1818 N Street, NW
Suite 410
Washington, DC 20036

July 7, 2023

Office of Science and Technology Policy (OSTP)
National Priorities for Artificial Intelligence Request for Information

Comments of Public Knowledge

Nicholas P. Garcia
Policy Counsel
nick@publicknowledge.org

Lisa McPherson
Senior Policy Analyst
lisa@publicknowledge.org

Meredith Filak Rose
Senior Policy Counsel
mrose@publicknowledge.org

Charlotte Slaiman
Competition Policy Director
charlotte@publicknowledge.org

I. INTRODUCTION & SUMMARY

Public Knowledge is a non-profit public interest organization that advocates for a world where people can use technology to realize their full potential, create authentic communities, and hold the powerful accountable. In our two decades of history advocating on behalf of the public, we have worked to promote freedom of expression, an open internet, and access to affordable communications tools and creative works. Drawing on our experience at the intersection of intellectual property, telecommunications, competition policy, and consumer protection, we hope to provide insight to the Office of Science and Technology Policy about how to align our national priorities for AI to serve the public interest and build a more creative and connected future for all.

Artificial intelligence has the potential to massively transform our relationship with information, communication, work, and even democracy itself. AI enabled tools can help us make sense of the massive amount of data generated by our increasingly digital world, enhance human faculties for reasoning and creativity, and fully or partially automate tasks ranging from drudgery to critical processes in many areas of work and life.

AI's greatest power, at the moment, is its potential. That potential, if mishandled, could turn destructive, or be monopolized by the powerful and deployed at the expense of the marginalized. Ensuring that AI technology truly enables broad human flourishing will require responsible, democratic stewardship, with principles rooted in long-standing values of equity and universal service. This includes addressing *both* long-term risks and known, ongoing harms; prioritizing equity and human rights at every stage of the development, deployment, and adoption process, and across the ecosystem as a whole; examining ways to make AI more equitable for creative workers *without* undermining bedrock principles of copyright law; engaging in ambitious, long-term strategic planning for AI's impacts on the workforce, including creative workers who serve as the proverbial canary in the coal mine; elevating the role of organized labor as a counterbalance against inequitable deployment of AI systems; promoting competition among AI development; ensuring universal equitable access to AI tools and technology; and establishing an AI "public option."

II. THE UNITED STATES MUST PRIORITIZE BOTH LONG-TERM PROBLEMS AND INTERVENTIONS TO ADDRESS IMMINENT AND ONGOING HARMS

The web of concerns about the rapid development of AI is both vast and complex, and spread across almost every policy field. Many of these problems predate AI, but are amplified by the scale and speed that the technology can provide. Some (but few) are genuinely new. We applaud OSTP for asking questions that acknowledge this broad range of challenges. In this moment of national attention—whether the changes continue to come fast (as they have in the past year), or whether they are still years away—we need to have long-term policy discussions without sacrificing our ability to discuss and address imminent and even ongoing harms.

The long-term vision of AI is often dominated by conversations around low-probability, high-impact events, such as the emergence of human-like artificial general intelligence (AGI), or the risk of a misanthropic (or orthogonally aligned[1]) AI superintelligence. Alongside this are debates about establishing sustainable, efficient regulation and governance. While it is true that we need to design

and establish these safeguards before the technology outstrips our ability to effectively govern it, these conversations must proceed in tandem with efforts to address existing and emerging harms.

A. AI Safety, Alignment, and the Possibility of AGI (Q1)

There is growing sentiment—fuelled by sudden leaps in generative AI capabilities, including the consensus that large language models (LLMs) exhibit unexpected emergent capabilities—that AGI, or even AI superintelligence, are real possibilities requiring consideration. The prospect of sapient machines has sparked cultural anxieties echoing those of the atomic age, as people attempt to reconcile the promise of tremendous scientific value with the specter of self-annihilation.[3] Advocates for "AI safety" and "AI alignment" frame irresponsible AI development as an extinction-level risk to humanity.

This perspective has deep and broad resonance. A recent statement on existential risk attracted hundreds of AI scientists and other notable signatories, including OpenAI CEO Sam Altman, Bill Gates, and Congressman Ted Lieu.[4] OpenAI has written about existential risk, and Sam Altman testified before Congress in support of regulation designed to set limits on these potentially catastrophic long-term risks posed by unconstrained or irresponsible AI development.[5] It has also driven the popular conception of AI, not only because of science fiction depictions of rogue AI systems (such as HAL from *2001: A Space Odyssey* or Skynet from *The Terminator*), but also because the idea's simplicity makes it easier to conceptualize than other, subtler forms of AI-driven harms.

The extreme stakes of these issues threaten to overshadow other elements of policy with regard to AI.[6] There is a need to mitigate existential risk and grapple with the society-shaking ramifications of AGI with human-like capabilities, even if these problems are low-probability. However, they cannot be the only priority. We must disentangle these high-impact/low-probability issues from ongoing, tangible, and imminent harms. In fact, dedicated government resources towards ongoing and imminent harms can support greater long term understanding of the technology as it continues to rapidly develop.

