

March 27, 2024

Stephanie Weiner
National Telecommunications and Information Administration
U.S. Department of Commerce
1401 Constitution Avenue NW, Washington, DC 20230

RE: Dual Use Foundation Artificial Intelligence Models With Widely Available Model Weights, Docket No. 240216-0052

Dear Ms. Weiner:

The below signed civil, human, and technology rights organizations thank the National Telecommunications and Information Administration (NTIA) for its request for comment on “Dual Use Foundation Artificial Intelligence Models With Widely Available Model Weights”¹ and its larger work to ensure that artificial intelligence is safe, trustworthy, and equitable. Although the emergence of foundation models with widely available model weights has spurred significant discourse about the potential risks they pose,² we urge NTIA to recognize the corresponding benefits of “openness” as it considers appropriate regulatory approaches. In particular, “openness”—a term covering data, model architecture, model weights, and source code in the AI context—can help advance transparency around models that impact civil rights, safety, and access to opportunity, where AI is already being deployed and already affecting people’s lives.

Both Advanced and More Rudimentary AI Technology Are Affecting Civil Rights Now, with the Potential for Large Foundation Models to Have Increased Impact in the Future

Computer-based decision-making processes, including statistical models, automated systems and other forms of artificial intelligence are already affecting individuals’ access to crucial life opportunities.³ These systems range from basic statistical models to foundation models including large language models. Examples of sectors where both legacy and foundation models are being used to make decisions that impact opportunities include:

- **Housing:** Tenant-screening algorithms are prone to errors and incorrectly include criminal or eviction records tied to people with similar names.⁴ Open Communities, a nonprofit housing advocacy organization, recently settled a lawsuit against a property management company and related organizations for the use of a chatbot, an application that can be driven by a foundation model, that systematically denied housing

¹ 89 Fed. Reg. 14059 (Feb. 26, 2024), [here](#) [hereinafter RFC].

² Sharon Goldman, *Why Anthropic and OpenAI Are Obsessed with Securing LLM Model Weights*, Venture Beat (Dec. 15, 2023), [here](#); cf. The Leadership Conference on Civil and Human Rights, et.al., *Civil Rights Principles in an Era of Big Data* (2020) [here](#).

³ Cf. 15 U.S.C. § 9401(3); The White House, *Blueprint for an AI Bill of Rights at 10* (2022), [here](#) (definition of “automated system”).

⁴ Kaveh Waddell, *How Tenant Screening Reports Make It Hard for People to Bounce Back From Tough Times*, Consumer Reports (Mar. 11, 2021), [here](#); Lauren Kirchner, *Can Algorithms Violate Fair Housing Laws?*, The Markup (Sept. 24, 2020), [here](#).

opportunities for individuals and families that use housing choice vouchers.⁵ Moreover, dynamic rental pricing powered by AI can be a barrier to housing for rental voucher recipients as rent prices potentially fluctuate above HUD fair market rent amounts. This can negatively impact low wealth groups who are disproportionately single, female-headed families with children and people with disabilities⁶ as well as people who live in rural communities.⁷

- **Employment:** Employers are using large language models,⁸ which are instances of foundation models to evaluate job applicants, and those technologies can unfairly advantage male candidates or de-preference first-generation college graduates and racial minorities.⁹ Other AI-driven hiring technology such gamified personality tests can be inaccessible to applicants with disabilities.¹⁰
- **Credit:** Credit scoring systems are algorithmic models that attempt to predict a borrower's risk and how well that person is likely to repay their debt obligations. These systems typically generate a numerical score used to help players in the financial services system determine the creditworthiness of a consumer. They are often used as part of a lender's decisions on underwriting and pricing. Algorithmic credit scoring disadvantages Black, Latinx, and Native American consumers who have historically had less access to traditional credit than white consumers.¹¹ As Federal Reserve Vice Chair of Supervision Michael Barr stated, "Artificial Intelligence...relies on the data that is out there in the world and the data...is flawed. Some of it is just wrong. Some of it is deeply biased...Information we have on the Internet is imperfect...if you train a Machine Learning device, if you train a Large Language Model on imperfect data, you're going to get imperfect results."¹²

⁵ *Open Communities Reaches Resolution in Case Alleging AI Discrimination*, Open Communities (Jan. 31, 2024), [here](#).

