Google

Building a responsible regulatory framework for AI

Building a responsible regulatory framework for AI

Google seeks to use science and technology to significantly improve the lives of people around the world. Artificial Intelligence (AI) is helping us do that in exciting new ways, solving problems for our users and customers. But we know that AI presents new challenges and critical questions. We are working to develop AI responsibly[1] in order to benefit both our users and the wider societies in which we operate.

Alongside industry, academia, civil society, and others, policy makers play a critical role in providing balanced, fact-based analyses of the opportunities and challenges presented by AI, reflecting views across diverse disciplines, perspectives, and walks of life. Government engagement in promoting responsible AI growth may increase societal acceptance and public confidence in AI adoption.

We have argued that AI is too important not to regulate. Thoughtful regulation can promote reliable, robust, and trustworthy AI applications. Of course, governments should not take a one-size-fits-all approach. There is no single definition or use case for AI; rather, it includes a broad range of current and future potential technologies and applications. In fact, widely-used applications such as Gmail, Google Search, and Maps have relied on AI for years. Regulations should be carefully and proportionately tailored to mitigate risks without unduly hindering innovation and limiting AI's promise. Stated differently, smart policy initiatives can not just minimize the negatives but also maximize the positives. Indeed, in many instances, given its remarkable promise across society, our biggest risk would be not to encourage the responsible use of AI to help us meet our biggest challenges.

While AI may present some novel challenges for regulators, including building up new expertise within government, they do not need to reinvent the wheel. Instead, they can draw on the same analytical processes regularly used to address complex questions in other domains.

Focusing on applications and outcomes rather than basic research

Regulatory reviews should focus on AI-enabled applications and the quality of their specific results. Regulating the underlying computer science at too early a stage risks not realizing the many benefits offered by AI applications. There is great value in learning from doing; many of AI's benefits may be realized only through research and experimentation. While it is not feasible to monitor all aspects of computer science, looking at the results of an AI process can provide key insights into whether it is meaningfully improving on existing systems, or creating new kinds of errors and issues. Investigating the science of AI can also help validate which externalities may require regulators' attention. Additionally, many "AI issues" are actually issues common to the operation of any complex software already used by retailers, banks, insurance companies, manufacturers, and others.

Maximizing use of sectoral experts and existing regulatory frameworks

We are increasingly seeing an immense diversity of AI applications[2] across almost all sectors of society—healthcare, financial services, transportation, energy, science, retail, agriculture, logistics, manufacturing, and beyond—and it is clear that they affect different sectors differently. That is why AI regulation is often better addressed through sectoral approaches that leverage existing regulatory expertise in specific domains, rather than one-size-fits-all horizontal approaches. Governments should therefore look first to existing regulatory experts, frameworks, and instruments that may encompass AI applications. Such sectoral experts typically will be well-positioned to assess context-specific uses and impacts of AI and to determine whether and how best to regulate them, although in some circumstances additional resources may be required, including internal technical AI expert capacity. For instance, health-focused agencies should be best positioned to evaluate the use of AI in medical devices. Similarly, energy regulators would have expertise in evaluating the use of AI in energy production and distribution.

Adopting a proportionate, risk-based framework

An effective risk-based approach to regulation (1) relies on robust opportunities for public participation and comment to identify detailed risks and benefits; (2) examines the pros and cons of a proposed regulatory approach to a specific AI application; and (3) consistent with general good regulatory practices,[3] works to tailor the regulation so that it is proportionate to the use case and the risks involved. Risk assessments are an opportunity to take into account both the probability and severity of the harm, including potential impacts on vulnerable populations. AI is not risk-free, but when developed and used responsibly, it can help reduce a vast array of risks inherent in everyday life.

Notably, it may be some of the very AI applications considered "high-risk" that are also "high value" to society. For instance, AI can not only help us address enduring challenges like disease, inequality, and traffic safety, but it may also give us tools to solve new and immediate threats like climate change. In assessing AI technology, regulators should consider the risk-substitution impacts of a specific AI application against the non-use scenario—i.e., evaluate the social, safety, and economic benefits of replacing human actions with AI-based actions, as well as the opportunity costs of non-use.

Like any system, including numerous human-based processes, AI systems are not perfect. They do, however, offer the opportunity to dramatically improve on current human-based decision making. Thus, the operational benchmark for AI systems should not be perfection, but instead, the processes they are replacing.

