Accelerating social good with artificial intelligence:

Insights from the Google AI Impact Challenge

Contents

Executive summary

Why we wrote this report

In fall 2018, we announced the Google AI Impact Challenge, an open call for proposals on how to use artificial intelligence (AI) to help address society's most pressing issues. Twenty organizations received a total of $25 million in grant funding from Google.org, coaching from Google's AI experts, credit and consulting from Google Cloud, and inclusion in a six-month Google Developers Launchpad Accelerator. In total, we received 2,602 submissions from six continents that together provide a view of the global AI for social good landscape. This publication shares insights gathered from these submissions, along with an extensive catalog of proposed projects, to help accelerate the way a variety of organizations can use AI to improve people's lives—and our planet.

Global momentum around AI for social good is growing—many organizations either are expressing interest in or are already using AI to address a wide array of societal challenges. We were impressed by the range of ideas we saw, and we're excited to share what we learned with the growing AI for social good community. As more social sector organizations recognize AI's potential, we all gain more high-impact opportunities to strengthen the emerging ecosystem.

Insights from our application review

As we reviewed the submissions and interviewed a portion of applicants, we uncovered trends about the state of AI for social good. While these trends are based only on a subset of organizations operating in this space, we believe that they point to significant opportunities for organizations seeking to use AI for social good.

Insight 1: Machine learning is not always the right answer.

Insight 2: Data accessibility challenges vary by sector.

Insight 3: Demand for technical talent has expanded from specialized AI expertise to data and engineering expertise.

Insight 4: Transforming AI insights into real-world social impact requires advance planning.

Insight 5: Most projects require partnerships to access both technical ability and sector expertise.

Insight 6: Many organizations are working on similar projects and could benefit from shared resources.

Insight 7: Organizations want to prioritize responsibility but don't know how.


The American University of Beirut is applying deep learning to agricultural data to help farmers in arid climates optimize irrigation systems.

In sharing these calls to action, we hope to inspire action and provoke conversation.

Opportunities

Nonprofits, academic institutions, social enterprises, funders, and technical partners—across both public and private sectors—all play a role in bolstering the developing AI for social good community. We also see an opportunity for strategic partnerships and collaboration among stakeholders, each of which brings a unique lens that can accelerate impact within this budding ecosystem.

As an extension of our insights, we have identified areas of opportunity for each stakeholder and have developed a set of proposed actions for those that are invested in advancing AI for social good.

In sharing these calls to action, we hope to inspire action and provoke conversation with illustrative examples. Of course, the most effective path forward will depend on many factors, such as geographic region, technical proficiency, resources at hand, the sector addressed, and organizational composition.

We called.
The world answered.

6
continents

119
countries

2,602
proposals

Launching a call to changemakers

In fall 2018, we launched the Google AI Impact Challenge, an open call to the world's nonprofits, social enterprises, and research institutions for proposals on how to use AI to help address society's most challenging, pressing, and important issues. Twenty organizations received a total of $25 million in grant funding from Google.org, coaching from Google's AI experts, credit and consulting from Google Cloud, and inclusion in a six-month Google Developers Launchpad Accelerator. In total, we received 2,602 proposals from organizations from 119 countries spanning six continents for projects addressing a wide range of sectors, from environmental to education to humanitarian aid.

The landscape it revealed

The proposals reflected the wide range of organizations already working with AI methods and capabilities to address social challenges and looking to accelerate and scale their work. We also saw many applicants taking their first steps to use AI. In fact, 55% of the not-for-profit organizations and 40% of the for-profit social enterprises that applied reported no prior experience with AI. Their applications point to the tremendous potential in this nascent space.

Most applicants proposed to take advantage of the open-source machine learning libraries that help bring AI from the world of academia to the beginnings of mainstream use. As awareness of AI capabilities grows and barriers to entry lower, it's likely we will see even more social sector organizations enter the AI space.

Submissions to the Impact Challenge, along with our observations from the review process, paint an initial picture of the global AI for social good landscape today. They also illuminate opportunities for stakeholders to help catalyze the impactful and responsible future use of AI.

By sharing what we learned from this process, we hope to contribute to the accelerating use of AI as a meaningful tool for improving people's lives and solving some of society's biggest problems.

A view of the AI for social good landscape

Nexleaf Analytics is using machine learning to predict vaccine degradation and efficacy.

A growing global belief in the social promise of AI

44

countries had 10 or more applicants

61

countries had at least 5 applicants

A view of the AI for social good landscape

The Google AI Impact Challenge was a wide-reaching global call for ideas using AI for social good. The volume, diversity, quality, and creativity of submissions revealed a growing belief in AI as a promising tool for tackling difficult social challenges.

Many kinds of organizations from around the world are interested in using AI for social good

The geographic diversity of the applications highlighted AI's potential for addressing both global and local issues in developed and developing nations alike. We received applications from nearly two-thirds of the world's countries, across six continents.

Thirty-three percent of the applications came from the United States, 19% from Europe, 16% from Asia, 14% from the Middle East and Africa, 9% from Latin America, 6% from Canada, and 3% from Oceania. Forty-four different countries had 10 or more applicants, and 61 had at least 5 (see the appendix for a full list of countries with 5 or more applicants). It's important to note that this geographic distribution is likely not fully representative of

More than half of the applications came from organizations with fewer than 25 employees. This points to AI as a potential force multiplier in smaller organizations.

the overall distribution in the sector, as the open-call announcement and application were in English and relied on local nongovernmental organization networks for dissemination. Nonetheless, it's a promising indicator of the global interest in applying AI for social good.

The submissions were balanced between nonprofits, social enterprises, and academic institutions. The size of these organizations was also diverse, ranging from entrepreneurial two-person nonprofits to leading research facilities with thousands of staff. More than half of the applications came from organizations with fewer than 25 employees, showing that AI is an accessible and powerful tool for both small and large organizations. This points to AI as a potential force multiplier in smaller organizations.

AI can be used to address a wide array of societal challenges

AI is flexible enough to extend beyond a single societal challenge or sector. In fact, proposals received as part of the AI Impact Challenge touched upon all 17 of the United Nations' Sustainable Development Goals.

Of the sectors represented, health-related applications were the most common, representing more than 25% of total submissions. Sixteen percent of applications sought to apply AI to environmental issues, followed by 12% for education, and 11% for economic development. Equality and inclusion, crisis response, and public and social sector topics also each received hundreds of submissions. The distribution of applications across sectors illustrates the versatility of AI capabilities and a broad-based interest within the social sector to leverage them.

For example, AI capabilities and techniques such as computer vision, deep learning, and natural language processing have the potential to enable new approaches to address social and economic issues. Across all major sectors mentioned above, computer vision was the most common AI capability chosen by applicants, with 41% referencing its use. Roughly a quarter of proposals sought to apply machine learning analytics, followed by nearly 20% each for structured deep learning and natural language processing.

41%

of applicants across all major sectors proposed to use computer vision, which was the most commonly referenced AI capability


WattTime is applying computer vision to satellite imagery to improve power plant emissions monitoring.

Common methods and capabilities leveraged by applicants

Core AI methods	
Rules-based solutions	Use explicitly stated rules to make decisions.
Machine learning solutions	Learn without explicit programming, using examples, to develop a model that can make decisions. • **Deep learning:** Using multiple layers of artificial neurons to create a network that can make a decision based on raw input. Applications of deep learning include computer vision and speech recognition.
AI-powered capabilities	
Audio processing	Hear, recognize, and process sound files and other auditory inputs. • **Speech recognition:** Using audio processing to translate human speech to text.
Computer vision	See, recognize, and process images, videos, and other visual inputs. • **Object detection:** Using computer vision to pick out and identify particular objects and/or physical properties. • **Image and video classification:** Using computer vision to understand and categorize or label visual inputs.
Machine learning analytics	Process and understand large volumes of data to identify patterns and make predictions.
Natural language processing	Process, decipher, understand, and generate human language. • **Sentiment analysis:** Using natural language processing to measure an author's or speaker's positivity or negativity.

