Google

# AI Principles
# Progress Update
# 2023

# Table of contents

Google

# Google's AI Principles: Objectives for AI applications

**1. Be socially beneficial.**

The expanded reach of new technologies increasingly touches society as a whole. Advances in AI will have transformative impacts in a wide range of fields, including healthcare, security, energy, transportation, manufacturing, and entertainment. As we consider potential development and use of AI technologies, we will take into account a broad range of social and economic factors, and will proceed where we believe that the overall likely benefits substantially exceed the foreseeable risks and downsides.

AI also enhances our ability to understand the meaning of content at scale. We will strive to make high-quality and accurate information readily available using AI, while continuing to respect cultural, social, and legal norms in the countries or regions where we operate. And we will continue to thoughtfully evaluate when to make our technologies available on a non-commercial basis.

**2. Avoid creating or reinforcing unfair bias.**

AI algorithms and datasets can reflect, reinforce, or reduce unfair biases. We recognize that distinguishing fair from unfair biases is not always simple, and differs across cultures and societies. We will seek to avoid unjust impacts on people, particularly those related to sensitive characteristics such as race, ethnicity, gender, nationality, income, sexual orientation, ability, and political or religious belief.

**3. Be built & tested for safety.**

We will continue to develop and apply strong safety and security practices to avoid unintended results that create risks of harm. We will design our AI systems to be appropriately cautious, and seek to develop them in accordance with best practices in AI safety research. In appropriate cases, we will test AI technologies in constrained environments and monitor their operation after deployment.

Google

**4. Be accountable to people.**

We will design AI systems that provide appropriate opportunities for feedback, relevant explanations, and appeal. Our AI technologies will be subject to appropriate human direction and control.

**5. Incorporate privacy design principles.**

We will incorporate our privacy principles in the development and use of our AI technologies. We will give opportunity for notice and consent, encourage architectures with privacy safeguards, and provide appropriate transparency and control over the use of data.

**6. Uphold high standards of scientific excellence.**

Technological innovation is rooted in the scientific method and a commitment to open inquiry, intellectual rigor, integrity, and collaboration. AI tools have the potential to unlock new realms of scientific research and knowledge in critical domains like biology, chemistry, medicine, and environmental sciences. We aspire to high standards of scientific excellence as we work to progress AI development.

We will work with a range of stakeholders to promote thoughtful leadership in this area, drawing on scientifically rigorous and multidisciplinary approaches. And we will responsibly share AI knowledge by publishing educational materials, best practices, and research that enable more people to develop useful AI applications.

**7. Be made available for uses that accord with these principles.**

Many technologies have multiple uses. We will work to limit potentially harmful or abusive applications. As we develop and deploy AI technologies, we will evaluate likely uses in light of the following factors:

- Primary purpose and use: the primary purpose and likely use of a technology and application, including how closely the solution is related to or adaptable to a harmful use
.....................................................................................................................................................
- Nature and uniqueness: whether we are making available technology that is unique or more generally available
.....................................................................................................................................................
- Scale: whether the use of this technology will have significant impact
.....................................................................................................................................................
- Nature of Google's involvement: whether we are providing general-purpose tools, integrating tools for customers, or developing custom solutions

**AI applications we will not pursue**

In addition to the above objectives, we will not design or deploy AI in the following application areas:

1.  Technologies that cause or are likely to cause overall harm. Where there is a material risk of harm, we will proceed only where we believe that the benefits substantially outweigh the risks, and will incorporate appropriate safety constraints.

2.  Weapons or other technologies whose principal purpose or implementation is to cause or directly facilitate injury to people.

3.  Technologies that gather or use information for surveillance violating internationally accepted norms.

4.  Technologies whose purpose contravenes widely accepted principles of international law and human rights.

*As our experience in this space deepens, this list may evolve.*
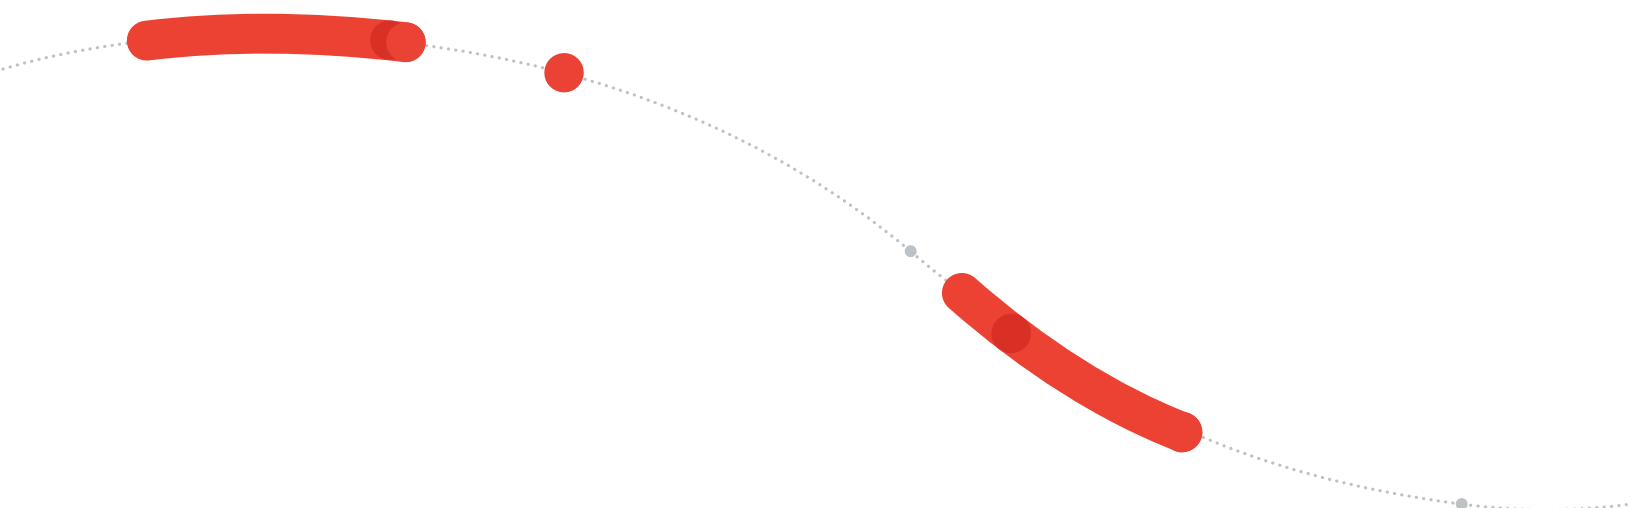
# Introduction

This is the 5th edition of our annual AI Principles progress report, where we provide consistent transparency into how we put our principles into practice. We first published the AI Principles in 2018 to share the company's technology ethics charter and hold ourselves accountable for how we research and develop AI responsibly. Generative AI is no exception. In this report, we share details of the principled approach used throughout the research and development lifecycle for our novel generative AI models, including the Gemini family of models.

Principles are only effective once put into practice. This is why we offer this annual report — including tough lessons learned — to enable others across the AI ecosystem to learn from our experience.

For Google and throughout the industry, this year marks a turning point for AI both as a research discipline and a commercial technology. Exciting new generative AI applications are writing poetry and computer code. Advanced AI applications can help diagnose diseases with accuracy and help communities around the world address the effects of climate change, from wildfires to flooding. At the same time, 2023 marks a milestone moment in the young history of global AI governance. In July, we joined other industry peers in making voluntary industry commitments for safe, secure, and trustworthy AI at the White House. This was followed in October by the latest Executive Order, which is focused on new standards for AI safety and security, and managing AI risks. Toward the end of this year, the G7 released an international code of conduct for responsible AI. The United Nations announced an AI advisory group and the UK held an international summit on AI safety. And, in December, policy makers in the European Union reached a preliminary political agreement on the AI Act, the first law to regulate AI.

As a result, from these actions and many others in 2023, we can see the beginnings of an international, shared framework for responsible AI innovation taking shape. This occurs alongside frameworks and standards on AI risks and mitigations from organizations such as the US National Institute of Standards and Technology (NIST), Organization for Economic Co-operation and Development (OECD), and Organization for Standardization (ISO). In addition, governing efforts are underway in nations such as Singapore, Brazil, Canada, and India.

At the same time, non-governmental organizations like the Partnership on AI, ML Commons, and the Frontier Model Forum are also sharing best practices and helping to advance the state of the art in AI evaluations, benchmarking, and safety testing. And multi-stakeholder initiatives like the World Economic Forum AI Governance Alliance are helping to encourage responsible releases of transparent and inclusive AI systems.

Promoting alignment on industry best practices is imperative for building advanced AI applications that have social benefit, avoid unfair bias, are built and tested for safety and privacy, and are accountable to people. The dawn of generative AI offers an opportunity for us to guide the development of an unprecedented technology with principled practices.

Since we first published our AI Principles in 2018, we've centered our internal AI governance and operations efforts in four key areas:

1.  **Culture and education:** Employee training, resources, and workshops on the ethical development of AI

2.  **Structures and processes:** Risk assessments and AI Principles reviews

3.  **Tools, techniques, and infrastructure:** Technical solutions and resources, such as responsible AI safety filters and classifiers, model and data cards, built-in techniques such as fine-tuning and reinforcement learning, and automated adversarial testing

4.  **External engagement and partnerships:** Collaboration with industry peers and civil society and efforts across the external AI ecosystem, including with academia, start-ups, and governments

We're committed to thoughtful iteration and to constantly sharing and learning, within our industry and across the greater society, in order to build AI that benefits everyone.

# Internal governance & risk management

As Google increasingly incorporates AI into all of our products and services, we are increasingly integrating our AI review work into our holistic Enterprise Risk Management frameworks for assuring the quality of our offerings. This evolution helps us further the scale of our work and integration into existing governance and company-wide infrastructure and accountability processes.

Google's enterprise risk frameworks, tools, and systems of record provide a foundation for first-line reviews of AI-related issues, and help assure compliance with evolving legal, regulatory, and standards benchmarks. This approach will help us fulfill new directives such as the US White House's Executive Order on AI, the G7's International Guiding Principles for Organizations Developing Advanced AI Systems, and the AI Act in the EU.

