McKinsey & Company | QuantumBlack AI by McKinsey

AI for social good: Improving lives and protecting the planet

This report is a collaborative effort by Medha Bankhwal, Ankit Bisht, Michael Chui, Roger Roberts, and Ashley van Heteren, representing views from McKinsey Digital.

May 2024

Table of contents

AI is already being used to further all 17 SDGs—from the goal of eliminating poverty to establishing sustainable cities and communities and providing quality education for all—and generative AI has opened new possibilities. As we look to the future, we see exciting potential for acceleration, with new tools and platforms putting ever-greater power in the hands of social entrepreneurs, public sector innovators, and private sector players to create effective solutions. But that power also brings with it the need to assure it is harnessed in trusted and responsible ways and that risks are monitored and managed actively to avoid unintended harms.

Six years ago, it was becoming clear that AI could play a major role globally in promoting not just productivity and economic growth but also social good. In a 2018 report, we outlined how AI capabilities, from natural language processing to sound recognition and tracking, could be used in about 170 use cases to benefit society[1]—to promote equality and inclusion, improve crisis response, protect the environment, and deliver impact in many more ways.[2]

Today's AI R&D is not just confirming our initial assessments but showing promise for further gains in the future. A series of improvements in AI techniques and progress on key enablers have substantially expanded the universe of problems that AI may be able to address. Much of this progress is centered on generative AI, which is enabling natural language interfaces; rapid language translation; synthesis of vast document repositories; creation of stories in text, images, and video; and much more.[3]

In this report, we take another look at how AI can become a key part of solutions to benefit people and the planet—and how it already has. One way to assess this is by mapping innovations and impact to the UN Sustainable Development Goals, or SDGs (see sidebar "Methodology," found at the end of the report). The SDGs comprise 17 goals and 169 targets that aim to improve lives around the world and protect the planet. But the UN's 2023 update on progress toward the SDGs indicates the world is on track to meet only 15 percent of SDG targets.[4] In real terms, this means that 2.2 billion people lack access to safe water and hygiene, and 3.5 billion lack access to safely managed sanitation[5]; roughly 3.3 billion people live in environments that are highly vulnerable to climate change[6]; and about 750 million people are facing hunger.[7]

Below, we illustrate the potential of AI to catalyze progress on these pressing social issues, and we highlight the challenges in the domains of data quality and governance, as well as access to AI talent (particularly for not-for-profits), that are hindering AI from scaling. We then outline some actions that stakeholders—including governments, foundations, universities, and businesses—could take to overcome these challenges. While the opportunities have associated risks, such as embedded biases and data privacy and security threats, thoughtful action could accelerate the deployment of AI-based solutions to advance progress on the SDGs and improve lives across the globe.

---

1 "Applying artificial intelligence for social good," McKinsey Global Institute, November 28, 2018. Additional research in 2023 yielded discovery of 13 more use cases piloted in 2018 that were not originally accounted for in our 2018 report, bringing the 2018 total up to about 170.

2 "'Tech for Good': Using technology to smooth disruption and improve well-being," McKinsey Global Institute, May 15, 2019; Amine Aït-Si-Selmi, Eric Hazan, Hamza Khan, and Tunde Olanrewaju, "Tech for Good: Helping the United Kingdom improve lives and livelihoods," McKinsey, July 31, 2020.

3 For more on generative AI, see "What is generative AI?," McKinsey, April 2, 2024.

4 *The Sustainable Development Goals report 2023: Special edition*, United Nations, July 10, 2023.

5 "The 17 goals," United Nations Department of Economic and Social Affairs, accessed April 24, 2024.

6 "Protecting people from a changing climate: The case for resilience," McKinsey, November 8, 2021.

7 "122 million more people pushed into hunger since 2019 due to multiple crises, reveals UN report," World Health Organization, July 12, 2023.

Chapter 1

How AI can accelerate progress toward reaching all of the SDGs

AI is not a magic bullet, and many risks need to be managed to harness its potential (see sidebar "Managing the risks of adopting AI"). But the universe of problems that AI can address is broad.

Current applications of AI are applicable to all the SDGs, including modeling proteins, screening drugs, designing vaccines, targeting aid and public services, solving supply chain problems such as route optimization for last mile delivery of food in remote geographies, forecasting the long-term impacts of climate change or giving early warning for natural disasters, and bringing expertise to frontline aid workers.

Additionally, adoption of generative AI could significantly increase and democratize access to new capabilities. AI tools now allow remote users to complete tasks that once required specific expertise, such as language translation, fact checking, identification of human or plant diseases, and identification of harmful online content. In a recent survey of more than 4,000 not-for-profits conducted by Google for Nonprofits, 75 percent of respondents said that generative AI had the potential to transform their marketing efforts by enhancing their translation and fact-checking capabilities.[8]

The experts we interviewed noted that AI could address or help solve social or environmental challenges in two circumstances: 1) when the AI solution could solve problems that bottleneck other efforts in the field—for example, a solution for water leakage in residential pipes requires predictions about the likelihood of leaks based on analysis of data such as pipe age and location; and 2) when data required for the model to work is (or will soon be) available and accessible.

To map the breadth of AI's applicability, we have developed a database of AI use cases, each of which highlights a type of meaningful problem whose solution could be enabled by one or more AI capabilities. At the time of our 2018 report, this database contained about 170 high-potential use cases. It now contains about 600—more than a threefold increase. This number is growing as more innovative uses come to light, as social impact leaders continue to experiment boldly, and as AI tools become more accessible and easier to use.

The number of real-life AI deployments has also increased significantly over the past six years. In 2018, only a small fraction of the about 170 use cases had been deployed. Today, about 490 of the 600 use cases, or more than 80 percent, have been implemented in at least one instance (Exhibit 1).[9]

Adoption of generative AI could significantly increase and democratize access to new capabilities. AI tools now allow remote users to complete tasks that once required specific expertise.

---

[8] *The Keyword*, "3 insights from nonprofits about generative AI," blog entry by Annie Lewin, March 28, 2024.

[9] Our library contains approximately 600 use cases, and our analysis of deployments is based on publicly available data. Neither is comprehensive or exhaustive, and both will continue to evolve.

Exhibit 1

About 600 AI-enabled use cases have the potential to support the UN Sustainable Development Goals.

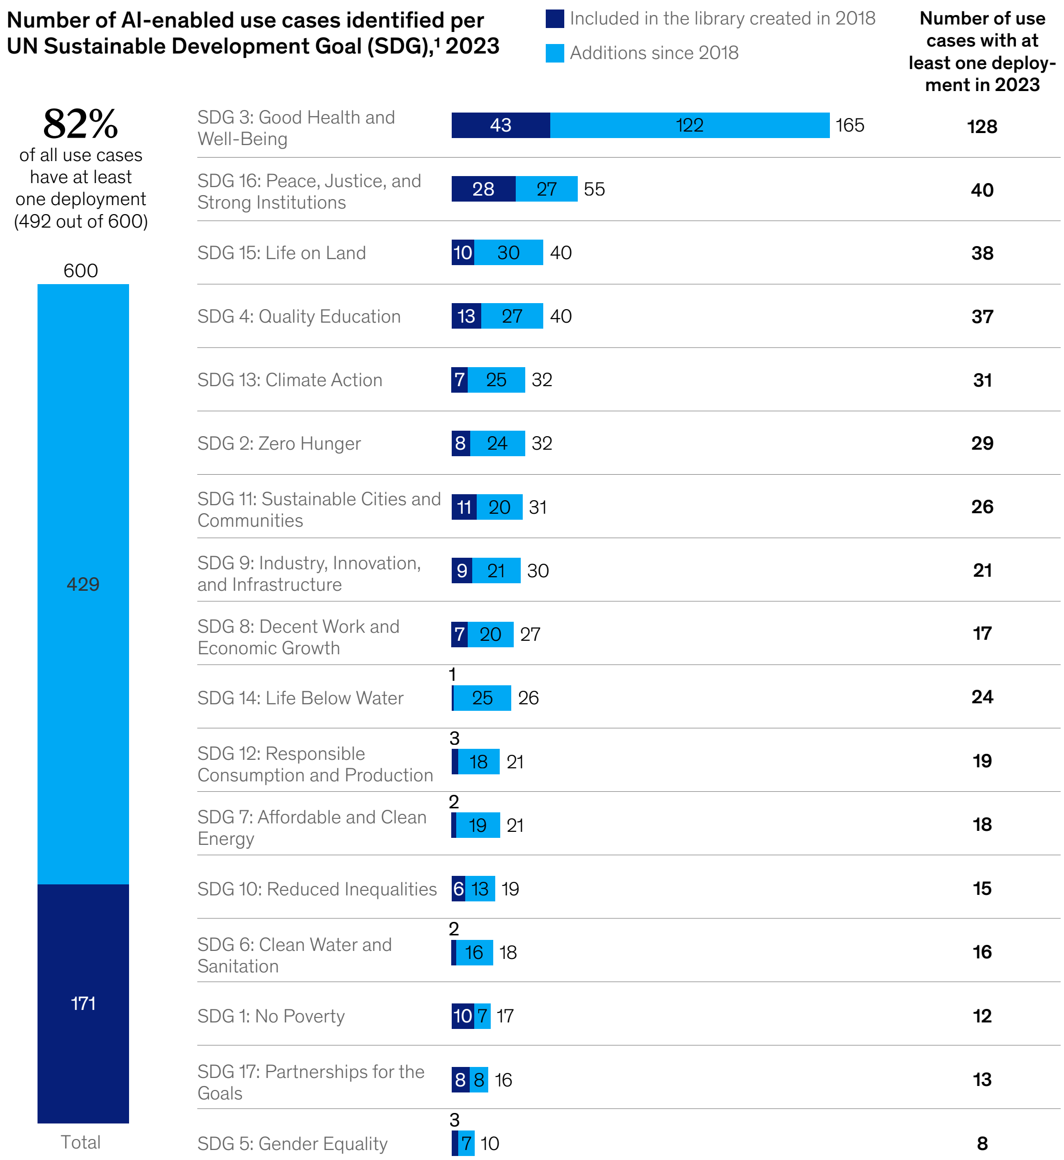

Note: Our library of 600 use cases and our analysis of deployments are based on publicly available data, are not comprehensive, and will continue to evolve; this library is a starting point and should thus not be treated as exhaustive. Many AI use cases are relevant for more than one SDG, which means that successful deployments of these use cases can spur progress on multiple fronts. Additional research in 2023 led to the discovery of 13 use cases piloted in 2018 that were not accounted for in our 2018 report.
[1]Each use case is mapped to primary UN SDGs only.
Source: AI for Sustainable Development Goals academy; Candid database 2018–23; *IRCAI global top 100 2022 report*, International Research Centre on Artificial Intelligence (IRACAI), 2022; *United Nations activities on artificial intelligence (AI) 2021*, International Telecommunication Union, 2021; United Nations Statistics Division; United Nations University Institute for Water, Environment and Health reports

McKinsey & Company

The experts we surveyed agreed that AI has particularly high potential to make a difference for five SDG goals: Good Health and Well-Being (SDG 3), Quality Education (SDG 4), Affordable and Clean Energy (SDG 7), Sustainable Cities and Communities (SDG 11), and Climate Action (SDG 13). In fact, 60 percent of not-for-profit AI for social good deployments were in these areas. Relative to their perceived AI potential, the goals for Zero Hunger (SDG 2), Life on Land (SDG 15), and Peace, Justice, and Strong Institutions (SDG 16) have many use case deployments, whereas Quality Education (SDG 4), Affordable and Clean Energy (SDG 7), and Climate Action (SDG 13) have fewer (Exhibit 2). We excluded Decent Work and Economic Growth (SDG 8); Industry, Innovation, and Infrastructure (SDG 9); and Partnerships for Goals (SDG 17) from the analysis of not-for-profit deployment, foundation grants, and private capital, because most projects can be tagged to these areas given their broad applicability.[10]

Exhibit 2

The number of not-for-profit deployments does not necessarily reflect the perceived AI potential for each Sustainable Development Goal.

