


AI for social good

How organizations can use artificial intelligence to build a better world

Contents

AI for social good can take many forms	4
In practice	7
A framework for action	10
Conclusion	19
Get in touch	19



Introduction


Artificial intelligence (AI), machine learning, and related technologies are swiftly and increasingly transforming the ways we live and work. Every day, organizations deploy powerful algorithms to generate important insights from vast amounts of data. This new understanding can then be used to drive better decision-making, unlock competitive advantage, increase efficiency, improve delivery of health care and other social services, and more. The possible applications and benefits of AI—for businesses, governments, and other organizations—are enormous. So how can AI be used to meet its potential to be a force for social good?

Often, the excitement and optimism about the use of AI and machine learning are tempered by skepticism and by wariness of the impact and/or misuse of those same technologies. Concerns about AI causing deepened economic and social inequalities should be taken seriously—as should claims that it will primarily benefit the technology and financial sectors.

Deloitte believes that AI can be—and more importantly, will be—a powerful force for social good. In fact, we're already using AI and related technologies to tackle some of the most significant problems facing society. And we're not alone.

In this report, we explore what it means to use AI to benefit society and we share insights drawn from our own experiences, as well as those of others in business, academia, and the not-for-profit sector. We hope to spark conversation within and between organizations about working together to harness AI and data to make a meaningful difference to society—and build a better world for all.

AI for social good can take many forms



There are numerous ways that organizations can use *AI for social good*: it's a flexible, multifaceted concept, and its execution is influenced and shaped by an organization's strategy, ambitions, capabilities, and resources. To understand AI for social good as an idea, we should begin by clarifying what we mean when we talk about AI and social good separately.

What do we mean by AI?


In popular usage, *AI* is a catch-all term for a number of related technologies, so it's worth defining the concepts used in this paper more precisely. AI describes computer systems that use data and algorithms to mimic or extend two categories of human capabilities: reasoning (including the ability to predict, recommend, optimize, and generate insights) and ability to interact with the world around us (this includes text, image, sound, and voice). *Machine learning* and *deep learning* are sometimes used as synonyms for AI, but it's more accurate to describe them as subsets of AI.

While "traditionally" AI has focused on creating machines that think and work like humans do and on automating tasks formerly done by humans, *augmented intelligence* takes a new approach. Augmented intelligence uses AI technologies to help people perform their work, by providing them with actionable, data-driven insights that support better decisions. Many *AI for social good* use cases are applications of augmented intelligence: AI augments humans with insights, allowing organizations and individuals to take actions that improve environmental, health, or other socially positive outcomes.

AI Defined Artificial Intelligence

Human functions that AI can mimic or extend

Reasoning	Interaction
Generate insights	Text
Predict	Image and video
Recommend	Sound and voice
Optimize	



Machine Learning

Algorithms that improve over time through exposure to more data

Deep Learning

A subset of Machine Learning that uses neural networks with massive amounts of data to learn



Artificial Narrow Intelligence (ANI)

An AI that can match human capabilities and ways of thinking in specified domains (e.g., chess, image recognition, speech recognition, optimization, scheduling)

Artificial General Intelligence (AGI)

An AI that matches human capabilities and ways of thinking in any domain

What do we mean by social good?


Sometimes defined as the common good, *social good* typically describes an item or action that benefits the general public, or at least a significant number of people in a particular community.

In this regard, the United Nations's Sustainable Development Goals (UNSDGs) provide a touchstone for many organizations. These 17 objectives are a core part of the 2030 Agenda for Sustainable Development, a blueprint for global peace and prosperity adopted by all UN member states in 2015. The UNSDGs include not only ending poverty and hunger, but also improving health and education, reducing inequality, creating economic opportunity, protecting our ecosystems, and tackling climate change (see below).

United Nations Sustainable Development Goals

1. End poverty in all its forms everywhere.
2. End hunger, achieve food security and improved nutrition, and promote sustainable agriculture.
3. Ensure health lives and promote well-being for all at all ages.
4. Ensure inclusive and equitable education and promote lifelong learning opportunities for all.
5. Achieve gender equality and empower all women and girls.
6. Ensure availability and sustainable management of water and sanitation for all.
7. Ensure access to affordable, reliable, sustainable and modern energy for all.
8. Promote sustained, inclusive and sustainable economic growth, full and productive employment, and decent work for all.
9. Build resilient infrastructure, promote inclusive and sustainable industrialization, and foster innovation.
10. Reduce inequality within and among countries.
11. Make cities and human settlements inclusive, safe, resilient and sustainable.
12. Ensure sustainable consumption and production patterns.
13. Take urgent action to combat climate change and its impacts.
14. Conserve and sustainably use the oceans, seas and marine resources for sustainable development.
15. Protect, restore and promote sustainable use of terrestrial ecosystems, sustainably manage forests, combat desertification, and halt and reverse land degradation and halt biodiversity loss.
16. Promote peaceful and inclusive societies for sustainable development, provide access to justice for all, and build effective, accountable and inclusive institutions at all levels.
17. Strengthen the means of implementation and revitalize the global partnership for sustainable development.