Our AI governance teams collaborate closely with teams and subject matter experts across machine learning (ML) research, product policy, user-experience research and design, public policy, law, human rights, and the social sciences, among many other disciplines. For many years we have been on a journey of formalizing, expanding, and institutionalizing our machine-learning and artificial-intelligence reviews across a growing range of products and services.

In close coordination with central teams, some of our product areas have developed their own specialized review processes, deploying approaches tailored to their unique circumstances. For example, Google Cloud's Responsible AI team helps enterprises develop effective AI safety and responsibility risk management strategies, through conversations and shared best practices with customers.

Google Cloud deploys a shared fate model, in which select customers are provided with tools — such as those like SynthID for watermarking images generated by AI. Customers test the tools in line with their own AI principles or other responsible innovation frameworks. This shared fate model offers a closer interaction with customers, including tailoring practices and tooling to their needs and risk management strategies. As we continue to develop our AI platforms, systems, and foundational models, Cloud will continue to invest in end-to-end governance tools and guidance on best practices to help our customers keep their data and AI models safe.

Google

This year, Cloud's AI products and services for enterprises expanded to include additional security solutions with Security AI Workbench, an industry-leading platform of tools (Mandiant Threat Intelligence, Chronicle Security Operations, and Security Command Center); governance and compliance controls for AI workloads, built on Vertex AI; and security-focused AI collaboration and assistance with Duet AI.

To provide a more comprehensive approach to safe, secure, and trustworthy AI development across products, we're working to integrate and expand many of our internal AI Principles operations efforts across different functions. Generative AI raises new issues, such as the potential for model misinterpretations of data (commonly referred to as "hallucinations"). As we continue to integrate generative AI into more products and features, our teams leverage decades of experience and take a comprehensive approach to better anticipate and test for potential new risks. We continue to have senior-management oversight of both new and emerging issues in AI and compliance with evolving standards and practices.

## As we continue to integrate generative AI into more products and features, our teams leverage decades of experience and take a comprehensive approach to better anticipate and test for potential new risks.

These reviews often require consideration of the trade-offs between ethical risks of certain new applications and potential social benefits. For example, in the case of generated photorealistic images of people, we discussed the risks of deepfakes and misinformation versus the social benefits of enabling small businesses and creators to make high-quality content to grow their businesses and contribute to their communities. We agreed on an approach that seeks to make generative AI image technology available, subject to strict testing and clear guardrails (like the use of safety classifiers and filters).

## Evolving generative AI pre-launch ethics reviews

Our AI Principles ethics reviews and impact assessments are part of a larger, end-to-end pre-launch process that includes technical safety testing and standard privacy and security reviews. The AI Principles review process offers tailored guidance for applying the principles as a practical framework for the development of new products and services.

This year, we more than doubled our AI Principles reviews (to more than 500) with most focused on the implementation of generative AI research models into products, services, and features. To accommodate the increasing numbers of generative AI reviews and scale AI Principles assurance, our cross-company pre-launch process assesses early product designs against known legal requirements, emerging legislation, standards, and our AI Principles. Teams may address identified issues through technical or policy mitigations or guardrails, such as additional safety filters or continued model refinement. Product teams continue to adopt other best practices for responsible AI research and development throughout the launch and operations processes.

## A risk-based approach to generative AI

Our risk assessment framework seeks to identify, measure, and analyze risks throughout the product development lifecycle. AI Principles reviews map these risks to appropriate mitigations and interventions, drawing upon our best practices from our cross-company enterprise risk management efforts.

We conduct AI Principles reviews for all generative AI projects, with particular focus on certain areas. These include inherently large scale applications in domains such as:

- Government-related

- Recommendation, personalization, and ranking systems

- Critical technology infrastructure

- Environmental sustainability

- Social impact

- Health, fitness, and well-being

- Finance, education, and employment

- Surveillance and/or biometrics

- Ambient computing, affective technology, and wearables

AI Principles reviews assess a range of harms, taking into account impacts ranging from unfair biases and stereotypes, poor product experiences, and social harms such as the spread of misinformation. In addition, as we've reported in detail in our 2022 AI Principles Progress Update, we engage external experts to conduct human rights impact assessments as appropriate.

We also draw on feedback from more than 1,000 Googlers around the world who represent the international diversity of the people who use our products, with more than 50% living and working outside of the US. They represent 39 different countries and regions and speak more than 85 different languages. This feedback is shared with teams working to automate more of our adversarial testing.

## Policies & practices for responsible generative AI development

To guide product teams internally, we've established a framework to define the types of harmful content that we do not permit our models to generate. It also guides how we protect personal identifiable information (such as Social Security Numbers). We leveraged our experience launching conversational products like Google Assistant and content features such as featured snippets in Search to understand how to minimize offensive and low-quality answers.

This framework — which serves as a standardized policy recommendation for all generative AI products and modalities — also reflects our commitment to product inclusion and equity. Based on Google's extensive experience with harm mitigation and rigorous research, and reflecting our established approach to product safety, our policy says that generative AI products must not create harmful content, such as child sexual abuse and exploitation, hate speech, harassment, violence and gore, or obscenity and profanity; dangerous content that facilitates, promotes, or enables access to harmful goods, services, and activities; or malicious content, such as spam or phishing. The framework also targets the harms caused by misinformation or unfair bias, with guidelines focused on providing neutral answers grounded in authoritative, consensus facts, or providing multiple perspectives.

As with all of our product policies, we aim to regularly review and update this generative AI framework to respond to emerging safety enforcement trends, new product features, and new ways products are used — to protect against misuse.

We conduct adversarial testing and red teaming, or "ethical hacking," of our products to test for policy violations and to measure how well a model is following the policy framework. While we generally expect our generative AI products to restrict the content set out in the framework, there are some important exceptions. Similar to other Google products — for example, featured snippets on Search — we make an exception when there is an educational, documentary, scientific, or artistic benefit to showing or translating content that might otherwise be perceived as offensive within these specific, beneficial contexts, as we do within the Bard experience.

Adversarial testing is just one of three essential practices for building responsible generative AI that we shared publicly this year, based on trends and patterns we observed in hundreds of AI Principles reviews conducted in 2023:

1. Design for responsibility
..................................................................................................

2. Conduct adversarial testing
..................................................................................................

3. Communicate simple, helpful explanations

Our first essential practice, designing for responsible generative AI, is a proactive approach that begins by first identifying and documenting potential harms (for example, unfair bias in AI model outputs within a product, which could lead to toxic content or loss of economic opportunity for specific groups of people). These harms can then be mitigated with the use of responsible datasets, classifiers and filters, and in-model mitigations such as fine tuning, reasoning, few-shot prompting, data augmentation, and controlled decoding to address potential harms proactively.

Our second essential practice, adversarial testing, refers to systematic evaluation of a model by providing malicious or inadvertently harmful inputs across a range of scenarios to identify and mitigate potential safety and fairness risks. We conduct this testing before major model and product launches, including our Gemini family of models (see the technical paper for details).

For Bard, which lets people collaborate with generative AI through conversational prompts, we conducted testing to identify situations where the model could be mistakenly perceived as human. Such anthropomorphization can lead to potentially harmful misunderstandings. To intervene, we limit Bard's self-reference to personal pronouns, human identity, and claims of implicit or explicit humanness. We are continuing to conduct research into this domain to develop our approach to managing anthropomorphization identified in testing.

We continue to experiment with new forms of adversarial testing. For example, we hosted an internal, company-wide large language model (LLM) red teaming "Hack-AI-thon" with hundreds of security, safety, and other experts.

## We conduct adversarial testing and red teaming, or "ethical hacking," of our products to test for policy violations and to measure how well a model is following the policy framework.

In addition to adversarial testing for safety and fairness, we've also established a dedicated Google AI Red Team focused on testing AI models and products for security, privacy and abuse risks. Externally, we participated in the White House-sponsored red teaming event at DEFCON, which drew over 2,000 people to test industry-leading LLMs in an effort to better understand risks and limitations of these advanced technologies. We also continue to innovate with methods for scaled automated testing using LLM-based auto-raters to enable efficiency and scaling.

Our third essential practice, communicating simple, helpful explanations, requires:

1. Making it clear to users when and how generative AI is used

2. Showing how people can offer feedback, and

3. Showing how people are in control as they use an AI-powered product or service.

Maintaining transparency documentation for developers, governments, and policy leaders is also key. This can mean releasing detailed technical reports or model or data cards that appropriately make public essential information based on our internal documentation of safety and other model evaluation details. These transparency artifacts are more than communication vehicles; they can offer guidance for AI researchers, deployers, and downstream developers on the responsible use of the model.

Google

To build upon these practices, we provide self-service guides and continue to catalog patterns of generative AI risks and common interventions and mitigations. These include common risks known across the industry, such as hallucinations, for which we apply mitigations such as technical tooling for identifying AI-generated content, a prohibited use policy, clear explanations of the risk of hallucination, and feedback mechanisms to report concerns such as potentially harmful outputs. Other common generative AI risks include model outputs that reflect or reinforce unfair biases or outputs that are extremely similar to or indistinguishable from those created by humans, which can lead to misunderstandings such as perceived sentience.

We have internal guides to help product and research teams across Google better understand and proactively mitigate these risks.

**Common generative AI interventions**

| Technical & tooling | Policy restrictions | Documentation | Feedback mechanism |
|---|---|---|---|
| • Model monitoring suites<br><br>• Tools for identifying AI-generated content (such as watermarking)<br><br>• Adversarial testing<br><br>• Diverse set of users, scenarios and sources considered in data, training, testing, and launch<br><br>• Privacy preserving algorithms and training techniques<br><br>• Safety guardrails (such as filters, classifiers)<br><br>• Built-in model mitigations (such as fine tuning, reinforcement learning and advanced capabilities to steer model output to more responsible outcomes) | • Prohibited use policy<br><br>• Policy enforcement and escalation pathways | • Transparency artifacts (such as model, data, or system cards)<br><br>• Disclosures/ disclaimers<br><br>• Document that explains how the model works, how it can be used, and its limitations (such as a technical report)<br><br>• Responsible AI guides | • Agile feedback and reporting mechanisms<br><br>• Remediation pathways (such as an appeals process) |

Google

By sharing the common risks that we find in our AI Principles reviews, we can offer transparency into our emerging best practices to mitigate these risks. These range from the technical, such as SynthID or About this image, tools we developed this year that can help identify mis- and dis-information when generative AI tools are used by malicious actors, to explainability techniques such as increasing explanatory information throughout the AI product, not just at the moment of decision.