⁶ See Claudia D. Solari et al., *Housing Insecurity in the District of Columbia*, Urban Institute (Nov. 16, 2023), [here](#); Sharon Cornelissen, *The Pandemic Aggravated Racial Inequalities in Housing Insecurities: What Can It Teach Us about Housing Amidst Crisis?*, Harvard Joint Center for Housing Studies (July 12, 2023), [here](#).

⁷ See Irina Ivanova, *Inflation is Hurting Rural Americans More Than City Folks - Here's Why*, CBS MoneyWatch (Dec. 2, 2021), [here](#).

⁸ Leon Yin et al., *OpenAI's GPT Is a Recruiter's Dream Tool. Tests Show There's Racial Bias*, Bloomberg (Mar. 7, 2024), [here](#).

⁹ Avi Asher-Schapiro, *AI is Taking Over Job Hiring, But Can It Be Racist?*, Thomson Reuters (June 7, 2021), [here](#); Jeffrey Dastin, *Insight - Amazon Scraps Secret AI Recruiting Tool that Showed Bias Against Women*, Thomson Reuters (Oct. 10, 2018), [here](#).

¹⁰ Lydia X.Z. Brown et al., *Center for Democracy & Technology, Algorithm-Driven Hiring Tools: Innovative Recruitment or Expedited Disability Discrimination?* (2020), [here](#).

¹¹ Emmanuel Martinez & Lauren Kirchner, *The Secret Bias Hidden in Mortgage-Approval Algorithms*, The Markup (Aug. 25, 2021), [here](#).

¹² See Federal Reserve Board of Governors Vice Chair Michael Barr, *Setting the Foundation for Effective Governance and Oversight: A Conversation with U.S. Regulators*, Responsible AI Symposium (Jan. 19, 2024), [here](#).

- **Governmental Benefits:** Algorithmic systems have been employed in a variety of governmental programs, including state Medicaid programs in Idaho and Arkansas, resulting in arbitrary denials of or reductions in benefits.¹³
- **Healthcare:** An algorithm widely used in hospitals is less likely to refer Black patients for care than equally ill white patients,¹⁴ resulting in 28.8% of Black patients being incorrectly deemed ineligible for additional care.¹⁵
- **Education:** Algorithmic proctoring tools fail to recognize students of color and flag “atypical” eye and body movements of students with disabilities as “cheating behaviors.”¹⁶
- **Child Welfare:** An ACLU and Human Rights Data Analysis Group audit of an algorithmic risk-scoring system used in child welfare determinations highlighted the ways in which the algorithm could potentially disproportionately flag Black families and families with disabilities for investigation, due in part to the data sources it relied on.¹⁷
- **Insurance:** Algorithms assign at least 10% higher car insurance premiums in minority zip codes than in majority-white zip codes¹⁸ and analyze individuals’ online activities to assign life insurance premiums.¹⁹ Additionally, the Casualty Actuarial Society (CAS) acknowledged that algorithmic bias can manifest in systems used in the insurance sector including underwriting, pricing, and claims models.²⁰
- **Criminal Legal System:** Algorithmic systems are used to predict the likelihood an incarcerated person will pose a risk on release, but those systems are punctuated by multiple decisions that incorporate existing disparities.²¹

These examples span the range of technology that may be deemed “artificial intelligence,” and the emergence of applications such as chatbots and facial recognition technology underscores that both rudimentary and the most sophisticated AI technologies are already affecting civil rights, safety, and access to opportunities.

¹³ Testimony of Ritchie Eppink, Hearing AI in Government Before the S. Comm. On Homeland Security & Government Affairs (May 16, 2023), [here](#); Colin Lecher, *What Happens When an Algorithm Cuts Your Health Care*, The Verge (Mar. 21, 2018), [here](#).

¹⁴ Heidi Ledford, *Millions of Black People Affected by Racial Bias in Health-Care Algorithms*, Nature (Oct. 24, 2019), [here](#).

¹⁵ Emily Sokol, *Eliminating Racial Bias in Algorithm Development*, Health IT Analytics (Dec. 26, 2019), [here](#).