Global distribution of applications

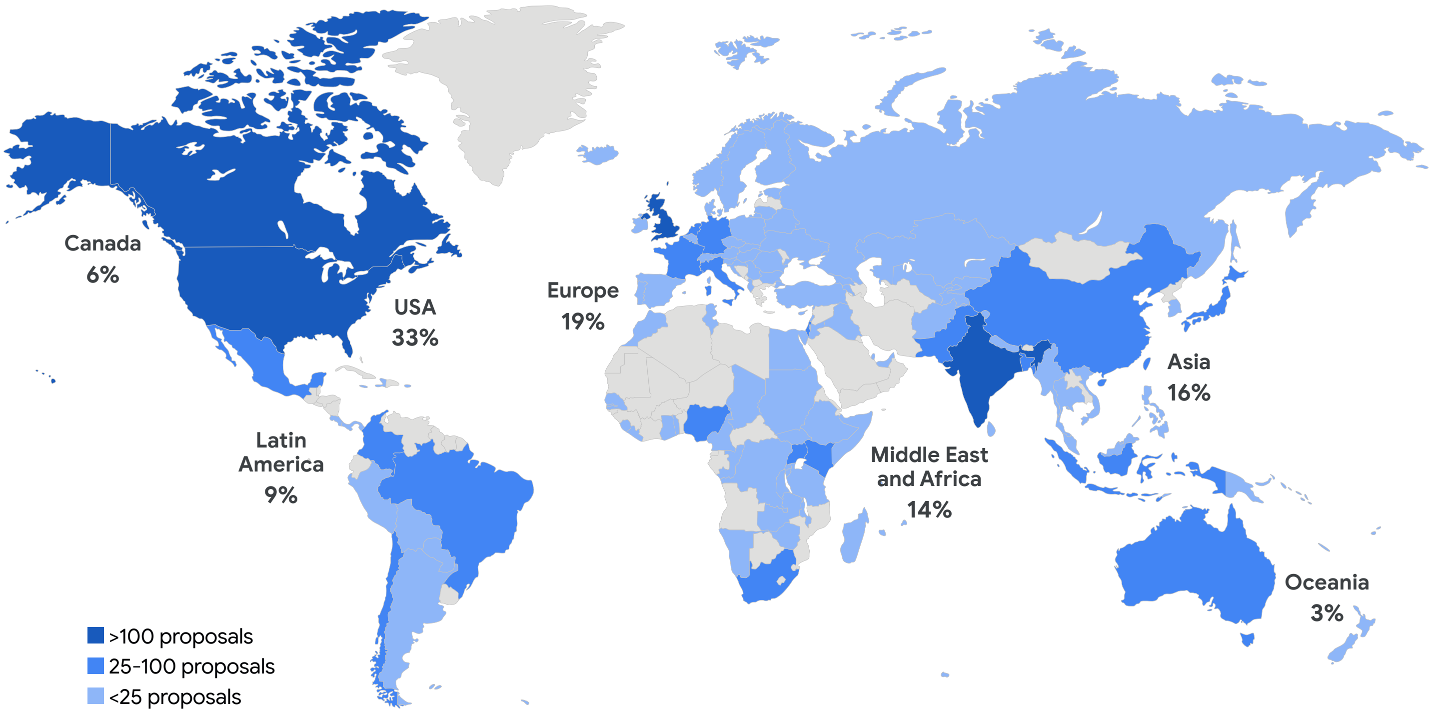

Issue areas addressed in applications

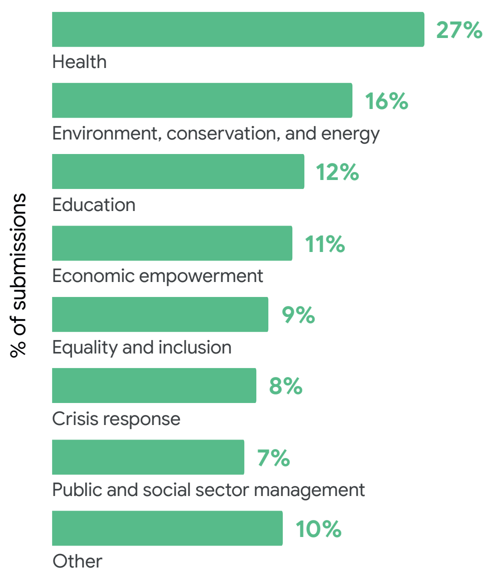

AI capability proposed in submission*

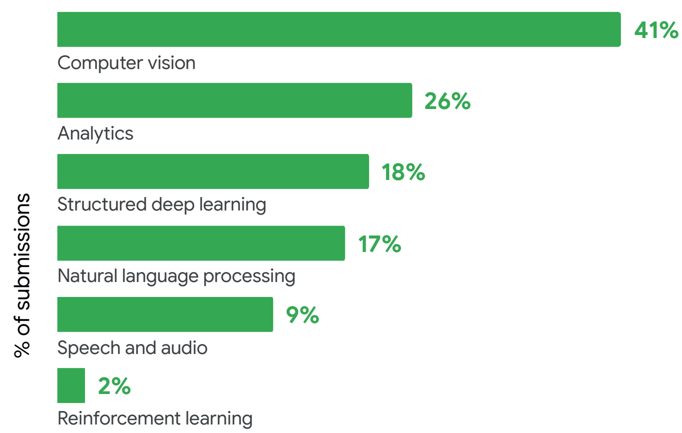

*Applicants could submit a proposal that leveraged more than one capability.

AI usage across applicants

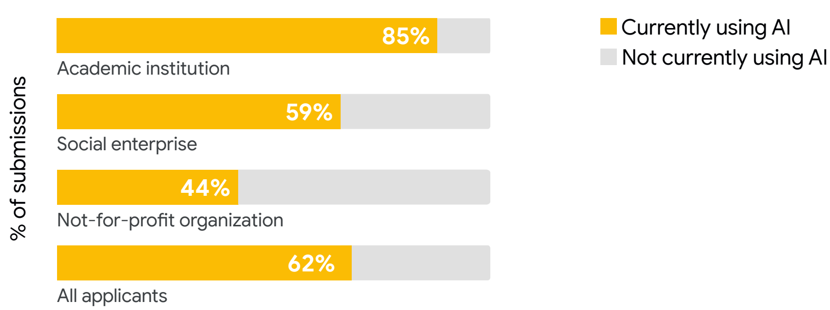


American University of Beirut (Lebanon)


Crisis Text Line, Inc. (USA)

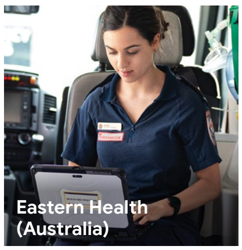
Eastern Health (Australia)


Fondation MSF (France)

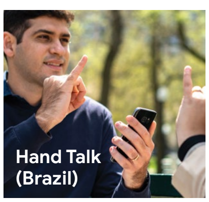
Hand Talk (Brazil)

Google AI Impact Challenge grant recipients

The Google AI Impact Challenge culminated in the selection of 20 grantees. Combined, these organizations received $25 million in grant funding from Google.org, coaching from Google's AI experts, credit and consulting from Google Cloud, and inclusion in a customized Google Developers Launchpad Accelerator program.