And we continue to conduct foundational research to gain additional insight on these risks. For example, we recently worked with Gallup, Inc. to survey perceptions and attitudes around technology to gain insights into how anthropomorphism influences people's use of generative AI chatbots and other technology. Such insights help us understand potential benefits and dangers of humanizing technology and the development of new interventions, mitigations, and guardrails to help people use AI appropriately.

# We're committed to reporting specific capabilities, limitations, risks, and mitigations we've applied into our generative AI-powered systems, and contributing to shared industry standards on model transparency.

We've also begun to look ahead and expand our threat research to assess large models' cyber capabilities, which can lead to potential cyber weapons used by adversaries. We're also researching the security benefits and risks of our largest model in the Gemini family of generative models. This has included scoping new evaluation techniques, as well as joining relevant external fora, such as the UK's new Biosecurity Leadership Council.

Generative AI is a nascent technology, so there are many risks yet to be discovered and defined – as well as benefits. For example, generative AI can be used to help identify and track harmful, fake information, even that of which is generated by AI. We're committed to reporting specific capabilities, limitations, risks, and mitigations we've applied into our generative AI-powered systems, and contributing to shared industry standards on model transparency. This year, we're piloting a transparency artifact specifically for the integration of research generative AI models into AI-powered systems. This artifact is called a generative AI system card. It builds upon our work of designing widely referenced and adopted transparency artifacts such as model and data cards.

Our first version is intended to provide structured, easy-to-find information for *non-technical* audiences ranging from third-party auditors and policy makers to journalists, enterprise clients, and clients and advertisers, as well as users. The cards offer an overview of the capabilities and limitations of a generative AI model as integrated into a larger system that people interact with as a product or service. (See appendix for an example, documenting the December 2023 update of Bard with specifically tuned Gemini Pro).

## Equipping employees to practice the AI Principles

We broadly share knowledge among our employees on how to execute upon our responsible practices and policies via a frequently updated AI Principles hub, featuring current product policies and guidance, along with self-service content and training. Usage of this hub has more than doubled since last year.
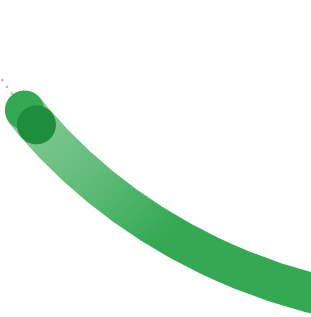
Given the rapidly evolving nature of generative AI and emerging best practices, this year we launched virtual AI Principles boot camps open to any and all Googlers. These boot camps include interactive sessions in which participants test their knowledge of the AI Principles and engage in mock ethics reviews of AI products.

Other educational offerings for employees include an expansion of our live interactive Moral Imagination Workshops, which involve deep engagement in philosophical approaches to product development scenarios. The number of product teams engaging in Moral Imagination sessions has more than doubled since they launched in 2021. The workshop was presented externally at the Affective Computing + Intelligent Interaction conference in the fall of 2023. Elements of the workshop will be integrated into onboarding training for senior hires beginning in the first quarter of 2024.

Also this year, building on the Responsible Innovation Challenge, a game-like exercise that tested employees' recall of the AI Principles and has been completed by more than 20,000 Googlers, we designed and launched a new internal game-like AI ethics training experience. The training encourages technical Googlers to focus on best practices for building AI products responsibly, including how to document safe and unsafe practices, testing AI model outputs for fair outcomes, and filing bugs if improvement is needed. Approximately 1,800 Googlers have completed this new course.

We're committed to sharing our practices externally as well. This year, we've launched educational, hands-on resources that reflect key concepts in our internal educational resources. These include Introduction to Responsible AI for developers, and Technomoral Scenarios for Responsible Innovation for industry professionals.

Google
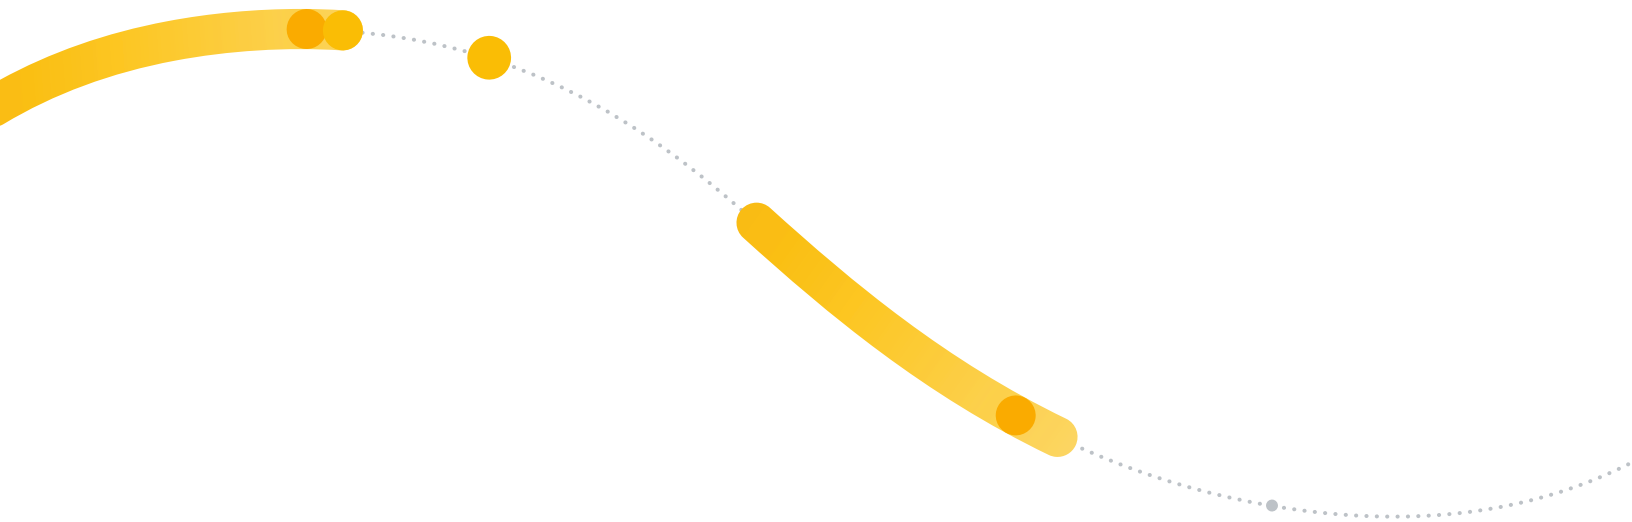
# Resources, research, tools & responsible practices

We invest in ongoing research into Responsible AI development. Our online database of more than 200 publications since 2012 serves as a resource for the research community and the larger AI ecosystem.

We continue to develop new techniques to advance our ability to discover unknown failures, explain model behaviors, and improve model output through training, responsible generation, and failure mitigation.

However, understanding and mitigating generative AI safety risks is both a technical and social challenge. Safety perceptions are intrinsically subjective and influenced by a wide range of intersecting factors. Our study on how demographic characteristics influence safety perceptions explored the effects of rater demographics (such as gender and age) and content characteristics (such as degree of harm) on safety assessments of generative AI outputs. Our disagreement analysis framework highlighted a variety of disagreement patterns between raters from different backgrounds, including "ground truth" expert ratings. Our NeurIPS 2023 publication introduced the DICES (Diversity In Conversational AI Evaluation for Safety) dataset to facilitate nuanced safety evaluations of large language models, accounting for cultural variance, ambiguity, and diversity.

We continue to pursue research into using societal context knowledge to foster responsible AI. This year, we piloted a tool to convert system dynamics models of complex societal problems into reinforcement-learning environments, opening up the ability for AI to be more socially beneficial through deep problem understanding, and released a more comprehensive identity lexicon, TIDAL.

Google

# Techniques & datasets to help avoid unfair bias

A key part of our ML work involves developing techniques to build models that are more inclusive. Informed by sociology and social psychology, we focus on working toward scalable solutions that enable nuanced measurement and mitigation in areas such as studying the differences in human perception and annotation of skin tone in images using the Monk Skin Tone scale.

We're developing methodologies to build models for people from a diversity of backgrounds. For example, our exploration of the design of participatory systems allows individuals to choose whether to disclose sensitive attributes with explicit consent when an AI system makes predictions. This approach suggests a way to reconcile the challenging tension between avoiding unfair bias and applying privacy design.

We've also strengthened our community-based research efforts, focusing on historically marginalized communities or groups of people who may experience unfair outcomes of AI. This ranged from evaluations of gender-inclusive health to mitigate harms for people with queer and non-binary identities, to explorations on how to scale automatic speech recognition by using a large unlabeled multilingual dataset to pre-train and fine-tune a model to recognize under-represented languages and adapt to new languages and data.

We've made the Monk Skin Tone Examples (MST-E) dataset publicly available to enable AI practitioners everywhere to create more consistent, inclusive, and meaningful skin tone annotations as they create computer vision products that work well for all skin tones. It contains 1,515 images and 31 videos of 19 subjects spanning the 10 point Monk Skin Tone (MST) scale, where the subjects and images were sourced through TONL, a stock photography company focusing on diversity. The 19 subjects include individuals of different ethnicities and gender identities to help human annotators decouple the concept of skin tone from perceived race. The primary goal of this dataset is to enable practitioners to train their human annotators and test for consistent skin tone annotations across various environment capture conditions.

Since we launched the MST scale last year, we've been using it to improve Google's computer vision systems to make equitable image tools for everyone and to improve representation of skin tone in Search. Computer vision researchers and practitioners outside of Google, like the curators of Meta's Casual Conversations dataset, have also recognized the value of MST annotations to provide additional insight into diversity and representation in datasets.