¹⁶ Lydia X.Z. Brown, *How Automated Test Proctoring Software Discriminates Against Disabled Students*, Center for Democracy & Technology (Nov. 16, 2020), [here](#); Todd Feathers & Janus Rose, *Students Are Rebelling Against Eye-Tracking Exam Surveillance Tools*, Motherboard (Sept. 24, 2020), [here](#).

¹⁷ Marissa Gerchick et al., ACLU, *The Devil is in the Details: Interrogating Values Embedded in the Allegheny Family Screening Tool* (2023), [here](#).

¹⁸ Julia Angwin et al., *Minority Neighborhoods Pay Higher Car Insurance Premiums Than White Areas With the Same Risk*, ProPublica (Apr. 5, 2017), [here](#).

¹⁹ Angela Chen, *Why the Future of Life Insurance May Depend on Your Online Presence*, The Verge (Feb. 7, 2019), [here](#).

²⁰ Ronda Lee, *AI Can Perpetuate Racial Bias in Insurance Underwriting*, Yahoo! Money (Nov. 1, 2022).

²¹ Marissa Gerchick & Brandon Buskey, ACLU, *ACLU Statement on the PATTERN Risk Assessment Tool* (2022), [here](#) [hereinafter ACLU PATTERN Statement].

Existing AI systems also play a critical role in preparing and enriching the data used to train foundation models, including large language models. From housing decisions to insurance decisions to decisions made by criminal legal systems built on data-driven algorithms, outputs from legacy AI systems often serve as inputs for large language models, and there are multiple instances where legacy AI systems are either integrated with foundation models or used to provide information needed by foundation models to function effectively. As both large language models and existing AI systems continue to evolve, we can expect even more innovative ways to leverage their combined capabilities. Hence, we urge the NTIA to also consider the role of non-foundation AI models in its efforts to ensure that dual use foundation models do not lead to irreparable roadblocks in the path to equitable, life-changing opportunities.

“Open” AI May Advance Regulatory Goals to Ensure Robust Auditing of and Transparency for Algorithmic Systems

Responsive to Questions 1 and 3.

Auditing and Transparency Are Crucial to Protect Civil Rights, Safety, and Access to Opportunity

In addressing AI’s risks for civil rights, safety, and access to opportunity, advocates, affected communities, and policymakers have championed a number of regulatory goals, including auditing and assessments, transparency, and explainability.²²

Audits and assessments of algorithmic systems for discriminatory impacts, fitness for its intended purpose, and unintended outcomes should take place prior to deployment of an algorithmic system and on an ongoing basis afterwards. Audits and assessments should occur in a real-world context that anticipates the conditions in which the algorithmic system will be used—including accounting for likely operators’ training and consulting with impacted communities. Assessments and audits should not only identify potential harms but also implement mitigations. An algorithmic system should be deployed for a particular use case only if its harms are outweighed by its benefits and the harms have been appropriately mitigated. Ongoing post-deployment monitoring should continue to assess risks and develop new mitigation techniques to reduce those risks. Ultimately, AI that is biased should not be used. Where the algorithmic system’s risks exceed an acceptable level and where mitigation is not practicable, the algorithmic system should be discontinued or decommissioned.

Crucially, to avoid conflicts of interests, audit and assessments should be conducted by independent third parties—who often do not have direct access to AI models or their underlying training data.²³ Assessments and auditing are critical if AI systems are to be trusted. In addition,

²² Other regulatory measures in addition to auditing and assessments, transparency, and explainability are also necessary to protect civil rights, safety, and civil liberties. See, e.g., Comments of the ACLU, Docket No. OMB-2023-0020 (2023), [here](#); Center for Democracy & Technology et al., Civil Rights Standards for 21st Century Employment Selection Procedures (2022), [here](#).

²³ The White House, Blueprint for an AI Bill of Rights at 20 (2022), [here](#).

audits and assessments should be based on NIST’s AI Risk Management Framework, a foundational tool for the White House Executive Order on Safe, Secure and Trustworthy Use and Development of AI.²⁴

Similarly, transparency and explanations about how an algorithmic system functions are essential for protecting civil rights, safety, and access to opportunity. Transparency and explanations should include:²⁵

- the system’s purposes;
- intended use cases;
- the factors and data it relies on;
- the data and factors’ relationship to the system’s intended purpose and each other;
- steps taken to ensure that the training data and samples are accurate and representative;
- measured disparate impact or discriminatory outcomes; and,
- corresponding mitigations.