Source: United Nations Department of Economic and Social Affairs, <https://sdgs.un.org/goals>



What AI for social good can be

This initiative can involve a range of related technologies and address a host of potential social issues. One organization may deploy its AI resources to search for solutions to some of the biggest challenges facing the world, while another may concentrate on providing insights that can lead to a consequential difference for a local community's most vulnerable people. New applications appear almost daily and vary significantly: AI models exist for predicting septic shock, preventing animal poaching, supporting student retention, and more.¹ There are countless opportunities for organizations to use AI and data to improve societal outcomes and make our communities, and even our world, healthier, cleaner, and more prosperous. In fact, AI can enable socially good outcomes that at one point might have been considered unfeasible, unaffordable, or virtually impossible.²

Discussions about the use of AI for social good often overlap with similar talks about its ethical use. Governments, industry groups, corporations, academics, and organizations such as the Partnership on AI and the Institute of Electrical and Electronics Engineers (IEEE)—both concerned with the future of AI and other technology that can benefit humanity—have proposed various principles or codes to promote the ethical use of AI. There is substantial agreement across bodies regarding these principles—including a widespread acceptance that AI and related technologies should be used for the common good, respect such values as fairness and privacy, and avoid harming people and otherwise undermining their rights.³ One useful synthesis of the various concepts incorporates ideas adapted from biomedical ethics.

In each instance, ethical AI use—including *AI for social good* use cases—should reflect the following:

- Beneficence: Will a particular AI use case bring about good?
- Non-maleficence: Will it avoid harm?
- Justice: Does it consider people fairly and impartially?
- Autonomy: Have the people affected given informed consent to participate?
- Explicability: Can we explain how the AI use case works—and who is accountable?⁴

In 2020, Google introduced its set of AI principles, which are largely consistent with these considerations. What's especially interesting is how the organization's doctrine explicitly prohibits a number of socially detrimental AI use cases: Google has pledged not to design or deploy AI for technologies that cause or are likely to cause overall harm, that are principally meant to cause or facilitate injury (e.g., weapons), that gather or use information for surveillance in violation of international norms, and whose purpose contravenes accepted principles of international law and human rights.⁵ By taking steps to prevent the use of AI for social harm, Google is working to ensure that all its AI use is, in some way, for social good.

For private-sector organizations in particular, AI for social good initiatives often fall under the umbrella of broader mandates to give back to the given companies' communities. This is relevant because most AI is being built in the private sector, observes Gillian Hadfield, who holds the Schwartz Reisman Chair in Technology and Society at the University of Toronto; much less is being built in the public and not-for-profit spaces. "When I think about AI for social good, I'm thinking about the ways it's focusing on areas where markets may not effectively drive investments in the development of social good," she says.

Peter Husar, VP of enterprise analytics strategy and planning at TD Bank Group, says, "Our AI initiatives obviously look at the impact on the bottom line. But when it comes to AI for social good, the profit aspect isn't considered—at least, not to the same extent. It's about the impact to society."

Profit may not be a driving factor in most AI for social good use cases, but that doesn't mean those same use cases can't deliver important financial or otherwise economic returns. For example, those projects can enable organizations to streamline processes, accelerate work, and better allocate scarce resources, reducing time and costs for relevant tasks and thus preserving budgets for investment in other priorities.

1. Luciano Floridi, Josh Cowls, et al. "How to Design AI for Social Good: Seven Essential Factors." *Science and Engineering Ethics* (2020) 26:1771–1796. <https://doi.org/10.1007/s11948-020-00213-5>

2. Ibid.

3. Jess Whittlestone, Rune Nyrop, et al. "Ethical and societal implications of algorithms, data, and artificial intelligence: a roadmap for research." Leverhulme Centre for the Future of Intelligence, University of Cambridge.

4. Nicolas Berberich, Toyooki Nishida, Shoko Suzuki, "Harmonizing Artificial Intelligence for Social Good." *Philosophy & Technology* (2020) 33:613–638. <https://doi.org/10.1007/s13347-020-00421-8>

5. Google, <https://ai.google/static/documents/ai-principles-2020-progress-update.pdf>

In practice

The Deloitte global network's commitment to driving societal change and promoting environmental sustainability is central to our corporate responsibility and sustainability efforts. By working in innovative ways with governments, non-profit organizations, and civil society, we're designing and delivering solutions that aim to contribute to a sustainable and prosperous future for all.

Deloitte Global's [WorldClass](#) initiative refines and extends these efforts, translating our aim to make an impact that matters into clear, measurable terms. Globally, millions of people are being left behind, unable to fulfill their aspirations and potential in the Fourth Industrial Revolution. At Deloitte, we believe that we make the greatest societal impact when our professionals use their knowledge and experience to help these and other people develop job skills, improve educational outcomes, and access opportunities to succeed in this new economy. Our WorldClass ambition is to support 100 million people in developing job skills, improving educational outcomes, and accessing opportunities to succeed by 2030. Aligning with others toward a common goal enables us to focus our resources and efforts to make a meaningful impact worldwide. One of the ways we're working is to achieve this goal in Canada is through our *AI for social good* projects. Across the country, we're using our AI technologies and the skills and experience of our people with our clients to foster improvements.

Large Agency of the Federal Government

An agency of the federal government invests billions worldwide in social initiatives, including children's education and women's programs. The agency had long used a manual system to record and track its spending on these programs, as well as to store the relevant data. They are now working with us to automate and streamline these processes using AI so that it can measure and report on its global social impact more quickly and accurately than ever before. This, in turn, should create more transparency and trust in how our tax dollars are being used.

Canadian Provincial Government

Using AI technologies, Deloitte teams have been working extensively with a provincial government in Canada to enable improvements in service levels and efficiency across a wide range of areas.

One improvement involves animal welfare. We used AI to analyze the migration patterns of animals such as moose and deer to reduce vehicle-animal collisions on the province's roads and highways. In addition to helping protect animals, this project also delivers important social and economic benefits to the residents: mitigating potential animal collisions helps to keep drivers safe and to reduce the need for costly repairs.

Another AI-related project involves analyzing population and demographic data to identify optimum locations for new schools in the province. While at first glance it may seem like a rather boring use of the technology, in fact it delivers real value as it helps to ensure the government's capital expenditures are optimized financially as well as socially.