American University of Beirut (Lebanon): Applying machine learning to weather and agricultural data to improve irrigation for resource-strapped farmers in Africa and the Middle East

Colegio Mayor de Nuestra Señora del Rosario (Colombia): Using satellite imagery to detect illegal mines, enabling communities and the government to protect people and natural resources

Crisis Text Line, Inc. (USA): Using natural language processing to optimize the assignment of texters in crisis to counselors, reducing wait times and maintaining effective communication

Eastern Health (Australia): Analyzing clinical records from ambulances to uncover trends and potential points of intervention to inform policy and public health responses around suicide

Fondation MSF (France): Detecting patterns in antimicrobial imagery to help medical staff in low-resource areas prescribe the right antibiotics for bacterial infections

Full Fact (UK): Developing trend-monitoring and clustering tools to aid fact checkers' analysis so that they can help contextualize the news and enable informed decisions

Gringgo Indonesia Foundation (Indonesia): Building an image-recognition tool to improve plastic recycling rates, reduce ocean plastic pollution, and strengthen waste management in under-resourced communities

Hand Talk (Brazil): Using AI to translate Portuguese into Brazilian Sign Language through a digital avatar, enabling digital communication for Brazilians who are deaf or have partial hearing loss

HURIDOCS (Switzerland): Using natural language processing and machine learning to extract and connect relevant information in case-related documents, allowing human rights lawyers to effectively research and defend their cases

Makerere University (Uganda): Tracking and predicting air pollution patterns via low-cost sensors in Kampala, Uganda, improving air quality forecasting and intervention

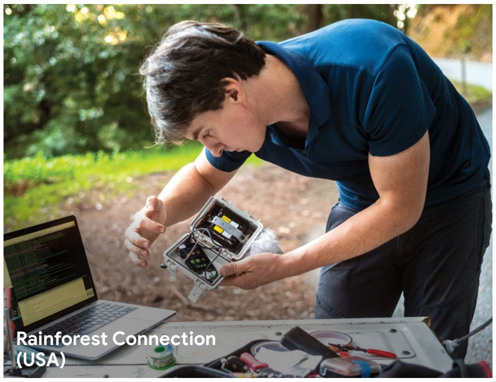
Rainforest Connection (USA)


Skilllab B.V. (Netherlands)

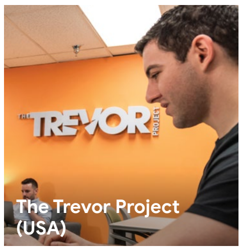

The Trevor Project (USA)


Wadhwani AI (India)


WattTime (USA)

New York University (USA): Partnering with the New York City Fire Department's analytics team to optimize response to its yearly 1.7 million emergencies, accounting for factors such as weather, traffic, and location

Nexleaf Analytics (USA): Building data models to predict vaccine viability throughout the cold vaccine supply chain and ensure effective delivery

Pennsylvania State University (USA): Using deep learning tools to better predict times of and locations at risk for landslides, creating a warning system to minimize the impact of natural disasters

Quill.org (USA): Using deep learning to provide low-income students with immediate feedback on their writing, enabling students to revise their work and quickly improve their skills

Rainforest Connection (USA): Using deep learning for bioacoustic monitoring and commonplace mobile technology to track rainforest health and detect threats

Skilllab B.V. (Netherlands): Helping refugees translate their skills to the European labor market and recommend relevant career pathways to explore

TalkingPoints (USA): Using AI to enable two-way translated parent-teacher engagement and coaching when language represents a barrier to communication

The Trevor Project (USA): Using natural language processing and sentiment analysis to determine an LGBTQ youth's suicide risk level to better tailor services for individuals seeking help

Wadhwani AI (India): Using image recognition to track and analyze pest-control efforts, enabling timely and localized intervention to stabilize crop production and reduce pesticide usage

WattTime (USA): Using image-processing algorithms and satellite networks to replace on-site power plant emissions monitors with open-source monitoring platforms

Insights from our application review

Makerere University is applying machine learning analytics to air pollution patterns in Kampala, Uganda.

We believe these trends point to significant opportunities for organizations and individuals to apply AI for social good.

Insights from our application review

As we reviewed the submissions and interviewed a portion of applicants, we uncovered trends about the current state of AI for social good (see the appendix for the application review criteria). Some trends confirmed our existing assumptions, while others were more surprising. Although these trends are based only on the subset of organizations that applied, we believe they point to significant opportunities for organizations and individuals to apply AI for social good—and ways we all can help amplify their impact.

Insight 1: Machine learning is not always the right answer.

Insight 2: Data accessibility challenges vary by sector.

Insight 3: Demand for technical talent has expanded from specialized AI expertise to data and engineering expertise.

Insight 4: Transforming AI insights into real-world social impact requires advance planning.

Insight 5: Most projects require partnerships to access both technical ability and sector expertise.

Insight 6: Many organizations are working on similar projects and could benefit from shared resources.

Insight 7: Organizations want to prioritize responsibility but don’t know how.

1 Machine learning is not always the right answer.

While we were excited by the number of organizations interested in exploring AI, throughout our evaluation process we identified opportunities for organizations to better understand what machine learning is and when it is an appropriate solution.

Some organizations submitted proposals that might be better implemented without machine learning by leveraging other methods that could result in faster, simpler, and cheaper execution. Other applicants underestimated the complexity of the work required, and in still other instances, machine learning is not yet sophisticated enough to make the proposed solution viable.

For example, several organizations proposed leveraging a deep learning platform to match individuals from an underserved population with the most appropriate legal knowledge and tools. While deep learning would be technically sufficient to achieve the desired outcome, similar results could be achieved faster and more cost-effectively through a rules-based system designed to recommend relevant content. Because rules-based systems are easier to understand and explain, the organizations could also help their intended audiences more quickly and effectively navigate the platform to find the most relevant resources.

HURIDOCS is using natural language processing of legal documents to help human rights lawyers effectively research and defend their cases.

INSIGHT 1: Machine learning is not always the right answer.

Opportunities

Implementers, policymakers, and funders need to be equipped to understand the potential value—and limitations—of AI, especially machine learning, to pressure test whether there is a faster, simpler, cheaper alternative to reach a proposed goal. Organizations and individuals with technical expertise play a critical role in providing educational resources and consulting with organizations on how best to leverage the right tools.

WHERE TO START:

Funders: Learn more about the main methods and capabilities of AI, responsible AI practices, and the right questions to ask about data and implementation to better evaluate the feasibility of proposals using AI.

Organizations using AI for social good: If technical expertise is needed to scope an AI project, reach out to organizations or individuals with that expertise to pressure test whether there is a faster, simpler, cheaper alternative.

Organizations with technical expertise: Host workshops and other forums to equip social sector organizations interested in using AI to assess whether AI is the right fit for their project or challenge.

Policymakers: Help boost public understanding of AI. For example, hold sessions with experts that are open to the public to facilitate more grounded and informed debate on key AI topics.

When can machine learning be useful?

When machine learning usually adds value	When machine learning may have limitations
• Predicting future events when there are many parameters and it's not easy to see a pattern • Personalizing user experience • Understanding natural language • Categorizing objects or behaviors • Detecting infrequent events that change over time	• Working with incomplete, limited, or inaccurate data • Maintaining predictability (e.g., placement of buttons in a user interface) • Offering complete "explainability" • Optimizing for speed and cost to market

Source: Google's People + AI Guidebook, https://pair.withgoogle.com

2 Data accessibility challenges vary by sector.

Access to reliable and meaningful data is a consistent barrier for social sector organizations interested in applying AI methods and capabilities. Data access strategies from applicants fell into one of five categories:

1. First-party data that they had already collected (e.g., electronic health records, student records)
2. Publicly accessible third-party data (e.g., census data)
3. Purchasable third-party data (e.g., satellite images)
4. Third-party data accessible through a partnership
5. A design to collect first-party data

We observed meaningful variation in access to data across sectors as well as in the types of common datasets proposed.