Transparency will entail making explanations, audits, and assessments not only publicly available but also *approachable* and *accessible*. Approachable transparency will require making documentation available in a format that is brief, easy to read, and in plain language, so that people may readily understand it, in addition to more technical documentation. Accessibility requires that explanations and transparency be available in a variety of formats that accommodate individuals’ disabilities and other barriers to access.

The Definition of “Open” AI Should Recognize Degrees of Openness and Corresponding Balancing of Risks and Benefits

“Open” AI may advance each of those regulatory goals.²⁶ As the NTIA recognizes, “openness” comes in a variety of degrees, from fully “closed” to fully “open.”²⁷ Moreover, the application of the concept of “openness” to complex AI models is not always clear.²⁸ As Gray Widder et al. recognized:

Some systems described as ‘open’ offer little more than an API or the ability to download a hosted model. In these cases, many question whether the term should be applied at all, or if it is ‘openwashing’ systems that should be understood as ‘closed’. Other more

²⁴ The National Fair Housing Alliance published a framework that can be used to audit any algorithmic system. The framework can help AI developers and deployers to identify sources of bias or risks in their algorithmic systems. The framework can be accessed [here](#), and used in conjunction with the NIST AI Risk Management Framework, [here](#).

²⁵ Cf. CFPB Acts to Protect the Public from Black-Box Credit Models Using Complex Algorithms, Consumer Financial Protection Bureau (May 26, 2022), [here](#).

²⁶ Nicol Turner Lee, Brookings, Written Statement to the U.S. Senate AI Insight Forum on Transparency, Explainability, Intellectual Property, & Copyright (2023), [here](#) (“Experts continuously emphasize the difficulty in understanding why an AI model behaves as it does That is why there is an increasing need and call for openness”).

²⁷ RFC at 14061.

²⁸ Ayah Bdeir & Camille François, *Introducing the Columbia Convening on Openness and AI*, Mozilla (Mar. 6, 2024), [here](#); Stefano Maffulli, *Open Source AI Definition: Where It Stands and What’s Ahead*, Open Source Initiative (Feb. 7, 2024), [here](#).

maximal versions of ‘open’ AI go further, offering access to the source code, underlying training data and full documentation, as well as licensing the AI system for wide reuse These precise distinctions matter, in part because the ideologies of ‘open’ and ‘open source’ are currently being projected onto ‘open’ AI systems even when they don’t fit—ideologies that were forged decades ago in the context of open source software at a time when the current tech industry was emergent, not immensely powerful.²⁹

Further compounding the complexity around “open” AI is the fact that it is not always easy to separate “openness” from the business interests of large AI developers, who may benefit from open innovation on their platforms and may later withdraw commitments to openness after the benefits have reached a critical mass, knowing that smaller developers are unlikely to have the resources necessary to independently compete.³⁰

In light of those considerations and multiple degrees of “openness,” we believe that NTIA should define “open” AI as a range of “AI systems in which one or many components [such as data, model weights, source code, and model architecture of the AI system] are offered transparently, reusably, or in ways that allow third parties to extend and build on top [of the components or systems].”³¹ In that definition, openness is built on three attributes:

- transparency, “the ability to access and vet source code, documentation and data”;
- reusability, “the ability and licensing needed to allow third parties to reuse source code and/or data”; and,
- extensibility, “the ability to build on top of extant off-the-shelf models, ‘tuning’ them for one or another specific purpose.”³²

That definition, however, should not be used in a vacuum. Instead, “openness” should always be discussed with reference both to the *components* of the AI system that are being made widely available—such as the model weights, architecture or coding, or training data—and the *attributes* that define the “open” system—namely, transparency, reusability, or extensibility. “Openness” varies wildly, depending on whether the components being made available are little more than an application program interface, documentation, the model weights, or the model’s code and training data. This approach to “openness” appropriately allows regulatory responses to the corresponding risks and benefits to be tailored appropriately, as described below.