We're also using AI to help this provincial government optimize its use of resources to keep the populace safe and minimize potential damage from climate-related disasters. Applying algorithms to data about changes in water levels in the province's lakes, rivers, and reservoirs allows for more accurate prediction of the likelihood and location of flooding, enabling authorities to take appropriate precautions.



Wildfire suppression is highly data-intensive: much information and analysis is required to schedule, manage, and direct appropriate firefighting resources. Still, the challenge of predicting fires, including their possible escalation, often means increased costs (e.g., from overtime pay), inefficient use of poorly allocated resources, and risk for greater fire damage when predictions fall short. To help, a Deloitte team built an AI model that better predicts realistic fire risk and behaviour based on likely short-term environmental conditions. A better understanding of the likely number, location, and timing of fires in a given time window enables provincial authorities to plan more effective responses, including by more accurately assessing the number of fire-suppression teams and the equipment likely to be required. Better fire prediction means improved resource utilization, which translates to potential savings of approximately \$5 million.

Deloitte's data-driven decarbonization tool

Companies around the world are looking for ways to meaningfully reduce their carbon footprints and do their part to combat climate change. However, many find it challenging to determine which emissions-reducing opportunities to pursue. Enter Deloitte's decarbonization tool.

Hundreds, if not thousands, of organizations have collected, analyzed, and supplied data that models potential impacts of climate change. We have incorporated into our decarbonization tool a number of these third-party data sets and methodologies—in particular, data about representative concentration pathways, which describes several scenarios that would allow the world to reach the common goals of net-zero emissions and/or to limit global temperature increases to 1.5°C by the end of the century. Aggregating these scenarios using the tool can allow organizations to measure the impact of their carbon-cutting efforts and compare these results to science-based targets. Our aim is to help companies do this by assessing their technical and financial parameters, as well as implications of their decarbonization options.

We have also been working with a major global retailer that's currently exploring a variety of decarbonization opportunities, from adopting electric vehicles and renewable-energy supplies to improving the energy efficiency of its facilities and other operations. Using our tool, we've been evaluating the technical and financial characteristics of decarbonization to understand their impact on the company and its emissions. We were then able to use our findings to conduct portfolio analyses to determine the optimal way the company can fund—and sequence—its pathway to net-zero emissions by its 2040 target.



Colombia's COVID-19 dashboard

Early in the COVID-19 pandemic, the Government of Colombia asked for the World Bank's help in developing a strategy to manage the health impact of the virus while minimizing the consequences of measures such as lockdowns on employment and productivity. Authorities wanted to know when, where, and how it would be safe to reopen—but they would first have to overcome the challenge of integrating more than two dozen fragmented data sets containing vital health and economic information.

The World Bank teamed up with academic partners at the University of Toronto's Dalla Lana School of Public Health and professionals and members of Deloitte Canada's AI practice to design the COVID-19 Safe Economic Reactivation Dashboard. This tool helps policymakers in 1,100 municipalities across Colombia monitor—in real time—virus spread, health-sector readiness and capacity, and the preparedness of 29 different economic sectors to operate safely at every stage of the pandemic.

Colombia's president unveiled the dashboard on national television in July 2020; by August 2021, the tool had more than 10,000 unique users, with more than 400 mayors, governors, and national policymakers trained on how to integrate dashboard insights into their decision-making. A mobile-friendly version of the dashboard is now available to the wider public, with other countries, such as Peru, having adapted the tool for their own needs.

BC Children's Hospital Foundation

This charitable organization devotes enormous amounts of time and energy to fundraising throughout the year via means such as marketing campaigns and special events. The foundation surmised that if it could gain sharper insights into donor demographics and activity—and combine these with economic-trends data—its fundraisers could better understand how to effectively turn occasional supporters into regular donors, increase overall donations, and reduce donor drop-off.


A Deloitte team offered pro bono assistance as part of our commitment to giving back to communities. With our support, the foundation is developing a process and algorithm that could analyze its trove of fundraising data and provide actionable insights into the demographic, timing, and economic factors influencing donor behaviour—and how the foundation could convert and retain donors overall.

AI Procurement in a Box

The World Economic Forum (WEF) notes that safeguarding the responsible use of AI technologies is a multi-stakeholder effort.

To help governments unlock the potential of AI in the public sector, the WEF's Centre for the Fourth Industrial Revolution—in collaboration with the UK government, Deloitte, and the US software company Splunk—created AI Procurement in a Box. This practical, multi-faceted initiative is intended to help policymakers and commercial teams rethink their approaches to AI procurement to more effectively and ethically adopt AI technologies in the public sector.

Since its launch in 2017, the project has been scaled to the United Arab Emirates, Bahrain, and South America (Brazil and Chile).