B. Expert and Adaptable Government Regulation (Q1, 2, 3, 12, 20)

Because AI technologies have enormous potential to be used for good—or to exacerbate inequality, undermine civil liberties, and supercharge information disorder[7]—oversight and accountability cannot be voluntary. Oversight of these powerful tools cannot be relegated to industry standards, corporate policies, or voluntary certifications. Government must take a leading role; we need expert regulators with the power to identify and prevent AI-based harms; provide accountability; and articulate policies and standards. The development of an innovative, trustworthy, and publicly beneficial AI sector is contingent on having well-informed and flexible regulators. To achieve this end, the United States should adopt a hybrid approach of reliance on our sector specific regulators—who are already deeply embedded in the domains that matter to us most—to avert immediate and anticipated harms, while also cultivating new expertise with a centralized AI regulator that can adapt with the technology to provide a broader view of the full ecosystem.

1. *Sector Specific Regulators*

As PK noted in a 2018 policy paper on AI regulation, AI will inevitably permeate the jurisdictions of numerous existing government agencies.[8] The twentieth century was marked less by sweeping legislative efforts than by a rise in the number and importance of executive and independent regulators. These regulators address the increasingly intricate and specialized policy challenges arising from an advanced economy. Consequently, there are already well-developed structures for tackling some of the potential risks. Existing sectoral regulators are generally already empowered to identify and deal with harms that flow from the use of AI. Whether it be issues of discrimination against marginalized communities, consumer protection, or health and safety, it must be the clear policy of the United States that the use of new technology like AI does not render an offense "novel" in a way that bypass existing safeguards and protections.

The White House has already taken strides down this road in the wake of the release of the Blueprint for an AI Bill of Rights.[9] An OSTP fact sheet described how numerous departments and agencies are applying their existing powers and expertise to the problems of their domain. For example, the Equal Employment Opportunity Commission (EEOC) and Department of Justice (DOJ) began updating guidance on anti-discrimination policies to account for algorithmic screening technologies,[10] the Consumer Financial Protection Bureau (CFPB) affirmed its longstanding rules about explanatory requirements for adverse financial decisions from creditors,[11] and the Department of Housing and Urban Development (HUD) provided guidance about how AI systems used to screen tenants may violate the Fair Housing Act.[12] The Federal Trade Commission (FTC), DOJ, CFPB, and the EEOC also issued a joint pledge to enforce their respective laws and regulations against harmful automated systems.[13] The FTC has gone even further, with Chair Lina Khan stating in a New York Times OpEd that the Commission's full competition and consumer protection authority will be used to make sure that "we continue to be the home of world-leading technology without accepting race-to-the-bottom business models and monopolistic control."[14] These actions highlight that, while the technology may be new, the potential harms or misuse cases are familiar. Discrimination, for example, is still discrimination whether it comes from an AI system or not. Placing a shiny technological veneer over prohibited practices cannot be allowed to thwart the clear jurisdiction of regulators.

This action is necessary, but not sufficient, and should not lull policymakers into complacency. While sector specific regulators are acting based on their existing authorities, they may not always have the resources they need to do their job effectively. Increased staffing, especially with allowances for more technologists, would bolster enforcement efforts. New rules and statutory authority may also be necessary to ensure that regulators can effectively address the issues at play.

This push to address new technology in familiar domains also underscores the value of enabling expert regulators to make reasoned evaluations of the pressing risks in their domain and develop an effective response. Embracing a general standard, like the "sensitive domains" standard in the Blueprint,[15] might seem more parsimonious but could lead to over- and under-inclusion that may hamper effective enforcement by overburdening a regulator unnecessarily in some areas, while inadvertently missing other contexts where the potential harm is significant but only a sector specific regulator might be aware of the issue. For example, communications networks and the paramount rights and values surrounding communications tools are critical to the healthy functioning of our democratic society. These values include free speech and free expression, secure and private

communications, digital equity, localism, access to news and information, universal service, accessibility, consumer choice and competition. Not only is this domain omitted from the Blueprint's initial list of sensitive domains while the Federal Communications Commission is holding public forums,[16] but the evolution of communications networks up the internet stack has resulted in multiple calls for an expert regulator for online digital platforms to address these key public interest protections on online platforms, including social media and messaging services.[17] The integration of AI in these Web 2.0 and 3.0 platforms is replete, and an expert regulator in this sector would inevitably involve research and/or oversight of AI and algorithmic decision making.

2. *The Need for An Expert AI Regulator*

The term "artificial intelligence" encompasses many different technologies, technical approaches, and use cases. While existing sector specific regulators are well-suited to taking on existing and emerging harms in their own domains, AI technology is too important, wide-ranging, and technically complex to rely on disjointed regulatory efforts. Therefore, Congress must either dramatically expand the jurisdiction of an existing agency, or create a new agency specifically charged to regulate AI.