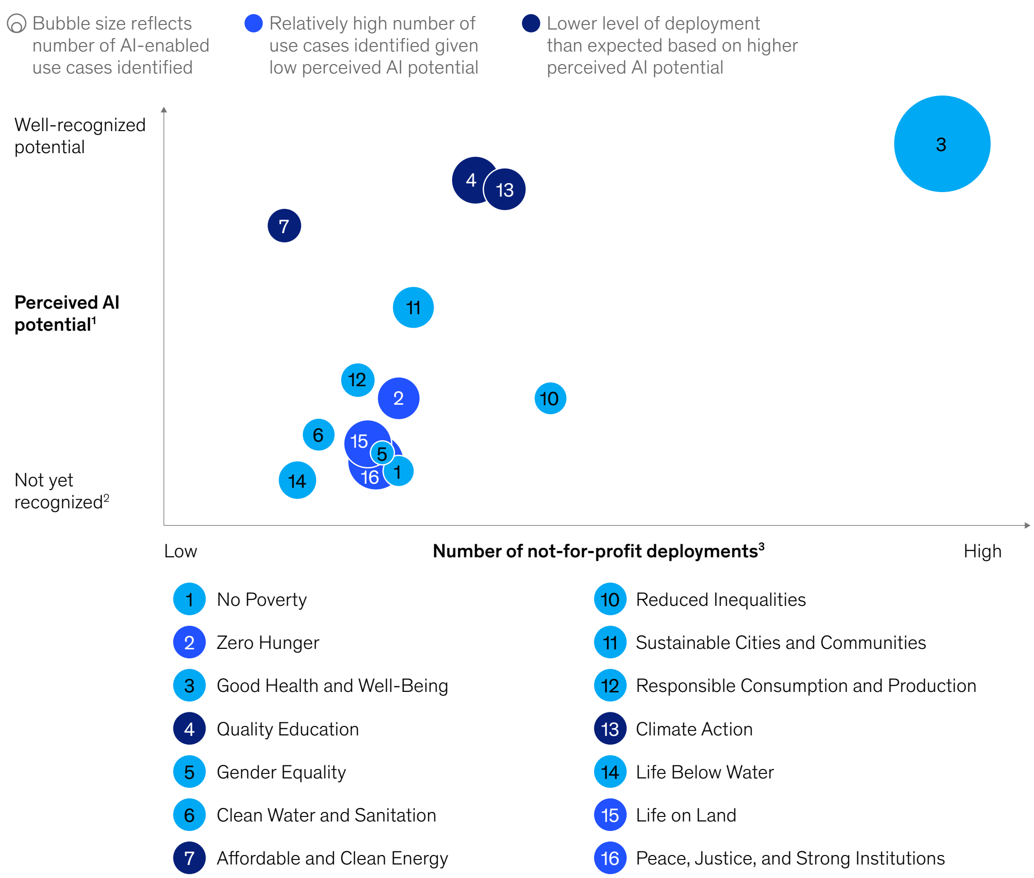

Note: We excluded Decent Work and Economic Growth (SDG 8); Industry, Innovation, and Infrastructure (SDG 9); and Partnerships for the Goals (SDG 17) from our analysis of not-for-profit deployment, grants, and private capital. This is because most projects can be tagged to these areas given broad applicability.
[1]AI potential determined through survey of ~60 experts representing 48 organizations (incl not-for-profits, foundations, technology companies, start-ups, academic institutions, and government) and 17 countries in response to the following question: "What are the top 5 Sustainable Development Goals (SDGs) in the list below where you think AI has the highest potential to accelerate progress toward the SDG targets?"
[2]There may be potential for AI use, but the surveyed experts are unaware of it at this point.
[3]Not-for-profit deployment determined from number of sample deployments in a collection of 1,121 AI applications largely deployed in not-for-profits.
Source: AI for Sustainable Development Goals academy; Candid database 2018–23; *IRCAI global top 100 2022 report*, International Research Centre on Artificial Intelligence (IRCAI), 2022; *United Nations activities on artificial intelligence (AI) 2021*, International Telecommunication Union, 2021; United Nations Statistics Division; United Nations University Institute for Water, Environment and Health reports

McKinsey & Company

[10] In our analysis of 600 use cases, each use case was tagged to a single primary SDG and SDG target.

Additionally, several SDGs that are behind on progress have relatively untapped AI potential. Consider the following examples:

— No Poverty (SDG 1): machine learning could be used to direct cash aid to those most in need or provide alternative credit scores to financially excluded individuals.

— Zero Hunger (SDG 2): AI could be used to help develop new crops, better select crop regions to minimize crop risks, and provide early warning for nutrition crises.

— Peace, Justice, and Strong Institutions (SDG 16): machine learning could be used to detect and curb the spread of misinformation, provide access to information that enables advocacy for policy change, and improve measurement of specific policy interventions.

Below, we explore potential and existing deployments in three of the SDGs with the most widely recognized potential: Good Health and Well-Being (SDG 3), Quality Education (SDG 4), and Climate Action (SDG 13). We also explore two SDGs where AI does not have widely recognized potential but has had an impact in select areas: No Poverty (SDG 1) and Zero Hunger (SDG 2).

We explore potential and existing deployments in three of the SDGs with the most widely recognized potential. We also explore two SDGs where AI does not have widely recognized potential but has had an impact in select areas.

Managing the risks of adopting AI

Risks are inherent to the use of AI. With generative AI (gen AI), risks include inaccurate outputs, biases embedded in the underlying training data, the potential for large-scale misinformation, and malicious influence on politics and personal well-being.[1] As we have noted in multiple recent articles,[2] AI tools and techniques can be misused, even if they were originally designed for social good.

Respondents to our survey of about 60 experts identified the top risks as impaired fairness, malicious use, and privacy and security concerns, followed by explainability (exhibit).[3] Respondents from not-for-profits expressed relatively more concern about misinformation, talent issues such as job displacement, and effects of AI on economic stability compared with their counterparts at for-profits, who were more often concerned with intellectual property infringement.

Exhibit

Experts say impaired fairness and malicious use are the top risks in using AI to address the Sustainable Development Goals.

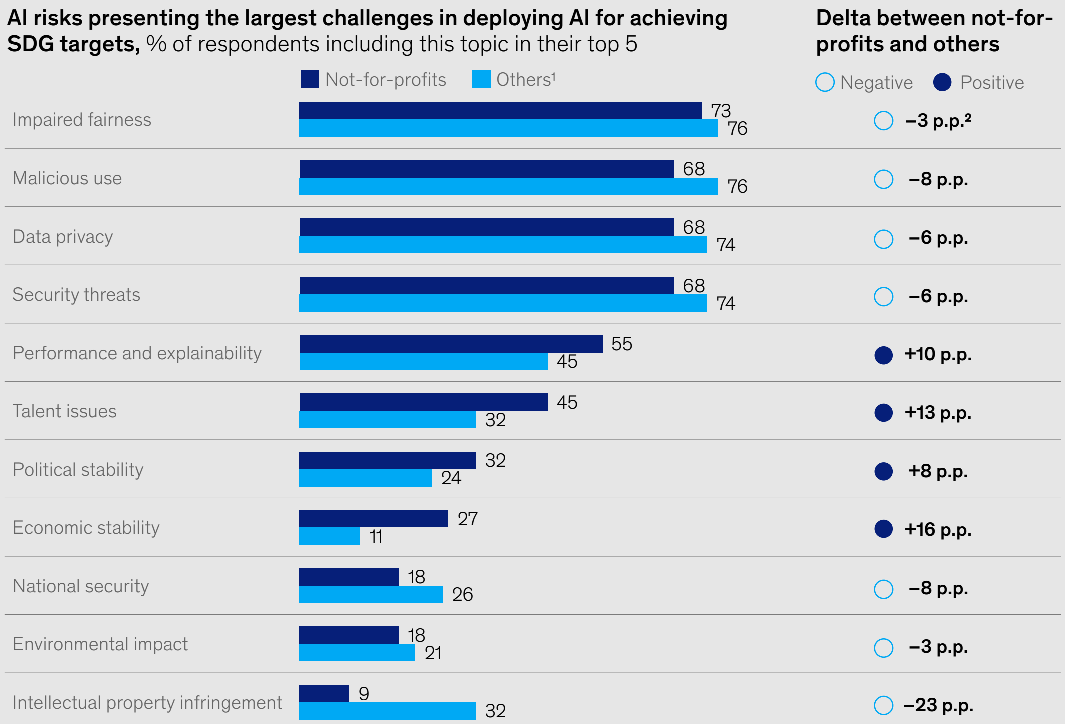

Note: "Impaired fairness" was framed as "bias and fairness" in the survey; "performance and explainability" was framed as "explainability"; "data privacy" and "security threats" were combined in the survey.
[1] "Others" includes for-profits, think tanks, academic institutions, and consultancies.
[2] Percentage points.
Source: Survey of ~60 experts representing 48 organizations (incl not-for-profits, foundations, technology companies, start-ups, academic institutions, and government) and 17 countries

McKinsey & Company

[1] "Implementing generative AI with speed and safety," *McKinsey Quarterly*, March 13, 2024.
[2] Ibid.; "The state of AI in 2023: Generative AI's breakout year," McKinsey, August 1, 2023; *New at McKinsey Blog*, "An inside look at how businesses are—or are not—managing AI risk," blog entry by Liz Grennan and Bryce Hall, August 31, 2023; "What is generative AI?," McKinsey, April 2, 2024.
[3] Our AI risks framework for social impact builds on McKinsey's gen AI risks framework (see "Implementing generative AI with speed and safety," *McKinsey Quarterly*, March 13, 2024). It includes additional categories such as political stability and environmental impact and excludes risks such as strategic risks that can be more relevant to for-profit enterprises.