The U.S. Department of Defense, Carnegie Mellon University, and CrowdAI are open-sourcing a labeled dataset containing

700K

before-and-after satellite images affected by eight different types of natural disasters.

Applicants in the crisis response, economic empowerment, and equality and inclusion categories were more likely to lack meaningful datasets. A significant number of crisis response submissions proposed to use accessible data sources such as satellite images and weather data but were limited by insufficient training data, as crisis events are not very common. For example, while satellite imagery is plentiful, there are not always enough labeled images of areas affected by a specific natural disaster to reliably train an image-classification model. The data challenges faced by economic empowerment and equality and inclusion proposals illustrate the difficulty in collecting large amounts of data from vulnerable populations that are often more transient, highly sensitive to privacy, and less likely to participate in the formal economy.

Applicants in health, environment, education, and the public and social sectors were more likely than other sectors to already have access to the necessary data. For these sectors, many of the data sources—for example, electronic health records in hospitals or student academic records in universities—are collected over the course of regular operations. The primary challenge for these organizations is acquiring access to the relevant data repositories. Without access to first-party data or public data sources, organizations need resources to purchase the respective data, such as satellite images, and/or time and assistance to broker the right partnership(s) for data access.

INSIGHT 2: Data accessibility challenges vary by sector.

Opportunities

In sectors where data already exists but is not easily accessible, organizations that own data have an opportunity to invest in data-sharing partnerships and responsible open-sourcing to allow other stakeholders to utilize this data (see Google's approach to open data*). In these cases, it will be important to consider privacy and security risks as well as potentially harmful use cases before sharing datasets broadly. In more data-sparse sectors, funders can help finance data collection. Funders and policymakers could leverage their resources and influence to support the collection and sharing of data, where appropriate.

WHERE TO START:

Funders: Fund data collection or responsible aggregation, ideally incentivizing sharing across organizations.

Organizations using AI for social good: Identify owned datasets that can be safely open-sourced or shared through data governance structures such as whitelists and data trusts.

Policymakers: Use data to help boost research and development in AI.

- Create a framework that incentivizes and facilitates the creation, sharing, and reuse of datasets relevant to priority fields, in a manner that respects user expectations of privacy.
- Make more public datasets available, especially in priority subject realms for innovation.

Example: To help organizations that want to use machine learning to assess building damage related to disaster recovery, the U.S. Department of Defense, Carnegie Mellon University, and CrowdAI are open-sourcing a labeled dataset containing 700,000 satellite images of buildings before and after they were affected by eight different types of natural disasters. This dataset could train AI to help first responders safely navigate through areas recently affected by a natural disaster and prioritize areas most in need of aid.

*https://www.blog.google/technology/ai/sharing-open-data/

Commonly used datasets by project type

Sector	Commonly used datasets	
Crisis response	• Satellite images	• Social media (e.g., Twitter, Instagram)
Economic empowerment	• Business transaction records • Census and socioeconomic data	• Satellite images (agricultural yield) • Weather data (agricultural yield)
Education	• Student and school records	• User data from learning mobile applications
Environment	• Air, water, bioacoustic sensor data • Satellite images	• Weather data
Equality and inclusion	• Census data • Issue-specific public databases (e.g., from Organization for Economic Cooperation and Development, United Nations)	• Legal cases and outcomes • Survey data
Health	• Electronic health records (including medical imaging)	• Historical infectious disease outbreak records
Public sector	• Procurement data from individual countries	• Satellite images

3 Demand for technical talent has expanded from specialized AI expertise to data and engineering expertise.

Many organizations simply need analysts to manipulate the data and engineers to run the datasets through pre-existing algorithms.

In recent years, new machine learning software libraries and other open-source tools have reduced the technical overhead required to implement machine learning. More than 70% of submissions, across all sectors and organization types, relied on existing AI frameworks (e.g., Caffe, cuDNN, TensorFlow, PyTorch).

These tools enable individuals without deep backgrounds in machine learning to use best-in-class algorithms and practices. As a result, organizations have less need to hire specialized AI experts to build and use machine learning systems. With the burden of algorithm design and development removed, most of the necessary technical work is focused around cleaning and formatting data, as well as identifying key features, tools, and tuning parameters. Many organizations simply need analysts to manipulate the data and engineers to run the datasets through pre-existing algorithms.

These engineers need access to enough AI domain expertise to understand which algorithms to use; how to select, structure, and test training data; and how to test and mitigate for robust, responsible systems. While many applicants recognized the need for someone to play this data engineer role, throughout our review process, we saw that even the most mature organizations underestimated the time and resources needed to prepare and maintain the data for algorithm use.

Additionally, many applicants that were new to AI needed a better understanding of the types of data roles required, and others found it difficult to compete with private sector organizations when hiring technical talent.

Pennsylvania State University is using computer vision to create a warning system for landslides.

INSIGHT 3: Demand for technical talent has expanded from specialized AI expertise to data and engineering expertise.

Opportunities

While data analysts and data engineers are more accessible than AI experts, many of the applications, particularly from nonprofits, still cited access to technical expertise as a critical bottleneck. In the near term, organizations with technical expertise and funding can provide resources to fill the expertise needs. In the longer term, training and educational resources from online courses and bootcamps can help ameliorate the data talent shortage.

WHERE TO START:

Funders: Provide funding to organizations to meet needs for data analysts and data engineers, and fill other gaps in technical expertise.

Organizations with technical expertise: Create formal opportunities for employees, particularly data analysts and data engineers, to volunteer their technical skills for social sector projects.

Example: Google, IBM, and Microsoft have all developed formal programs to provide full-time technical expertise to nonprofit organizations. Smaller companies may not be able to dedicate full-time resources but could offer as-needed consults with nonprofit technical teams.

Policymakers: Create incentives and programs for technical talent to support organizations that want to use AI for social good.

- Offer grants to support the development and provision of AI-oriented vocational training for people employed or seeking jobs in priority sectors.
- Offer training grants to encourage people from diverse backgrounds to learn about AI in order to bring fresh perspectives and wider community representation.
- Partner with industry in priority sectors to establish apprenticeship schemes to train the next generation of data scientists.

It would be beneficial to engage potential users early and often throughout the development process to understand user needs and how AI might help meet them.

4 Transforming AI insights into real-world social impact requires advance planning.

Many applicants had deep AI expertise and access to meaningful datasets but needed a clear implementation plan to translate insights and analysis into social impact.

For example, several organizations aimed to use satellite imagery to map in real time the damage caused by natural disasters such as earthquakes, fires, and floods, but they did not have a plan to share the model's insights with real-world programs that could act on these learnings to save lives. While it is easy to see the potential value of such a model—for example, improved warning systems and better distribution of aid—without properly defining a path toward implementation, that value cannot be fully realized. Moreover, it would be beneficial to engage potential users early and often throughout the development process to understand user needs and how AI might help meet them.

Academic organizations, in particular, should consider this challenge, as many of the proposals were led by professors focused on research and technical development. The American University of Beirut, for example, received a grant from the AI Impact Challenge for its work to save water by optimizing irrigation schedules through machine learning. The team has already developed relationships with local agricultural associations and intends to work with these partners to distribute the schedules to farmers in the form of a user-friendly mobile application.

Gringgo Indonesia Foundation is building an image-classification tool to improve waste management.

INSIGHT 4: Transforming AI insights into real-world social impact requires advance planning.

Opportunities

Before an AI system is developed, organizations need to plan for how they will operationalize it for impact and get meaningful feedback from users and beneficiaries. Identifying experts and partnering with affected communities to test and iterate ideas will lead to more impactful implementation of AI-powered solutions and avoid interesting research projects that may otherwise fail to reach their potential impact.