A framework for action

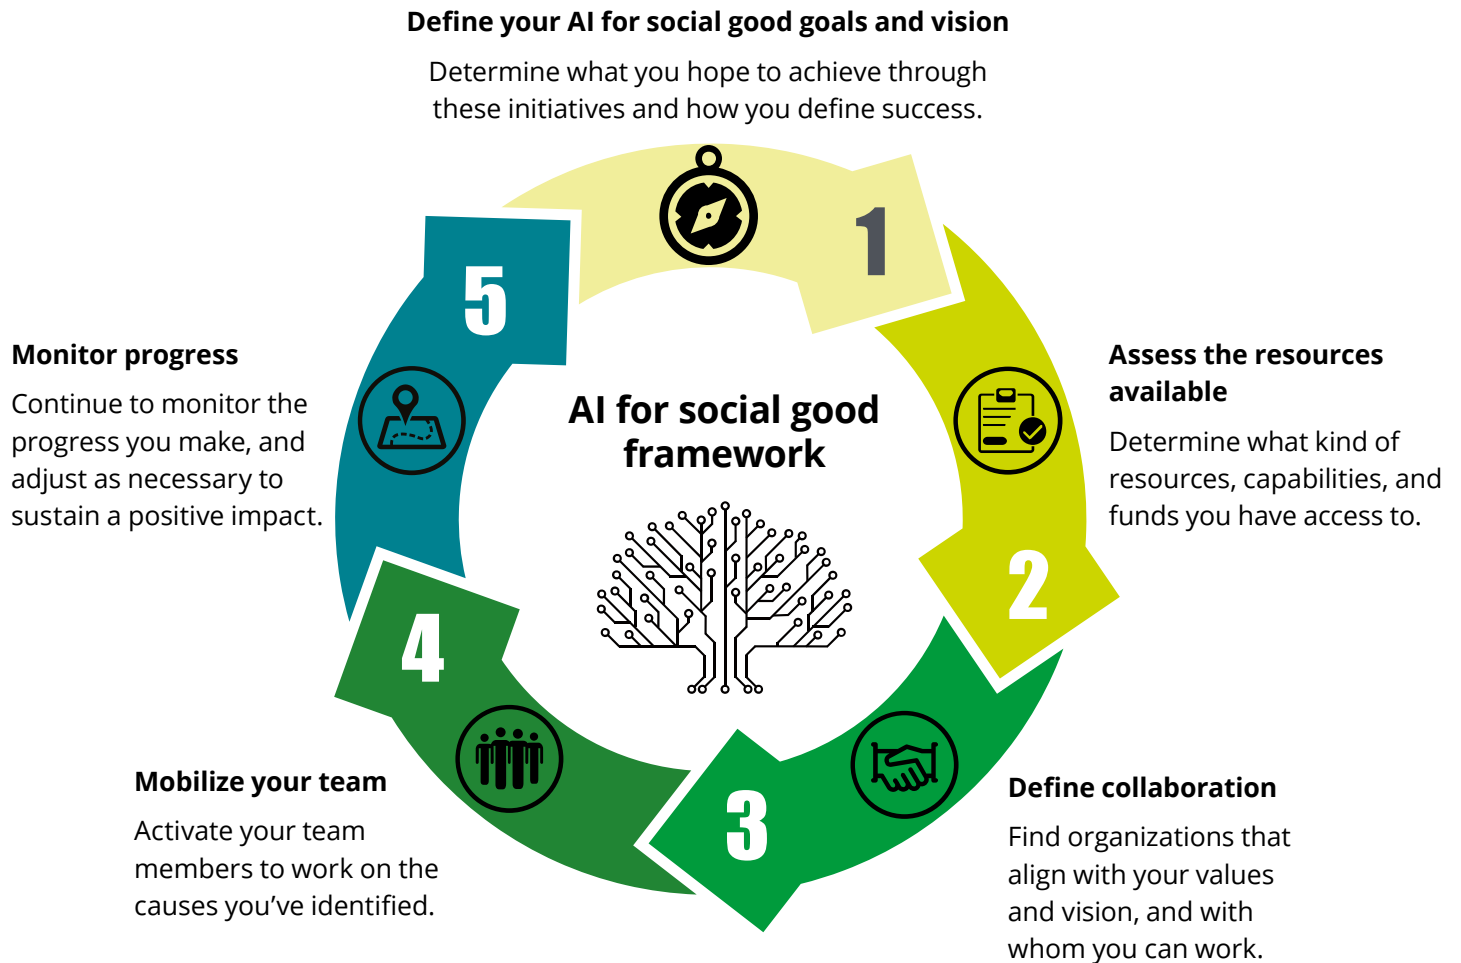
Industry insights on how to make the world a better place through AI

As more organizations invest in AI technologies and talent to further their own ambitions through data-driven insights, more opportunities are created to allow them to use their AI resources for social good. But how should they proceed?

Based on the experiences of our own people—and on conversations with representatives at other organizations that are active in the AI for social good space—we have developed a framework for thinking about such projects. Companies eager to use their AI investments to bring about positive change in their communities and contribute to making the world a better place can use this system as a starting point for their own efforts.



Applying the framework



Source: The Institution of Engineering and Technology, UK, and Deloitte Center for Financial Services analysis

The framework comprises five stages:

We spoke with a number of leaders at a variety of corporations, academic institutions, and other organizations for insights into each of these stages. Their comments, advice, and explanations are included in the next sections.



1. Define your AI for social good goals and vision

These initiatives should be rooted in a firm understanding of an organization's goals and its vision in undertaking such efforts. Some examples of successes follow.

At the AI for Good Foundation, the driving ambition is to apply new work on AI to global challenges—and have a real impact, says CEO James Hodson. But the foundation doesn't try to define what social good looks like; instead, it relies on the parameters set out in the 17 UNSDGs.

"We wanted to have an external framework that included lots of cultural perspectives," notes Hodson, adding that the UNSDGs offer universal priorities that most people can agree define actions that "have a positive impact on humanity and the ecosystems upon which we rely."

Similarly, the AI for Good Foundation strives to ensure that AI is subordinate to the social-good aspect of a project. AI isn't the end goal or even the only solution, Hodson says. Instead, the foundation focuses on achieving the UNSDGs by looking at how AI can support work already under way in the field. At the outset of a potential opportunity, then, researchers explore what role—if any—AI can and should play, the networks or communities that would need to be established, and the expected impact.

TD Bank Group has a long history of giving back to its communities. The TD framework captures this community-minded spirit by expressing the organization's vision of being "the better bank," as well as its purpose: to enrich the lives of its bank customers, colleagues, and communities. More recently, it unveiled its TD Ready Commitment to help create conditions that could allow everyone a chance to succeed in a changing world. The organization set a target of \$1 billion in community giving by 2030, focusing on work that supports change, nurtures progress, and helps make the world a better, more inclusive place.

