

# Google's comments to NITI Aayog

## Toward Responsible #AIforAll - Working Document

Google welcomes the opportunity to provide input to NITI Aayog's deliberations around the development of Responsible AI. We appreciate the thoughtful approach laid out in the [working document](#), balancing the importance of advancing research and adoption of artificial intelligence technologies with the need to foster trust and ensure that AI is developed and used to help people.

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, including those from marginalised communities. The challenge is to ensure that any interventions in the development and use of AI are suitably tailored and appropriately balanced so as to have the desired effect with minimal unwanted repercussions.

### Executive Summary

As a leader in the development and use of AI, Google has recognized both -- the enormous benefits and the issues that AI raises. Google's [AI Principles](#), released in 2018, guide the ethical development and use of AI in its research, products, and services. Google is also investing in responsible AI research, developing new tools, techniques and practices to address key responsible AI challenges and working closely with governments and civil society stakeholders around the world in this regard.

Below are Google's key recommendations in response to NITI Aayog's working document.

- 1. Adopt proportionate, flexible and risk-based approaches to AI:** Smart 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 of AI. There is no practical "one-size-fits-all" approach to regulation, so governments must take account of relevant costs and benefits and consider non-regulatory responses as well.
- 2. Promote fairness, with particular attention to marginalised communities in India:** Investing in research to better understand Indian cultural conventions around fairness and how they manifest in technology will be essential to building robust technology for Indian users. Key sub-groups that require attention include caste, gender and gender identity,

religion, ability and class. Google is currently undertaking research on ML fairness in India to understand how safeguards to protect against counter-abuse, along with appropriate privacy, transparency and informed consent measures can help mitigate harms to marginalised communities in India.

- 3. Balance transparency requirements with practical constraints:** Transparency is essential to managing risk, empowering users and fostering trust in AI systems. Documenting the processes followed for development and training is important. However, providing explanations at scale remains a challenge because the detail of what is needed varies significantly from sector to sector and across audiences. Therefore, a sensible compromise is needed that balances the benefits of using complex AI systems against the practical constraints that different standards of explainability would impose. If it is not feasible to provide the desired level of explainability in a given instance, the same goals could be achieved by placing strict guardrails on an AI system's use.
- 4. Allow flexible self-assessment guides for AI risk assessments:** It is reasonable to expect organisations to carry out and document risk assessments prior to the launch of complex and large scale AI systems, with deeper analysis of products and services that are deemed to present a higher risk. We support the approach of documenting clear goals for fairness, accuracy, and explainability that are appropriate for a given technology and application, and ensuring they are reflected in the design and deployment of the system throughout its lifecycle. However, it is important to allow for flexibility in the self-assessment process, since the precise nature of the trade-offs varies depending on each specific case.
- 5. Multiple ethics review bodies promote flexibility and accountability:** There is no one-size-fits-all solution for understanding the ethical considerations of different AI applications, and no one ethical committee can fully understand and address these issues. Rather than a single ethics committee structure, establishing review bodies that are empowered to bring in relevant internal and external perspectives and tailor their processes for different AI applications can provide both flexibility and accountability.
- 6. Promote access to AI technologies and community-driven innovation:** It is important to make AI open and available to a wide range of stakeholders. For that reason, Google's seventh AI principle enshrines its commitment to building and sharing beneficial AI with the world and calls for open access to AI technologies that accord with the AI principles. Openness and access to AI technologies enables community-driven innovation and creative problem solving at the local level, and is a critical driver of the rapid development of AI research globally. We suggest that the government incorporate a similar principle in the development of national AI ethics principles to promote greater community participation.


