Google

March 12, 2020

Mr. Russell Vought
Acting Director
Office of Management and Budget
Washington, DC 20503

Dear Mr. Vought,

Google welcomes the opportunity to provide comments in response to the Office of Management and Budget's *Request for Comments on a Draft Memorandum to the Heads of Executive Departments and Agencies, "Guidance for Regulation of Artificial Intelligence Applications."*[1] In particular, we support the memorandum's message that artificial intelligence (AI) and machine learning (ML) hold extraordinary potential to improve our lives and that the responsible development of AI technology should be encouraged, not discouraged. Moreover, we agree that federal agencies must avoid creating unnecessary barriers to AI development and use by taking a thoughtful, context-specific approach when regulating AI that carefully assesses risks and weighs costs and benefits.

The Promise of AI

Google's mission is to organize the world's information and make it universally accessible and useful. In pursuing this mission, Google is prioritizing investment in advanced technologies such as AI/ML. These technologies make Google's core products and services much more useful to the public, including Android, Assistant, Cloud, Gmail, Maps, Photos, Pixel, Search, YouTube, and many more.

Google is also creating tools to ensure that everyone can access AI, including researchers and developers, entrepreneurs and businesses of all sizes, academics, nonprofits, and governments. Wider accessibility is how AI will have its biggest impact and how society can reap its full promise. Critical to this approach is open-sourcing AI tools through systems such as TensorFlow,[2] a framework which makes machine learning faster, smarter, more flexible, and available to a wider community of developers. Google Cloud brings this technology to the enterprise world, offering a range of AI-powered products and solutions, from pre-built APIs for computer vision and natural language processing to end-to-end solutions that are helping to transform sectors such as financial services, retail, healthcare, and beyond.

AI is making it easier for people to go about their daily lives, from managing household and workplace tasks to breaking down language barriers, while also transforming organizations across both the private and public sectors. But the future potential is far greater. AI provides new hope for addressing the world's toughest problems, from rethinking transportation to

---

[1] Federal Register Doc. 2020-00261, https://www.whitehouse.gov/wp-content/uploads/2020/01/Draft-OMB-Memo-on-Regulation-of-AI-1-7-19.pdf.

[2] https://www.tensorflow.org/.

advancing scientific discovery to tackling complex environmental challenges. That's why we're conducting research that advances the state-of-the-art in the field, applying AI across many domains. For instance, we've recently pioneered an AI model that can help doctors spot breast cancer in mammograms with greater accuracy, and another that makes immediate, hyperlocal rainfall forecasts more accurately than existing methods to help everything from agricultural productivity to resilience against flooding.

<u>Using AI Responsibly</u>

We believe that AI will be overwhelmingly beneficial for society. But it is critical that the technology is developed and used to help people — that it is socially beneficial, fair, accountable, and works for everyone. AI must be developed and used responsibly in ways that build trust. And potential harms must be acknowledged and mitigated, or prevented.

As a leader in the development and use of AI, Google has recognized both the enormous benefits and the issues that AI raises, and has wrestled with those issues in the context of our own operations. To that end, in 2018 we established principles governing our development and use of Google AI applications, best practices to share in our work with communities outside of Google, and programs to operationalize our efforts.

Our [AI Principles](#)[3] guide the ethical development and use of AI in our research, products, and services. These guidelines help us avoid unfair bias, rigorously review for safety, design with privacy top-of-mind, and make the technology accountable to people. They also specify areas where we will not design or deploy AI, such as where human rights might be negatively impacted. But principles that simply remain on paper are meaningless. So we've developed business processes to put them into action, such as requiring vigorous testing of Google's AI decisions for fairness and conducting independent assessments of new products against our principles.

We have gone even further and are leading efforts to push for the responsible and socially beneficial use of AI applications developed outside of Google. We believe that all organizations creating AI tools should adopt guiding principles and robust internal review processes. That is why we regularly share our [recommended responsible AI practices](#)[4] with the broader AI community, other companies, nonprofit organizations, and academia. Google Cloud has also created commercial tools that enable enterprise organizations to incorporate these practices into their own AI applications. For example, we have introduced tools like [Explainable AI](#)[5] to help our customers better test and understand the outputs of their models. Additionally, we have invested in scaling frameworks like [Model Cards](#)[6] to increase transparency and comprehension around the proper use and limitations of AI models.

What's more, we've worked closely with governments and civil society stakeholders around the world to seize opportunities and address challenges related to AI. This includes stakeholders at all levels of government here in the United States. Google has provided

---

[3] https://ai.google/principles/.
[4] https://ai.google/responsibilities/responsible-ai-practices/.
[5] https://cloud.google.com/explainable-ai.
[6] https://modelcards.withgoogle.com/about.

technical expertise about AI, insight into responsible approaches to AI, and advice on how to turn AI ethical principles into practice.