PK's 2018 policy paper on AI regulation outlined five key advantages and essential features of a expert AI authority: (1) bolstering sector-specific regulators and confronting overarching policy challenges raised by AI; (2) protecting public values in government procurement and implementation of AI; (3) attracting AI practitioners to civil service, and building centralized and durable AI expertise within government; (4) identifying major gaps in the laws and regulatory frameworks that govern AI; and (5) coordinating strategies and priorities for international AI governance.[18] These features have only become more critical in the five years since the paper was published. While the paper predates the latest developments in AI technology, its discussion of the importance of the role of an expert AI regulator has only increased in relevance.

While we encourage OSTP to review the 2018 paper in full, a few key points from that paper are directly responsive to questions in the instant proceeding. Questions 24 through 28 ask about how government should approach the implementation of AI technologies for government services and missions.[19] An important function of an expert AI regulator would be to coordinate standards for government adoption and implementation of AI systems. An expert AI regulator could set, and oversee, standards for the development, acquisition, and ongoing assessment of AI systems across the government. This would create uniformity, ensure that public values are respected, and enhance public trust in government AI systems.

In the intervening five years, it has become even more apparent that an expert federal regulator with rulemaking and enforcement authority is not only necessary, but would in fact be warmly welcomed by key stakeholders including industry leaders like OpenAI,[20] Google,[21] and Microsoft.[22] Moving beyond an intra-governmental coordinating role and into proactive regulation requires both an explicit public interest mandate, and broad authority to allow for agency flexibility.

III. EQUITY, RIGHTS, AND SOCIAL JUSTICE

As AI accelerates many elements of our society, we run the risk of magnifying existing disparities and inequalities. To ensure that AI systems enable us to dismantle, rather than reinforce, systems of oppression and marginalization, they must be designed, built, and deployed with intention and oversight.

A. Mitigating and Eliminating Bias and Discrimination (Q1, 2, 10, 12)

AI systems built using machine learning techniques derive their abilities from exposure to vast repositories of (predominantly human-generated) data. This content is inevitably colored by the bias of its human sources, and too often reflects pervasive social injustices such as racial bias, misogyny, and violence. If unchecked, the biases present in these datasets are replicated and amplified within AI systems, leading to outputs that place the lives and well-being of marginalized individuals in jeopardy.[23] While there may be technical methods of minimizing these problems, the answer to AI bias is not, and cannot be, purely technological. Bias is no mere design bug; it is deeply embedded in the fabric of our society. As a result, responsible AI development requires a more systemic approach that upholds inclusion, representation, and accountability throughout the ecosystem. Affirmative engagement with historically marginalized and underrepresented people and communities is necessary to build inclusive data sets and ensure that AI systems are designed with the needs of diverse communities in mind.

While efforts to counteract bias and discrimination created or enabled by AI systems must be backed by enforceable accountability measures, there is no single intervention point for "solving" bias in AI systems. From flawed datasets to irresponsible deployment practices, combating bias in AI systems will require ongoing intervention throughout the entire ecosystem.[24]

B. Protecting Privacy (Q1, 2)

The lack of a comprehensive federal privacy law is a gaping hole in our increasingly data-driven world, and data-hungry AI systems risk accelerating the erosion of digital privacy rights. We have already seen LLMs leak personally identifiable information[25] and chat histories.[26] But those data breaches are not the only reason AI would benefit from strong privacy protections. Establishing comprehensive protections for privacy, especially related to the collection and use of personal data, would have a profound impact on the AI ecosystem with regard to the data that goes into training models, and about whom those models are able to make inferences. Giving people greater control over who has access to data about them and how it is used would ensure that the influence of AI systems over people's lives is more transparent and volitional. Restrictions on data brokers and data minimization requirements like those in the proposed American Data Privacy and Protection Act (ADPPA)[27] would provide people long-overdue protection of their privacy and autonomy in an increasingly AI-saturated future.

C. Protecting and Enhancing Democracy from Information Disorder (Q7, 15)

One of the key challenges posed to democracy by AI systems is their ability to further distort the integrity of our news environment. More (and increasingly credible) disinformation will lead to continued declines in citizens' trust in news and other democratic institutions, as well as having the

potential to create harm and spark violence. Disinformation narratives, whether of domestic origin or foreign, also prevent people—including policymakers—from solving our most pressing problems.

Chat GPT, as one example, has already been described as "the most powerful tool for spreading misinformation that has ever been on the internet."[32] OpenAI researchers have conveyed their own concerns that their systems could be misused by "malicious actors … motivated by the pursuit of monetary gain, a particular political agenda, and/or a desire to create chaos or confusion."[33] Chat GPT and other generative artificial intelligence systems can compound the challenges in our information environment in at least three ways: increasing the number of parties that can create disinformation narratives, making them less expensive to create, and making them more difficult to detect. Traditional cues that alert researchers to false information, like language and syntax issues and cultural gaffes in foreign intelligence operations, will be missing. This isn't just about AI "hallucinations" – researchers have already proven that clean, convincing news articles, essays and television scripts can be purposefully created using AI. Image generators, like Stability AI's Stable Diffusion, may undermine the classic entreaty to "believe your own eyes" to determine what is true and what is not.