Impaired fairness. Algorithmic systems can inherit biases from their creators or from the data sets on which they are trained. When these algorithms are deployed in decision-making capabilities, these biases can reinforce preexisting prejudices and social inequalities, with potentially negative impacts on marginalized communities. One organization, Data Science for Social Good, builds bias detection tools that allow developers to audit data science systems for bias and equity.[4]

Malicious use. Malicious use includes creating and disseminating false information or fake content, scams, phishing attempts, hate speech, and activities that harm individuals and national security. A 2022 UN report found that misinformation had been used to incite hatred against marginalized groups and to prevent civilians from finding humanitarian corridors during conflicts, such as the one in Ukraine.[5] According to a recent report by the World Economic Forum, "growing misinformation and disinformation could further increase vaccine hesitancy, which has already led to the re-emergence of locally eradicated diseases."[6] The Global Disinformation Index uses models based on large language models (LLMs) to detect disinformation with the goal of tracking news sites supported by hostile states.[7] Similarly, Full Fact is an independent fact checking organization that deploys a range of AI and machine learning methodologies to detect and curb the proliferation of misinformation across the evolving landscape of information-spreading platforms.

Data privacy and security threats. Many of the UN Sustainable Development Goals use cases require access to health or financial data of vulnerable populations. While organizations are well aware of the harm that could result from breaches in their data systems, many social enterprises have resource constraints that may limit their ability to use the latest cybersecurity capabilities. Several organizations have developed data privacy guidelines, tool lists, and custom security frameworks for not-for-profits with limited resources.[8]

Performance and explainability. Many AI solutions employ complex algorithms that can make it difficult to identify the data or logic used to arrive at a decision. This is particularly relevant for gen AI solutions, which may provide inaccurate or toxic answers. Explainable AI models have several advantages for not-for-profits: they may make it easier to verify the correctness and fairness of results, to assign credit to data providers, and to assign accountability for model outcomes. The Allen Institute for AI recently released a platform for comparing large text data sets to measure the prevalence of toxic, low-quality, duplicate, or personally identifiable information used to train various LLMs.[9]

To mitigate the risks of AI, organizations must first understand and prioritize the risks they are most likely to face, both from inbound AI threats such as disinformation and from developing and deploying their AI solutions. While risks such as data privacy may be addressed through traditional software tools, emerging risks, such as bias in systems driven by LLMs, may require the development of new monitoring systems and guardrails.

---

[4] "The bias and fairness audit toolkit for machine learning: Aequitas," Center for Data Science and Public Policy, accessed April 24, 2024.
[5] *A/77/288: Disinformation and freedom of opinion and expression during armed conflicts - Report of the Special Rapporteur on the promotion and protection of the right to freedom of opinion and expression*, Office of the High Commissioner for Human Rights, United Nations, August 12, 2022.
[6] *The global risks report 2023: 18th edition*, World Economic Forum, 2023.
[7] "What we do," Global Disinformation Index, accessed April 24, 2024.
[8] "Online privacy for nonprofits: A guide to better practices," Electronic Frontier Foundation, accessed April 24, 2024; "Learn," Digital Defense Fund, accessed April 24, 2024; website of SOAP, accessed April 24, 2024; "Frontline policies," Open Briefing Ltd, accessed April 24, 2024.
[9] Akshita Bhagia et al., "What's in my big data?," arXiv:2310.20707, March 2024.

Existing deployments related to SDG 3: Good Health and Well-Being

SDG 3 aims to promote well-being and ensure people live healthy lives.[11] Specific targets for this SDG include reducing maternal mortality; fighting communicable diseases such as AIDS, tuberculosis, and malaria; and establishing universal access to sexual and reproductive care, family planning, and education.

AI is now well integrated into many medical research pipelines. Key AI applications in this area include protein modeling, genome sequencing, computerized tomography (CT) analysis, vision support, and vaccine design. Health is relatively accessible for AI work compared with many SDGs: the field is technology-forward, data availability is high (relative to other SDGs), and health outcomes are frequently measurable. Yet major opportunities remain to support SDG 3 targets that have received less attention, such as treating neglected communicable diseases and preventing substance abuse.

Sample use case: Addressing maternal and newborn health in Kenya. Jacaranda Health provides AI-enabled solutions that improve the quality of care for women with the goal of reducing the number of maternal deaths in Kenya. For example, PROMPTS is an SMS exchange that sends personalized messages to women, empowering them to seek care. An accompanying free digital healthcare platform uses natural language processing to categorize user questions in real time and connects those who need urgent care with a help desk agent. More than two million new and expectant mothers have enrolled with PROMPTS. Mothers who use the services are 20 percent more likely to attend more than four prenatal visits; women who adopt the service are also twice as likely to use postpartum family planning services as women who do not.[12] Jacaranda Health shares feedback from PROMPTS users with governments and facilities to improve their services.[13]

Sample use case: Addressing maternal and newborn health in India. More than 1.3 million women in India have died in pregnancy or childbirth over the past two decades, mostly from preventable causes.[14] ARMMAN was founded in 2008 to address systematic problems that prevent at-risk women from accessing care.[15] The organization developed numerous interventions, including mMitra, an automated voice messaging system that delivers key information on preventive care. These messages have a high correlation with positive health outcomes, such as improved rates of taking iron supplements and better knowledge of family planning. However, 40 percent of women drop out of the program before giving birth.[16] ARMMAN has resources to call some women and encourage them to stay in the program. The organization partnered with Google Research India to develop an AI-based prediction model for this intervention that selects women to receive service calls. The solution is a resource optimization model based on a restless multi-armed bandit approach to optimize resource allocation in a changing world. In a randomized controlled trial, dropout rates were 32 percent lower for women called according to the algorithm than women called using a round robin control group method.[17] Using mMitra, ARMMAN has reached roughly 3.6 million women in nine states, many of whom would likely have dropped out without the AI-targeted intervention. ARMMAN has now developed a similar AI model for use with Kilkari, a voice technology program that brings time-sensitive care information to families.[18]

11 "3: Good Health and Well-Being," Global Goals, accessed April 24, 2024.
12 *PROMPTS*, Jacaranda Health, 2023.
13 "Impact at a glance," Jacaranda Health, accessed April 25, 2024.
14 R. Begum et al., "Trends in maternal mortality in India over two decades in nationally representative surveys," *British Journal of Obstetrics and Gynaecology*, March 2022, Volume 129, Number 4.
15 "ARMMAN: About us," LinkedIn, accessed April 25, 2024.
16 *Google Research Blog*, "Using ML to boost engagement with a maternal and child health program in India," blog entry by Milind Tambe and Aparna Taneja, August 24, 2022.
17 Aparna Hegde et al., "Field study in deploying restless multi-armed bandits: Assisting non-profits in improving maternal and child health," *Proceedings of the AAAI Conference on Artificial Intelligence*, June 2022, Volume 36, Number 11.
18 "Kilkari," ARMMAN, accessed April 25, 2024.

Sample use case: Predicting the structure of proteins to aid drug discovery. DeepMind developed AlphaFold 2 in 2020 and AlphaFold 3 in 2024 to tackle a challenge that had plagued scientists for more than 50 years: the protein-folding problem. This problem involves three related puzzles, as defined by a National Library of Medicine paper: What is the folding code? What is the folding mechanism? And can we predict the native structure of a protein from its amino acid sequence?[19] AlphaFold2 is an attention-based deep learning system that predicts protein structures with a higher degree of accuracy than was previously possible. The DeepMind team released a database of more than 200 million protein structure predictions that is now widely used in structural biology research.[20] A million researchers have accessed the protein structure database since its launch, using the predictions to solve real-world problems, including developing treatments for neglected diseases and fighting antibiotic resistance.[21] AlphaFold 3 extends beyond proteins to include a wide range of biomolecules impacting life sciences and medical research, agriculture, materials sciences, and more.

Existing deployments related to SDG 4: Quality Education

SDG 4 aims to ensure inclusive and equitable quality education and promote lifelong learning opportunities.[22] Targets include establishing free primary and secondary education, ensuring equal access to quality preprimary education, and achieving universal literacy and numeracy.

One in 20 school-age children from low-income countries has internet access at home, while nearly nine in ten from high-income countries do.

AI algorithms are already being used in this space, such as predictive tools that help identify a student's likelihood of completing high school or dropping out and that enable at-risk students to get early intervention and support.[23] AI can be used to create more-inclusive educational platforms for young children, teenagers, adults, and people with disabilities; increase student enrollment; and formulate lesson plans for teachers—including creating materials tailored to students' unique development areas and interests.

Yet implementation has proved challenging, partly because of limited infrastructure—including internet access and data records—in developing countries. Roughly one in 20 school-age children from low-income countries has internet access at home, while nearly nine in ten from high-income countries do.[24] Parents are unable to engage with schools using digital platforms due to factors such as digital literacy and internet access, so use cases that focus on parent engagement are not yet an option.

Sample use case: Enabling people who are nonverbal or experiencing learning disabilities to communicate. Livox uses intelligent algorithms and machine learning to adapt content for students with a variety of disabilities, including verbal, motor, cognitive, and visual.[25] The Livox interface adapts to the student's needs, and its software tracks improvements in visual, auditory, cognitive, and behavioral function, making it easier for teachers to monitor students' progress. More than 25,000 people with disabilities have used this service, which is available in 25 languages.

[19] Ken A. Dill et al., "The protein folding problem," *Annual Review of Biophysics*, June 2008, Volume 37.
[20] "AlphaFold: Protein structure database," EMBL's European Bioinformatics Institute, accessed April 25, 2024.
[21] Oana Stroe, "Case study: AlphaFold uses open data and AI to discover the 3D protein universe," EMBL, February 9, 2023.
[22] "4: Quality Education," Global Goals, accessed April 25, 2024.
[23] *IDeas Blog*, "Rebuilding the Educate Girls machine learning model," blog entry by Sid Ravinutala, April 29, 2019.
[24] "How many children and young people have internet access at home? Estimating digital connectivity during the COVID-19 pandemic," UNICEF, December 2020.
[25] "About us," Livox, accessed April 25, 2024.