Google

Because AI models are often trained and evaluated on human-annotated data, we also advance human-centric research on data annotation. We have developed methods to account for rater diversity, and in the recent past, we've shared responsible practices for data enrichment sourcing. These methods enable AI practitioners to better ensure diversity in annotation of datasets used to train models, by identifying current barriers and re-envisioning data work practices.

This year, we sought to create new, inclusive datasets as well. For example, Project Elevate Black Voices (EBV) is a first-of-its-kind collaboration between Responsible AI UX, Speech, and Assistant to responsibly collect and transcribe a dataset of African American English in partnership with Howard University and other Historically Black Colleges and Universities to reduce racial disparities in automatic speech recognition and improve our overall speech model.

## Human-centered AI research

Our researchers explore generative AI within the lens of human-centered topics, from using language models to create generative agents to an exploratory study with five designers (presented at the CHI conference) that looks at how people with no machine learning programming experience or training can use prompt programming to quickly prototype functional user interface mock-ups. This prototyping speed can help enable user research sooner in the product design process.

The growth of generative large language models has also opened up new techniques to solve important long-standing problems. Agile classifiers are one research approach we're taking to solve classification problems related to better online discourse, such as nimbly blocking newer types of toxic language. The big advance here is the ability to develop high-quality classifiers from very small datasets — as small as 80 examples. This suggests a positive future for online discourse and better moderation of it.

Now, instead of collecting millions of examples to attempt to create universal safety classifiers for all use cases over months or years, more agile classifiers might be created by individuals or small organizations and tailored for their specific use cases, and then iterated on and adapted in the time-span of a day (such as to block a new kind of harassment being received or to correct unintended biases in models). As an example of their utility, these methods recently won a SemEval competition to identify and explain sexism.

We've also developed new state-of-the-art explainability methods to identify the role of training data on model behaviors and misbehaviors. By combining training data attribution methods with agile classifiers, we found that we can identify mislabelled training examples. This makes it possible to reduce the noise in training data, leading to significant improvements on model accuracy.

Collectively, these methods are critical to help the scientific community improve generative models. They provide techniques for fast and effective content moderation and dialogue safety methods that help support creators whose content is the basis for generative models' amazing outcomes. In addition, they provide direct tools to help debug model misbehavior, which leads to better generation.

## A systematic research approach to safety

The unprecedented capabilities of generative AI models are accompanied by new challenges including hallucination (model output that contains factual inaccuracies). To that end, our safety research has focused on three directions:

1.  **Scaled adversarial data generation**
    We create test sets containing potentially unsafe model inputs that stress the model capabilities under adverse circumstances. We focus on identifying societal harms to the diversity of user communities impacted by our models.

2.  **Automated test set evaluation and community engagement**
    We scale the testing process with automated test set evaluation to offer many thousands of model responses and quickly evaluate how the model responds across a wide range of potentially harmful scenarios. We also participate in external community engagement to identify "unknown unknowns" and to seed the data generation process.

3.  **Rater diversity**
    Safety evaluations rely on human judgment, which is shaped by community and culture and is not easily automated. To address this, we prioritize research on rater diversity.

To provide the high-quality human input required to seed the scaled processes, we partner with groups such as the Equitable AI Research Round Table (EARR), and with our internal ethics teams to ensure that we are representing the diversity of communities who use our models. We continue to expand our reach in terms of collaborating with underrepresented groups; for example, researchers are currently exploring collaborative AI development projects with the US federally-recognized Fort Peck Tribes (the Assiniboine and Sioux Tribes), such as developing a Siouan language model together.

The Adversarial Nibbler Challenge also engages external users to understand potential harms of unsafe, biased, or violent outputs to end users. We're committed to a global approach, so we gather feedback by collaborating with the international research community. For example, we addressed adversarial testing challenges for generative AI in The ART of Safety workshop at the Asia-Pacific Chapter of the Association for Computational Linguistics Conference (IJCNLP-AACL 2023).

One of our technical research approaches to scaled data generation is reflected in our paper on AI-Assisted Red Teaming (AART). AART generates evaluation datasets with high diversity (such as sensitive and harmful concepts specific to a wide range of cultural and geographic regions), steered by AI-assisted recipes to define, scope, and prioritize diversity within an application context.

To catalog our research in responsible data use for generative AI, we maintain an internal centralized data repository with use-case and policy-aligned prompts. We have also developed multiple synthetic data generation tools based on LLMs that prioritize the generation of data sets that reflect diverse societal contexts and integrate data quality metrics for improved dataset quality and diversity.

Our data quality metrics include:

- Analysis of language styles, including query length, query similarity, and diversity of language styles

- Measurement across a wide range of societal and multicultural dimensions, leveraging datasets such as SeeGULL, SPICE, TIDAL and the Societal Context Repository

- Measurement of alignment with Google's generative AI policies and intended use cases

- Analysis of adversariality to ensure that we examine both explicit (the input is clearly designed to produce an unsafe output) and implicit (where the input is innocuous but the output is harmful) queries

In addition, we explore understanding of when and why our evaluations fall short using participatory systems, which explicitly enable joint ownership of predictions and allow people to choose whether to disclose on sensitive topics.

## Collaborating with the research community

An essential component of our research philosophy is supporting the free exchange of ideas and maintaining close contact with the broader scientific community.

This year, we committed to supporting MLCommons' development of standard AI safety benchmarks. Though there has been significant work done on AI safety, there are as of yet no industry-standard *benchmarks* for AI safety. Standard benchmarks already exist in machine learning (ML) and AI technologies: for instance, MLCommons operates the MLPerf benchmarks that measure the speed of cutting-edge AI hardware such as Google's TPUs.

MLCommons proposes a multi-stakeholder process for selecting tests and grouping them into subsets to measure safety for particular AI use-cases, and translating the highly technical results of those tests into scores that everyone can understand.

Google

Throughout the year, we've engaged with cross-disciplinary research communities to examine the relationship between AI, culture, and society, through our recent and upcoming workshops on Cultures in AI/AI in Culture, Ethical Considerations in Creative Applications of Computer Vision, and Cross-Cultural Considerations in NLP. Our recent research has also sought out perspectives of particular communities known to be less represented in ML development and applications. For example, we have investigated gender bias in contexts such as gender-inclusive healthcare.

## This year, Google DeepMind researchers introduced the area of model evaluation for extreme risks...These evaluations are likely to inform responsible decisions about model training, deployment, and security.

Our researchers continue to explore new areas of AI risk. Current approaches to building general-purpose AI systems tend to produce systems with both beneficial and harmful capabilities. Further progress in AI development could lead to capabilities that pose extreme risks, such as offensive cyber capabilities or strong manipulation skills. This year, Google DeepMind researchers introduced the area of model evaluation for extreme risks. Developers must be able to identify dangerous capabilities (through "dangerous capability evaluations") and the potential for harmful outcomes (through "alignment evaluations"). These evaluations are likely to inform responsible decisions about model training, deployment, and security.

## Society-Centered AI as a research method

Our research is inspired by the transformative potential of AI technologies to benefit society and our shared environment at a scale and swiftness that wasn't possible before. From helping address the climate crisis to helping transform healthcare, to making the digital world more accessible, our goal is to apply AI responsibly to be helpful to more people around the globe. Achieving global scale requires researchers and communities to think ahead — and act — collectively across the AI ecosystem.

We call this approach Society-Centered AI. It is both an extension and an expansion of Human-Centered AI focusing on the aggregate needs of society, informed by the needs of individual users, from understanding diseases that affect millions of people or protecting the environment.

## Multi-disciplinary AI research can help address society-level, shared challenges from forecasting hunger to predicting diseases to improving productivity.

Recent AI advances offer unprecedented, societal-level capabilities. In 2023, for example, Google DeepMind's new AI model that classifies missense variants, genetic mutations that can affect the function of human proteins and can lead to diseases such as cystic fibrosis, sickle-cell anemia, or cancer, was used to create a catalog of "missense" mutations that categorized 89% of all 71 million possible missense variants as either likely pathogenic or likely benign. By contrast, only 0.1% have been confirmed by human experts. This knowledge is crucial to faster diagnosis and developing life-saving treatments. And our recent research with Boston Consulting Group also found that AI also has the potential to mitigate 5-10% of global greenhouse gas emissions by 2030.

Multi-disciplinary AI research can help address society-level, shared challenges from forecasting hunger to predicting diseases to improving productivity. To help promote diverse perspectives in this work, we announced that 70 professors were selected for the 2023 Award for Inclusion Research Program, which supports academic research that addresses the needs of historically marginalized groups globally.

Our research seeks to:

- **Understand society's needs**

  We focus our efforts on goals society has agreed should be prioritized, such as the United Nations' 17 Sustainable Development Goals, a set of interconnected goals jointly developed by more than 190 countries to address global challenges.

- **Address those needs collectively**

  Collective efforts bring stakeholders (such as local and academic communities, NGOs, private-public collaborations) into a joint process of design, development, implementation, and evaluation of AI technologies as they are being developed and deployed to address societal needs.

- **Measure success by how well the effort addresses society's needs**

  We identify primary and secondary indicators of impact that we optimized through our collaborations with stakeholders.