“Open” AI Will Further Civil Rights, Safety, and Access to Opportunity

Each degree of openness may further civil rights goals of transparency and explainability, especially when bolstered by additional protections as necessary. The White House’s AI Bill of

²⁹ David Gray Widder et al., *Open (For Business): Big Tech, Concentrated Power, and the Political Economy of Open AI*, SSRN at 4 (2023), [here](#).

³⁰ *Id.* at 3, 12; Devin Coldewey, *Why Elon Musk’s AI Company ‘Open-Sourcing’ Grok Matters— and Why It Doesn’t*, Tech Crunch (Mar. 18, 2024), [here](#) (“If you’re not already in possession of, say, a dozen Nvidia H100s in a six-figure AI inference rig, don’t bother clicking that download link.”); Benj Edwards, *Elon Musk’s xAI Releases Grok Source and Weights, Taunting OpenAI*, Ars Technica (Mar. 18, 2024), [here](#).

³¹ Gray Widder et al., *supra* note 29, at 3.

³² *Id.* at 3-4.

Rights recognized the need for independent auditors to “be given access to the system and samples of associated data,”³³ and varying degrees of openness will aid that access.

The request for comment is “particularly focused on the wide availability, such as being publicly posted online, of foundation model weights,”³⁴ pursuant to the charge in Section 4.6 of President Biden’s Executive Order on “Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence.”³⁵ Even this degree of openness can significantly aid transparency and explainability, especially if samples of associated data are provided. As others have observed, “Widely available model weights enable external researchers, auditors, and journalists to investigate and scrutinize foundation models more deeply,” including to assess harms to marginalized communities, by better understanding the relationship among the parameters evaluated by the model, especially in the context of sample data used to derive the weights.³⁶

“Openness” might also refer to openness around a model’s underlying training data, source code, and architecture.³⁷ Many of the discriminatory harms perpetuated by AI can arise from the data used to train a model or the human judgments represented as the model’s code and architecture.³⁸ For example:

- The federal Bureau of Prison uses the “PATTERN” risk assessment tool to purportedly predict an individual’s likelihood of recidivism—but instead relying on historic data regarding recidivism *per se*, the model was developed with data regarding rearrest or return to custody following release, which more accurately reflects racially biased policing practices than an individual’s risk of committing a crime.³⁹ Similarly, PATTERN’s thresholds for low- and high-risk scores reflect the value judgments of the model’s creators.⁴⁰
- An algorithm used in one Pennsylvania county to score families for investigation similarly relied on training data that potentially biased the algorithm’s predictive outcomes.⁴¹ The algorithm relied on existing government databases, including county child welfare, juvenile probation, and behavioral health records. Problematically, those databases reflect the lives of those who have more contact with government agencies and systems shaped by historical and ongoing discrimination—not necessarily those who pose greater “risk” to their children—potentially over-representing some groups of people and underrepresenting others.

In each of these cases, the availability of the training data, production data, model weights, and information about the design of the model were integral in assessing its potential impacts.

³³ The White House, Blueprint for an AI Bill of Rights at 20 (2022), [here](#).

³⁴ RFC at 14062.

³⁵ Executive Order 14110 of October 30, 2023, 88 Fed. Reg. 75191 (November 1, 2023), [here](#).

³⁶ Sayash Kapoor et al., *On the Societal Impact of Open Foundation Models*, arXiv at 4-5 (2024), [here](#).

³⁷ Andreas Liesenfeld et al., *Opening Up ChatGPT: Tracking Openness, Transparency, and Accountability in Instruction-Tuned Text Generators*, arXiv at 1-3 (2023), [here](#).

³⁸ *E.g.*, ACLU PATTERN Statement, *supra* note 21; Ivey Dyson, *How AI Threatens Civil Rights and Economic Opportunities*, Brennan Center (Nov. 16, 2023), [here](#).

³⁹ ACLU PATTERN Statement, *supra* note 21, at 4-5.

⁴⁰ ACLU PATTERN Statement, *supra* note 21, at 7 n.21.

⁴¹ Marissa Gerchick et al., ACLU, *The Devil is in the Details: Interrogating Values Embedded in the Allegheny Family Screening Tool* (2023), [here](#).

Voluntary commitments to openness can aid in future assessments and audits of AI's impact on civil rights, safety, and access to opportunities. AI systems used in such consequential ways should be able to withstand scrutiny.