WHERE TO START:

Funders: Ensure AI for social good projects have a clear path to impact and have the necessary resources and plans to engage users and the sector early in the design process.

Organizations using AI for social good:

- For organizations focused on research or developing AI systems, invite implementers to research workshops and actively seek to work with organizations that can apply your research for real-world impact in tackling societal challenges.
- For organizations aiming to both create and implement the technology, develop your AI systems and implementation plan with frequent user testing and feedback from target beneficiaries and organizations working with these populations.

Policymakers: Elevate the most important social and environmental challenges that need attention, and offer sectoral expertise to organizations developing solutions.

- Offer subsidies to support investment in the physical infrastructure underpinning AI in regions where it is lacking (e.g., discounts on electricity, faster CapEx depreciation).
- Look for tangible ways to make it easier for organizations to access AI capabilities, including through cloud-based services (e.g., by providing more flexible rules around data localization).
- Encourage universities to include training on applying AI across their curriculums, beyond engineering, so that the next generation of graduates to enter the industry is well equipped.

5 Most projects require partnerships to access both technical ability and sector expertise.

The most compelling proposals demonstrated fluency in both social sector and technical expertise. Applicants submitting such proposals fell into three broad categories:

1. Technologically savvy nonprofits and social enterprises that have in-house technical talent
2. Academic institutions that were willing to own their project and its implementation post-development
3. Partnerships between nonprofits with deep sector expertise and academic institutions or technology companies with the technical ability to shape and execute the AI portion of the project

On average, we saw that organizations with established partnerships had the most comprehensive applications and were best positioned to implement AI for social good.

Given the nascency of AI for social good work, few organizations fell into the first two categories. On average, we saw that organizations with established partnerships had the most comprehensive applications and were best positioned to implement AI for social good.

To accelerate impact, funders, technical partners, and policymakers should help equip organizations to be fluent in both AI and the sectors they are working to affect. Still, there may always be organizations with deep social sector expertise that will need to work with technical partners. Partnerships also face challenges, because nonprofits, technical partners, and academic institutions approach issues from different perspectives and may have different working cultures, priorities, and goals.

To provide an example of a mutually enriching partnership, one organization with a long history of working directly with refugees identified a need for better distribution of humanitarian aid. However, it lacked the technical expertise to create a machine learning forecasting model to inform the flow of aid from refugee camps with excess to those that were often short. This organization partnered with a company that specializes in applying advanced data analytics and AI tools for a variety of industries to create a prototype that improved forecasting of refugee supply needs. In return, the nonprofit provided the subject-matter expertise and distributed the findings to the relevant stakeholders.

INSIGHT 5: Most projects require partnerships to access both technical ability and sector expertise.

Opportunities

We believe partnerships between organizations with deep sector expertise and organizations with technical expertise offer the most actionable near-term opportunity to develop and operationalize the use of AI for social good. However, partnerships are not easy to find and develop. Common forums where organizations interested in using AI for social good could share missions and needs with technical experts may help facilitate connections.

WHERE TO START:

- **Funders:** Host gatherings for grantees interested in using AI for social good to share current needs, and facilitate mutual introductions that may lead to strategic partnerships.
- **Organizations using AI for social good:** Have a clear understanding of your own strengths and limitations related to applying AI and developing potential partner profiles.
- **Policymakers:** Create frameworks and forums to foster cross-sector collaboration on AI research.

6 Many organizations are working on similar projects and could benefit from shared resources.

The applications we received highlighted many overlapping problem areas and solutions (see the catalog of common project submissions for examples). Given the obstacles we outline above around data accessibility and technical expertise, we believe that greater openness and collaboration can help jump-start initiatives by allowing resource-constrained organizations to draw from a shared pool of relevant knowledge and tools.

This kind of collaboration is by no means easy. Organizations may be hesitant to share their data and technology for a number of reasons, such as competition for a limited pool of grant funding. Even when there is a willingness to share knowledge, resources are needed to coordinate and organize this knowledge for maximum accessibility and to identify and mitigate privacy and security concerns. Nonetheless, we believe that promoting a culture of responsible openness can help accelerate programs that will help individuals and make the world better.

30

applications proposed using AI to identify and manage agricultural pests.

For example, we received more than 30 applications from all around the world proposing to use AI to identify and manage agricultural pests. For these applicants, sharing images of pests, relevant model code, and best practices could help improve the accuracy of models and reduce time spent on duplicated work for all. This allows the individual projects to focus on the unique value they are bringing to the space, be it their relationships with farmers in a certain region or understanding of pests unique to a certain area.

For tangentially related tasks, such as forecasting pest damage in upcoming agricultural seasons, sharing functions and tools (for example, by constructing a public dataset that contains detailed information about pest diet, migration patterns, or treatment options) will benefit all organizations working on these issues.

INSIGHT 6: Many organizations are working on similar projects and could benefit from shared resources.

Opportunities

A more open ecosystem requires both willingness to share responsibly and ease of access to these shared resources. Stakeholders can contribute to the open ecosystem in three ways: (a) share their own knowledge across multiple stages of AI implementation, (b) incentivize other organizations with knowledge to open-source, and (c) ensure that there are systems in place to make shared knowledge easily accessible for organizations it could benefit.

WHERE TO START:

Funders:

- Join together with key stakeholders in the AI for social good ecosystem to create a third-party body to develop open-sourcing best practices and aggregate open-source AI for social good projects.
- Require grantees to responsibly open-source funded projects (as Google.org asks of our grantees).

Organizations using AI for social good: Invest in responsible open-sourcing to share intellectual property (e.g., models and web and mobile applications), and share these investments with existing sector associations.

Policymakers: Increase access to publicly funded research and hold domain-specific research conferences for organizations to share their work on AI for social good.

Collaboration and open-source opportunities across the stages of AI implementation

Data standards: Organizations that seek to leverage tangential datasets may develop agreed-upon standards for quality, format, and reporting to help facilitate data interchange.

Datasets: Organizations can responsibly share their underlying datasets with other organizations and, when possible, publicly accelerate training and accuracy of machine learning models.

Algorithms and models: Relevant algorithms and models, and specifications on their design and capabilities, may also be shared with other organizations under open-source licenses or through APIs.

Web and mobile applications: Many applications proposed web or mobile applications as a means to implement their models' insights. Open-sourcing these applications can provide a helpful reference point around best practices for other organizations.

7 Organizations want to prioritize responsibility but don't know how.

Any organizations that propose the use of AI need to understand how to implement AI responsibly and the potential risks of applying AI. We found that applicants varied significantly in their understanding of AI responsibility. Some had detailed plans to test datasets and model outputs for bias, protocols for handling personally identifiable information, and risk-mitigation plans for data security. Others focused on the promise of the technology and acknowledged the importance of responsibility but had no plans to ensure that their AI projects would be developed responsibly.

Understanding and managing responsible AI use is both difficult and nonnegotiable.

Understanding and managing responsible AI use is both difficult and nonnegotiable. As part of the review process, we evaluated all shortlisted applications for fit with the Google AI principles.

These state that applications of AI should do the following:

1. Be socially beneficial
2. Avoid creating or reinforcing unfair bias
3. Be built and tested for safety
4. Be accountable to people
5. Incorporate privacy design principles
6. Uphold high standards of scientific excellence
7. Be made available for uses that accord with these principles

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

- 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
- Weapons or other technologies whose principal purpose or implementation is to cause or directly facilitate injury to people
- Technologies that gather or use information for surveillance violating internationally accepted norms
- Technologies whose purpose contravenes widely accepted principles of international law and human rights

We found that many organizations needed to more carefully consider risks related to creating or reinforcing unfair bias, incorporating privacy design principles, and mitigating risk of harmful use or misuse. In addition to the unfair bias and privacy challenges that any AI project might face, applicants using data from historically disadvantaged populations face particular challenges that require careful attention. For example, many applicants highlighted the need to check for unfair bias in models that were trained on a more general population and grappled with how to respect privacy by gathering consent from low-digital-literacy or transient groups.