7. **Encourage inter-disciplinary and participatory design processes:** The government should encourage AI developers to work directly with marginalized communities that may be impacted by their systems. Further, stakeholders across socio-economic strata and disciplines should be involved in the AI research process, including promoting greater industry-academia collaboration.
8. **Incentivise fundamental research on responsible AI:** The government should enable researchers, both globally and in India, to pursue fundamental research on responsible AI. Funding through grants and scholarships, open datasets, APIs, and benchmarks are necessary for researchers to advance basic research. The government should also convene research conferences and responsible AI workshops to help socialise India's own models of fairness into the development of global responsible AI frameworks.
9. **Develop a rich data ecosystem through voluntary data sharing:** In many cases, limited access to India-specific datasets creates challenges to developing AI tools suited for the local context. Such data sets are crucial to the development of robust AI solutions for Indians, including solutions in natural language processing and preventative healthcare. We urge policymakers to promote voluntary data sharing mechanisms with adequate privacy, intellectual property protections and business confidentiality safeguards.
10. **Promote cross-border data flows to enable global access to AI tools:** Enabling AI innovation in India and building a rich data ecosystem requires deepening India's links to the global data ecosystem. Data flows are an essential component of digital trade, economic progress, scientific research and the development of new technologies. Data flows benefit Indian companies in sectors such as healthcare, finance and manufacturing by helping them significantly reduce IT costs, leverage AI technologies and access global markets. Most importantly, data flows are empowering individuals across the world -- making information available to anyone, anywhere in the world, and helping to reduce the digital divide.

## Detailed Comments

### Google's approach to responsible AI

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](#) 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.

Our review process is an important part of putting Google's AI Principles into practice, helping to ensure they are used to guide the responsible development of Google's AI-related research, product development, and sales decisions. In addition, Google is committed to respecting the rights enumerated in the Universal Declaration of Human Rights and its implementing treaties, as well as upholding the standards established in the United Nations Guiding Principles on Business and Human Rights.

We are also investing in Responsible AI research, developing new tools, techniques and practices to address key responsible AI challenges:

- Fairness - aiming to avoid unfair bias in our AI models and products through innovations in data (e.g. [Facets](#), [Crowdsourcing](#)), model constraints (e.g. for YouTube recommendation engine - [weighting per user rather than per watch time](#)), and testing (e.g. [Profair](#))
- Explainability and accountability - ability to understand how ML models work and the data they are trained on (e.g. [Model Cards](#)) and explain decisions and predictions made by a machine learning model (e.g. [Cloud Explainable AI](#), [TensorFlow Model Analysis](#), Explanation snippets e.g. Quick Access in Google Drive)
- Privacy - ensuring state-of-the-art privacy protections with ML models (e.g. our [open-source RAPPOR technology](#), [Federated learning](#), and [Deep neural networks with differential privacy guarantees](#))

- Security - building safety and security into our ML systems. This is still a new field and the majority of work is in the research phase (e.g. [Adversarial Logit Pairing \(ALP\)](#), [CleverHans](#), development of [adversarial training](#), [Concrete problems in AI safety](#))

**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](#) 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](#) to help our customers better test and understand the outputs of their models. Additionally, we have invested in scaling frameworks like [Model Cards](#) to increase transparency and comprehension around the proper use and limitations of AI models. Importantly, we want these tools to be available to developers outside of Google as well. In June, we released [Responsible AI with TensorFlow](#) to enable third party ML developers to leverage Google's Responsible AI tools and techniques.

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 from governments around the world, including India. **Google has provided technical expertise about AI, insight into responsible approaches to AI, and advice on how to turn AI ethical principles into practice**, including sharing specific feedback with NITI Aayog in the formulation of the '[National AI Strategy](#)'. Google is also part of relevant standard setting bodies like the Bureau of Indian Standards - Artificial Intelligence Sectional Committee.

## The role of regulation

As our CEO Sundar Pichai [wrote](#) in January, 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 is that we unnecessarily hamper its use in areas where it can provide significant benefits to society. 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. As the working document notes, for many AI uses, such as AI-powered medical devices, existing legislation and

regulatory frameworks already cover potential harms of AI systems. However, in certain instances, governments will need to provide additional clarity on when and how existing regulations apply to AI applications. 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.**

## Perspectives on the AI Ecosystem in India

### ML Fairness for Indian communities

Strengthening India’s data ecosystem, in particular by building data sets for automated decision-making systems, have the potential to enable fairer and more inclusive approaches compared to ad hoc rules or human judgements. At the same time, any unfairness in such systems can have wide scale negative impact on certain communities. Therefore, it is critical that we work towards systems that are fair and inclusive for all.