Appropriately Tailored Regulatory Responses May Mitigate the Risks of Both Open and Closed AI

Despite the benefits of openness for civil rights, safety, and access to opportunity, open AI—like “closed” AI—may still pose risks, and regulatory responses should be tailored accordingly. As others have observed, “[E]ven the most open of ‘open’ AI systems do not, on their own, ensure democratic access to or meaningful competition in AI, nor does openness alone solve the problem of oversight and scrutiny.”⁴² The same harms that might be posed by *any* AI system, such as discriminatory hiring decisions or the automated rejection of housing choice vouchers, may be posed by open AI. Consequently, any system that might plausibly—or implausibly—qualify for the label “open”⁴³ should not be automatically exempted from fundamental regulatory protections.

For example, the use of open source AI by employers to evaluate job applicants raises the same concerns as the use of “closed” AI of discrimination and a poor fit for purpose. Employers that deploy that AI to make hiring decisions—and the vendors that develop or fine-tune it for that purpose—should subsequently be subject to requirements, such as those to audit and evaluate their systems both by themselves and independent third parties.⁴⁴ Although “open” AI may substantially advance transparency and explainability, it should not serve as a broad exemption from regulation. Similarly, the release of data sets collected by a developer to train models should be subject to the same laws as any release of information, such as privacy protections for personal information contained in the data set.⁴⁵

In imposing obligations on the deployers and developers of AI systems, policymakers should also be cognizant of the position of open foundation models in the wider AI ecosystem.⁴⁶ To the extent it is appropriate to impose obligations on the developers of open foundation models, they should be carefully tailored to those developers’ role in the open ecosystem without diminishing the benefits of open AI.⁴⁷ In the end, both developers and deployers of AI systems have a shared responsibility to ensure that the systems they design, develop, and deploy are fit for purpose—that they work and are not biased.

⁴² Gray Widder et al., *Open (For Business)*, *supra* note 29, at 7.

⁴³ Gray Widder et al., *Open (For Business)*, *supra* note 29, at 2 (“Having a grasp on the multiplicity of definitions for ‘open’ and ‘open source’ AI, alongside what such systems do and do not enable, also matters because classifying a system as ‘open’ or ‘open source’ could have downstream effects on how it is regulated.”).

⁴⁴ See ACLU et al., EEOC Priorities on Automated Systems and Technological Discrimination at 7-9 (2024), [here](#) (“[T]he EEOC should issue guidance that clarifies when digital platforms and software vendors can themselves be liable for the discriminatory functioning of their tools.”); CDT et al., *supra* note 22, Civil Rights Standards.

⁴⁵ Cf. CDT et al., *supra* note 22, Civil Rights Standards at 30-31 (requiring retention of records for audits to comply with privacy laws).

⁴⁶ Sayash Kapoor et al., *On the Societal Impact of Open Foundation Models*, arXiv at 9 (2024), [here](#).

⁴⁷ *Id.*

"Open" AI Could Help Mitigate Market Concentration In Emerging AI Markets and Democratize Its Benefits, Including for Marginalized Communities.

Advanced AI models like foundation models require vast amounts of training data and high-powered compute, meaning that large existing tech companies already have a strong advantage.⁴⁸ There are already great fears that the AI revolution will further entrench existing Big Tech companies instead of advancing a more traditional cycle of disruption that comes from new challengers. Existing incumbents are increasingly investing in and licensing the technology to newer AI companies, creating concerns about potential conflicts.⁴⁹

The potential promise of “open” AI is that it may allow increased competition and customization of AI models, disrupting the potential concentration developing in the advanced AI market. However, this competition will only exist if “open” AI models are able to be hosted at the scale necessary for success. Additional competition concerns arise from the high-powered compute required to train or run advanced AI models like large language models. Because few companies can afford to build or run their own infrastructure, most AI companies are restricted to a handful of major commercial cloud computing vendors to develop or host their AI models.⁵⁰ Each of those major commercial cloud computing vendors is also developing proprietary closed AI models,⁵¹ which they will sell alongside their commercial cloud computing hosting services. Currently, the major commercial cloud computing vendors allow other AI models, including “open” AI models, to be hosted on their cloud computing services.⁵² But there is no requirement for any major commercial cloud computing vendors to allow “open” AI models to be hosted on their services, and the potential for self-preferencing may make the use of non-native AI models more difficult or expensive. Extra attention should be paid to the commercial cloud computing vendors and other infrastructure, and the ways that “open” AI may enable competition in AI markets.⁵³