Organizations also varied in their ability to evaluate potentially harmful use cases and identify ways to mitigate intentional or unintentional misuses of their technology. The organizations that did proactively identify these concerns proposed limiting the use of their models through governance structures, such as whitelists and data trusts.

Ultimately, organizations were interested in finding ways to implement technology responsibly and raised thoughtful questions. We hope that more social impact organizations using AI will help advance the conversation around responsibility, not only for the social sector, but also more broadly.

INSIGHT 7: Organizations want to prioritize responsibility but don't know how.

Opportunities

Ultimately, everyone has a critical role in ensuring responsibility.

WHERE TO START:

Funders: Ensure projects are vetted for potential responsibility concerns and that there is an ongoing mechanism for reviewing concerns with grantees.

Organizations using AI for social good:

- Have a clear idea of the responsibility guidelines to be followed (Google's AI principles are just one example of many).
- Where possible, make transparent modeling decisions and use transparent data collection methods to allow others to pressure test for responsible use of your technology.
- Engage a diverse set of stakeholders, including affected populations, to discuss potential risks and mitigations.
- Evaluate model performance across different dimensions that may highlight areas of unfair bias (e.g., different demographics).
- Develop a risk-mitigation plan for potential areas of harmful use or unintentional misuse.

Policymakers: Promote constructive governance frameworks and build responsible AI expertise in government bodies.

- Make government a role model for responsibly embracing AI.
- Encourage industry to share best practices and promote codes of conduct.
- Promote ethics training for government-funded researchers using machine learning.

Wadhwani AI is using image classification to detect crop pests and inform local pest-control efforts.

Laying a foundation for growth

At Google, we are incredibly excited about the prospect of discovering how to better use AI responsibly and to support the growing AI for social good ecosystem to bring substantial, lasting quality-of-life improvement to people across all communities. It is our hope that the findings and insights in this report will help spark new progress and collaboration in this nascent space. Together, we have an opportunity to lay a foundation for AI to support social good actors in their missions.

Catalog of common project submissions

The following catalog highlights some of the common project types we saw from applicants. We share it with the hope that this information will help inform potential areas of collaboration and investment that can benefit projects at scale.

- Crisis response
- Economic empowerment
- Education
- Environment
- Equality and inclusion
- Health
- Public sector

COMMON PROJECT SUBMISSIONS

Crisis response

Common project types	AI capabilities used	Data sources
Natural and man-made disasters		
Optimizing natural disaster relief efforts and funding	• Machine learning analytics	• Emergency response organization internal data (e.g., emergency calls) • Satellite images
Predicting earthquakes, floods, wildfires, and other catastrophic events	• Deep learning • Image classification • Object detection	• Data on historical occurrences • Emergency response organization internal data • Public international/ regional data • Satellite images • Weather data
Real-time monitoring of extent and severity of damage	• Image classification • Natural language processing — sentiment analysis • Object detection	• Satellite images • Social media (e.g., Twitter, Instagram)

COMMON PROJECT SUBMISSIONS

Economic empowerment

Common project types	AI capabilities used	Data sources
Agricultural quality and yield		
Conserving and measuring soil health	• Deep learning • Machine learning analytics	• Sensor data • Weather data
Forecasting, identifying, and managing crop damage (pests, disease, etc.)	• Image classification • Machine learning analytics • Object detection	• Beneficiary-generated data (e.g., use of mobile application to take pictures of crops)
Guidance on the right crop to grow	• Deep learning • Machine learning analytics	• Sensor data • Weather data
Obtaining estimates for irrigation, fertilizer, and pesticide usage	• Deep learning • Machine learning analytics	• Sensor data • Weather data
Financial inclusion		
Improving access to credit for financially excluded individuals and communities	• Deep learning • Machine learning analytics	• External transaction data (e.g., utility payments, mobile phone usage) • Financial organization internal transaction data • Historical borrower repayment rates
Skills supply and demand matching		
Identifying local job market needs to improve youth employability training	• Deep learning • Natural language processing	• Individual résumés • National labor statistics • Public local job advertisements • Skill ontologies (e.g., ESCO)

COMMON PROJECT SUBMISSIONS

Education

Common project types	AI capabilities used	Data sources
Access and completion of education		
Creating learning tools and emotional training for individuals with autism spectrum disorder	• Facial recognition • Speech processing	• Faces and poses in images and videos • User-generated data from mobile applications
Identifying at-risk students	• Deep learning • Machine learning analytics • Natural language processing	• Student and district records
Mapping out school locations in developing regions	• Image classification • Natural language processing • Object detection	• Census data • National and regional poverty, employment, health outcomes • Publicly available articles and news about specific schools • Publicly available United Nations Statistics Division data • Satellite images
Supporting interactive language learning, especially through mobile applications	• Machine learning analytics • Speech processing	• Existing language curriculum • User-generated data from mobile applications
Maximizing student achievement		
Career coaching	• Machine learning analytics • Natural language processing	• Skills ontology • Student records • Survey data
Providing adaptive learning applications for individual students	• Deep learning • Machine learning analytics • Natural language processing	• Student records • User-generated data from mobile applications
Teacher and administration productivity		
Grading and feedback to improve skills	• Deep learning • Machine learning analytics • Natural language processing	• Past graded papers, projects, homework

COMMON PROJECT SUBMISSIONS

Environment

Common project types	AI capabilities used	Data sources
Climate change		
Estimating causal effects of air pollution (e.g., health outcomes, severe storms)	• Deep learning • Machine learning analytics	• Insurance claims data • Publicly available community demographic data • Publicly available emissions data (e.g., from the Environmental Protection Agency) • Weather and storm data (e.g., Federal Emergency Management Agency maps)
Estimating urban air pollution	• Deep learning • Image classification • Object detection	• Air sensor data • Publicly available emissions data (e.g., from the Environmental Protection Agency) • Satellite images • Traffic flow data
Tracking concentrated methane emissions	• Deep learning • Image classification • Object detection	• Air sensor data • Hyperspectral imaging • Publicly available data on known methane leaks • Satellite images
Conservation		
Identifying illegal fishing vessels	• Image classification • Object detection	• Field monitoring videos and images • Satellite images
Identifying illegal mining	• Image classification • Object detection	• Satellite images
Monitoring ecological communities (e.g., endangered animals and habitats)	• Audio processing • Deep learning • Image classification • Object detection	• Bioacoustic sensor sound data • Field monitoring videos and images
Providing land-cover classification for precision conservation (e.g., water, deforestation, coral reefs)	• Deep learning • Image classification • Object detection	• Satellite images
Waste management		
Identifying and monitoring plastic debris (water and land)	• Image classification • Object detection	• Existing street images (e.g., Google Earth's Street View) • Satellite images • User-generated pictures and videos through mobile application

COMMON PROJECT SUBMISSIONS

Equality and inclusion

Common project types	AI capabilities used	Data sources
Accessibility and disabilities		
Assisting in environment-sensing for the visually impaired	• Image classification • Object detection	• Images and videos of everyday objects
Translating voice and text to sign language	• Audio processing • Natural language processing	• Videos of people speaking
Fair analysis in criminal proceedings		
Identifying bias and stereotypes	• Deep learning • Machine learning analytics	• Historical court data, including case outcomes
Human exploitation		
Predicting human trafficking patterns related to recruitment and transactions	• Deep learning • Machine learning analytics	• Arrest and indictment records • Federal, state, and local human trafficking cases • Police narratives • Publicly available U.S. Department of Labor enforcement data
Migrants and refugees		
Forecasting demand for aid within refugee camps	• Deep learning • Machine learning analytics	• Census data • International Organization for Migration datasets (e.g., displacement tracking) • Survey data • Weather data
Matching jobs to skills	• Deep learning • Natural language processing	• Individual résumés • National labor statistics • Public local job advertisements • Skill ontologies (e.g., ESCO)
Supporting the immigration and asylum processes with multilingual chatbots	• Audio processing • Natural language processing	• Survey data • Unstructured text and audio data from current requests