India’s diverse communities by socio-economic strata require a deep engagement with various sub-groups, their lived experiences and harm mitigation strategies. At a data and model fairness level, indicators like names, locations, and photos can carry proxies and stereotypes. A model that works for a privileged caste, urban user may end up harming a Dalit, rural user because of the encoded bias in the ML models.

**AI is in emergent stages in India, and its impact is not yet fully understood.** Google has undertaken research to examine ML Fairness in India, acknowledging assumptions that might get challenged or generalize from fairness practices enacted in the west. **Key sub-groups that require attention to reduce potential harms include Caste, Gender and Gender Identity, Religion, Ability, and Class. Lack of safeguards, including counter-abuse and privacy measures, radical transparency, and informed consent in Indic languages have been key issues found to exacerbate harms of these vulnerable groups.** Datasets are central to fairness optimizations of sub-groups. But human-data practices in India like frontline workers collecting data, shared usage of devices, and about half the population lacking access to Internet-enabled devices means that datasets must be understood differently, as they may not represent people or phenomena accurately.

Ensuring fairness can be a difficult task because AI models learn from data collected from the real world, and so an accurate model may learn or even amplify problematic pre-existing biases in the data based on race, gender, religion or other characteristics. Even with the most rigorous and cross-functional training and testing, it is a challenge to ensure that a system is fair in all situations. In addition, defining unfair bias is not always simple, and notions of fairness differ across cultures

and societies. Fairness is often multi-dimensional, and optimising for one measure of fairness may require trading off another.

**Investing in research to better understand Indian cultural conventions around fairness, and how they manifest in technology, will be essential to building robust technology for Indian users.**

## Enabling a rich data ecosystem through voluntary sharing

Indian data is crucial to the development of robust AI solutions for Indians. AI could offer powerful solutions to some of India's key challenges, for example using Natural Language Processing (NLP) to address the diversity of languages across India, or using AI analytics in preventive and wellness-oriented healthcare, to counter suboptimal primary healthcare delivery in parts of the country. However, in many cases, limited access to India-specific datasets creates challenges to developing AI tools to address these challenges.

Over the years, private enterprises have voluntarily made large data sets available, which helps empower other organisations to innovate and develop programs that produce socio-economic value. For example, firms make data sets available to local governments and other partners to [inform sustainable mobility projects](#) and to help in [urban planning](#). Businesses have also made [open data sets](#) and [tools](#) available to third parties, free of charge, so they can leverage the power of emerging technologies like AI/ML and [cloud technologies](#). These tools help produce valuable insights around [image processing](#), [video annotations](#), [natural language processing](#) and [search trends](#) around the world. It is important to note that these tools and datasets have been made freely available to third parties after a careful review of the sensitivity of the data involved and to ensure conformity with intellectual property laws.

**Enabling voluntary data sharing, including broadening access to government data sets, could support indigenous innovation and AI development to address India's unique challenges.** The government should also explore building tools to enable organisations and individuals to voluntarily upload de-identified datasets. It is important that these voluntary data sharing frameworks provide adequate privacy, intellectual property protections and business confidentiality safeguards. Key privacy and data protection principles, including collection and purpose limitation should be followed -- for example, open datasets should have clear metadata, provenance and intended use (for eg. open datasets should be collected for the purpose of ML only and not re-used across survey data analytics and feature-rich ML). Clarifying policies around data collection and security, supporting proactive efforts to build rich, robust, privacy-respecting data sets to enable AI applications for Indian users, particularly users from marginalized and underrepresented communities, and fostering partnerships between the private sector, academia, civil society and the government to share data responsibly, could strengthen the AI ecosystem in India and support new, high value-add AI applications.