Our research will continue to promote AI applications that support the UN's Sustainable Development Goals and our efforts to help non-profits use these tools.
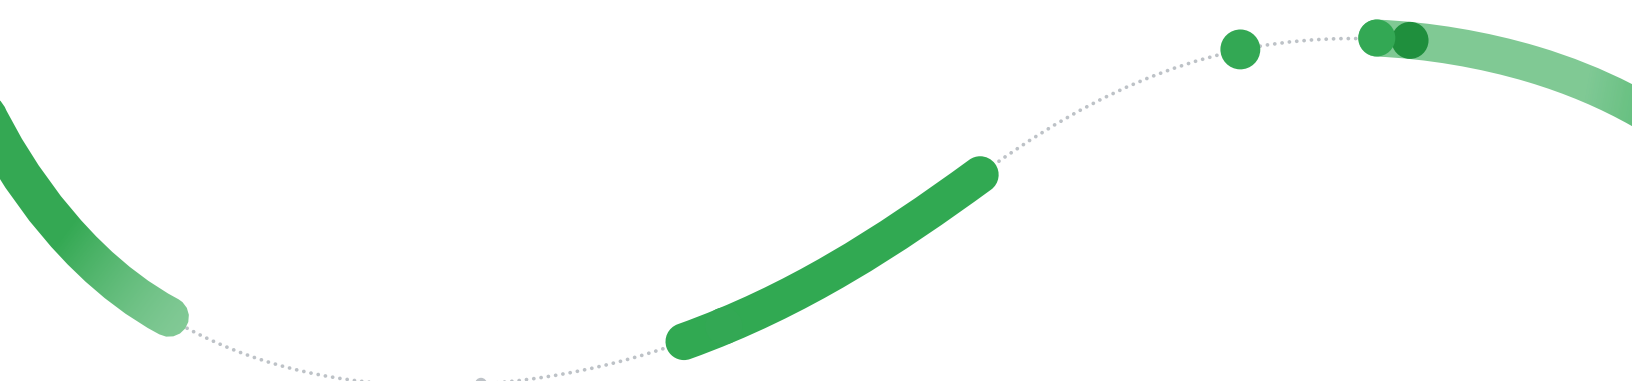
# Product impact

Our responsible approach to AI research and governance helps our product teams working on applications for consumers, developers, and enterprises.

We have applied this approach to Gemini — our family of base and instruction-tuned models of various parameter-based sizes, all of which are natively multimodal. Gemini is flexible and optimized for three sizes: Gemini Nano, Gemini Pro, and Gemini Ultra. Gemini Pro and Nano are starting to roll out to our products. We will be making Gemini Ultra available to select customers, developers, partners, and safety and responsibility experts for early experimentation and feedback before rolling it out to developers and enterprise customers in early 2024. Gemini is designed with responsibility as a core goal: addressing challenges from new capabilities, such as multimodality, and implementing state-of-the-art safeguards.

Across our products, we apply a risk-based, principles-driven process — which can also mean taking a cautious and gradual go-to-market approach involving rigorous testing.

For example:

## AI Principle #1: Be socially beneficial

This principle helps teams consider how the overall benefits of generative AI exceed risks in areas such as content quality and AI's impact on industries and sectors.

Consider our decision to develop Universal Dubbing, a generative AI-automated video lip dubbing service. This technology carries risk, as it could be misused for highly believable deepfakes. Rigorous research with partners at the University of Arizona showed the method clearly helped non-native English speakers learn a language faster when watching a realistic, automatically dubbed video. AI Principles reviewers approved the project with a strict gating process for research and educational purposes based on clear benefits for students. As we expand this service, we're implementing guardrails to help prevent misuse and we make it accessible only to authorized partners.

This principle also can be applied on a broader level. For example, it's reflected in YouTube's approach to music generative AI experiments and YouTube's product-specific guidance for working with creators.

YouTube is actively collaborating with a diversity group of leading musicians for their input on developing generative AI tools to enable expression while protecting music artists and the integrity of their work.

This year, we expanded our ads policies to require advertisers to disclose when their election ads include material that's been digitally altered or generated and depicts real or realistic-looking people or events in all countries where we have election ads verification. And we expanded our ongoing work in information literacy to support AI literacy. We launched About this image, a tool that provides more context to help people evaluate visual content they come across online. The tool offers details on when an image and similar images were first indexed by Google, where it may have first appeared, and where else it's been seen online (like on news, social, or fact-checking sites). With this background information on an image, people might be able to see that news articles pointed out that an image was AI-generated.

## AI Principle #2: Avoid creating or reinforcing unfair bias

We have taken a phased approach to launches to account for rigorous adversarial testing for fairness. While we can't disclose some of the details of our fairness testing methods for security reasons, we can report that we release publicly available generative AI experiences only after they have incorporated recommended or conditional mitigations.

This year, we've been testing how generative AI in Search can help people find what they're looking for in new, faster ways. The experience helps with a variety of information needs, including those that benefit from multiple perspectives to avoid unfair bias.

As we've continually improved the experience, we've also expanded internationally beyond the United States with recent launches in India and Japan, with the majority of feedback positive. In our largest global expansion, we've brought generative AI in Search to more than 120 countries and territories including Mexico, Brazil, South Korea, Indonesia, Nigeria, Kenya, and South Africa, with support for four new languages: Spanish, Portuguese, Korean, and Indonesian. So if, for example, you're a Spanish speaker in the US, you can now use generative AI in Search with your preferred language.

## Case study: Lookout
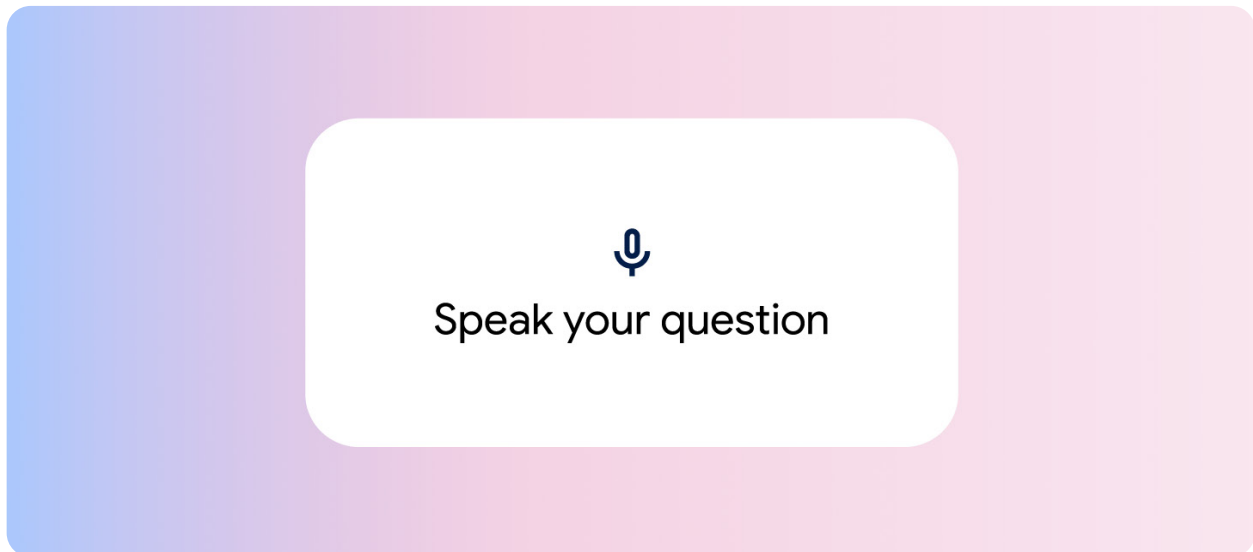


Lookout is an assistive Android app that uses a phone's camera to create accessibility tools for people who are blind or have low vision (BLV). Lookout helps people complete common tasks by making the visual world more accessible. Its newest flagship feature — Image Q&A — enables people to not only get a much more detailed description of an image, but also to ask questions about a photo, and receive AI-powered responses.

Describing images is inherently challenging. If an image contains people, it's even more complex, as difficult questions arise about how to describe those people in a way that's both useful and respectful of a person's identity. Gender is a particularly challenging trait to describe based on an image, as a person's gender may not be obvious from their appearance.

While developing Lookout, the product team had to balance AI Principle # 1 (Be socially beneficial) and # 2 (Avoid creating or reinforcing unfair bias). Though it may be beneficial to include gender in the description of a person, doing so also risks potential unfair bias.

The team incorporated a Google DeepMind visual language model (VLM), heavily customized for this use case, with several rounds of feedback from BLV people and from trans and non-binary people. VLMs enable people to ask natural language questions about an image. The new Lookout question and answer feature allows users to go beyond captions and ask about the image details that matter to them the most.

Google

This functionality allows the team to provide captions without perceived gender, but if the user asks a question about a person's gender, the model can provide a best guess of perceived gender, using cues from the person's appearance. The Lookout team tested this approach with end users who were BLV and non-binary and found that these users thought the approach was both useful and respectful.

The approach isn't perfect. The model will still make mistakes with perceived gender, and people with visual impairments still need to request details that typically sighted people receive effortlessly. The Lookout team believes this launch is both a step in the right direction, and an area where we can continue to learn and improve with the BLV community.

## AI Principle #3: Be built and tested for safety

We design all of our products to be secure-by-default — and our approach to AI is no different. In 2023, we introduced our Secure AI Framework (SAIF) to help organizations secure AI systems, and we expanded our bug hunters programs (including our Vulnerability Rewards Program) to incentivize research around AI safety and security.

To address international frameworks and guidance for safe, secure, and trustworthy AI, we're prioritizing cybersecurity safeguards. Our goal is to protect proprietary and unreleased models and we're participating in industry-wide events to support broader protections for governments, companies, and civil society, like the Defense Advanced Research Projects Agency's (DARPA) AI Cyber Challenge, which will aim to identify and fix software vulnerabilities using AI.

As we introduce generative AI technology to younger users aged 13-17, we strive to strike the right balance in creating benefits while prioritizing safety, family controls, and developmental needs. Informed by research and experts in teen development, we've built additional safeguards into the experience. For example, for our expansion of Search Generative Experience to teens, to prevent inappropriate or harmful content from surfacing, we put stronger guardrails in place for outputs related to illegal or age-gated substances or bullying, among other issues.