COMMON PROJECT SUBMISSIONS

Health

Common project types	AI capabilities used	Data sources
Diagnosis		
Detecting diseases (e.g., cancer, neurodegenerative conditions) using CT and MRI images	• Image classification • Object detection	• Electronic health records
Detecting diseases using images from mobile cameras (e.g., skin conditions)	• Image classification • Object detection	• User-generated data from mobile application
Mental health		
Directing those with mental health questions to appropriate resources via chatbot	• Natural language processing — sentiment analysis	• Unstructured text and audio data from current requests
Triaging individuals at most imminent risk of suicide	• Natural language processing — sentiment analysis • Speech processing	• SMS conversations • Social media • Speech recordings with at-risk individuals
Outbreaks and epidemics		
Forecasting outbreaks and spread of infectious disease (e.g., malaria, dengue fever)	• Deep learning	• Publicly available records of past outbreaks (e.g., HealthMap) • Spatial data on population density and community demographics • Spatiotemporal disease data
Prevention		
Developing patient risk profiles (e.g., chronic diseases, pregnancies)	• Deep learning • Machine learning analytics	• Electronic health records
Treatment		
Personalizing treatment plans	• Machine learning analytics	• Electronic health records
Predicting surgical outcomes to improve surgery standardization	• Machine learning analytics	• Electronic health records • Hospital surgical outcome records

COMMON PROJECT SUBMISSIONS

Public sector

Common project types	AI capabilities used	Data sources
Citizen services		
Aiding citizen inquiries and customer service for government services	• Natural language processing	• Unstructured text data from current requests
Corruption		
Monitoring public procurement processes for corruption	• Machine learning analytics • Natural language processing	• Procurement data from individual countries
Land use		
Forecasting urban growth to plan for sustainable land use	• Image classification • Object detection	• Satellite images

Resources

Machine learning is not always the right answer

Using AI for Social Good: This guide helps nonprofits and social enterprises understand the types of problems their organizations can solve with machine learning, learn how to identify and prepare data sources, and review guidelines on how to develop and utilize machine learning responsibly. https://ai.google/education/social-good-guide

Machine Learning Crash Course with TensorFlow APIs: A self-study guide for aspiring machine learning practitioners. https://developers.google.com/machine-learning/crash-course

PAIR guidebook: Originally written for user experience (UX) professionals and product managers as a way to help create a human-centered approach to AI on their product teams, this guidebook is useful to anyone in any role wanting to build AI products in a more human-centered way. https://pair.withgoogle.com

Responsible AI

Building Responsible AI for Everyone: Reliable, effective, user-centered AI systems should be designed following general best practices for software systems, together with practices that address considerations unique to machine learning. Google's top recommendations are located here, with additional resources for further reading. https://ai.google/responsibilities/responsible-ai-practices

Responsible Development of AI: In this white paper, we provide an overview of our approach to responsible use and development of AI and share recommendations for government policy frameworks. https://ai.google/static/documents/responsible-development-of-ai.pdf

Facets tool: Facets contains visualizations to aid in understanding and analyzing machine learning datasets. Get a sense of the distribution of features in your datasets using Facets Overview, or explore individual data points using Facets Dive. https://pair-code.github.io/facets

What-If tool: The What-If tool makes it easy to efficiently and intuitively explore a model's performance on a dataset. Evaluate your model's performance as you manipulate features on individual data points, and explore various optimization strategies. https://pair-code.github.io/what-if-tool

Data sources

Google Cloud Platforms datasets: The Google Cloud Public Datasets program hosts copies of structured and unstructured data to make it easier for users to discover, access, and utilize public data in the cloud. These datasets are hosted for free. https://console.cloud.google.com/marketplace

Kaggle datasets: Explore, execute, share, and comment on code for any open dataset with the in-browser analytics tool, Kaggle Kernels. You can also download datasets in an easy-to-read format. https://www.kaggle.com/datasets

Google Earth Engine catalog: Earth Engine's public data archive includes more than 40 years of historical imagery and earth science datasets, updated and expanded daily. https://developers.google.com/earth-engine/datasets/catalog

Dataset Search: Google's Dataset Search tool allows users to search across thousands of dataset repositories on the web. Once you've found a relevant dataset, Dataset Search will direct you to the repository or provider where the data is hosted. https://toolbox.google.com/datasetsearch

Data Commons: Data Commons attempts to synthesize a single Open Knowledge Graph from different data sources. It links references to the same entities (cities, counties, organizations, etc.) across different datasets to nodes on the graph so that users can access data about a particular entity aggregated from different sources without data cleaning or joining. https://www.datacommons.org

Appendix

Funders

How can I help support AI for social good?

Funders play a vital role in bolstering the developing AI for social good community. But it's not always easy to know where to start.

As part of our work, we've identified some ways that funders can make a difference:

- Learn more about the main methods and capabilities of AI, responsible AI practices, and the right questions to ask about data and implementation to better evaluate the feasibility of proposals using AI.
- Fund data collection or responsible aggregation, ideally incentivizing sharing across organizations.
- Provide funding to organizations to meet needs for data analysts and data engineers, and fill other gaps in technical expertise.
- Ensure AI for social good projects have a clear path to impact and have the necessary resources and plans to engage users and the sector early in the design process.
- Host gatherings for grantees interested in using AI for social good to share current needs and facilitate mutual introductions that may lead to strategic partnerships.
- Join together with key stakeholders in the AI for social good ecosystem to create a third-party body to develop open-sourcing best practices and aggregate open-source AI for social good projects.
- Require grantees to responsibly open-source funded projects (as Google.org asks of our grantees).
- Ensure projects are vetted for potential responsibility concerns and that there is an ongoing mechanism for reviewing concerns with grantees.

Organizations

How can I help support AI for social good?

Organizations that use or are interested in using AI for social good and organizations with technical expertise in AI both play a vital role in bolstering and developing the community. But it's not always easy to know where to start.

As part of our work, we've identified some ways that organizations can make a difference:

Organizations interested in using AI for social good

- If technical expertise is needed to scope an AI project, reach out to organizations or individuals with that expertise to pressure test whether there is a faster, simpler, cheaper alternative.
- Identify owned datasets that can be safely open-sourced or shared through data governance structures such as whitelists and data trusts.
- For organizations focused on research or developing AI systems, invite implementers to research workshops and actively seek to work with organizations that can apply your research for real-world impact in tackling societal challenges.
- For organizations aiming to both create and implement the technology, develop your AI systems and implementation plan with frequent user testing and feedback from target beneficiaries and organizations working with these populations.
- Have a clear understanding of your own strengths and limitations related to applying AI and developing potential partner profiles.
- Invest in responsible open-sourcing to share intellectual property (e.g., models and web and mobile applications), and share these investments with existing sector associations.
- Have a clear idea of the responsibility guidelines to be followed (Google's AI principles are just one example of many).
- Where possible, make transparent modeling decisions and use transparent data collection methods to allow others to pressure test for responsible use of your technology.
- Engage a diverse set of stakeholders, including affected populations, to discuss potential risks and mitigations.
- Evaluate model performance across different dimensions that may highlight areas of unfair bias (e.g., different demographics).
- Develop a risk-mitigation plan for potential areas of harmful use or unintentional misuse.