Although tools are emerging to detect and authenticate whether content is created with AI, and some of the creators of AI systems offer their own, they are imperfect and may be outpaced by developments in the technology itself. It is unlikely these tools would win a technological arms race with motivated generators of disinformation. And some solutions that have been proposed to mitigate disinformation, like digital media literacy, will not be effective when the outputs are so convincing.

IV. PLANNING FOR THE AI ECONOMY

AI has the power to reshape enormous swaths of our economy. We commend OSTP for taking a wide-ranging, proactive perspective on how to address these challenges. To put it simply, we must be able to walk and chew gum at the same time. There are several areas that we think deserve particular attention as policy moves forward: copyright and the future of creative work; competition and openness; universal access to AI tools and datasets; and developing a "public option" for AI.

A. Copyright and the Future of Creative Work (Qs 12, 16,18, 19, 20, 29)

Generative AI has been a significant source of anxiety among creative professionals. Broadly speaking, AI concerns largely focus on two areas: its intersection with (and potential impact on) copyright law, and its effects on the creative labor market. Although AI raises some limited doctrinal issues around copyright law, it also represents a policy opportunity for lawmakers to structure and direct the use of these tools in the market. The most fundamental question is also the most difficult one: will AI displace creative workers? Sound policy must balance the potential upsides of these new tools against the economic precarity of existing creative workers. Policy—and organized labor—has a role to play in ensuring that creative workers are *beneficiaries* of generative AI, rather than its victims.

1. *Copyright (Q19, 29)*

The recent public fascination with AI has been driven in no small part by GAI's ability to mimic human creativity.[34] Despite the high profile intersections between AI and creative work, copyright by

itself is not a useful policy tool for addressing broader concerns about the job market and creative labor. There are, however, discrete areas where copyright policy and GAI intersect in as-yet-unsettled ways.

a. Copyright and training data

When technology impacts creative fields, the discussion almost invariably starts with copyright law. Many developers use copyrighted material in the process of training their GAI models, leading to anger at what some commentators see as copyright infringement. However, not all unauthorized use of copyrighted work is infringement; fair use provides a limit on infringement liability in many cases.[35] Most uses of copyrighted work for GAI training likely fall under the fair use umbrella.[36] However, it is worth noting that the fair use inquiry is intensely fact-specific, and a fair use finding may depend on (among other variables) the specific learning model at issue and how it was trained. Moreover, it is possible that the adoption of synthetic (as opposed to human-generated) works to train AI may eventually render this a self-limiting problem.[37]

Creators, however, should be given tools to proactively prevent their works from being included in AI training data sets. Though logistically difficult, we feel this is a ripe field for further discussion; many groups are already attempting to solve the "consent problem."[38] We would, however, caution against any system which, as a primary or second-order effect, limits other forms and purposes of web crawling or scraping, such as historical web archiving or search indexing. Any resulting system must also account for the potential privacy risks posed by some forms of metadata, as well as accommodate other potential fair uses of scraped material.

b. Copyright eligibility for AI-enabled work

Human authorship is a statutory prerequisite for copyright protection. There is broad stakeholder consensus around cases on both ends of the AI spectrum: a work made entirely by artificial intelligence would not be eligible for copyright protection, but a work made by a human author with minimal, non-substantial AI assistance would be. Most cases, however, will fall somewhere in the middle, with no clear answer in doctrine or practice. The Copyright Office has taken the position that any *elements* of a work generated by artificial intelligence are unprotected by copyright.[39] While doctrinally sound, this raises endless enforcement questions: how can the Office (or other users, who may wish to borrow or re-use unprotectable elements of a work) tell what is and is not human-authored? Should there be labeling requirements at the registration stage, or permanent watermarking requirements that signal to future users what is and is not AI generated? How would litigants prove which parts of a work are attributable to a human author, and which to an AI? What should be the role of metadata, which can potentially pose privacy risks to both creators and end users? Rather than attempt to solve this on purely doctrinal grounds, we encourage OSTP to examine the use of registrability as a policy tool. Registration of a copyrighted work provides numerous benefits to the author, chief among them the ability to bring an infringement suit in federal court. Withholding registration for AI-enabled works would be a useful tool to discourage GAI in disfavored contexts.

Nor is GAI's legal impact limited to copyright law. GAI imitation of an artist's likeness is squarely covered by "right of publicity" or name/image/likeness (NIL) laws. However, these currently exist as a patchwork of state-level regulations. A federal right of publicity law could be a strong tool to combat both economic (unfairly capitalizing on a public figure's likeness) and disinformation harms posed by

GAI. There is also growing consensus among creative stakeholders around the need for such a federal law; as such, we encourage the OSTP to further investigate this route.

Finally, it is important to note that copyright law does not protect artistic style. This is by design; protection for a "style" would prove unworkably broad, severely curtail freedom of expression, and be fundamentally un-administrable either by the courts or by the current registration apparatus. The administration should resist calls to extend copyright law to broad stylistic traits.