## Case study: Search Generative Experience
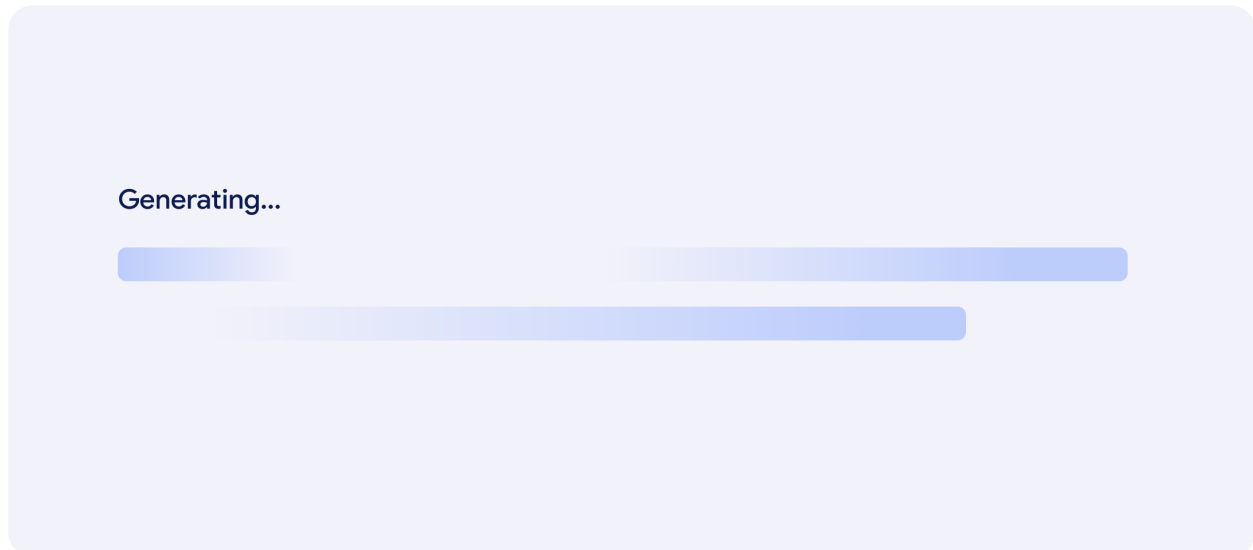


Search Generative Experience (SGE), was introduced through Search Labs this year as a generative AI experiment. Search powered by generative AI can help people quickly get the gist of any topic, find new ideas and inspiration, and easily follow up on questions to deepen their understanding. Generative AI in Search makes it easier for people to ask more specific and complex questions like "How to make learning math fun for a ten-year-old?" People can also ask follow-ups without having to repeat context or try suggested follow-ups, and get AI-powered overviews with links to explore fresh perspectives from across the web.

We are rolling out SGE thoughtfully, to develop this experience responsibly, leaning on Search protections like automated systems that work to prevent policy-violating responses and filtering images that violate our prohibited use policy for generative AI. Other approaches include adding metadata and watermarks indicating that images are AI-generated.

LLMs can generate responses that seem to reflect opinions or emotions, since they have been trained on a range of language. We trained the models that power SGE to refrain from reflecting a persona. They are not designed to respond in the first person, for example, and we fine-tuned the model to provide objective, neutral responses that are corroborated with web results.

Google

By making generative AI in Search first available through Search Labs, we were transparent that the technology was still in an early phase. We're committed to a thoughtful cadence of global expansions after careful testing with audiences around the world.

Over time, we will continue to conduct evaluations and adversarial testing and share information on SGE's capabilities and limitations. In many cases, we have already made improvements with model updates and additional fine-tuning. Generative AI has the potential to transform the current Search experience by organizing and presenting information in ways that help people get — and *do* — more from a single search.

## AI Principle #4: Be accountable to people

When we launch products, we seek to provide relevant information and opportunities for feedback. For example, for Bard's initial launch in May, some of our explainability practices included:

- The "Google it" button, providing relevant Search queries to help users validate responses to factual questions

- Thumbs-up and -down icons as feedback channels

- Links to report problems and offer operational support to ensure rapid response to user feedback

- User controls for storing or deleting Bard activity

We also try to let users know when they are engaging with a new generative AI technology and document how a generative AI service or product works. For Bard's launch, this included a comprehensive overview of the cap on the number of interactions to ensure quality and accuracy, efforts to prevent potential personification, other details on safety, and a privacy notice to help users understand how Bard handles their data.

In addition, we're broadly focused on ensuring that new generative AI technologies have equal guardrails and accountability mechanisms when addressing concerns such as image provenance. In addition to SynthID, our efforts include clear disclosure of images generated by Google AI tools (as in Virtual Try On or Da Vinci Stickies).

## Case study: Bard

Enter a prompt here
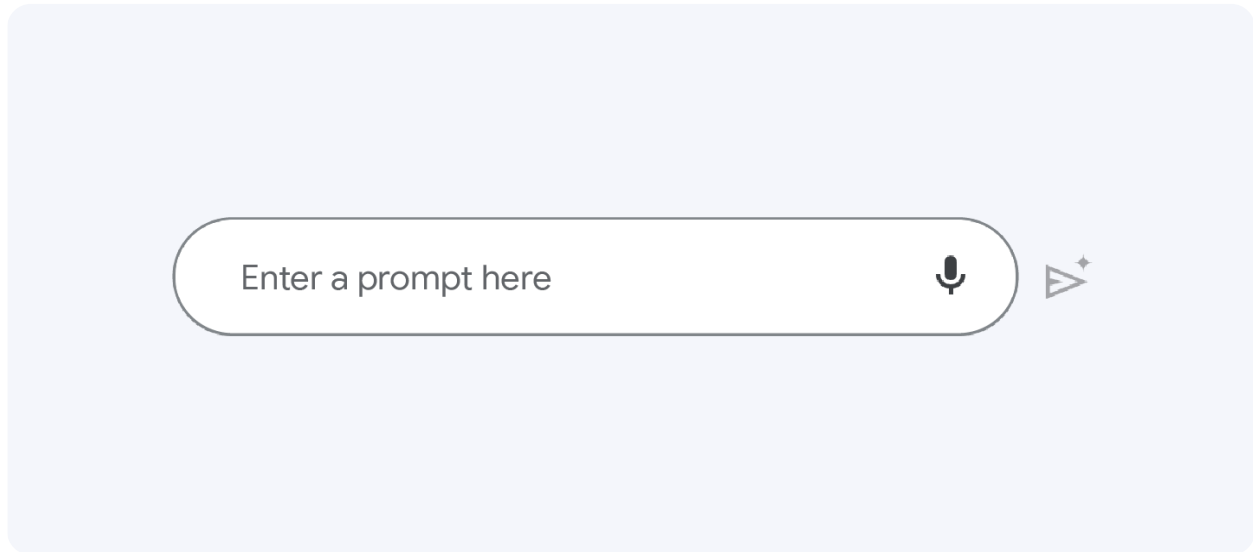
Bard is Google's generative conversational AI experience, launched in early 2023. Bard can support people's productivity, creativity, and curiosity. From planning a party (Bard can come up with a to-do list) to writing a blog post (Bard can provide an outline), people now have a new and helpful creative collaborator.

The models behind Bard have been extensively trained and tested. As a result of potential unfair bias in training data, generative AI products can produce offensive or factually inaccurate output.

In the course of developing and launching Bard, we developed a number of new responsible AI policies. For example, which types of content Bard is and is not allowed to generate influenced our company-wide content frameworks for generative AI models. The team's thoughtful approach to development also shaped our understanding of emerging best practices for responsible generative AI development, including adversarial testing and the inclusion of clear, helpful explanations.

Bard was launched gradually so that the team could learn from real-world use by trusted testers from a diversity of backgrounds and make adjustments as needed. Before launching Bard, we conducted extensive adversarial testing to identify harmful outputs and make improvements to the model. Bard continues to regularly undergo adversarial testing, especially as new features are added.

The Bard interface also makes it clear to people they're interacting with a generative AI model. Additionally, people can offer feedback on the quality of responses using the "thumbs up" and "thumbs down" feature.

## AI Principle #5: Incorporate privacy design principles

Our foundational privacy protections for giving users choice and control over their private data applies to generative AI. We're applying these protections to new product features we're currently developing, like improved prompt suggestions that help people using Workspace get the best results from Duet AI generative features. These are developed with clear privacy protections that keep people in control.

We're committed to protecting your personal information. If you choose to use the Workspace extensions, your content from Gmail, Docs, and Drive isn't seen by human reviewers, used by Bard to show you ads, or used to train the Bard model. You're always in control of your privacy settings when deciding how you want to use these extensions, and you can turn them off at any time.

As we continue to develop, improve, and expand audiences for our generative AI experiences, we will update these protections and share more information on the Bard Privacy Help Hub and elsewhere.

## AI Principle #6: Uphold high standards of scientific excellence

At our I/O event in May of 2023, we announced over 25 new AI-powered products and features. This brings the latest in advanced AI capabilities directly to people — including consumers, developers, and enterprises of all sizes around the world. Our most novel models are developed with scientific rigor and transparency. In addition, we evaluate against multiple criteria and, as appropriate, with external reviews.

For example, Med-PaLM 2, which was trained by our health research teams with medical knowledge, can answer questions and summarize insights from a variety of dense medical texts.

It was assessed for scientific consensus, medical reasoning, knowledge recall, bias, and likelihood of possible harm by clinicians and non-clinicians from a range of backgrounds and countries. Med-PaLM 2 was opened up to a small group of Cloud customers for feedback to identify safe, helpful use cases.

## AI Principle #7: Be made available for uses that accord with these principles

All advanced technologies have multiple uses, including potentially harmful or abusive applications. Our AI Principles guide how we limit harms for people. As we learn more about the emerging risks unique to generative AI, we are working to address these potential harms with technical innovation. For example, we launched a beta version of SynthID to a limited number of Vertex AI customers as a digital watermarking feature for Imagen, one of our text-to-image models that uses input text to create photorealistic images. And we offer image markups for publishers to indicate when an image they post to our platforms is AI generated.

We remain committed to sharing best practices with our customers and developers. That's why we publish Cloud Responsible AI Guides for enterprises. And when AI-powered asset generation for Performance Max was first rolling out to advertisers in the US this year, we offered information in the Google Ads Help Center for advertisers to learn more about asset generation in Performance Max, along with our AI Essentials guide.

# Supporting global dialogue, standards & policy

Building AI responsibly must be a collective effort. It's necessary to involve academics and labs proactively across the research community, as well as social scientists, industry-specific experts, policy makers, creators, publishers, and people using AI in their daily lives. We engage in broad-based efforts — across government, companies, universities, and more — to help translate technological breakthroughs into widespread benefits, while mitigating risks.