Organizations with technical AI expertise

- Host workshops and other forums to equip social sector organizations interested in using AI to assess whether it's the right fit for their project or challenge.
- Create formal opportunities for employees, particularly data analysts and data engineers, to volunteer their technical skills for social sector projects.

Policymakers

How can I help support AI for social good?

Policymakers play a vital role in bolstering the developing AI for social good community. But it's not always easy to know where to start.

As part of our work, we've identified some ways that policymakers can make a difference:

- Help boost public understanding of AI (e.g., invite the public to attend sessions with experts to facilitate more grounded and informed debate on key AI topics).
- Create a framework that incentivizes and facilitates the creation, sharing, and reuse of datasets relevant to priority fields, in a manner that respects user expectations of privacy.
- Make more public datasets available, especially in priority subject realms for innovation.
- Create incentives and programs for technical talent to support organizations that want to use AI for social good.
  - Offer grants to support the development and provision of AI-oriented vocational training for people employed or seeking jobs in priority sectors.
  - Offer training grants to encourage people from diverse backgrounds to learn about AI in order to bring fresh perspectives and wider community representation.
  - Partner with industry in priority sectors to establish apprenticeship schemes to train the next generation of data scientists.
- Elevate the most important social and environmental challenges that need attention, and offer sectoral expertise to organizations developing solutions.
  - Offer subsidies to support investment in the physical infrastructure underpinning AI in regions where it is lacking (e.g., discounts on electricity, faster CapEx depreciation).
  - Look for tangible ways to make it easier for organizations to access AI capabilities, including through cloud-based services (e.g., by providing more flexible rules around data localization).
  - Encourage universities to include training on applying AI across their curriculums, beyond engineering, so that the next generation of graduates to enter the industry is well equipped.
- Create frameworks and forums to foster cross-sector collaboration on AI research.
- Increase access to publicly funded research, and hold domain-specific research conferences for organizations to share their work on AI for social good.
- Promote constructive governance frameworks, and build responsible AI expertise in government bodies:
  - Make government a role model for responsibly embracing AI.
  - Encourage industry to share best practices and promote codes of conduct.
  - Promote ethics training for government-funded researchers using machine learning.

List of countries that submitted five or more applications

>100 proposals

Canada
India
United Kingdom
United States

25–100 proposals

Australia
Bangladesh
Brazil
Chile
China
Colombia
France
Germany
Indonesia
Israel
Italy
Japan
Kenya
Mexico
Netherlands
Nigeria
Pakistan
South Africa
Uganda

<25 proposals

Argentina
Austria
Belgium
Bulgaria
Cameroon
Czechia
Denmark
Egypt
Finland
Ghana
Hong Kong
Ireland
Lebanon
Malaysia
Morocco
Nepal
New Zealand
Panama
Peru
Philippines
Poland
Portugal
Romania
Russia
Serbia
Singapore
Slovakia
Slovenia
South Korea
Spain
Sweden
Switzerland
Tanzania
Thailand
Turkey
Ukraine
United Arab Emirates
Vietnam

AI for social good projects tackling the United Nations' Sustainable Development Goals

UN SDGs	Examples of AI for social good projects
1 NO POVERTY	Applying machine learning on satellite imagery and survey data to extract socioeconomic indicators and generate visualizations and predictions of poverty in areas without survey data
2 ZERO HUNGER	Applying machine learning to satellite images, localized food prices, and conflict data to predict and address severe acute childhood malnutrition
3 GOOD HEALTH AND WELL-BEING	Using machine learning analytics on vaccine transit data, vaccine potency data, temperature data in vaccine refrigerators, equipment metadata, and other datasets to predict viability for common vaccines at every point in the supply chain
4 QUALITY EDUCATION	Applying natural language processing on conversations between educators and non-English-speaking parents as well as machine learning analytics on parent and student profiles to construct a multilingual family engagement platform and deliver personalized resources that help parents digitally connect with teachers regardless of language
5 GENDER EQUALITY	Leveraging natural language processing methods, including topic modeling, psycholinguistic feature modeling, and audio signal processing on voice recordings and chat transcripts from crisis call hotlines for women to escalate calls with high risk of intimate partner violence
6 CLEAN WATER AND SANITATION	Leveraging machine learning analytics on thermal satellite images, weather data, and farmer-supplied agriculture data to estimate evapotranspiration and help farmers optimize the exact amount of water needed to irrigate crops
7 AFFORDABLE AND CLEAN ENERGY	Utilizing machine learning analytics and image recognition on satellite images, power grid outlines, and relevant socioeconomic information to determine optimal resource allocation for electricity infrastructure in developing countries
8 DECENT WORK AND ECONOMIC GROWTH	Using machine learning analytics on publicly available skills and occupational data to map an individual's skill set captured through a guided assessment directly to relevant occupations
9 INDUSTRY, INNOVATION AND INFRASTRUCTURE	Using computer vision on Google Street View and LIDAR images to help residents to assess defensible space and identify flammable vegetation around their homes

AI for social good projects tackling the United Nations' Sustainable Development Goals (cont.)

UN SDGs	Examples of AI for social good projects
10 REDUCED INEQUALITIES	Applying machine learning to image and text data to develop a dictionary of norm-revealing phrases that can be used as an alternative credit-scoring mechanism to make consumer lending more accessible to low-income individuals
11 SUSTAINABLE CITIES AND COMMUNITIES	Applying machine learning analytics on incident dispatch data and other correlative data (weather, anomalous events, city demographics, etc.) to build a predictive model around emergency response times for first responders in urban areas
12 RESPONSIBLE CONSUMPTION AND PRODUCTION	Leveraging deep neural net and image-recognition technology on garbage and waste management images to help automate classification and sorting of recyclable items at waste management facilities
13 CLIMATE ACTION	Applying machine learning analytics and computer vision on emissions data, satellite imagery of power plants, weather conditions, and grid conditions to monitor power plant emissions
14 LIFE BELOW WATER	Using image-recognition technology on waste facility images and data to help increase recycling rates and reduce ocean plastic pollution
15 LIFE ON LAND	Applying machine learning analytics and audio recognition on live audio streams from rainforests and other ecosystems to help locals derive insights and help root out any ecological threats
16 PEACE, JUSTICE AND STRONG INSTITUTIONS	Using natural language processing and machine learning methods on legal and judicial documents (e.g., laws, jurisprudence, victim testimonies, and resolutions) to extract relevant information and empower human rights advocates
17 PARTNERSHIPS FOR THE GOALS	Leveraging machine learning analytics on developing country economic indicators and population survey data to create a repository to help inform public and private sector entities as well as the government to better target assistance and measure return on investment

Google AI Impact Challenge application review criteria

The submissions were reviewed by over 200 technical and issue experts within and outside of Google. Each submission was evaluated across five criteria associated with high-potential, high-impact projects.

Potential for impact	The project has a clear path to significant, real-world social impact and is grounded in research and data about the problem and solution.
Appropriate use case for AI	It is clear how using AI uniquely allows the project to have impact (i.e., introduces a solution that otherwise would not be possible). There is a plan to deploy an AI model and a strong link between the AI model's deployment and project goals.
Feasibility	The team has a well-developed, realistic plan to execute on the proposal, including access to a meaningful dataset as well as to domain and technical expertise and a robust implementation plan. The proposal correctly identifies the biggest risks and how they can be mitigated.
Scalability	The project can scale its impact or will serve as a model for other efforts to advance the field.
Responsibility	The project is consistent with the principles laid out in Google's AI principles, making use of recommended technical practices in the responsible AI practices. Potential ethical concerns are recognized, and mitigation strategies are proposed.

Google AI Impact Challenge
September 10, 2019