The Role of Regulation

As our CEO Sundar Pichai wrote[7] recently, smart government approaches to regulation will play an important role in building trust and ensuring that AI technology is used responsibly, while also encouraging innovation. Regulatory frameworks should consider safety, explainability, fairness, and accountability to promote development of the right technologies in the right ways. Governments should also take a proportionate, risk-based approach by balancing potential harms with the social and economic benefits that will be created by AI. One of the biggest risks related to AI might be that we unnecessarily hamper its use in areas where it is urgently needed. So any regulatory framework should be flexible enough to evolve with this dynamic technology space.

Smart regulatory policy can provide broad guidance across many sectors while allowing for tailored risk-management solutions for individual AI applications in specific contexts. For some AI uses, such as AI-powered medical devices, existing regulatory frameworks are good starting points. In other instances, governments may need to update regulations or even create new regulatory frameworks. In all cases, regulators must take account of relevant costs and benefits and consider non-regulatory responses. There is no reasonable or practical "one-size-fits-all" approach.

Our Comments on the Draft Memorandum

OMB's draft principles (outlined below) represent a good foundation for advancing AI innovation while also protecting vital public interests. The principles will likewise help prevent conflicting regulatory approaches to AI from developing across the federal government. Going forward, we encourage OMB to develop supplemental guidance for regulatory agencies as experience leads to best practices on regulatory design and impact assessment, and to promote interagency discussions around key learnings. We also support efforts to ensure that agencies have adequate resources and in-house technical expertise to address AI as it continues to evolve.

1. Public Trust in AI

Google supports OMB's efforts to bolster public trust in AI. Indeed, this is a core objective that should unite technologists, businesses, policymakers, and citizens. The potential of this technology will not be fully realized if its development is held back by unfounded fears and misunderstanding of AI. Public trust is best achieved if AI technology is developed responsibly and transparently. Accordingly, our AI principles encourage architectures with privacy safeguards, sufficient transparency and accountability over the use of data, and appropriate human direction and control. Google urges agencies to contribute to public efforts to demystify AI technology and highlight its societal benefits.

---

[7] https://www.ft.com/content/3467659a-386d-11ea-ac3c-f68c10993b04.

2. Public Participation

Rulemaking related to AI technologies should include robust and fulsome opportunities for public participation and comment, which will allow diverse perspectives to be incorporated and help to avoid unintended consequences. These opportunities should begin at the earliest stages of the rulemaking and continue after the notice is published. Comment periods should be sufficiently long to allow stakeholders of all sizes to prepare useful comments, and supplemental comment should be requested where important questions are raised that warrant further consideration. Agencies should go above and beyond in terms of stakeholder outreach and engagement, including consideration of creative deep-dive formats (e.g. "citizen juries"[8]), expert roundtables, and more. To maximize opportunities for public participation, we encourage OMB to add additional language in the memo to this effect.

3. Scientific Integrity and Information Quality

Google agrees that AI policy decisions must be made based on scientifically sound analysis and the highest quality information. Technological innovation is rooted in the scientific method and a commitment to open inquiry, intellectual rigor, integrity, and collaboration. That is why our own AI principles emphasize the importance of aspiring to high standards of scientific excellence, including drawing on scientifically rigorous and multidisciplinary approaches.

For example, when a research team develops a ML model for an application, such as automated lip reading, and writes a paper documenting the model's unprecedented accuracy and the team's evaluation process, we recommend a review aligning the research with potential socially beneficial uses and the estimated scale and severity of potential harms before sharing the paper's findings via publication or presentation. In the case of a lip-reading AI application, for example, the review could determine that the model has benefits for people with hearing or speech impairments and therefore is socially beneficial; the same review could also determine a potential harm such as use of the lip-reading model for nefarious surveillance purposes. They could apply frameworks such as international human rights to assess trade-offs, identify mitigations, and determine whether the model's benefits outweigh its harms, before deciding to share the research.

Accordingly, Google encourages OMB to emphasize that AI knowledge should be shared responsibly, where practicable, through publication of educational materials, best practices, and research, as well as, when appropriate, free open-source tools that can enable more people to develop useful AI applications the right way. OMB should also urge agencies to invest in information quality for any open datasets released to the public for the purpose of training AI systems.