Sample use case: Bolstering girls' enrollment in school. Educate Girls is a not-for-profit organization that works to educate girls in India's rural and educationally underresourced areas. The organization uses a machine learning model to reduce the operational cost of locating girls who are not attending school. Before developing this model, Educate Girls staff members had to travel from village to village to gather the required data, which they would then manually compile and analyze to identify areas where their services could have the most impact. The machine learning model uses census data, which is manually cleaned and updated where necessary, and district-level out-of-school data to recommend target areas, allowing Educate Girls staff to reach a greater number of prospective students faster and target interventions more accurately.[26] Educate Girls aims to enroll 1.6 million girls—or 40 percent of the population of out-of-school girls—into grades one through ten.[27]

Existing deployments related to SDG 13: Climate Action

SDG 13 focuses on combating climate change and its impacts, including strengthening resilience and adaptive capacity to climate-related disasters and integrating climate change measures into policies and planning.[28]

AI can be used to analyze large climate data sets and model the impact of specific variables, improve the yield of agriculture, and reduce emissions from transportation and industrial processes, to name a few applications. Not-for-profit deployments are lower than perceived potential for this SDG. In the past few years, AI has been used to provide detailed climate information to improve climate change education and awareness, and to track emissions and improve the sustainability of operations across industries. It has also been used to strengthen resilience efforts through the creation of early-warning systems, forecasts for natural disasters, and up-to-date information for disaster relief.

Sample use case: Conserving forests by stopping illegal deforestation. Global Forest Watch (GFW) assesses forest conditions and monitors destruction using satellites, computer vision, and deep learning. Governments, civil society organizations, and companies can use this data to make informed decisions about sustainable sourcing and conservation. For example, Friends of the Earth Nigeria, a civil society organization, used GFW's forest cover and land concession data to monitor deforestation and corporate land acquisition. This evidence was crucial in informing policy making for the government of Nigeria, which resulted in nearly 14,000 hectares (35,000 acres) of land being given back to affected communities (a result that also benefits SDG 15: Life on Land).[29]

Sample use case: Improving flood forecasting. Google provides actionable flood forecasts to governments, local aid organizations, and people at risk by combining an AI model that forecasts the amount of water flowing in a river with another model that predicts which areas will be affected and how severely.[30] The Flood Hub offers flood forecasts up to seven days in advance of a risk, as well as providing real-time alerts. Flood Hub currently predicts riverine flooding and helps protect livelihoods in more than 80 countries up to seven days in advance, including in data-scarce and vulnerable regions.[31] Between October and December 2023, Google provided the International Committee of the Red Cross (ICRC) with inundation risk maps and daily flood forecasts in 20 locations suffering from riverine flooding made worse due to the El Niño effect in Somalia. This information enabled the ICRC to target its humanitarian efforts toward key activities such as publicizing risks and floodproofing facilities.[32]

[26] "Rebuilding the Educate Girls machine learning model," April 29, 2019.
[27] *IDeas Blog*, "Educate Girls: Improving learning outcomes for millions of children in India," blog entry by Ben Brockman et al., May 20, 2021.
[28] "13: Climate Action," Global Goals, accessed April 25, 2024.
[29] *Global Forest Watch Blog*, "Big data is all around. How do we harness it to drive the change we need?," blog entry by Andrew Steer, April 18, 2017.
[30] "Flood forecasting," Google Research, accessed April 25, 2024.
[31] Deborah Cohen et al., "Global prediction of extreme floods in ungauged watersheds," *Nature*, March 20, 2024, Volume 627.
[32] *The Keyword*, "How AI flood forecasting can help communities in need," blog entry by Brigitte Hoyer and Moriah Royz, February 5, 2024.

Existing deployments related to SDG 1: No Poverty

Even for SDGs with a lower perceived potential for AI and few deployments, there are successful impact initiatives that can be scaled elsewhere. SDG 1, which aims to end poverty in all its forms globally, is one such example.[33] Targets for this SDG include eradicating extreme poverty; implementing social protection systems; and ensuring equal rights to ownership, basic services, technology, and economic resources. Making progress on this SDG requires both capital and the political will to provide social protection and distribute resources; AI has helped improve the efficiency of existing distribution systems.

Sample use case: Distributing aid during the COVID-19 pandemic. The government of Togo partnered with UC Berkeley, GiveDirectly, Innovations for Poverty Action, and local telecom providers to provide aid to Togolese citizens whose livelihoods were most disrupted by the COVID-19 pandemic. The partnerships enabled the government to pull together diverse data sources, such as geospatial analytics and mobile phone metadata, and use machine learning to accurately estimate poverty and improve aid targeting. Roughly 140,000 Togolese were remotely identified and paid. The AI approach helped reach 8 to 14 percent more eligible recipients than the rule-based alternative approach being considered at the time.[34]

Existing deployments related to SDG 2: Zero Hunger

SDG 2 aims to end hunger, achieve food security and improved nutrition, and promote sustainable agriculture through specific targets such as ending all forms of malnutrition, doubling the productivity and incomes of small-scale food producers, and promoting sustainable food production and resilient agricultural practices.[35] AI has already been shown to improve agriculture processes and productivity, food quality (including nutrition), and storage and distribution processes.[36]

Sample use case: Educating farmers on approaches to improve yield. Outgrow is an agricultural and farmer engagement platform in India that connects farmers with high-value products and services, including automated soil labs and microclimatic weather devices.[37] Using the AI-based platform, farmers can predict and detect crop disease in real time, learn how and when to irrigate, and connect with agricultural experts in multiple Indian languages. Outgrow's platform supports more than 18 crops and provides services to more than 200,000 farmers.

Sample use case: Preserving plant biodiversity and enabling the discovery of resilient genes. Future Seeds, a global agricultural innovation hub, uses predictive models to identify habitats that host important crop biodiversity.[38] The organization's Colombia-based facility uses robotics, drones, and AI to help scientists identify traits that can help crops cope in extreme weather conditions, including water stress. With the goal of preserving key specimens and supporting food security, Future Seeds already has the largest collection of cassava, beans, and tropical forages in the world.[39]

[33] "1: No Poverty," Global Goals, accessed April 25, 2024.
[34] "Study: AI targeting helped reach more of the poorest people in Togo," GiveDirectly, July 28, 2021.
[35] "2: Zero Hunger," Global Goals, accessed April 25, 2024.
[36] Lyndah Chiwazi et al., "Artificial intelligence applications in the agrifood sectors," *Journal of Agriculture and Food Research*, March 2023, Volume 11.
[37] "Agri services," WayCool, accessed April 25, 2024.
[38] "Future Seeds: A ground-breaking genebank to help climate-proof food systems," Alliance of Bioversity International and CIAT, accessed April 25, 2024.
[39] Ibid.

Chapter 2

How funding for AI initiatives supports SDGs

Funding toward harnessing AI for the SDGs aligns directionally with the five highest-potential areas identified by respondents to our survey of experts. Grant and private capital funding centers on Good Health and Well-Being (SDG 3), Quality Education (SDG 4), Affordable and Clean Energy (SDG 7), Sustainable Cities and Communities (SDG 11), and Climate Action (SDG 13). In addition, about 40 percent of private capital investments into the 20,000 AI companies we analyzed contributed directly or indirectly toward progress on SDGs.[40]

About 40 percent of private capital investments into the 20,000 AI companies we analyzed contributed directly or indirectly toward at least one of the 17 SDG thematic areas.

Some SDGs are well funded by both grants and private capital; the analysis shows Good Health and Well-Being (SDG 3) receives seven times more grant funding than other SDGs and the most private capital, which corresponds with its high number of not-for-profit deployments. Meanwhile, others are funded predominantly by one source. Consider Affordable and Clean Energy (SDG 7), Sustainable Cities and Communities (SDG 11), and Climate Action (SDG 13), which are primarily funded through private capital, and Quality Education (SDG 4), which receives more grant funding. Private sector capital and grant funding could seem to complement each other, with one source addressing funding gaps left by the other. However, some SDGs—such as Life Below Water (SDG 14) and Gender Equity (SDG 5)—are not funded significantly by either source.

Funding is not always in line with where the experts we surveyed see potential. Consider the relatively low private capital funding for Quality Education (SDG 4). Even for Affordable and Clean Energy (SDG 7) and Climate Action (SDG 13), more than 50 percent of private capital investment went toward autonomous vehicles for improving energy efficiency and reducing emissions. This suggests further potential for private entities to deploy capital toward SDG themes where AI has high potential (Exhibit 3).

An analysis of the location of grant recipients' headquarters from a database of US-majority foundations reveals that from 2018 to 2023, only 10 percent of grants allocated toward AI initiatives that address one or more of the SDGs went to organizations based in low- or middle-income countries.[41] While organizations may have impact outside of the countries where they are headquartered, 60 percent of experts responding to our survey agreed that AI efforts today do not focus enough on benefiting lower-income countries (as opposed to higher-income or developed countries), where the need and impact can be the highest.

Analysis of private capital shows that 36 percent of 9,000 companies addressing SDGs are headquartered in the United States, but these companies received 54 percent of total funding. We also found that while 20 percent of 9,000 companies addressing SDGs are headquartered in lower- or middle-income countries, they received a higher proportion (25 percent) of total funding. One reason for this is that Chinese companies receive a high proportion of investment.[42] The remaining developing countries in the sample received only 3 percent of funding while representing 7 percent of the sample.

[40] SDGs 8, 9, and 17 were excluded from this analysis due to the general applicability of AI to those SDGs.
[41] This analysis is not holistic, because it analyzed grants only from Candid, which focuses primarily on US-based foundations. Additionally, Candid's database takes 1.5 to 2.0 years to update because of a lag from IRS data. See "AI for Sustainable Development Goals," Candid Foundation Directory, 2018–2023.
[42] China is in the middle-income category. See "The World Bank in China," World Bank, updated April 2024.

Exhibit 3

Grant and private capital funding tends to flow toward Sustainable Development Goals where AI has the most recognized potential.