In addition, the National Artificial Intelligence Research Resource Task Force emphasized the importance of public access to resources for building AI in its 2023 report. By making datasets, computing power, and educational materials openly available, researchers outside of well-funded corporations or labs gain the tools necessary to develop competitive AI solutions. This broader participation can prevent a small number of companies from controlling the field's

⁴⁸ Amba Kak, Sarah Myers West & Meredith Whittaker, *Make No Mistake—AI Is Owned by Big Tech*, MIT Technology Review (Dec. 5, 2023), [here](#); Parmy Olson, *Big Tech Has a Troubling Stranglehold on Artificial Intelligence*, Bloomberg (June 20, 2023), [here](#).

⁴⁹ *FTC Launches Inquiry into Generative AI Investments and Partnerships*, Federal Trade Commission (Jan. 25, 2024), [here](#).

⁵⁰ Amba Kak and Sarah Myers West, *AINow, 2023 Landscape: Confronting Tech Power at 6* (2023), [here](#).

⁵¹ *Azure OpenAI Service*, Microsoft, [here](#) (last visited Mar. 25, 2024); Krystal Hue, *Amazon Dedicating Team to Train Ambitious AI Model Codenamed 'Olympus'*, Reuters (Nov. 8, 2023); *Overview of Generative AI on Vertex AI*, Google, [here](#) (last visited Mar. 25, 2024).

⁵² *Amazon Bedrock*, Amazon, [here](#) (last visited Mar. 25, 2024).

⁵³ Comments of the Center for American Progress, Solicitation for Public Comments on the Business Practices of Cloud Computing Providers, Docket No. FTC-2023-0028 (June 21, 2023), [here](#).

direction and potentially lead to more innovation and accessibility in the long run.⁵⁴ As seen in other technological contexts, diffusing market concentration, especially over gateway or bottleneck facilities, can increase the diversity of voices, including for marginalized communities.⁵⁵

Conclusion

We thank the NTIA and the Administration for their continued focus on mitigating the potential civil rights harms of AI, including bolstering transparency and accountability. If you have any questions about these comments, please do not hesitate to contact us at cvenzke@aclu.org and reponsibleai@nationalfairhousing.org.

Sincerely,

American Civil Liberties Union
Center for American Progress
Leadership Conference on Civil and Human Rights
National Fair Housing Alliance

⁵⁴ As part of the National Artificial Intelligence Initiative Act of 2020, Congress established the National Artificial Intelligence Research Resource (NAIRR) Task Force to "investigate the feasibility and advisability of developing" the NAIRR as a national AI research cyberinfrastructure, and "to propose a roadmap detailing [how the NAIRR] should be established and sustained." National Artificial Intelligence Research Resource Task Force, Strengthening and Democratizing the U.S. Artificial Intelligence Innovation Ecosystem at 4 (2023), [here](#).

⁵⁵ *E.g.*, Comments of the Leadership Conference, ACLU, et al., Amendment of Section 73.3555(e) of the Commission's Rules, MB Docket No. 17-381 (Mar. 19, 2018), [here](#) ("When there are more owners, it is more likely that a woman or person of color, or a member of any other underrepresented group, can purchase a station."); Comments of the Leadership Conference, ACLU, et al, 2010 Quadrennial Review, MB Docket No. 09-182 (Mar. 5, 2012), [here](#); Comments of the ACLU, 2002 Biennial Regulatory Review, MB Docket No. 02-277 (May 22, 2003), [here](#) ("Government action should be exercised to promote greater competition and thus to encourage diversity of views. Extreme care should be taken by the Commission to see that as a practical matter, no monopoly in the presentation of news and opinion is created."); Comments of the ACLU, Framework for Broadband Internet Service, GN Docket No. 10-127 at 6 (July 15, 2010), [here](#); *cf.* Sanda Fulton, *We Still Need Diversity and Minority Ownership in Our Media*, ACLU (Jan. 19, 2012), [here](#).