"TD believes data, advanced analytics, and AI can help accelerate the bank's ability to help communities succeed and thrive in years to come," says TD Bank Group's Peter Husar. "Data, analytics, and AI have to be embraced not just by industry, but by everyone so that our society can advance and we can make things better everywhere." Accordingly, the TD Ready Commitment supports investments outside the organization TD to drive the use of AI and analytics and thus help society at large.

TD's Layer 6—the bank's in-house AI team—eagerly seized the opportunity to use its technology and talent for socially beneficial ends as part of the commitment. The team uses two principal criteria for choosing *AI for social good* projects: potential impact and available data. Initiatives that benefit significant numbers of people and involve data sets large enough to justify machine learning are ideal for their purposes. When the pandemic struck, Layer 6 worked with the Vector Institute—a Toronto-based organization focused on AI research that fosters economic growth—to create a dashboard featuring high-level aggregated metrics that could help the Institute assess and develop insights into the ensuing economic impact of the global crisis. The team also collaborated with Mount Sinai and SickKids hospitals, using computational biology and machine learning to investigate the properties and proteins of the COVID-19 virus in order to better understand what triggers its mutations.

For technology giant TELUS, AI for social good projects represent an opportunity to prove that AI—and the data it relies on—can be used in responsible, trustworthy ways. "Technology innovations can transform society and the economy, and solve many of our social problems," says Jesslyn Dymond, the company's director of data ethics and governance. "It's about spotting the opportunity in technology in a way that adds value and mitigates any risks or concerns that come with AI."

In addition to publishing its commitment to responsible AI use on its [website](#), TELUS also created its [Data for Good program](#). This outlines the company's pledge to apply its fountain of critically useful data to socially beneficial pursuits, all the while ensuring it does so ethically, transparently, and with regard to preserving users' privacy. The data it compiles can help generate valuable insights into mobility patterns. Not only are such statistics advantageous to understanding and reducing the spread of COVID-19, but they can also be helpful in reducing traffic congestion and improving traffic safety. "When we can make data available in a way that respects privacy, it can deliver really useful insights to inform actions that can have tremendous impact on social good," says Dymond.

"For Microsoft, a top goal is to make AI accessible to every person and organization on the planet," says John Kahan, the company's vice president and chief data analytics officer. This objective is grounded in the company's mission to enable and empower users to do more with technology.

"We aim to accelerate the pace of innovation by bringing together philanthropic work and connectivity advances, and by providing resources—and, in some cases, funding—to make technology more accessible around the world," says Kahan. Microsoft's AI for Good program (not to be confused with the similarly named AI for Good Foundation) has a target of US\$165 million in philanthropic investments over the next five years, using cloud and AI technologies to overcome environmental challenges, advance humanitarian

objectives, improve health outcomes worldwide, preserve and enrich shared cultural heritage, and support the billion people around the world who live with disabilities.

While Microsoft has a clear vision for its AI for Good, it defers to its non-profit and other partners to determine their own goals for any given initiative. “As long as we’re solving their problems and doing it as fast as possible using technology and AI, then we’re achieving what we need to in the world,” says Kahan.

Biotech firm Roche Canada is clear about its vision and goals for its AI for social good initiatives, says Fanny Sie, the head of the company’s [AI Centre of Excellence](#). “We look at impact. We look at the way we can improve people’s lives so that they can live their *best* lives. That’s not just about longevity: it’s about people’s quality of life in the community, each and every day,” she says.

Roche established its AI centre in November 2020 after joining forces with startups, multinationals, academic researchers, and health-care sector professionals that spring to aggregate and curate critical data—and after these combined efforts quickly resulted in vital insights that the group realized could be shared with pandemic task forces worldwide. This AI division uses the same collaborative model to break down silos between public and private organizations in multiple sectors to bring swift solutions to patients.

Google aims to enable the real-world impact of AI, focusing its AI for social good efforts on five pillars: accessibility, public health, crisis response, nature and society, and climate change. The company also works alongside non-profit partners through its Google.org Impact Challenges and Fellowship programs.

“We believe that artificial intelligence can provide new ways of approaching problems and meaningfully improving people’s lives,” says Brigitte Hoyer Gosselink, the corporation’s director of product impact. “Our work on AI for social good is focused on maximizing social gains from AI, and we share a vision with the organizations we work with that AI can help humans solve some of the toughest problems we’re facing.”



2. Assess the resources available

Once an organization establishes its vision for *AI for social good* efforts and then either sets out or adopts specific goals to achieve these ends, it’s important to assess available resource, including computing power, data, talent, and financial support.

Today, more than half of the world’s existing AI resources are found in the high-technology and financial sectors, observes Microsoft’s Kahan, while non-profit organizations are home to less than 5% of these global assets. “The world’s asking organizations with few resources, comparatively, to solve world-scale challenges,” Kahan says. “Therefore, Microsoft and those with [more substantial AI] resources have an obligation to do something. This is why we created the AI for Good program.”

“Those with [more substantial AI] resources have an obligation to do something.”

– John Kahan, Microsoft

Ensuring AI talent is given the time and space it needs to take part in AI for social good programs is equally important, observes TD Bank Group’s Husar. The bank’s Analytics for Social Good program is always on: at any given time, TD is involved with approximately 15 related projects, each with a specific non-profit. The program has more than 500 volunteers, with eight to 10 of them committing two to four hours per week each over the course of a year to do analytics work for their chosen non-profit. The initiative allows staff to fulfill their passion for giving back by applying the skills and expertise they use on the job to their preferred volunteer project.

