models,” arXiv, November 28, 2023, http://arxiv.org/abs/2311.16867.

149 Aohan Zeng et al., “GLM-130B: An Open Bilingual Pre-trained Model,” arXiv, October 5, 2022, http://arxiv.org/abs/2210.02414.

150 Team GLM et al., “ChatGLM: A Family of Large Language Models from GLM-130B to GLM-4 All Tools,” arXiv, June 18, 2024, doi:10.48550/arXiv.2406.12793.