Potential for AI to have an impact on Sustainable Development Goals (SDGs) vs grant funding and private capital funding across SDGs (2018–23)

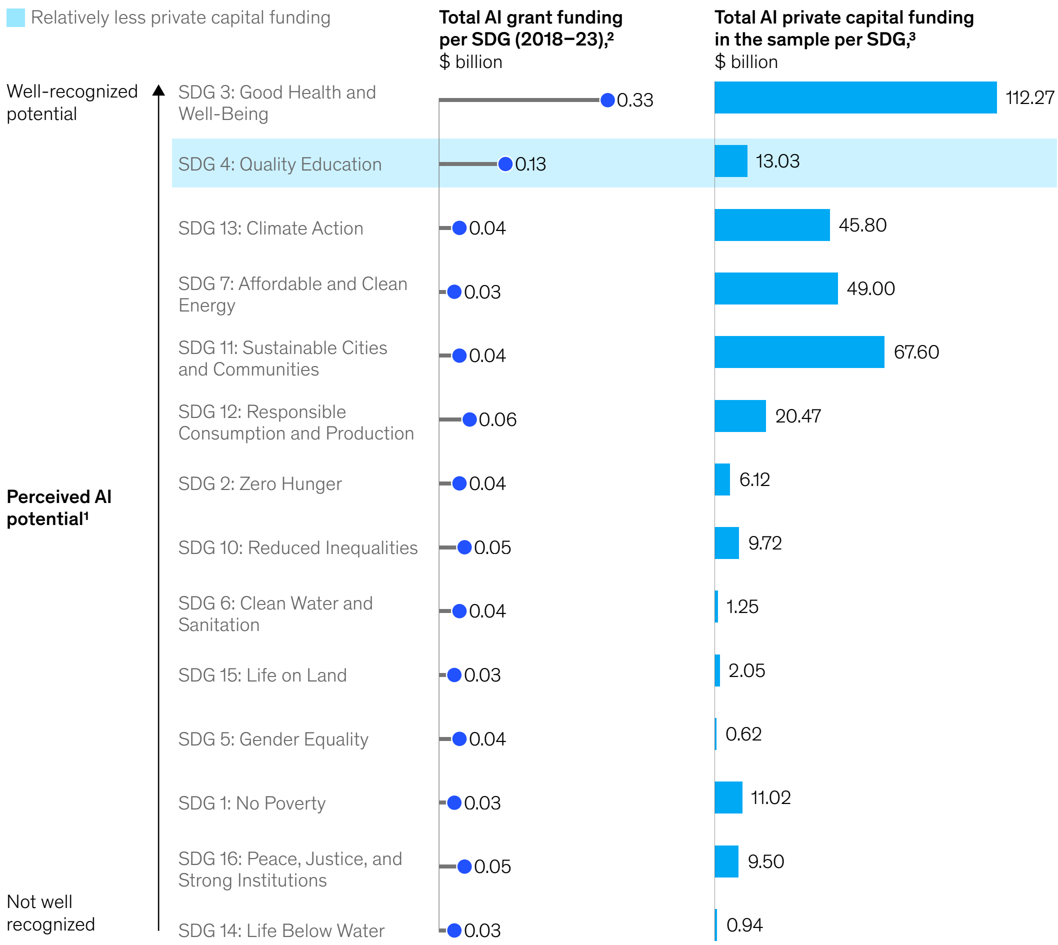

Note: We excluded Decent Work and Economic Growth (SDG 8); Industry, Innovation, and Infrastructure (SDG 9); and Partnerships for the Goals (SDG 17) from our analysis of not-for-profit deployment, grants, and private capital. This is because most projects can be tagged to these areas given broad applicability. This analysis is not holistic; Candid's Foundation Directory is focused on US-based foundations and typically sees a time lag since it is based on IRS data. Grants may be double counted where projects are tagged to multiple SDGs.
[1]Determined through survey of ~60 experts representing 48 organizations (incl not-for-profits, foundations, technology companies, start-ups, academic institutions, and government) and 17 countries in response to the following question: "What are the top 5 SDGs in the list below where you think AI has the highest potential to accelerate progress toward the SDG targets?"
[2]Grants filtered by search terms "AI," "ML," "artificial intelligence," and "machine learning." Foundations outside the US made 200 of the 1,159 grants, accounting for 35% of total funding. Grant making countries include Australia, Brazil, Belgium, Canada, China, India, Norway, Spain, Switzerland, the United Kingdom, and the United States.
[3]Data has not been reviewed by PitchBook analysts.
Methodology: Analysis of approximately 20,000 companies, sourced from PitchBook, that have a clear focus on AI, machine learning, and big data in their products or services and have raised a funding round since 2020. Based on their description, companies were tagged to one or more SDGs only if they addressed any of the relevant thematic areas. The total lifetime funding of each tagged company was used to determine the total funding per SDG. The companies (and their subsequent funding) tagged to more than 1 SDG were counted separately for each.
Source: Foundation Directory, Candid, 2018–23; PitchBook

McKinsey & Company

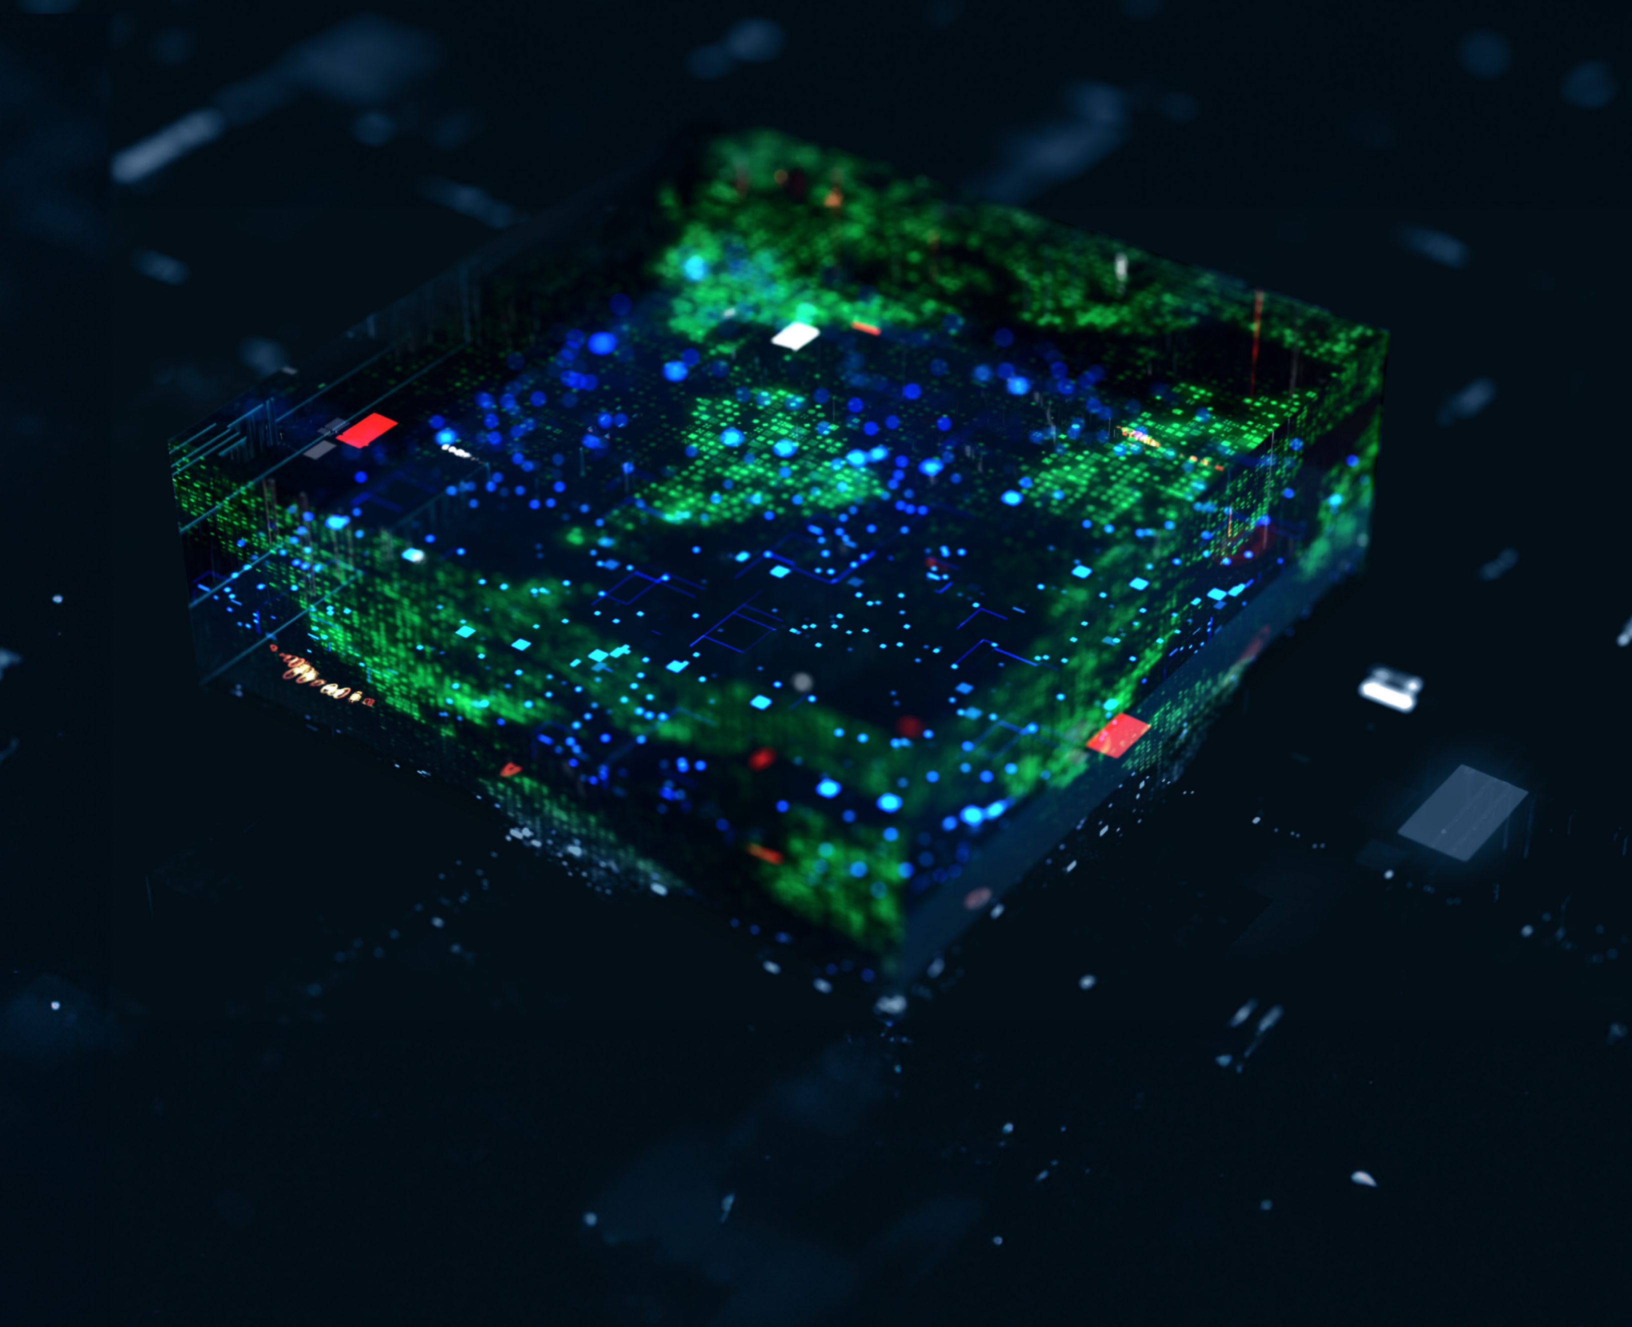

Chapter 3

Challenges of scaling AI for social good

A recent report from Stanford University and Project Evident found that AI already has a considerable presence in the social and education sectors: 48 percent of funders and 66 percent of not-for-profit respondents say their organization uses some type of AI, and about 80 percent of funders and not-for-profits say that their organization would benefit from using more AI.[43] Fully realizing the potential of AI for social good will require AI solutions to be deployed on a larger scale—for instance, moving from regional in-country pilots to cross-country or cross-continent efforts.