## Promoting cross-border data flows

Enabling AI innovation in India and building a rich, robust, and privacy-respecting data ecosystem is not just about building up data domestically, but about deepening India's links to the global data ecosystem. Data flows contribute significantly to economic growth and digital trade. A McKinsey report states that global data flows contribute \$2.8 trillion to annual trade, and contribute a larger share of the increase in global GDP compared to the global trade in physical goods. Another report by AIMA estimates that greater openness to data flows could result in a fourteen-fold increase in value to India's domestic sectors, reaching \$512 billion by 2030. A recent report by ICRIER-IAMAI estimates cross border data flows increase India's total volume of good trade by ~\$43 billion annually. **In summary, data flows are an essential component of digital trade, promotes economic progress, scientific research and the development of technologies. Data flows also benefit Indian companies in sectors such as healthcare, finance and manufacturing by helping them significantly reduce IT costs, leverage AI technologies and access global markets. Most importantly, data flows are empowering individuals across the world -- making information available to anyone, anywhere in the world, and helping to reduce the digital divide.**

On the other hand, data localisation could increase privacy and security threats, restrict digital trade and harm consumers by limiting choice, reducing universal access to information and increasing overall costs of consumption. It is concerning that India's [Personal Data Protection Bill](#) imposes strict data localisation requirements for 'sensitive personal data' (requirement to store a copy in India) and for 'critical personal data' (processing is allowed only in India).

As long as the government is able to meet its sovereign objectives through data transfer mechanisms, data localisation measures are unnecessary and disruptive to business and innovation. For instance, **global data transfer mechanisms like international certifications and bilateral agreements can ensure compliance with local laws, even when data is transferred outside the jurisdiction**, and also help address users' grievances.

## Participatory design with communities

**The government should encourage AI developers to work directly with marginalized communities that may be impacted by their systems.** Techniques from Participatory Design can help with co-designing responsible AI systems, problem statements, representative datasets, model behaviour, error handling, and explainable outputs. Communities and civil society should have a say in the systems. Research ethics and protocols may need to be developed to advise developers in conducting research with communities. Further, **stakeholders across socio-economic strata and disciplines should be involved in the AI research process, including promoting greater industry-academia collaboration.**



## Developing a Responsible AI research ecosystem

**The government should enable researchers to pursue fundamental research on responsible AI. Funding through grants and scholarships, open datasets, APIs, and benchmarks are necessary for researchers to advance basic research.** India lies at a strategic intersection of cutting edge technology and diversity of sub-groups, with great promise for responsible AI research that can transcend beyond India to other countries. India's own models of fairness, such as reservations and indigenous justice systems, should be further explored in model optimizations. India's numerous value systems can be further incorporated into fairness decision-making. Linguistic, demographic, and cultural diversity can lead to interesting technical problems, such as de-biasing models as they translate from a well covered Indic language to an endangered Indic language. Research conferences and responsible AI workshops can lead to high-quality venues for knowledge sharing on these topics.

## Specific responses to the working document

We appreciate NITI Aayog's two-pronged approach to evaluating AI's risks, looking at both direct and indirect impacts of technology at the system-level and societal level. We also believe it is important to look at solutions and risk mitigation strategies through a similar lens. In addition to direct policies to manage risks that arise in AI development and deployment, for example by clarifying how existing regulatory authorities apply to the use of AI, policymakers should also look to societal solutions to some of AI's challenges. For example, while we believe that AI will ultimately create more new jobs and industries than it displaces, **ensuring that people have access to education and new skills can help manage the impact of AI on jobs** and minimize the pain of that transition for displaced workers, without creating undue barriers to innovation and deployment of AI in traditional industries that can unlock significant gains in productivity and competitiveness.