With TD’s Layer 6 team, 20% to 30% of each staff member’s time is dedicated to research-related activities. While some focus on AI explainability and/or trustworthiness, others concentrate on fields such as health care. Since 2019, Layer 6 researchers have published more than 10 peer-reviewed articles on how to apply AI and machine learning to solve problems in clinical or public-health settings.

As an organization, TELUS is committed to giving back to local communities; Data for Good is the realization of that commitment to social good through technology. Grassroots organizations don’t always have significant technology investments; some may even still rely largely on information compiled and stored on paper. And even when data exists, getting access to it can be difficult, as organizations must navigate data-privacy and security concerns. But with TELUS’s massive data sets—including the company’s vast array of valuable information on mobility patterns—a variety of social-good use cases are afforded access to invaluable resources they otherwise never could have had.

At Deloitte, there are a number of programs that enable our specialists to use their AI knowledge and skills to give back to their communities. Our people are encouraged to volunteer their talents through firm-sponsored events such as Impact Day—which, in FY2021, saw nearly 900 of us engage for a combined total of almost 11,000 hours. Additionally, Deloitte’s Community Advisory Projects (CAP) program, which provides strategic advisory services to a wide array of non-profit organizations, hosts events that offer the firm’s AI specialists unique opportunities to drive real-world impacts. Our CAP program work includes support for the BC Children’s Hospital Foundation (as noted previously), as well as for the Ontario Trillium Foundation to create a natural language processing model that could increase the rate at which that organization could derive insights from its applications. In FY21, 135 of our people contributed more than 5,000 cumulative hours across 25 CAP engagements.

At Google, hundreds of staff are directly involved in AI for social good programs, says the company’s Hoyer Gosselink. Initiatives include in-house teams working on flood forecasting and Google.org Fellowship groups collaborating with organizations such as Doctors Without Borders/Médecins sans frontières and Full Fact.

“We have a mix of formally run programs focused on specific topics and issues, where we’ve identified a particularly high potential for Google’s work to make an impact,” says Hoyer Gosselink. “We also have grassroots programs designed to support Googlers in pursuing their personal vision for a better world,” including paid time off for volunteering.





3. Define Collaborations

Collaborations are hallmarks of *AI for social good* projects. Using AI to tackle social, environmental, and related issues requires a range of expertise, including in talent, data, and technology—a team’s members could run the gamut from data modellers and academic experts to local field staff on the ground.

Such projects involve many very smart people, notes Roche Canada. “But the synergy of bringing all of these different perspectives together is really important to create impactful, tangible outcomes in the community. Machine-learning specialists need biologists, biologists need clinicians, and so on.”

The AI for Good Foundation focuses a great deal of energy on establishing connections that can lead to productive, effective partnerships down the road. “We’re an organization that builds bridges between communities,” says its CEO, James Hodson. The foundation teams with researchers to facilitate empirical work in their specialties, supplying them with data sets and connecting them with communities where they could make an impact. The organization also helps NGOs and governments of all levels—from the global Organisation for Economic Co-operation and Development (OECD) and the UN to authorities at the municipal level—develop policy frameworks and identify large-scale opportunities to introduce AI. The foundation also works with NGOs and other not-for-profits to implement these and other initiatives to ensure that the real-world impacts of their AI for social good projects are measured—and that findings are shared with all stakeholders.

“The synergy of bringing all of these different perspectives together is really important to create impactful, tangible outcomes in the community.”

– Fanny Sie, Roche Canada

TELUS partners with the Vector Institute. Its AI for good focus “is incredibly well-aligned with our goals and priorities at TELUS,” says the company’s data ethics and governance director. As previously noted, TD’s Layer 6 team also boasts a strong connection with Vector: its chief AI officer, Tomi Poutanen, was one of Vector’s founders. Additionally, Layer 6 has extensive ties with academic institutions across the AI community; when the need arises, these relationships facilitate rapid collaborations.

The right partnerships can also remove some of the potential obstacles for AI for social good projects. For example, in its collaborations with ICES (formerly the Institute for Clinical and Evaluative Sciences), Layer 6 doesn’t have to deal with data issues: ICES controls and regulates any data used, and then devotes significant energy to ensuring privacy issues are addressed and data quality is high. So Layer 6 can be confident that their ensuing analyses won’t disclose sensitive or otherwise personal information.

Still, selecting the right partners can take time, notes Susan Athey, a professor of economics technology at the Stanford Graduate School of Business. “Things need to be at the right point of maturity,” she says. “You need the infrastructure; if there’s only one server, for example, there’s only going to be so much you can do.”

Athey prefers to focus on building longer-term partnerships at her Golub Capital Social Impact Lab at Stanford. Once data-sharing agreements are in place, the lab collaborates with a new partner on a series of simpler projects before moving on to more complex ones. Over time, as teams grow more embedded, they can develop more sophisticated projects whose scopes reach beyond those initially proposed in order to expand and vary the group’s overall impact.

Understanding what motivates potential partners can be especially important to nurturing fruitful relationships, she notes. For example, organizations should be wary of academic teams that might be particularly interested in using their customers or constituents as lab rats for research. And as AI for Good Foundation’s Hodson observes, many large companies and other for-profit organizations might see AI for social good projects primarily as self-marketing tools. Moreover, for these organizations, the true value of partnerships under the guise of social good may lie in gaining access to technical talent. Governments, too, notes Hodson, might in some cases seek to benefit from these initiatives by regarding them as opportunities to be seen as technology leaders. To that end, the AI for Good Foundation considers everyone’s motivations and investments to help ensure that AI advances and innovations continue to be welcome, promoted, and in demand.

According to Susan Athey of the Stanford Graduate School of Business, selecting the right partners can take time; understanding what motivates them can be especially important to nurturing fruitful relationships.