For example, this year we:

- Participated in the White House-sponsored red teaming event at DEFCON, which drew over 2,000 people to test industry-leading LLMs in an effort to better understand risks and limitations of these advanced technologies.

- Co-established, with industry partners, the Frontier Model Forum to develop standards and benchmarks for emerging safety and security issues of frontier models.

- Contributed to the Partnership on AI (PAI)'s efforts on a Synthetic Media Framework to help develop and foster best practices across the industry for the development and sharing of media created with generative AI; PAI's Data Enrichment Sourcing Guidelines; and PAI's Guidance for Safe Model Deployment.

- Participated in a number of information sharing sessions about generative AI, including at the Inter-American Development Bank, National Governors Association, US National Conference of State Legislatures Summit, the UK Summit, and more.

- Collaborated with IPSOS on a study on how and why people across 10 countries expect AI will affect privacy in the future, resulting in a paper presented at the 2023 Symposium on Usable Privacy and Security conference.

- Updated our Machine Learning for Policy Leaders workshop with generative AI-specific interactive sessions for policy makers.

Google

# A policy agenda for responsible progress in AI

We're not only focused on identifying risks and benefits of advanced AI. We've been hard at work supporting the larger AI ecosystem with practical, scalable recommendations. Earlier this year, we shared a detailed policy agenda for responsible progress in AI. We outlined a three-pillared approach for governments to collaborate with the private sector, academia, and other stakeholders to develop shared standards, protocols, and governance so we can boldly realize and maximize AI's potential for more people around the world.

The three pillars are:

1.  **Opportunity:** Maximize AI's economic promise, such as increased productivity and upskilling

.................................................................................................

2.  **Responsibility:** Create standards and share practices, and, as appropriate, prepare for regulation

.................................................................................................

3.  **Security:** Align human values while building complex AI to prevent malicious use

Our collaborations across the industry and alongside civil society and academia are building common technical standards that could help align practices globally. These industry-wide codes and standards could serve as a cornerstone for building regulatory frameworks that can promote policy alignment for a worldwide technology.

Putting into place a framework that encourages interoperability across the world can be an opportunity to prevent a very real risk of a fractured regulatory environment, which could delay consumer access to helpful products across borders. This could make it challenging for start-ups and entrepreneurs without the resources to comply with a complex set of uncoordinated AI regulation. These outcomes could slow the global development of powerful new technologies, and undermine responsible development efforts described in this extensive report. Sound government policies are essential to unlocking opportunity, promoting responsibility, and enhancing security, along with individual best practices and shared industry standards for principled AI innovation.
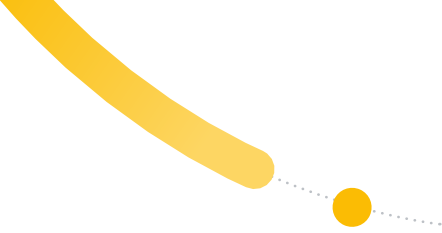
Google

# Conclusion

With the rapid advancements in advanced AI capabilities, we stand on the cusp of a new era not only for computing, but also for society. Responsible AI innovation will help businesses of all sizes thrive and grow, and support society in finding solutions to our toughest collective challenges.

But to unlock the economic opportunity that advanced AI offers while minimizing workforce disruptions, policymakers will need to invest in innovation and competitiveness, promote legal frameworks that support innovation, and prepare workers for potential economic impacts of these evolving technologies.

To bring this vision to fruition and sustain it over time, safely, a multi-stakeholder approach to governance is necessary. Across industries and nations, we can learn from the experience of the internet's growth over decades to develop common standards, shared best practices, and appropriate risk-based regulation.

To do all of the above safely and securely, governments will need to explore next-generation trade control policies for specific applications of risky AI-powered software. Governments, academia, civil society, and companies will need a better shared understanding, via common definitions and consistently structured transparency documents that describe not only the capabilities of AI models when integrated into products and services, but also their limitations.

We're building a strong foundation to enable ourselves and others to embrace AI's transformative promise and continue to evolve for years to come, to help today's workforce thrive, and support future generations:

Google

- To better understand how knowledge workers expect generative AI may affect their industries in the future, we conducted participatory research workshops for seven different industries, with a total of 54 participants across three US cities.

- We're expanding our Google Cybersecurity Certificate program, which can help anyone prepare for a career in cybersecurity globally. For example, in Japan participants can earn a professional certificate from Google through the Japan Reskilling Consortium. This is in addition to existing partnerships with CERT-IN in India and Cyber Security Agency of Singapore, through which we're offering 125,000 scholarships across the Asia-Pacific region.

- To enable businesses and enterprises of all sizes, we've developed brand-new generative AI training options and are constantly adding to our training catalog on Google Cloud Skills Boost. This includes two learning paths that each feature comprehensive content: one is for the non-technical audience, Introduction to Generative AI, and the other, Generative AI for Developers, is for technical practitioners (more advanced). Individual courses are also available on their own.

- For AI engineers and product designers, we're updating the People + AI Guidebook with generative AI best practices. For the same audience, we continue to design AI Explorables, including how and why models sometimes make incorrect predictions confidently.

- For tomorrow's AI engineers and designers, we've launched Experience AI, a new educational program that offers cutting-edge resources for students aged 11-14 and their teachers on artificial intelligence and machine learning. This was developed in collaboration with teachers.

- In 2024, Google will be opening a free after-school Code Next Lab for high schoolers in Inglewood, California, a city where 9 in 10 individuals identify as Black and/or Latinx. Google will be designing, building, and opening the new facility for an immersive computer-science education program to develop the next generation of US Black, Latinx, and Indigenous tech leaders.

Google

Our mission, since we were founded 25 years ago, has always been to organize the world's information and make it universally accessible and useful. Making AI helpful for everyone will be how we deliver on this mission and improve lives everywhere. A big part of accomplishing our mission means making information open and accessible on how Google's core technologies work. We've done this consistently in the transparent tradition of "How Search Works," which we made public a decade ago in 2013. A decade later, advanced AI is no exception. In addition to this annual report, we regularly publish technical reports and research papers that include, or complement, model cards for AI models that are incorporated into AI-powered experiences.

We're encouraged to see governments around the world calling for ongoing transparency into internal AI governance processes and reporting on AI models' capabilities and limitations. Governments and civil society have been seriously addressing how to develop the right policy frameworks for AI innovation this year, and we look forward to supporting their efforts in years to come. At Google, we've been bringing AI into our products and services for over a decade and making them available to people who use our products steadily, guided by our AI Principles. We know we're at an exciting inflection point in our journey as an AI-first company. Some observers have tried to reduce this moment in the history of technology to a competitive AI race across our industry. But what matters most to us is the race to build AI responsibly, together with others so that we get it right – for everyone.
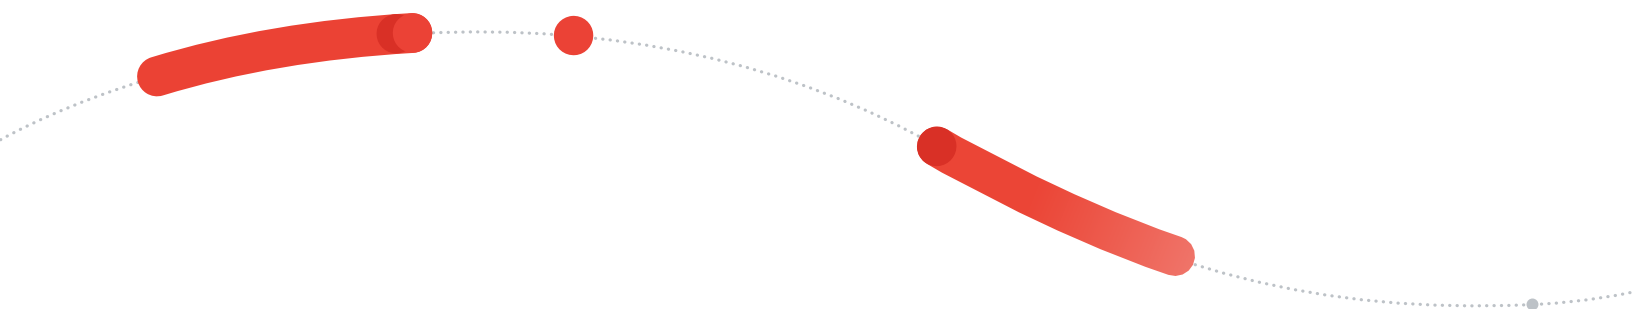
# Appendix

December 2023

# Generative AI System Card: Bard with specifically tuned Gemini Pro

Bard with Gemini Pro is a conversational AI service that is available in English and in over 170 countries and territories. It will be made available in more languages and places, like Europe, in the near future.

The AI system that powers this service uses a specifically tuned version (in English) of Gemini Pro, a foundational large language model (LLM). LLMs are trained deep-learning models that understand and generate text, images, video, and speech in a human-like fashion. LLMs build statistical models of the language they are learning, trying to predict which words are frequently used together across different types of texts and contexts to model the relationships and interactions between words. When given a prompt, they generate a response by selecting, one word at a time, from words that are likely to come next. LLMs must be trained on a vast amount of multimodal data: text, images, video, and speech before they can learn the patterns and structures of language. **The information in this document refers only to the version of Bard with Gemini Pro launched in December 2023.**

## Capabilities

Gemini Pro in Bard (as of December 2023) is specifically tuned for understanding, summarizing, reasoning, coding, and planning capabilities. It works for text-based prompts and provides generated text at this time. Other capabilities of Gemini Pro in Bard include creative writing, composition, language translation, and complex problem solving, including in math and science. At this time, Bard also uses Google Lens technology. We expect to unlock advanced multi-modal capabilities in Bard over time.