2. *The future of creative work (Q18, 20)*

Creative workers already face an unstable and unfavorable labor market. The market for creative work tends toward monopsony; buyers and publishers are highly concentrated, and wield significant price-setting power against a large (and largely diffuse) sea of independent suppliers.[40] In many creative industries—including fiction writing and music—these workers are considered independent contractors by default, with no ability to organize or collectively bargain. They experience higher-than-average unemployment rates and lower-than-average incomes.[41] To these workers, generative AI tools pose a significant risk of economic displacement, as certain kinds of complex creative work—such as video game asset design[42] and book cover art[43]—are being "streamlined" by AI tools in lieu of human creators.

In the short- to medium-term, generative AI provides the greatest benefit to larger, concentrated purchasers of creative work. One need look no further than the Copyright Office roundtables; while major trade organizations have offered uncharacteristically mild statements about AI and its role in the creative process, independent creators have raised the alarm about potential job displacement. In addition to enhancing productivity, AI allows major labels, publishers, and other major entertainment companies the ability to flood an already-crowded market with low-cost alternatives;[44] "sure thing" bets, such "new" music from deceased artists;[45] or cheap sequel scripts to existing franchises.[46] Any regulations designed at intervening in the market must ensure that those productivity gains are enjoyed primarily by the creative workers, rather than being used to solely enhance the balance sheets of employers, as has happened in the past.[47]

Collective labor has a significant role to play in shaping how market forces utilize AI, both in creative fields and more generally across the market. The ongoing strike by Writers' Guild of America West is an useful illustration of how workers concerned about strategic use of AI could result in a significantly worse deal for their members. By negotiating contract terms that clearly delimit the scope, scale, and context of AI usage by studios, WGA has sought to craft a solution that will allow creative workers to share in the benefits of AI, while protecting them from the worst of its potential harms.

AI is, as noted above, a wildly diverse tool with a range of applications we cannot fully imagine today. Its use and role in the economy will continue to grow and evolve. The government has a critical and central role to play; it is rarely, however, the "first to know" of new developments. Because of this, we believe that the best "first responders" to the evolving role and use of AI are the workers whose jobs are directly impacted. We believe that organization among workers should be widely encouraged not only for its own sake, but as a way to check irresponsible growth and mismanagement of a tool that has the ability to fundamentally reshape the economy.

B. Competition and Innovation (Q 29)

Safeguarding competition to preserve innovation and user choice will be essential to guard against the kind of consolidation that has plagued online platforms. Gatekeeper power allows Big Tech platforms to control the marketplace where they also compete. The adoption of a new technology can sometimes be an inflection point where there's an opportunity for new companies to gain share and unseat dominant incumbents. However, through acquisitions, partnerships, monopoly leveraging, or in some cases, healthy competition, incumbents have so far been able to weather previous inflection points like the shift to mobile. We should watch closely to ensure that today's tech incumbents don't use necessary inputs like data and computing power to pick winners and losers in AI. This might look like withholding access to computing power (at any price) unless a potential AI competitor gives an equity stake in their company to the cloud provider.

There are significant barriers to entry and incumbency advantages in AI, so we may see increasing pressures towards a consolidated market structure. As that happens, market participants will benefit greatly from non-discrimination protections like those in the American Innovation and Choice Online Act. For example, LLMs in particular may become a powerful tool for recommending products and services, and ensuring that models remain fair will be essential to protecting consumers and preserving competitive markets.

Finally, the open-source AI sector has an important role to play in ensuring vigorous competition and maintaining innovation and should be protected from efforts to enclose AI research. One of the most computing-intensive parts of the process is pre-training, when the model is created, but much of the important work happens in training, which doesn't require as much computing power. Open-source models can obviate the need for that compute-intensive pre-training stage and give potential competitors a head start, freeing them from needing special relationships with a cloud provider to access huge amounts of computing power. At least one such open-source model is already available, with more undoubtedly on the way. This may facilitate entry into the AI sector by smaller competitors and supercharge innovation and competition in AI by lowering capital costs and helping new entrants bypass existing incumbents in the cloud space.

C. Universal Access, Public Support for AI and a Public Option for AI (Q10, 12, 16, 29)

As with other transformative technologies like electricity, telephone, and (most recently) internet access, we must champion the value of universal access to AI technology. This could take many forms, depending on how AI technology continues to develop as a tool, but universal access should serve as a north star in guiding AI policy, when relevant. Like power and communications technology before it, AI technology will serve society best when it is shared broadly, empowering people and their communities in new ways.

The United States should support AI research to secure the transformative promise of AI technology and ensure that AI systems align with our values. The United States should also consider models for how it might develop a public option for AI: a publicly accessible AI foundation model funded, provisioned, and governed by a democratically accountable entity in the public interest.

The United States intends to commit considerable resources towards AI development.[48] The National Artificial Intelligence Research Resource "presents a unique and critical opportunity to 'design

in' the standards for responsible AI research practices and governance processes that uphold our priority to develop and harness these groundbreaking technologies in a manner that reinforces our Nation's democratic values and Americans' personal freedoms."[49] Acting on the NAIRR Task Force's recommendations would allow the United States to support AI research while ensuring AI development aligns with our national priorities and principles.