Providing broad guidance on good AI processes, while allowing flexibility for tailoring application-specific rules

There are common underlying governance issues[4], some unique to AI, that regulators should be prepared to address in collaboration with wider civil society and AI practitioners. These include explainability standards, approaches to appraising fairness and bias, safety considerations, and transparency and accountability frameworks. At the same time, because AI is a dynamic, constantly evolving, and diverse technology, AI regulations should be flexible across application and adaptable over time. For example, using AI to identify cancerous tumors and using AI to sort photos pose very different operational and ethical questions, even though they use the same AI vision technologies. Similarly, the risks associated with using AI facial recognition technology for public surveillance are far greater than using it to unlock a phone. Flexibility can also be promoted through the use of periodic retrospective review of regulation to assess performance, as well as partnerships with expert third party organizations to evaluate whether regulation remains fit-for-purpose or requires modification.

Acknowledging that there are tradeoffs that require careful consideration

Regulators will want to consider explainability, fairness, safety, and accountability, and also the tradeoffs they often present. For example, using less identifying data protects users' privacy, but can lead to unintended consequences, such as outcomes that are biased against underrepresented populations. Similarly, restricting data sets to data that has been specifically approved by a given jurisdiction can reduce the geographical and cultural diversity of resulting AI applications. Making algorithms simpler to make them more explainable can also make them less accurate and precise. Human intervention can increase accountability or reduce the accuracy of results by introducing errors, depending on the context. Sectoral regulators evaluating individual applications and use cases through a risk-based framework will be best positioned to design the appropriate risk management solution.

Building up international standards

Internationally recognized standards can serve as the basis for robust self- and co-regulatory regimes, as guideposts for regulators, and as regulatory standards themselves if incorporated by reference. Because such standards are based on a broad and deep foundation of expertise from a wide variety of industry and civil society perspectives, they can be flexible and nimble in a way that static regulation cannot, evolving over time as the technologies innovate and change. Global regulatory buy-in can support the development and adoption of effective standards for AI.

Promoting interoperability in AI governance

Given the cross-border nature of the digital economy, global AI regulatory frameworks and technical standards need to operate across nations and regions. Increased global alignment on AI regulation, including in the context of trade, helps to facilitate the adoption, use, and interoperability of AI technologies across different jurisdictions. By taking coordinated (even if not identical) approaches, regulators can avoid adopting measures that inhibit cross-border research or disproportionately impact AI applications created in other countries. Finally, cross-border regulatory coherence and cooperation is critical to helping governments jointly develop and deploy AI to address global challenges related to public health, humanitarian assistance, climate change, and disaster response.

Considering a broad range of approaches

Policy makers should also explore more flexible opportunities to foster the responsible development and use of AI. These include pilot and other experimental programs, software sandboxes, behavioral nudges, sector-specific policy guidance and frameworks, and self- and co-regulatory frameworks. Such approaches, potentially including some degree of government oversight, can offer greater flexibility for risk management of AI applications, especially in dynamic and rapidly evolving fields.

Beyond AI governance – boosting AI adoption

Beyond regulation, other government actions will play a key role in helping us realize the many benefits that AI offers to society. For instance, governments can serve as role models for responsibly embracing AI, showcasing best practices, prioritizing opportunities to utilize AI, and demonstrating its many use cases. Policy makers can help to boost public confidence and understanding of AI by providing balanced, fact-based information about both the opportunities and the challenges presented by AI. And regulators should consider facilitating public access to and the use of government data sets, properly scrubbed and anonymized, to support the design, development, deployment, and operation of AI applications. Finally, governments should encourage increased investment in AI by actively facilitating and funding AI research and educational efforts in both the private and public domains. In short, Google believes that these foundational principles of good AI regulatory design, in combination with government leadership in AI use, will enable regulators to minimize risk to the public while at the same time advancing innovation and harnessing the substantial social, safety, and economic benefits of AI.

Appendix: Additional AI resources

- **Google: Advancing AI for everyone:** https://ai.google/
- **Google: Building Responsible AI for everyone:** https://ai.google/responsibilities/
- **Google: Responsible Development of AI:** https://ai.google/static/documents/responsible-development-of-ai.pdf
- **Google: Perspectives on Issues in AI Governance:** https://ai.google/static/documents/perspectives-on-issues-in-ai-governance.pdf
- **Artificial Intelligence at Google:** Our Principles: https://ai.google/principles/

Endnotes

1 https://ai.google/principles/

2 https://blog.google/technology/ai/

3 See, e.g., OECD Best Practice Principles for Regulatory Policy, available at https://www.oecd-ilibrary.org/governance/oecd-best-practice-principles-for-regulatory-policy_23116013

4 https://ai.google/static/documents/perspectives-on-issues-in-ai-governance.pdf

Google