4. Risk Assessment and Management

Agencies must take a holistic, risk-based approach to AI. Like all other technologies, AI is not risk-free. But one of AI's greatest promises is that it will help reduce a vast array of risks inherent in everyday life, from traffic accidents and injuries sustained on the job, to life-threatening illnesses, to the secondary effects related to the emission of pollutants. In

---

[8] https://www.epa.gov/international-cooperation/public-participation-guide-citizen-juries.

assessing AI technology, reviewers should consider the risk-substitution impacts of a specific AI application against the non-use scenario. For example, commercial drones equipped with AI-powered analytics can be used to improve inspections of critical infrastructure, such as wind turbines mounted hundreds of feet in the air. Here we see the power of AI applications to advance social good — economic efficiency, environmental benefits, and safety. This is risk substitution in action, as one considers the human safety benefits of replacing human inspections with AI/drone-based inspections (see video example [here](#)[9]), as well as the opportunity costs of non-use. We urge agencies to be transparent and consistent as they weigh such considerations to ensure that innovation is not hampered by regulatory uncertainty or overreach.

5. *Benefits and Costs*

As noted above, Google agrees that cost-benefit analysis must be undertaken thoughtfully. Advances in AI will have transformative impacts in a wide range of fields, including healthcare, security, agriculture, energy, transportation, manufacturing, and entertainment. As agencies consider potential responses to AI technologies, they should be careful to take into account the broad range of social and economic factors and benefits related to those technologies, not solely a specific application's immediate impacts or hypothetical downsides. In this section of the memo, OMB should point agencies to the "Non-Regulatory Approaches to AI" section, urging them to consider novel approaches to understanding the impact of a given application before attempting to regulate it. This should include pilot programs and experiments, which will provide agencies with additional information to make an informed decision, rather than evaluating risks and costs based only on speculative *ex ante* forecasts of an application's use.

6. *Flexibility*

AI technologies will allow us to make significant — and fast — safety, efficiency, and productivity advances in industry and throughout society. Google thus agrees that any government responses must be flexible in nature, not rigid or overly prescriptive. Because AI is a continuously evolving technology space, regulations should be designed to accommodate future innovation. A challenge for regulators will be to draft definitions that are sufficiently flexible to account for this inevitable change without being so vague and overbroad as to inject unnecessary uncertainty. OMB should consider how U.S. and international standard-setting bodies can be useful partners with government to assist in designing flexible regulatory and non-regulatory approaches and supporting their evolution over time.

7. *Fairness and Non-Discrimination*

AI algorithms and datasets can reflect, reinforce and amplify, or reduce unfair biases and discrimination. We recognize that defining unfair bias is not always simple, and notions of fairness differ across cultures and societies. Fairness is often multidimensional, and optimizing for one measure of fairness may require trading off another measure of fairness.

Google's AI principles highlight the need to avoid unjust impacts on people, particularly those related to sensitive characteristics such as race, ethnicity, gender, nationality, income, sexual

---

[9] https://www.youtube.com/watch?v=lF-u7j1x0C0&feature=youtu.be.

orientation, ability, and political or religious beliefs. ML fairness is an emerging area of AI research in which we are heavily invested, and we have launched relevant open-source tools, including a What-If Tool[10] that empowers developers to visualize biases, Fairness Indicators[11] that help Cloud users check ML model performance against defined fairness metrics, and an ML Fairness Gym[12] for building model simulations that explore the potential long-run impacts of ML-based decision systems in social environments.

We applaud OMB for ensuring that this is a focus and a priority. Moreover, OMB should consider urging agencies to find ways to promote best practices in this space, rather than merely assessing AI technologies on the basis of fairness and non-discrimination criteria. In this regard, organizations in the private sector and academia can be helpful partners.

8. *Disclosure and Transparency*

Google agrees that transparency is important, as we have highlighted prominently in our own AI principles. Google has pioneered the research[13] and implementation[14] of Model Cards, a framework for providing practical information about the performance and limitations of AI models in order to help developers make better decisions about what models to use for what purpose and how to deploy them responsibly.

However, it is important for agencies to recognize that what constitutes appropriate disclosure and transparency will vary by industry, application, and even specific use cases and audiences. What's more, agencies must consider whether and how transparency features in certain circumstances conflict with privacy interests. Accordingly, we urge agencies to ensure that any government actions related to disclosure and transparency take a common-sense approach, are context-specific, and are not overly broad.

9. *Safety and Security*

We strongly believe that AI applications should be built and tested for safety. This is a foundational element of Google's AI principles, and we apply strong safety and security practices to avoid unintended results that create risks of harm. We also design our AI systems to be appropriately cautious and seek to develop them in accordance with best practices in AI safety research. In many cases, we will test AI technologies in constrained environments and monitor their operation after deployment.