Policymakers should also consider how public computational resources, datasets, and expert oversight would enable not just public-private partnerships or academic research, but also present the opportunity to develop publicly owned and operated AI systems. The idea of a "public option" for AI has been promoted and discussed by various experts as an opportunity to put these powerful technologies squarely in the public sector, rather than rely on profit-driven private actors.[50] Rooted in the notion that the government could exercise stewardship over a shared data commons, or through its investment in publicly-funded computational resources, a public option for AI would provide a competitive alternative to private models.

Public foundational models or AI systems could be leveraged, not just for research, but to ensure that the public has the benefit of access to potent AI technologies and that AI is built and deployed consistent with democratic values. For example, companies developing Large Language Models (LLMs) "make decisions with huge consequences for democracy, but little democratic oversight. We don't hear about political trade-offs they are making. Do LLM-powered chatbots and search engines favor some viewpoints over others? Do they skirt controversial topics completely? Currently, we have to trust companies to tell us the truth about the trade-offs they face. A public option LLM would provide a vital independent source of information and a testing ground for technological choices with big democratic consequences."[51]

V. Conclusion

The field of AI is rapidly evolving, and holds both promise and potential risks. The decisions we make now regarding our priorities are crucial for shaping its future. We thank the Office of Science and Technology Policy for the opportunity to comment in this proceeding to help shape these priorities to place the United States on a path to a safe, responsible, diverse, and prosperous relationship with AI technology.

Endnotes

1. *See* Nick Bostrom, *The Superintelligent Will: Motivation and Instrumental Rationality in Advanced Artificial Agents*, Minds and Machines, 22(2), 71-85 (2012), available at https://nickbostrom.com/superintelligentwill.pdf (describing how goals and intelligence are orthogonally related, allowing for the possibility of AI and humans having divergent goals without a need for direct opposition).

2. Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama et al., *Emergent abilities of large language models,* arXiv preprint arXiv:2206.07682 (2022), https://arxiv.org/abs/2206.0768 .

3. Center for Humane Technology, The AI Dilemma (Mar 24, 2023), https://www.humanetech.com/podcast/the-ai-dilemma.

4. *Statement on AI Risk*, Center for AI Safety (last accessed Jul. 7,2023), https://www.safe.ai/statement-on-ai-risk.

5. Sam Altman, Planning for AGI and Beyond, OpenAI (24 Feb 2023), https://openai.com/blog/planning-for-agi-and-beyond; *Oversight of A.I.: Rules for Artificial Intelligence: Hearing Before the Subcomm. On Privacy, Technology, & the Law of the S. Comm. on the Judiciary*, 118th Cong. 1 (2023) (statement of Sam Altman, CEO, OpenAI), available at https://www.judiciary.senate.gov/imo/media/doc/2023-05-16%20-%20Bio%20&%20Testimony%20-%20Altman.pdf.

6. Natasha Tiku, *How elite schools like Stanford became fixated on the AI apocalypse*, Washington Post (5 Jul 2023), https://www.washingtonpost.com/technology/2023/07/05/ai-apocalypse-college-students/ .

7. Claire Wardle & Hossein Derakhshan, *INFORMATION DISORDER: Toward an interdisciplinary framework for research and policy making* (Sep. 27, 2017), Council of Europe Report DGI(2017)09, https://rm.coe.int/information-disorder-report-version-august-2018/16808c9c77.

8. Ryan Clough, *The Inevitability of AI Law & Policy: Preparing Government for the Era of Autonomous Machines*, Public Knowledge (Oct 2018), https://publicknowledge.org/policy/the-inevitability-of-ai-law-policy-preparing-government-for-the-era-of-autonomous-machines/ ("2018 PK AI Policy Paper").

9. *Blueprint for an AI Bill of Rights*, White House Office of Science and Technology Policy, available at https://www.whitehouse.gov/ostp/ai-bill-of-rights/ (last accessed Jun. 12, 2023) ("WH Blueprint").

10. Artificial Intelligence and Algorithmic Fairness Initiative, Equal Employment Opportunity Commission, https://www.eeoc.gov/ai (last accessed Jun. 12, 2023); U.S. EEOC and U.S. Department of Justice Warn against Disability Discrimination (May 12, 2022), Equal Employment Opportunity Commission, https://www.eeoc.gov/newsroom/us-eeoc-and-us-department-justice-warn-against-disability-discrimination.

11. Adverse action notification requirements in connection with credit decisions based on complex algorithms, Consumer Financial Protection Circular 2022-03 (26 May 2022), https://www.consumerfinance.gov/compliance/circulars/circular-2022-03-adverse-action-notification-requirements-in-connection-with-credit-decisions-based-on-complex-algorithms/.

12. Justice Department Files Statement of Interest in Fair Housing Act Case Alleging Unlawful Algorithm-Based Tenant Screening Practices, Department of Justice Office of Public Affairs (9 Jan 2023),

13. FTC Chair Khan and Officials from DOJ, CFPB and EEOC Release Joint Statement on AI, Federal Trade Commission (25 Apr 2023), https://www.ftc.gov/news-events/news/press-releases/2023/04/ftc-chair-khan-officials-doj-cfpb-eeoc-release-joint-statement-ai.