## AI principles

NITI Aayog's principles-based approach to AI ethics, and the seven principles laid out in the working document, closely parallel our approach and many of Google's own AI Principles, particularly around risk management and human-centered development. In addition to managing risk, **our AI Principles also enshrine our commitment to building and sharing beneficial AI with the world.** Google's first AI Principle emphasizes the importance of proactively investing in developing AI that is socially beneficial. We have partnered with the Indian government to We have built ML applications to address important societal challenges in India, including [flood forecasting](#) and [diabetic retinopathy](#).

We're also collaborating with leading academic AI researchers and nonprofit organizations to address various social, humanitarian and environmental challenges with AI. Recently, we issued

grants for six research projects led by organizations from India, with expertise in areas such as computer vision, natural language processing, and other deep learning techniques. The research projects cover areas such as wildlife conservation, maternal healthcare, water management and local languages.

We also believe it is important to make AI open and available to a wide range of stakeholders. Our seventh AI principle calls for open access to AI technologies that accord with the AI principles. **Openness and access to AI technologies enables community-driven innovation and creative problem solving at the local level, and is a critical driver of the rapid development of AI research globally.** That is why our seventh AI Principle calls for making this technology as widely available as possible for beneficial uses, while managing the risk of this technology being misused. In practice, we empower developers and organisations to build and use AI using a variety of tools that are readily deployable on the cloud, including the [AI platform](#), resources available on [AI Hub](#) and specialised tools for enterprises using [TensorFlow](#).

## Transparency and Explainability

Transparency is essential to managing risk, empowering users, and fostering trust and enabling adoption of AI systems. **Documenting the processes followed for development and training is important**, and Google encourages a strong emphasis within its teams on transparency of an AI system's performance during validation. For example, providing information about how well it performs for evaluation datasets against key metrics; providing an indication of the frequency and cost weighting that were assigned to different kinds of errors (e.g. false negatives/false positives); and if relevant, how it compares to existing human-performance benchmarks are all important.

However, **it is vital that transparency requirements remain sufficiently flexible to account for a wide variety of contexts and delivery formats.** Should required documentation be too expansive it could undermine privacy or trade secrets, or increase the risk that bad actors can manipulate the system. Google strongly cautions against making it mandatory to share the precise data used, or to reveal full details about AI models or the underlying code, as that could risk undermining business confidentiality and enable adversarial gaming of the system.

AI's greatest value is seeing patterns in complex situations that are beyond human comprehension — thus (by definition) such AI systems will not be fully explainable in a way that a person can grasp. Even if the source code were shared in such a situation (an extreme form of algorithmic transparency which Google does not support) it would not help, as it would still be too complex to fathom even for experts. However, **it is a fallacy that AI systems are black boxes. With enough effort and the right tools, it is possible to get some insight into why any AI system behaves in a certain way.**

The problem is that explainability is costly, either in terms of technical resources or in terms of trade offs with other goals like model accuracy (if more accurate but harder-to-explain techniques have to be foregone). Tailoring explanations to be meaningful to a range of audiences is also difficult and time intensive. While there has been much progress in tools to support developers, such as Google's recently launched Explainable AI tool for Cloud AI customers, **providing explanations at scale remains a challenge because the detail of what is needed varies significantly from sector to sector and across audiences.**

Fortunately, just as not everyone needs to be an expert mechanic to get a driving licence and trust that a car is safe to drive, nor are explanations always necessary when using AI systems. In considering the level of explainability demanded in a specific instance, it is worth comparing the standards applied to current (non-AI) approaches. For example, an oncologist may struggle to explain the intuition that leads them to believe they fear a patient's cancer has recurred. In contrast, an AI system in the same circumstance may be able to provide biomarker levels and historical scans from 100 similar patients as a reference, even if it remains a struggle to fully grasp how the data are processed to predict an 80% chance of cancer. There is a risk that innovative uses of AI could be inadvertently precluded by demanding that AI systems meet a "gold standard" of explainability that far exceeds that required of established non-AI (including human-based) approaches. **A sensible compromise is needed that balances the benefits of using complex AI systems against the practical constraints that different standards of explainability would impose.**

Finally, it's important to acknowledge that explainability is seldom an end in itself, but rather a means of providing accountability and boosting trust. **If it is not feasible to provide the desired level of explainability in a given instance, the same goals could be achieved by placing strict guardrails on an AI system's use** — e.g., rigorous ongoing testing, or triggering human review if the probability of accuracy falls below a certain threshold, using interfaces that allow meaningful consideration of an AI system's output while mitigating the risk of confirmation bias.