“Generally, two kinds of expertise are required to implement any successful AI for social good project: knowledge of AI and knowledge of the problem,” notes Google’s Hoyer Gosselink. His company strives to find, support, and have its work guided by partners who are close to the challenges to be addressed. “We find that the most successful projects happen when we bring together diverse expertise across domains and geographies to collaborate and create greater outcomes,” says Hoyer Gosselink. “Working on the world’s hardest problems is not something any organization or individual can do alone.”

Still, substantial persuading is sometimes needed before organizations will agree to use AI as part of their solutions, notes Nihar Dalmia, a Deloitte partner in the firm’s Omnia AI division. “People are often scared of AI because they’re afraid of what they don’t know about it or what they’ve heard about it, or because they don’t know how to use it. If you can make it easier for potential partners to understand AI—and show how it acts as a tool that complements and supports the work people are already doing—it makes collaboration much, much easier.”





4. Mobilize your team

Your project's overall vision and goals have been established, partnerships have been formed, and the necessary resources have been marshalled, from data to technology to talent. The next step is to mobilize your team for action.

Connecting *AI for social good* projects with existing commitments to community action—such as Deloitte's *WorldClass* program and the TD Ready Commitment—can help drive teams. But the intrinsic benefits of these AI initiatives often serve as powerful motivating factors themselves. Husar of the TD Bank Group believes such projects help with talent retention. "People want to work on this sort of thing. They get excited about it," he says. "We're seeing people volunteering more than they normally would, because this is giving back *and* using the skills they have and love to use. And they can see the impact of their contributions more directly." In addition, he says, because their engagement in these programs exposes TD volunteers to various challenges facing society, they can gain new perspectives and appreciation for a diversity of thought outside their usual workplace bubbles of technology, banking, and finance.

Deloitte's Vipul Patil was already volunteering with local initiatives in the Vancouver area when the chance arose to help the BC Children's Hospital Foundation. "I jumped at the opportunity," he recalls. Not only did the Deloitte CAP engagement allow him to support the foundation, but he also got to work alongside "great minds"—that is, his teammates. "I don't look at it as extra time spent away from client work. I look at it as my own personal time that I'm spending to make BC better."

The University of Toronto's Gillian Hadfield says that a willingness to dig in and work across organizational barriers is essential to effective teaming. "You need to get people involved who will roll up their sleeves and collaborate really closely and learn about how the various systems work," she says. In health-care-related AI projects, for example, success depends on "close, close collaboration" among hospital administrators, doctors, nurses, patients, and machine-learning engineers.

"You need to get people involved who will roll up their sleeves and collaborate really closely."

– Gillian Hadfield, University of Toronto

"AI for social good projects transcend companies and competitors," remarks Microsoft's Kahan. He points to the COVID-19 pandemic as an example: "People got together to share data and help solve problems. Competitive barriers dropped and technology firms—as an industry—started working with health-care professionals and policymakers to solve the problems we faced."

Sometimes, however, institutional barriers can keep team members from participating as fully as they might like. As the AI for Good

Foundation's CEO points out, many academics and researchers have a tremendous desire to apply their expertise to global societal challenges, but they typically face significant pressures and incentives to focus on publishing their research, not on applying it to solve real-world problems. Creating opportunities for researchers to use social-good projects for their own work can thus help advance the larger cause. Hodson and the foundation also encourage universities to measure faculty by their social-good impacts versus their publication rates alone.

Support is also essential to mobilizing and sustaining AI for social good projects. "It's about finding the right champion, someone whose interests or values are aligned with the proposed framework or approach," says Jayant Narayan, the project lead of the World WEF's Global AI Action Alliance. "Finding that first champion helps get the project going, and then that first project becomes a tool, a scaling mechanism, to start other projects."

"Having very strong top-level executive sponsorship behind AI for social good initiatives is extremely important," says Deloitte's Nihar Dalmia. That's especially true in times of economic uncertainty, when social good can easily be seen as a secondary goal. "The best examples are those where there's a clear agenda set by the CEO, the board, and other executives to make AI for social good part of the culture and brand of the organization. It needs to be ingrained into the culture."



5. Monitor progress

Once an *AI for social good* project is under way, it's important to track its progress; the team can then adjust as necessary to achieve and sustain beneficial results.

According to Microsoft's Kahan, to effectively measure the impact of a given social-good project, a holistic view is needed: it's important to understand the partners and their perspectives, as well as the nature of the problem the team is trying to solve, so that the group can select metrics that are relevant to the challenge at hand.

"KPIs matter," says Stanford's Athey. However, choosing relevant metrics can be more challenging than expected. For example, she says, a team developing an educational app designed to promote sustainable agricultural practices might be able to easily measure log-ins and identify which stories are being read, but those metrics can't determine whether farmers are as a result changing their practices and enjoying long-term benefits. "It's easy to train AI to get people to click more," notes Athey. "It's harder to make sure those clicks lead to long-term outcomes."

"It's easy to train AI to get people to click more. It's harder to make sure those clicks lead to long-term outcomes."

– Susan Athey, Stanford Graduate School of Business

“A willingness to ‘dig in’ and work across organizational barriers is essential to effective teaming.”

– Gillian Hadfield, University of Toronto

“One of the ways to build confidence in AI is to use AI for social good. And there’s growing recognition that AI for social good is AI for economic good, as well.”

– Nihar Dalmia, Deloitte Canada

“Sometimes the key impact [of AI for social good projects] is really straightforward: Did we keep someone out of hospital today?”