According to 72 percent of respondents to our expert survey, most efforts to deploy AI for social good to date have focused on research and innovation rather than adoption and scaling. Fifty-five percent of grants for AI research and deployment across the SDGs are $250,000 or smaller, which is consistent with a focus on targeted research or smaller-scale deployment rather than large-scale expansion (Exhibit 4).

Exhibit 4

Fifty-five percent of foundation grants to support AI research and deployment across the Sustainable Development Goals are $250,000 or less.

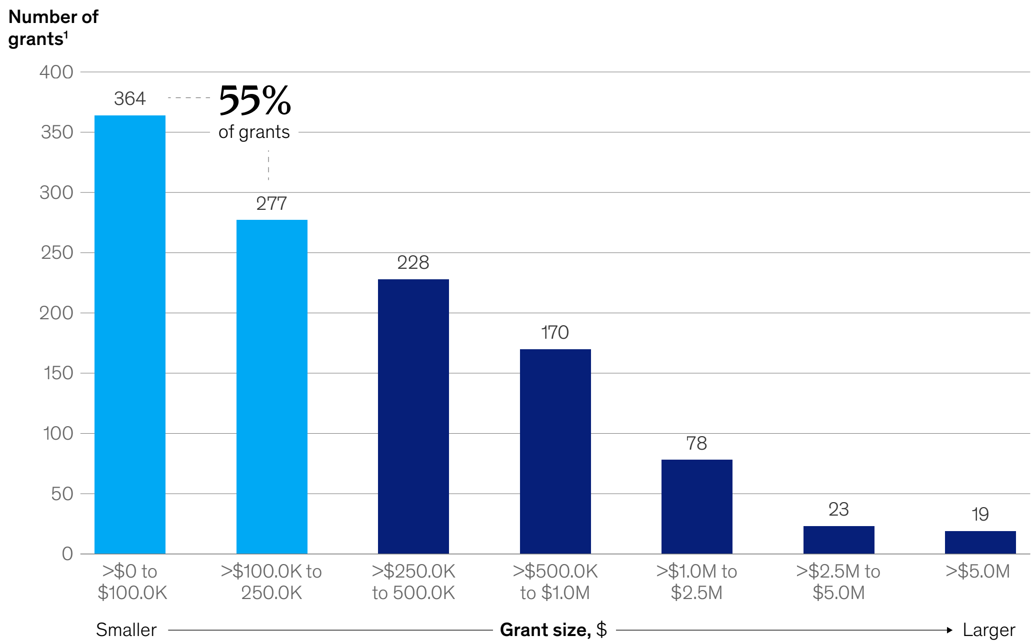

[1]This analysis is not holistic; Candid's Foundation Directory is focused on US-based foundations and typically sees a time lag since it is based on IRS data. Foundations outside the US made 200 of the 1,159 grants, accounting for 35% of total funding. Grant-making countries include Australia, Brazil, Belgium, Canada, China, India, Norway, Spain, Switzerland, the United Kingdom, and the United States.
Source: Foundation Directory, Candid, 2018–23, grants filtered by search terms "AI," "ML," "artificial intelligence," and "machine learning"

McKinsey & Company

[43]Kelly Fitzsimmons et al., *Inspiring action: Identifying the social sector AI opportunity gap*, Stanford Institute for Human-Centered Artificial Intelligence working paper, February 2024.

Organizations that aim to deploy AI in support of the SDGs can face many challenges to scaling. In our 2018 report,[44] we outlined bottlenecks broadly facing the industry to deploying AI for social good. These are largely still the same challenges today. The three factors cited most often by respondents to our expert survey as impeding AI-driven progress toward SDG targets are data availability, accessibility, and quality; AI talent availability and accessibility; and organizational receptiveness and change management (Exhibit 5). While these issues are common to any organization undertaking AI development, not-for-profits and other social good organizations face additional unique challenges. Below, we dig into what's changed in these areas since 2018 and why they are still in many cases the biggest hurdle in deploying AI on a large scale.

Exhibit 5

Experts say the top challenges to scaling AI for social good relate to data, talent, and change management.

Top 5 AI risks that present the largest challenges in deploying AI for achieving the SDG targets,
% of respondents including this topic in their top 5

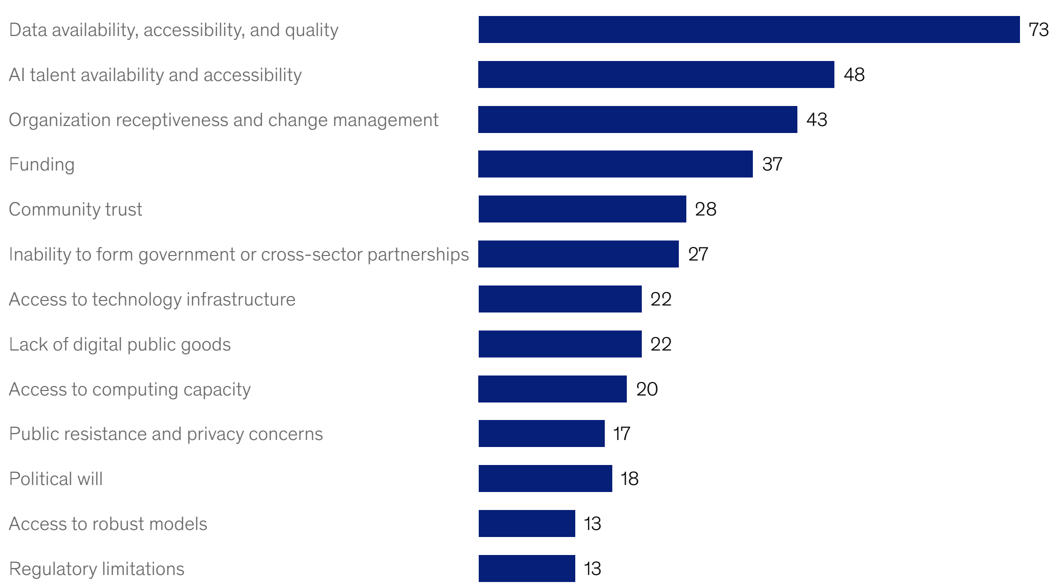

Source: Survey of ~60 experts representing 48 organizations (incl not-for-profits, foundations, technology companies, start-ups, academic institutions, and government) and 17 countries

McKinsey & Company

[44] "Applying artificial intelligence for social good," November 28, 2018.

Data availability, accessibility, and quality
Relevant data sets for SDG-related topics are challenging to create or curate because they require data from less-developed regions, in languages that are less widely spoken, and from more-vulnerable demographic groups (which are the intended beneficiaries of many SDG initiatives). Where rich data sets do exist—such as those owned by telecom, satellite, or financial companies—they may be privately owned or expensive and therefore inaccessible to small organizations. Accessible data can have quality issues: entries may be missing, data may be poorly organized, or data may be in hard-to-access formats such as paper records and spreadsheets.

In short, many social organizations simply do not have sufficient data or must manually clean significant amounts of data, a process that is too resource-intensive for many organizations.

AI talent availability and accessibility
The supply of talent in AI-related fields has increased globally. According to Stanford HAI's *AI Index report 2024*, the penetration of AI skills significantly increased from 2015 to 2023. The countries with the highest penetration rates as of 2023 were Germany, India, and the United States; penetration was 2.2 times higher in the United States than in the rest of the world.[45] Talent is thus very unevenly distributed. As a result, governments, social enterprises, and not-for-profits—especially those in low- and middle-income countries—have limited access to the skills they need to innovate with AI. In addition, there is growing competition from the private sector. For example, applied AI job postings across sectors grew by 6 percent from 2021 to 2022.[46] Furthermore, SDG initiatives may require AI skills in combination with knowledge in other areas, such as sustainable development, specific organizational contexts, or local culture; this combination can be especially hard to come by. In a recent survey of more than 4,000 not-for-profits by Google for Nonprofits, 40 percent of respondents said nobody in their organization was educated in AI.[47]

Organizational receptiveness to AI applications
AI can help frontline workers determine which services or support to provide for vulnerable populations, such as a pregnant mother in a village, a survivor of human trafficking, or an unemployed individual or refugee looking to enter the workforce. But organizations can implement AI tools only if their frontline workers are receptive. Thus, organizations may need to adapt processes to accommodate new ways of working,[48] such as using generative AI (gen AI) tools.

Not-for-profit organizations may have additional reasons to be less receptive to AI applications, including the following:

— *Applications may require time- and resource-intensive training.* Without such training, workers may misunderstand the applications and become frustrated or fail to adopt them.

— *Processes and operating models may need updating.* To integrate AI applications, organizations usually need to simultaneously adjust processes and operating models to capture the benefits of these applications. These changes may be challenging and time-consuming.

— *Explainability may need to be treated as a priority from the start.* Many AI applications rely on recommendations or predictions, and explainability of the AI—that is, the ability to explain model outputs appropriately and accurately—is key to getting users to trust and engage with the applications.

[45] *Artificial Intelligence Index report 2024*, Stanford Institute for Human-Centered Artificial Intelligence, 2024.
[46] Michael Chui, Mena Issler, Roger Roberts, and Lareina Yee, *Technology Trends Outlook 2023*, McKinsey, July 20, 2023.
[47] "3 insights from nonprofits," March 28, 2024.
[48] Eric Lamarre, Alex Singla, Alexander Sukharevsky, and Rodney Zemmel, "A generative AI reset: Rewiring to turn potential into value in 2024," *McKinsey Quarterly*, March 4, 2024.