## Self-assessment guide

When developed and used responsibly, AI can help reduce a vast array of risks inherent in everyday life. However it is important to be cognisant of the risks and challenges that AI itself may pose to individuals and society as a whole, including potential impacts on vulnerable populations. Thus, **prior to any launch it is reasonable to expect organisations to carry out and document risk assessments**, with deeper analysis of products and services that are deemed to present a higher risk.

The working document's self-assessment guide for responsible AI is a helpful tool for outlining key considerations for companies and organizations when assessing risk in AI products and applications.

In particular, **we support the approach of documenting clear goals for fairness, accuracy, and explainability that are appropriate for a given technology and application, and ensuring they are reflected in the design and deployment of the system throughout its lifecycle.**

**But it is important to allow for flexibility in the self-assessment process.** What constitutes responsible behavior can vary greatly across different AI applications and use contexts. While a detailed assessment process and continuous monitoring will be appropriate for sensitive or high-risk AI applications, for many systems that present few risks, for example translation services, where a less complex approach may be more appropriate.

While our assessments vary by context—AI work at Google ranges from publishing research to launching commercial APIs—we assess each project against the full set of Principles including applications we will not pursue, taking into account the scale and likelihood of potentially beneficial and/or negative outcomes. **The precise nature of the trade offs varies depending on each specific case under review, however there are almost always competing factors.**

## AI ethics committees

Incorporating diverse perspectives in AI governance is essential, and creating dedicated processes and bodies to assess risk and evaluate AI solutions is important to translate ethical principles into action. At Google, consultation is the primary mechanism used for understanding AI's risks and the tradeoffs between ethical principles, determining what mitigations are possible, and where the ultimate balance should be struck.

However, **there is no one-size-fits-all solution for understanding the ethical considerations of different AI applications**, and no one ethical committee can fully understand and address the wide array of AI systems and use cases in this dynamic field. **Rather than a single ethics committee structure, establishing review bodies that are empowered to bring in relevant internal and external expertise and perspectives and tailor their processes for different AI applications can provide both flexibility and accountability.** An ecosystem of civil society, researchers, judiciary, state, and industry is necessary for AI ethics to function. Incentives, mechanisms for audits, transparent disclosure of purpose and impacts, and funding are needed for stakeholders to audit AI systems and their impacts.

**At Google, a central review team serves as a hub for project teams seeking expert support in carrying out a formal AI Principles review.** It also plays a vital role in calibrating decisions across the company, and consolidating a library of case studies as a reference for future decision-making, from which patterns can be identified and precedents established, and which are revisited regularly as new solutions and learnings emerge.

**But some product areas have set up their own dedicated review teams to address the product-specific issues that they face.** For example, Google Cloud has developed its own governance processes to supplement the central Google governance structure to address questions that arise in third party enterprise AI deals that differ from consumer-facing Google products.

**In addition, review teams are empowered to bring in relevant domain experts in machine-learning fairness, security, privacy, human rights, and other areas to help identify issues and to carry out deep-dive specialist evaluations.** External advisors (e.g., human rights experts, ethicists, etc.) are also brought in to projects where deemed appropriate. For example, whenever relevant, we work with external experts to conduct human rights impact assessments of new products well before launch (e.g., BSR's formal review of our Celebrity Recognition tool, which informed not only the product's design but also the policies around its use). In addition, because any robust evaluation of AI needs to consider not just technical methods but also social context(s), review teams also seek a wider spectrum of perspectives, including social scientists, ethicists and Google's employee resource groups.

\*\*\*\*\*