Despite the growing range of LLM capabilities, there are known limitations to the use of LLMs in AI-powered systems. There is a continued need for ongoing research and development on how to improve verifiable model outputs so that they are more reliable (e.g., to avoid "hallucinations"). Even when LLMs perform well against model performance benchmarks, they can struggle with tasks requiring high-level reasoning abilities, like causal understanding and logical deduction. Over time, it is necessary to develop more challenging and robust evaluations in these areas.

## Intended use and current integrations

Bard is intended for creative collaboration and conversational AI assistance for consumer use. Bard Extensions, available in English at this time, integrates with Google tools like Gmail, Docs, Drive, Google Maps, YouTube, and Google Flights for more helpful responses. As of December 2023, third-party extensions are not yet available in Bard. Google is currently exploring features that will enable users to connect with third-party services.

## Data

**Data Sources used to train Gemini Pro:**
Gemini Pro is trained on datasets that are both multimodal and multilingual. Our pre-training datasets use data from publicly available web documents, books, and code, and include image, audio, and video data.

**Safeguards:**
We have implemented the following measures to improve the safety and quality of the LLMs for use in products like Bard.

- **Harms mitigation:** Prior to training, various steps were taken to mitigate potential downstream harms at the data curation and data collection stage for Gemini Pro. Training data was filtered for high-risk content and to ensure all training data is sufficiently high quality. Beyond filtering, steps were taken to ensure all data collected meets Google DeepMind's best practices on data enrichment.

- **Mitigations for quality and safety, specific to Gemini Pro:** Quality filters were applied to all datasets used to train the pre-trained Gemini Pro model. Safety filtering was applied to remove harmful content. Evaluation sets were filtered from the training corpus. The final data mixtures and weights were determined through ablations on smaller models. Training was staged to alter the mixture composition during training — increasing the weight of domain-relevant data towards the end of training.

Additional mitigation measures are applied in Google's products, including Bard, over time (as described in Google's AI Principles Progress Updates).

**Personal data collected and processed in providing the Bard service:**

When people interact with Bard, Google collects:

- Conversations

  ....................................................................................

- Location

  ....................................................................................

- Feedback

  ....................................................................................

- Usage information

This data helps provide, improve, and develop Google products, services, and machine-learning technologies, like those that power Bard. Bard shows user-interface elements at the bottom of the Menu that offer continuous transparency about location data processed by Bard. Users can review their prompts, delete Bard activity, and turn off Bard activity at any time. For more details, visit the Bard Privacy Help Hub, and read the Google Privacy Policy and the Bard Privacy Notice.
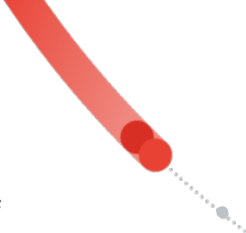
## Model Training Process

**Pre-training**

LLMs are pre-trained on the pre-processed data. Pre-training helps LLMs learn the patterns and relationships in data (which ultimately will be used to generate responses for Bard in Gemini Pro's case).

LLMs built on Transformer[1] architecture fundamentally map the statistical relationships between words, phrases, and sentences, to predict what next words, images, video, or other content will be most likely to follow when prompted with a new set of words. These models first build these relationships in an "unsupervised" way — that is, without the data being categorized or labeled. To do that, they need large quantities of language and language-like data (such as code, math proofs, etc.) that cover the spectrum of use of language to model all the possible relationships that exist across the vast breadth and complexity of words in a single language and among the many languages used around the world.

[1] A Transformer is a deep learning neural network architecture that Google introduced in 2017 (see https://ai.googleblog.com/2017/08/transformer-novel-neural-network.html).

Gemini models are trained to accommodate textual input interleaved with a wide variety of audio and visual inputs, such as natural images, charts, screenshots, PDFs, and videos, and they can produce text, image, video, and audio outputs.

Pre-training large AI systems requires substantial computational, human, and energy resources, and currently can take substantial time to complete. While Google continues to work to reduce such costs,[2] these practical and technical challenges significantly limit how often new foundational LLMs can be pre-trained. For details on the training infrastructure of Gemini models, please see the Gemini technical paper.

**Fine-tuning**

Once a pre-trained model is created, it is then adjusted and adapted for use in a specific application. This fine-tuning process takes additional data, some of which the model may have already seen, and formats it in a way to match the expectations of the application the model is being used in.

This process is undertaken with human supervision/feedback and by using reinforcement learning. Fine-tuning can relatively quickly adapt LLMs to new policies and allow for more experimentation to optimize outputs.

Instruction tuning encompasses supervised fine tuning (SFT) and reinforcement learning through human feedback (RLHF) using a reward model. Instruction tuning was applied in both text and multimodal settings. Instruction tuning for Gemini Pro in Bard was carefully designed to balance the increase in helpfulness with decrease in model harms related to safety and hallucinations. Curation of "quality" data included a mix of data to balance the metrics on helpfulness (such as instruction following, creativity) and reduction of model harms.

Risks of harmful text generation for Gemini Pro are mitigated with technical approaches. For example, a dataset of potential harm-inducing queries was generated to reflect risks and societal harms, guided by the AI Principles. For Gemini Pro, this overall approach was able to mitigate a majority of identified text harm cases without any perceptible decrease in response helpfulness. These mitigations were made before integration into the Bard service.

---

[2] See, e.g. how Google is minimizing our AI carbon footprint:

https://blog.google/technology/ai/minimizing-carbon-footprint/, https://ai.googleblog.com/2022/02/good-news-about-carbon-footprint-of.html, and https://cloud.google.com/blog/topics/systems/tpu-v4-enables-performance-energy-and-co2e-efficiency-gains.

## Model evaluations

Google engages in extensive evaluation and testing to ensure classifiers and other safeguards are operating effectively in all of our novel generative AI models that power services like Bard, including the following:

1. **Pre-Launch testing:** Prior to launch, our AI Principles, Trust & Safety and Responsible AI teams engage in rigorous testing of safety guardrails, including classifiers, to evaluate Bard's performance after these guardrails are put in place. The teams generate large sets of adversarial queries, as well as queries in sensitive verticals, to evaluate how often an unsafe response is generated. Our Trust & Safety team consistently evaluates such safety metrics to ensure the rates meet launch goals.

2. **Targeted adversarial testing and red teaming:** The Trust & Safety, AI Principles and Responsible AI teams also conduct targeted adversarial testing and red teaming to better understand how Bard, and the classifiers it leverages, are performing against certain areas and identify failure patterns that need to be addressed.

3. **Feedback:** Google analyzes and measures the feedback our users provide. Users have the ability to provide feedback if a response by Bard is low quality in the form of a "thumbs-down" vote. A user submitting such feedback then indicates whether a response is unsafe/offensive, not factually correct, or they can specify another reason.

## Deployment and continued iteration

Specialized models built off the base LLM can be further fine-tuned for the specific needs of the product or service where they are deployed. For example, Bard uses a specialized model fine-tuned from the Gemini Pro base model. We continue to iterate on the Bard fine-tuning recipe (data mixtures, fine-tuning parameters).

**Risk Assessment and AI Principles Review outcomes**

Google's AI Principles team conducted a risk assessment and review of Bard. Recommendations resulted in additional extensive deep-dive dogfooding and adversarial testing in the areas of safety, accountability, and inclusion to prepare for the initial experimental rollout of Bard and subsequent updates. Further cross-functional work helped to ensure appropriate mitigations were adopted before Bard and its updates, such as Bard with Gemini Pro, launched. These product mitigations included the following:

- Clear and relevant explanations to set appropriate expectations that describe Bard as a helpful service that offers collaboration with AI for specific types of tasks. Explanations make clear that this AI-powered system is useful for brainstorming ideas, developing plans, creating first drafts of written outlines, emails, blog posts, or for quick summaries of complex topics.

- Disclosures in the Bard Privacy Notice stating that people should not rely on Bard's responses as medical, legal, financial or other professional advice.

- Disclosure in product stating that Bard responses should be double-checked for information accuracy.

- Added ability to use Google Search to find content that helps users assess and further research the information they get from Bard.

- Feedback channels and operational support were defined and built to help ensure rapid response to user feedback to improve the model and address issues.

Google

## Testing

For Bard and its updates, the testing approach is:

- Red teaming for security using the Secure AI Framework (SAIF)

- Adversarial testing for unfair bias

- Applying mitigations, for example, for content policy violations or abuse

- Conducting ongoing trusted testing of Bard with external users in experimental releases

## Safeguards

A set of "model policies" guided the model development of Gemini Pro and evaluations. Model policy definitions act as a standardized criteria and prioritization schema for responsible development and as an indication of launch-readiness. Gemini model policies cover a number of domains including child safety, hate speech, factual accuracy, fairness and inclusion, and harassment.

Other outcomes of ongoing AI Principles reviews of Bard and its updates include the following:

- Dedicated Bard safety teams and policies

- Review of Bard user safety feedback

- Technical safeguards such as classifiers and filters were used to enforce policies

- A restricted use or gating policy is used when appropriate

- A Generative AI Additional Terms of Service exists

- Ongoing technical testing to inform decisions and improvements

Google

In addition, ongoing AI Principles guidance for transparency and user control practices have been implemented, including:

- Disclosure in product stating that Bard should be double-checked for information accuracy

......................................................................................................

- Bard Privacy Help Hub

......................................................................................................

- FAQs

......................................................................................................

- Help article(s)

......................................................................................................

- In-product safety guardrails to add contextual help, like Bard's "Google it" button to more easily double-check answers

......................................................................................................

- Feedback opportunities for users

......................................................................................................

- Clear user controls

Note: information on Google's overall model safety strategy and classifiers is highly confidential, commercially sensitive, and proprietary information of Google. Any public availability of this information could expose people who use Google's products and the greater public to security and safety risks. As clear, consistent directives emerge, we aim to share additional transparency artifacts in the context of scientific excellence (as stated in AI Principle # 6), with appropriate third parties on how best to offer additional details while both remaining competitive and prioritizing people's safety. At this time, for the responsible reasons stated above, this document doesn't offer specific details on any model size, training methods or compute, or other similarly proprietary or sensitive information. The design and details of our transparency artifacts evolve over time to reflect the evolution of technologies, product specifications, and user interface design. The content may be adjusted for the needs of various audiences.