— *Organizations may have concerns about AI-related risks to vulnerable people.* As discussed above, risks such as bias, malicious use, and privacy concerns can discourage organizations from using AI. These risks can have particularly negative impacts on the vulnerable people not-for-profits work with, such as pregnant mothers and survivors of human trafficking.

— *Widespread alignment may be necessary.* Implementation at scale may require collaboration and adoption across many stakeholders, including funders, governments, frontline organizations, and tech companies. Any one organization that is less receptive can impede the progress of the whole stakeholder group.

The emergence of simpler gen AI tools offers hope. Gen AI tools are easy to use, and they make it easier to develop other software, which may expand AI use among organizations with limited data science capabilities. Gen AI is also expected to unlock value from open data by giving low-resource organizations, such as some not-for-profits or organizations in low-income countries, access to improved user experiences.[49] About 70 percent of respondents to our experts survey believe that gen AI will have a net positive impact on SDG progress by enhancing productivity, personalization, and language translation; democratizing access to knowledge and its dissemination; and increasing awareness that can potentially lead to policy change. However, these benefits are yet to be realized, and much awareness building and upskilling work remains ahead. In a Google for Nonprofits survey of more than 4,000 not-for-profits, 50 percent of respondents said they plan to provide gen AI training to the communities they serve.[50]

Additional barriers to scale

In addition to the three challenges outlined above, many other obstacles could hinder progress. For example, many people still have limited access to the internet.[51] Although compute costs are lower today than in 2018, they are still unaffordable for many not-for-profits, and while free tools are emerging to build solutions, the high costs to run recent gen AI models at scale may widen the gap. Grants and political will are also required to create incentives for widespread adoption of many applications.

[49] For a deeper look at open data, see “Open data: Unlocking innovation and performance with liquid information,” McKinsey Global Institute, October 1, 2013.
[50] *Nonprofits and generative AI*, Google, accessed April 25, 2024.
[51] *Measuring digital development: Facts and figures 2022*, ITU, 2022.

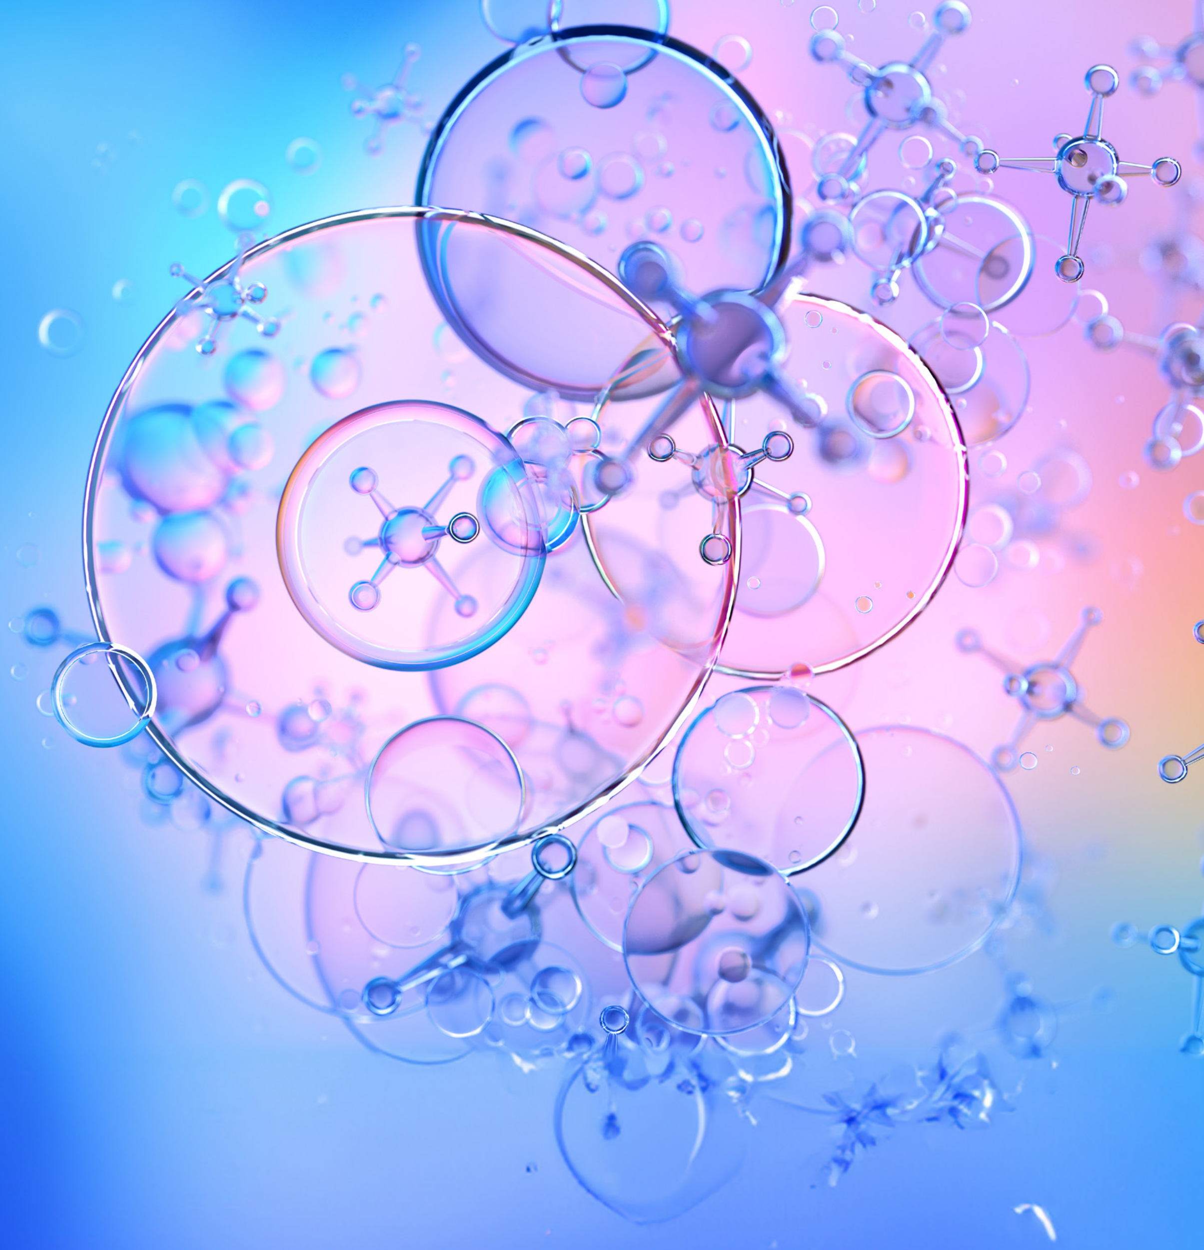

Chapter 4

How stakeholders can accelerate the deployment of AI for social good

In this section, we outline six complementary, cross-cutting approaches that have the potential to mitigate one or more of the challenges above. These actions are not mapped one-to-one to the scaling challenges outlined above, but each could help address at least one challenge. Mission-driven organizations, governments, foundations, universities, ecosystems of developers, businesses, and other stakeholders can implement these approaches to support the use of AI for social good.

1. Forming partnerships to accelerate impact

Partnerships can fulfill many needs of mission-driven organizations and alleviate many of the challenges discussed above. AI applications may be developed by researchers or tech companies but require a frontline nongovernmental organization (NGO) to adopt them for impact; partnerships can bring these groups together to get solutions to the front line faster. Partnerships can also supplement data science or engineering talent and allow groups to share or develop data; codevelop or gain access to critical infrastructure, models, and applications; and promote awareness and trust through larger stakeholder communities.

Some stakeholders have already explored partnerships that provide expanded reach, data, and resources.

Partnering for reach. Khushi Baby's digital health platforms have reached more than 17,000 medical centers in India through increasingly large partnerships with village and state governments. These partnerships have proved crucial to reaching local patients.[52]

Partnering for data. AirQo, a not-for-profit in Uganda, partners with city environment management authorities to measure and build data sets for air quality, enabling the authorities to formulate and implement actions to improve air quality.[53]

Partnering for resources. PATH develops solutions to address pressing health issues. It has partnered with Amazon Web Services for computing credits and with Microsoft's AI for Good team to identify new opportunities to use AI for public health.[54]

Both public and private sector organizations can build innovation ecosystems that bring together stakeholders, generate ideas for AI solutions that target existing issues, and reduce barriers to impact.

Disha, a UN-led initiative, brings together foundations, tech and telecom companies, AI ethics centers, academic institutions, and data providers to jointly build responsible AI solutions. Initial products include a solution built in partnership with a telecom provider to serve multiple NGOs in Asia on disaster resilience and recovery efforts by estimating people movement and poverty incidence using mobile data. Coalitions such as Disha can enable the deployment of more robust, standardized, and shared AI products to support millions of people globally.

[52] "Maternal & Newborn Health: Khushi AI," MIT Solve, accessed April 25, 2024.
[53] "AirQo," Climate and Clean Air Coalition, UN Environment Programme, accessed April 25, 2024.
[54] "PATH – AWS Imagine Grant winner," YouTube, December 9, 2019; Juan M. Lavista, Brian Taliesin, William B. Weeks, "Using artificial intelligence to advance public health," *International Journal of Public Health*, 2023, Volume 68.

2. Supporting the development of digital public goods

Digital public goods can include models, software, standards, content, and data intended to contribute to sustainable development. Their availability can also simplify the development of novel AI for SDG solutions, alleviating challenges related to talent, funding, and compute.

Organizations that produce SDG-relevant data sets and models may be able to make them accessible—or partially accessible—for clearly defined public interest initiatives. Existing AI for SDG initiatives have demonstrated that data collected by governments, telecom companies, technology companies, utilities providers, and content creators, as well as AI models themselves, can be harnessed for sustainable-development efforts. For example, Google DeepMind's GraphCast is a ten-day global weather forecasting system that can be used to predict extreme weather events and prevent adverse outcomes for millions of people. The underlying open-source model is being used by weather agencies such as the European Centre for Medium-Range Weather Forecasts.[55]

Governments, technology companies, and funders can sponsor efforts to organize and curate existing digital public goods, reducing technical barriers to developing AI solutions for SDGs. Open-data initiatives have begun to provide large-scale improvements to availability and accessibility by improving training data, sharing data sets, and improving model performance. One such initiative is Data4SDGs, a network of more than 700 organizations that has brokered more than 100 data partnerships to support goals related to climate, health, social inclusion, education, and more.[56] In another example, Data Commons provides access to data sets from publicly available, reliable sources such as the United Nations' Intergovernmental Panel on Climate Change.[57] Users can ask questions in their own language and get immediate responses, accompanied by visuals.