We urge agencies to consult widely with the private sector as questions of safety and security are explored and to stay up-to-date on emerging practices. This is another area where novel, non-regulatory approaches to AI applications must be considered. In this section, OMB should reinforce that agencies examine public-private collaboration (e.g., hackathons, challenges, and joint experiments) to promote safety and security. It is also an area ripe for partnership with

---

[10] https://pair-code.github.io/what-if-tool/.

[11] https://ai.googleblog.com/2019/12/fairness-indicators-scalable.html.

[12] https://ai.googleblog.com/2020/02/ml-fairness-gym-tool-for-exploring-long.html.

[13] https://arxiv.org/abs/1810.03993.

[14] https://modelcards.withgoogle.com/about.

U.S. and international industry standard-setting organizations (e.g., NIST, ISO, IEEE), consistent with the principles and policies set forth in <u>OMB Circular A-119</u>.[15]

10. Interagency Coordination

Google supports robust interagency coordination to prevent piecemeal and inconsistent government actions related to the development and use of AI applications. This should include fora and processes to allow agencies to share experience, learnings, and technical and regulatory know-how and capacity. Likewise, we strongly urge OMB to prioritize *international* coordination related to AI-related regulatory and non-regulatory actions, consistent with <u>E.O. 13609</u>.[16] The United States should strive with other nations that share our democratic and societal values to align regulatory approaches so that innovation is not held back by a confusing global patchwork of variable AI standards. The <u>OECD Principles on AI</u>[17] are an example of how countries and civil society can work together to reach consensus in this space.

Non-Regulatory Approaches to AI

As noted above, Google supports the OMB memorandum's direction that agencies examine non-regulatory approaches to AI, such as voluntary consensus standards, pilot programs, sandboxes, and sector-specific policy guidance. AI is a dynamic, quickly evolving technology space that will impact every sector of society. To strike the right balance, agencies should carefully consider, in the specific context of a particular AI application, options aside from a regulatory action. In some cases, this may warrant close monitoring to see how a particular application evolves and to further evaluate its societal impacts. In all of these circumstances, various non-regulatory approaches may be valuable tools to simultaneously promote adequate protections while also preserving the benefits of open innovation. They also offer the opportunity to harness the expertise of private sector technologists and to identify previously unforeseen opportunities for driving responsible innovation.

Reducing Barriers to the Deployment and Use of AI

Google strongly encourages reduction of barriers related to the responsible deployment and use of AI. OMB's emphasis on access to federal data is important. Consistent with legal, privacy, and security considerations, we urge agencies to proactively identify new opportunities for opening datasets to the public for use in development of AI applications. It is essential that federal agencies lead by example in adopting AI solutions that generate value for society and improve the functioning of government.

Facilitating adoption of AI, however, extends beyond immediate regulatory considerations. It should be considered within the broader context of the Administration's work in areas such as IT modernization, including migration from legacy infrastructure to the cloud. Reducing

---

[15] https://www.whitehouse.gov/wp-content/uploads/2017/11/Circular-119-1.pdf.

[16] https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/inforeg/inforeg/eo_13609/eo13609_05012012.pdf.

[17] https://www.oecd.org/going-digital/ai/principles/.

barriers for AI thus entails building bridges and encouraging adoption of new technology and IT infrastructure, without which AI cannot be deployed at scale in a meaningful way.

Agency Plans to Achieve Consistency with this Memorandum

Executive Order 13859 requires that implementing agencies submit plans to OMB for achieving consistency with the memorandum. Google urges OMB to ensure that agencies conduct outside consultations on their implementation plans, including with the public, business, and all other interested stakeholders. Moreover, we request that agencies publish and seek public comment before finalizing those individual agency plans.

* * *

Google appreciates the opportunity to comment on the memorandum and welcomes any questions, feedback, or opportunities for further discussion.

Respectfully submitted,
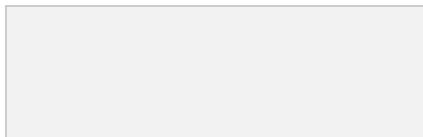
Kent Walker
Senior Vice President, Global Affairs
Google

CC:
Mr. Paul Ray
Administrator
Office of Information and Regulatory Affairs

Mr. Dominic Mancini
Deputy Administrator
Office of Information and Regulatory Affairs

Mr. Michael Kratsios
Chief Technology Officer
Office of Science and Technology Policy

Ms. Lynne Parker
Deputy Chief Technology Officer
Office of Science and Technology Policy