14. Lina Khan, *We Must Regulate A.I. Here's How*, NY Times (3 May 2023), https://www.nytimes.com/2023/05/03/opinion/ai-lina-khan-ftc-technology.html.

15. WH Blueprint, *Definitions*, available at https://www.whitehouse.gov/ostp/ai-bill-of-rights/definitions/ (last accessed Jun. 12, 2023).

16. *FCC & NSF to Host an Artificial Intelligence Workshop*, Federal Communications Commission (20 Jun 2023), https://www.fcc.gov/document/fcc-nsf-host-artificial-intelligence-workshop

17. *See, e.g.*, S. 1671, 118th Cong. (2023), https://www.congress.gov/bill/118th-congress/senate-bill/1671?s=3&r=11; Adam Connor & Erin Simpson, *How to Regulate Tech: A Technology Policy Framework for Online Services*, Center for American Progress (Nov. 16, 2021), https://www.americanprogress.org/article/how-to-regulate-tech-a-technology-policy-framework-for-online-services/ ; Tom Wheeler, Phil Verveer, and Gene Kimmelman, New Digital Realities; New Oversight Solutions, Shorenstein Center (Aug 2020), https://shorensteincenter.org/new-digital-realities-tom-wheeler-phil-verveer-gene-kimmelman/ ; Harold Feld, *The Case for the Digital Platform Act: Market Structure and Regulation of Digital Platforms*, Public Knowledge & Roosevelt Institute (8 May 2019), https://rooseveltinstitute.org/publications/the-case-for-the-digital-platform-act-market-structure-and-regulation-of-digital-platforms/ .

18. 2018 PK AI Policy Paper at 2,3.

19. OSTP RFI, Questions 24-28.

20. Cecilia Kang, *OpenAI's Sam Altman Urges A.I. Regulation in Senate Hearing,* NY Times (16 May 2023), https://www.nytimes.com/2023/05/16/technology/openai-altman-artificial-intelligence-regulation.html.

21. Kent Walker, *A Shared agenda for responsible AI progress*, The Keyword (28 Mar 2023), https://blog.google/technology/ai/a-shared-agenda-for-responsible-ai-progress/ ("When it comes to AI, we need both good individual practices and shared industry standards. But society needs something more: Sound government policies that promote progress while reducing risks of abuse…As we've said for years, AI is too important not to regulate — and too important not to regulate well.").

22. Brad Smith, *How do we best govern AI?*, Microsoft (May 25, 2023), https://blogs.microsoft.com/on-the-issues/2023/05/25/how-do-we-best-govern-ai/ ("There will then be a need to develop new law and regulations for highly capable AI foundation models, best implemented by a new government agency.").

23. *See, e.g., Artificial Intelligence: Societal and Ethical Implications: Hearing Before the H. Comm. on Science, Space, & Technology*, 116th Cong. 1 (2019) (statement of Joy Buolamwini, Algorithmic Justice League), available at https://www.congress.gov/116/meeting/house/109688/witnesses/HHRG-116-SY00-Wstate-BuolamwiniJ-20190626.pdf ; Leonardo Nicoletti and Dina Bass, *Humans are Biased. Generative AI is Even Worse*, Bloomberg (9 Jun 2023), https://www.bloomberg.com/graphics/2023-generative-ai-bias/.

24. *See, e.g.*, Comments of Public Knowledge, National Telecommunications and Information Administration, AI Accountability Policy Request for Comment, Docket No. 230407-0093 (12 Jun 2023) at 13-14, available at https://publicknowledge.org/policy/ntia-ai-accountability/ (PK NTIA AI Accountability Comments).

25. Jon Porter, *ChatGPT bug temporarily exposes AI chat histories to other users*, The Verge (21 Mar 2023), https://www.theverge.com/2023/3/21/23649806/chatgpt-chat-histories-bug-exposed-disabled-outage

26. Andrew Tarantola, *OpenAI says a bug leaked sensitive ChatGPT user data*, Engadget (24 Mar 2023), https://www.engadget.com/openai-says-a-bug-leaked-sensitive-chatgpt-user-data-165439848.html.

27. H.R.8152, 118th Cong. (2023), https://docs.house.gov/meetings/IF/IF00/20220720/115041/BILLS-117-8152-P000034-Amdt-1.pdf .

28. Carole-Jean Wu et al, *Sustainable AI: Environmental Implications, Challenges and Opportunities* (2021), https://arxiv.org/pdf/2111.00364.pdf.

29. See, e.g., Environmental Health Disparities and Environmental Justice, Nat'l Institutes of Environmental Health Sciences, https://www.niehs.nih.gov/research/supported/translational/justice/index.cfm (last accessed 6 Jul 2023).

30. OSTP RFI, Question 8.

31. It is also worthwhile to note that AI can be a valuable tool for helping to *combat* climate change. *See* Harold Feld, *AI Policy and the Uncanny Valley Freakout*, Public Knowledge (29 Jun 2023), https://publicknowledge.org/ai-policy-and-the-uncanny-valley-freakout/ (noting that AI can "help us identify

everything from the impact of solar fluctuations to effective techniques for managing environmental resources to offset global climate change.")