However, initiatives to share data, code, and best practices remain underdeveloped. One researcher who focuses on data science for social good told us, "The coordination is poor, and we are not learning from each other. We are not sharing best practices. As a result, solutions are being invented from scratch in different parts of the world, leading to a waste of resources."

While open-source communities do exist—and are tracked by the Digital Public Goods Alliance[58]—there is a large gap in ready-to-customize AI assets and tools across most SDGs. It may be beneficial for those who consume these goods to create communities around them and build on one another's work.

[55] "GraphCast," Google DeepMind, 2023; *Google DeepMind*, "GraphCast: AI model for faster and more accurate global weather forecasting," blog entry by Remi Lam, November 14, 2023.
[56] "Our impact," Global Partnership for Sustainable Development Data, accessed April 25, 2024.
[57] Muhammad J. Amjad et al., *Data Commons*, September 12, 2023.
[58] "Registry," Digital Public Goods Alliance, accessed April 25, 2024.

3. Bolstering quality and usability of data

Stakeholders could improve the scalability of data-driven AI for SDG applications, while reducing risks and improving output quality, by creating high-quality data sets (either publicly available or through select access management). It is particularly critical to support data collection in resource-poor contexts and for disadvantaged populations, and to create incentives for organizations with rich data sets to make them available for social good applications.

Not-for-profit organizations that do not have—or cannot access—the data they need for a particular purpose may benefit from getting creative. For example, Rainforest Connection's Guardian Platform uses solar-powered acoustic streaming devices—connected via Global System for Mobile Communication (GSM)—to gather a continuous recording of forest soundscapes. This audio record is then transmitted to the cloud, where AI tools analyze it in real time to detect illegal deforestation activities, enabling immediate intervention.[59] Rainforest Connection has worked with more than 100 NGOs and local conservation agencies to deploy sensors in more than 35 countries. In addition, tracking "exhaust data," such as the location and navigation route of maritime vessels, could reveal unregulated and illegal fishing activities.

4. Expanding the pool of AI for SDG talent

Increasing the AI for SDG talent pool will require both near-term and longer-term initiatives. In the near term, academic institutions or companies with technical talent can make their own talent available to low-resource organizations on a part- or full-time basis. In the longer term, investments in training or relevant degrees can help grow the talent base, particularly in underrepresented regions or within vulnerable communities.

For example, several large companies—including Google, Microsoft, and Two Sigma—have used their data science talent for social endeavors by loaning or seconding talent to other organizations or allocating time for them to work on social good projects.[60] DataKind, a global not-for-profit, works to make these loaning efforts possible by pairing global data science talent with projects.[61] Additionally, Data.org launched the Capacity Accelerator Network (CAN), with the goal of training one million data practitioners around the world by 2032.[62]

Many of the experts we spoke to emphasized that AI talent needs to have a deep understanding of SDG fields or target regions. Not-for-profit educational initiatives have begun expanding the pool of talent in social good fields. In addition, the Google DeepMind scholarship program provides financial support to students from underrepresented groups pursuing graduate studies in AI-related and adjacent fields. The program is available at universities around the world, including partner institutions in Africa.[63]

[59] "Guardian Platform," Rainforest Connection, accessed April 25, 2024.
[60] "Our work," Google.org, accessed April 25, 2024; "Expand opportunity," Microsoft, accessed April 30, 2024; "Data for good: A corporate perspective," Two Sigma, accessed April 30, 2024.
[61] "Our approach," DataKind, accessed April 25, 2024.
[62] "Capacity Accelerator Network (CAN)," Data.org, accessed May 10, 2024.
[63] "About: Education," Google DeepMind, accessed April 25, 2024.

5. Taking an inclusive, user-centric approach
AI solutions can be developed before an application has been tested with the end user. Human-centric AI is a process by which the user experience is developed in tandem with AI solutions, with frequent user testing including during early deployment. This approach addresses challenges related to low organizational receptiveness and the need for change management to adapt to AI, specifically gen AI. Incorporating users from the beginning can build trust with communities and help ensure users are receptive to solutions. It also helps organizations avoid spending development time and costs on solutions that do not fully meet user needs.

For example, Rainforest Connection started because forest rangers were spending more time fighting loggers than restoring biodiversity. So the organization recycled cell phones into autonomous, solar-powered listening devices that could remotely monitor and detect logging activity in rainforests. The devices, called Guardians, are hidden high in trees for better cell service and access to sunlight for power. They use machine learning models to detect logging sounds, such as chain saws and trucks, around the clock. While the solution worked in theory, the need to place the devices high in trees was a barrier to installation. The organization worked with biodiversity scientists to enhance adoption by helping local community members install the devices. Today, Rainforest Connection monitors more than one million acres in 106 protected reserves in 35 countries.[64]

6. Creating a 'business model' (where applicable)
Fundraising can be a barrier for organizations, particularly when it comes to the increased funding required to move from initial pilots to at-scale solutions. Respondents to our survey of experts identified funding as the fourth-most-critical challenge to scaling AI for SDG impact. Some not-for-profit organizations have created business models that supplement or offset fundraising requirements by establishing continual-revenue streams—either through direct technology or by providing platforms for local entrepreneurship.[65]

For example, Thorn, a not-for-profit that builds tools to defend children from sexual abuse, offers both pro bono and fee-based services to a mix of not-for-profits, government agencies, and technology platforms.[66]

Scientific breakthroughs have increased the effectiveness of AI at pattern recognition, prediction, and creation. This progress has coincided with a rapidly growing number of successful AI deployments, but there are still challenges to scaling their use for addressing the SDGs. Realizing this potential will require stakeholders to collaborate more closely to ensure access to adequate talent, robust data solutions, and AI applications and models—and to ensure these models are more open-sourced or scalable across user geographies around the world and can therefore meet people at the point of need.

By collaborating to find ways to put AI to work at scale for social good, mission-driven organizations, governments, foundations, universities, ecosystems of developers, and businesses can help solve some of the world's most challenging and intractable problems. They can help thwart human trafficking, ensure girls and children all over the world receive the education they deserve, protect forests from illegal deforestation, support the health and safety of pregnant women and newborns, and so much more. If these things aren't worth fighting for, what is?

[64] Interview with Rainforest Connection leadership, August 2023.
[65] Many opportunities in social good focus on addressing market failures where there might be fewer and less-compelling incentives to build for-profit business models.
[66] Interview with Thorn leadership, August 2023.

Methodology

We set out to understand AI's potential impact on the UN Sustainable Development Goals (SDGs) and assess how private and charitable and public capital is flowing toward the SDGs with the highest potential for AI impact. We collected several kinds of data to do so:

Use case library and not-for-profit deployments. We created a library of 600 use cases in which AI could be used to advance the SDGs and more than 1,000 not-for-profit deployments of AI toward these goals. To build this data set, we used a range of publicly available sources made available by not-for-profits, foundations funding their activities, and the United Nations.

Survey and interviews. We surveyed about 60 experts from 17 countries and 48 organizations (including not-for-profits, foundations, technology companies, start-ups, academic institutions, and government) to learn their views related to SDGs where AI has the highest potential, AI risks, and barriers to scaling. Additionally, we interviewed more than 50 executives and senior leaders from not-for-profits, tech companies, and foundations globally.

Foundation grants. We analyzed more than 1,000 foundation grants filtered by the search terms "AI," "ML," "artificial intelligence," and "machine learning" from 2018 to 2023.[1] Grants went to recipients in countries including Australia, Brazil, Belgium, Canada, China, India, Norway, Spain, Switzerland, the United Kingdom, and the United States.

Analysis of private capital investment. Based on data from PitchBook, we analyzed approximately 20,000 companies that have a clear focus on AI, machine learning, and big data and have gone through a funding round since 2020. Based on their descriptions, companies were tagged to one or more SDGs. The total lifetime funding of all tagged companies was used to determine the total funding per SDG. Because funding can be applicable to multiple SDGs, when the value of SDG funding is calculated, it exceeds the total amount invested when considered without SDG tagging. Roughly 35 percent of companies were headquartered in the United States.

[1] Analysis focuses on grants from mostly US-based foundations, and there is typically a time lag since the data is based on IRS data. But the analysis is representative of distribution across SDGs as well as scale of funding. See "AI for Sustainable Development Goals," Candid Foundation Directory, 2018–2023.

Acknowledgments

Medha Bankhwal is a consultant in McKinsey's San Francisco office, where **Michael Chui** is a McKinsey Global Institute partner; **Ankit Bisht** is a partner in the Dubai office; **Roger Roberts** is a partner in the Silicon Valley office; and **Ashley van Heteren** is a partner in the Amsterdam office.

The authors wish to thank the following McKinsey colleagues for their contributions to this report: Ibrahim Almubarak, Oliver Bevan, Damien Bruce, Jake Bryant, Enrico Cremonese, Ayush Dharap, Doria Du, Ankit Fadia, William Frohlich, Liz Grennan, Martin Harrysson, Eric Hazan, Martin Hirt, Rasmus Kastoft-Christensen, Kelly Kochanski, Jessica Lamb, Matej Macak, Clara Marquardt, Christopher McGrillen, Mark Patel, Saurabh Sanghvi, Daan Schellekens, Aakanksha Srinivasan, Eline van Es, Wilbur Wang, Hannah Weilert, Olivia White, Lola Woetzel, and Peter Zipparo.

This independent initiative is based on our own research, the experience of our McKinsey colleagues more broadly, and McKinsey's research collaboration with Google.org and AI researchers at organizations including Google DeepMind and Google Research. Many others informed our research. We wish to thank: Caitlin Augustin and Jake Porway, DataKind; Engineer Bainomugisha, AirQo; Nicole Dunn and Shannon Farley, Fast Forward; Nick Cain and Vilas Dhar, Patrick J. McGovern Foundation; Andrew Dudfield, Full Fact; RV Guha and Prem Ramaswami, Data Commons; Joe Huston, GiveDirectly; Zia Khan, the Rockefeller Foundation; Rebecca Portnoff, Thorn.org; Heejae Lim, TalkingPoints; Neha Madhiwala, ARMMAN; Gavin McCormick, Climate TRACE and WattTime; Danil Mikhailov, Data.org; Yongyi Min, UN Stats; Aditya Nayan and Shekar Sivasubramanian, Wadhwani AI; Bourhan Yassin, Rainforest Connection; and Leid Zejnilovic, Data Science for Social Good Foundation.

Noble Intelligence is McKinsey's social responsibility initiative to apply our data and AI capabilities toward social impact initiatives globally.

McKinsey & Company
May 2024
Copyright © McKinsey & Company

www.McKinsey.com

@McKinsey
@McKinsey
@McKinsey