– Fanny Sie, Roche Canada

AI for social good often relies at least in part on economic metrics—after all, delivering social good isn’t incompatible with delivering financial benefits. However, the economic impacts of these types of projects don’t always align with the way business leaders think about economic benefits. The AI for Good Foundation, for example, has collaborated with economists at leading US and European universities to understand the potential economic impacts of their shared social-good projects, considering metrics other than return on capital. Instead, teams seek to measure how their programs improve people’s lives in ways that are only indirectly linked to economics. A team might ask: How does our project improve crop yields/water usage/energy generation? These metrics measure meaningful contributions to a community’s well-being; but their economic benefit, while real, is secondary.

At Google, a project’s success is based on its on-the-ground impact, as well as on building a robust ecosystem in which the social sector, academics, and technology companies can come together to solve big problems, says Brigitte Hoyer Gosselink. “We want to maximize the good AI can do in our society by ensuring both that we deliver end-to-end results and that more organizations have access to the benefits of AI and other advanced technologies,” she says.

Google often looks at the specific impact AI has made on time and cost savings, as well as on the quality of a partner organization’s primary outcomes overall. But the company also considers metrics related to the success of individual projects. For example, its collaboration with the India-based non-profit Wadhvani AI—which works to help developing countries’ underserved communities, including those in agricultural areas—increased smallholder profits by more than 20%, while its initiative to support Full Fact’s human fact-checkers with AI processes led to a 1,000% increase in the number of facts the organization could verify each day.

Ultimately, AI for social good projects are about their impacts on the target population, says Sie from Roche Canada. These can include, for example, empowering patients to make informed health decisions, providing insights that allow businesses and public organizations to build better health-care systems, and informing government policy decisions that affect resource deployment. “Sometimes the key impact is really straightforward: did we keep someone out of hospital today?” says Sie.



Conclusion

Artificial intelligence and its related technologies—such as machine learning—have incredible potential to transform how we live and work. The powerful algorithms that can help businesses make better decisions, drive revenue, improve efficiencies, and secure new competitive advantages can also make lasting contributions to the public good. AI can help governments, businesses, non-profits, and other organizations around the globe create a more equitable, more inclusive, more sustainable world.

Using AI for social good is a collaborative effort. Achieving optimal outcomes and bringing about positive change requires organizations, including businesses, academic institutions, and governments, to work across barriers to combine their processing power and share expertise, knowledge, and relevant data. It can be done; indeed, it *is* being done today, by Deloitte and many other businesses.

It's time for organizations to consider their contribution to the greater social good—and to maximize these efforts with AI initiatives.

Acknowledgements

We would like to thank the following individuals for sharing their time and insights with us as we prepared this report:

- Brigitte Hoyer Gosselink, Director, Product Impact, Google
- Philip Nelson, Director Software Engineering, Google
- Fannie Sie, Head of Roche's AI Centre of Excellence
- John Kahan, Vice President and Chief Data Analytics Officer Emeritus at Microsoft
- James Hodson, AI for Social Good Foundation
- Jesslyn Dymond, Director of Data Ethics & Governance, Telus
- Susan Atfey, Professor of Economics Technology, Stanford Graduate School of Business
- Gillian Hadfield, Schwartz Reisman Chair in Technology and Society, University of Toronto
- Ginger Zielinskie, Chief Strategy Officer, Data.org
- Jayant Narayan, Project Lead, World Economic Forum's Global AI Action Alliance
- Jeremy Veillard, Senior Health Specialist in the Latin America and Caribbean Region of the Health, Nutrition and Population Global Practice at the World Bank
- Nathan Steeghs, National Climate Lead, Deloitte Canada
- Preeti Shivpuri, Data and AI Governance Leader, Deloitte Artificial Intelligence
- Riphay Al-Hussein, Senior Manager, Omnia AI, Deloitte Canada's Artificial Intelligence Practice
- Vipul Patil, Manager, Strategy, Analytics and M&A, Omnia AI, Deloitte Canada

Get in touch

Partners

Nihar Dalmia
National Government and
Public Services Industry
Leader
Deloitte Artificial Intelligence
ndalmia@deloitte.ca

Stefan Popowycz
AI Insights & Engagement
Portfolio Leader
Deloitte Artificial Intelligence
spopowycz@deloitte.ca

Contributors

Audrey Ancion
Leader, AI Institute Canada
Deloitte Artificial Intelligence
aacion@deloitte.ca

Aisha Greene
Manager AI Institute
Deloitte Artificial Intelligence
aigreene@deloitte.ca

Nick Murray
Intern, AI Strategy Analyst
Deloitte Artificial Intelligence

Anamjit Singh Sivia
Manager, Omnia AI
Deloitte Artificial Intelligence
asivia@deloitte.ca

Deloitte.

Deloitte provides audit and assurance, consulting, financial advisory, risk advisory, tax, and related services to public and private clients spanning multiple industries. Deloitte serves four out of five Fortune Global 500® companies through a globally connected network of member firms in more than 150 countries and territories bringing world-class capabilities, insights, and service to address clients' most complex business challenges.

Deloitte LLP, an Ontario limited liability partnership, is the Canadian member firm of Deloitte Touche Tohmatsu Limited. Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee, and its network of member firms, each of which is a legally separate and independent entity www.deloitte.com/about for a detailed description of the legal structure of Deloitte Touche Tohmatsu Limited and its member firms.

Our global Purpose is making an impact that matters. At Deloitte Canada, that translates into building a better future by accelerating and expanding access to knowledge. We believe we can achieve this Purpose by living our shared values to lead the way, serve with integrity, take care of each other, foster inclusion, and collaborate for measurable impact.

To learn more about Deloitte's approximately 330,000 professionals, over 11,000 of whom are part of the Canadian firm, please connect with us on [LinkedIn](#), [Twitter](#), [Instagram](#), or [Facebook](#).