32. Tiffany Hsu and Stuart A. Thompson, *Disinformation Researchers Raise Alarms About A.I. Chatbots*, NY Times (updated 20 Jun 2023), https://www.nytimes.com/2023/02/08/technology/ai-chatbots-disinformation.html

33. Irene Solaiman et al, *Release Strategies and the Social Impacts of Language Models*, OpenAI (Nov 2019), https://arxiv.org/ftp/arxiv/papers/1908/1908.09203.pdf.

34. *See* Harold Feld, *supra* n 31.

35. For a general overview of fair use, *see* Rich Stim, *What is Fair Use?*, Stanford Libraries, https://fairuse.stanford.edu/overview/fair-use/what-is-fair-use/ (last accessed 3 Jul 2023).

36. *See* Nicholas Garcia, *Generative AI is Disruptive, But More Copyright Isn't the Answer.* Public Knowledge (11 May 2023), https://publicknowledge.org/generative-ai-is-disruptive-but-more-copyright-isnt-the-answer/ .

37. Adam Zewe, *In machine learning, synthetic data can offer real performance improvements*, MIT News (3 Nov 2022), https://news.mit.edu/2022/synthetic-data-ai-improvements-1103.

38. *See, e.g.*, Spawning AI, https://spawning.ai/ (last accessed 5 Jul 2023).

39. Copyright Registration Guidance: Works Containing Material Generated by Artificial Intelligence, 88 Fed. Reg. 16,190 (16 Mar 2023) (37 CFR pt 202), https://www.copyright.gov/ai/ai_policy_guidance.pdf.

40. *See* United States v Bertelsmann SE & Co., et al, Civil Action No. 21-2886-FYP (D.D.C. Nov 14, 2022).

41. Randy Cohen, Artists in the U.S. Workforce 2006-2020, Americans for the Arts (Mar 2021),

42. Shannon Liao, *A.I. May Help Design Your Favorite Video Game Character*, NYTimes (22 May 2023), https://www.nytimes.com/2023/05/22/arts/blizzard-diffusion-ai-video-games.html.

43. Daniel Piper, *This AI-generated book cover is causing controversy*, Creative Bloq (16 May 2023), https://www.creativebloq.com/news/ai-book-cover; Lidna Codega, *Tor Tried to Hide AI Art on a Book Cover, and It Is a Mess*, Gizmodo (16 Dec 2022), https://gizmodo.com/tor-book-ai-art-cover-christopher-paolini-fractalverse-1849904058.

44. Amanda Hoover, *AI-Generated Music Is About to Flood Streaming Platforms*, Wired (17 Apr 2023), https://www.wired.com/story/ai-generated-music-streaming-services-copyright/.

45. Mark Savage, *Sir Paul McCartney says artificial intelligence has enabled a 'final' Beatles song*, BBC (13 Jun 2023), https://www.bbc.com/news/entertainment-arts-65881813.

46. Jake Coyle, *A.I. is one of the main reasons that Hollywood writers are on strike: 'Too many people are using it against us and using it to create mediocrity'*, Fortune (5 May 2023), https://fortune.com/2023/05/05/writers-strike-hollywood-ai-scripts/.

47. *See, e.g.*, Lawrence Mishel, *Growing inequalities, reflecting growing employer power, have generated a productivity–pay gap since 1979*, Economic Policy Institute (2 Sep 2021), https://www.epi.org/blog/growing-inequalities-reflecting-growing-employer-power-have-generated-a-productivity-pay-gap-since-1979-productivity-has-grown-3-5-times-as-much-as-pay-for-the-typical-worker/; Michael Brill et al, *Understanding the labor productivity and compensation gap*, Beyond the Numbers: Productivity, vol. 6, no. 6 (U.S. Bureau of Labor Statistics, June 2017), https://www.bls.gov/opub/btn/volume-6/understanding-the-labor-productivity-and-compensation-gap.htm.

48. *See, e.g.*, *Strengthening and Democratizing the U.S. Artificial Intelligence Innovation Ecosystem: An Implementation Plan for a National Artificial Intelligence Research Resource* at ii, National Artificial Intelligence Research Resource Task Force (Jan 2023), https://www.ai.gov/wp-content/uploads/2023/01/NAIRR-TF-Final-Report-2023.pdf.

49. *Id*. at ii.

50. Ben Gansky, Michael Martin and Ganesh Sitaraman, *Artificial Intelligence Is Too Important to Leave to Google and Facebook Alone,* New York Times (10 Nov 2023), https://www.nytimes.com/2019/11/10/opinion/artificial-intelligence-facebook-google.html; Bruce Schneier, Henry Darrell, and Nathan E. Sanders, *How Artificial Intelligence Can Aid Democracy,* Slate (21 Apr 2023), https://slate.com/technology/2023/04/ai-public-option.html; Ezra Klein & Alondra Nelson, *Transcript: Ezra Klein Interviews Alondra Nelson*, NY Times (11 Apr 2023), https://www.nytimes.com/2023/04/11/podcasts/ezra-klein-podcast-transcript-alondra-nelson.html.

51. Schneier et al, *supra*.