

Advancing Equitable AI in the US Social Sector

How thoughtfully designed AI-powered applications can drive stronger and more equitable outcomes for nonprofits.

By **Kelly Fitzsimmons** | Mar. 12, 2024




(Illustration by Raffi Marhaba, The Dream Creative)

Artificial intelligence is rapidly evolving and reshaping how people and organizations think and behave, across many sectors and around the world. In the United States, companies like Netflix and Amazon have been leveraging AI for years to tailor recommendations and provide virtual assistance to customers, while research institutions and AI labs like **DeepMind** are using it to accelerate medical research and fight climate change.

Nonprofit organizations, however, have been less involved in this moment of technological innovation. To some degree, this makes sense. The nonprofit sector faces widespread challenges, including a lack of investment in research and development and a shortage of staff with AI expertise, that other sectors don't. But it also needs to change. AI's impact on society—how people work and live—will only increase over time, and the social sector can't afford to not engage with it. Indeed, nonprofits have an important role to play in its development. When designed and implemented with equity in mind, AI tools can help close the data gap, reduce bias, and make nonprofits more effective. Funders, nonprofit leaders, and AI experts need to move quickly and in alignment with one another to advance equitable AI in the social sector.

Using AI to Support Equity

Of course, while AI offers many exciting opportunities, its potential to cause serious harm is well-known. Developers train AI algorithms using data culled from across society, which means biases are baked in from the start. For example, financial service providers often use AI to make lending decisions, but the financial industry in the United States has a long history of systematic discrimination against women and communities of color, including redlining and inequitable appraisal and underwriting policies. Because algorithms used to make lending decisions are trained on historical data that reflect the intentional disadvantaging of certain zip codes, occupations, and other proxies associated with race or gender, they can



The Global Pursuit of Equity

This article series, devoted to advancing equity, looks at inequities within the context of seven specific regions or countries, and the

perpetuate unfair lending practices and financial inequities if left unaddressed. Even well-intentioned nonprofits could easily design flawed AI applications that result in unintended and damaging consequences. An organization providing seed funding to social entrepreneurs, for instance, could train AI using biased financial data and end up working against its mission of advancing wealth equity by mistakenly preferencing certain populations.

ways local innovators are working to balance the scales and foster greater inclusion across a range of issue areas.

There's also a general fear that rapid improvements in AI's ability to perform administrative, analytical, and creative tasks could render many professions and even entire industries obsolete. The digital media company BuzzFeed **recently began using AI to generate content** for its site; it's not hard to imagine a scenario where a nonprofit with a limited budget might decide to cut its marketing team and rely on the AI-powered language model ChatGPT instead. In addition, the lack of regulation and financial incentives for organizations to consider the equity aspects of AI makes it less of a priority. Leading tech companies like Google and Microsoft, for example, **have cut ethical AI staff in recent years**.

These are legitimate concerns that organizations must strive to address. However, when developed thoughtfully and with equity in mind, AI-powered applications have great potential to help drive stronger and more equitable outcomes for nonprofits, particularly in the following three areas.

1. Closing the data gap. A widening data divide between the private and social sectors threatens to reduce the effectiveness of nonprofits that provide critical social services in the United States and leave those they serve without the support they need. As Kriss Deiglmeir wrote in **a recent *Stanford Social Innovation Review* essay**, “Data is a form of power. And the sad reality is that power is being held increasingly by the commercial sector and not by organizations seeking to create a more just, sustainable, and prosperous world.” AI can help break this trend by democratizing the process of generating and mobilizing data and evidence, thus making continuous research and development, evaluation, and data analysis more accessible to a wider range of organizations—including those with limited budgets and in-house expertise.

Take **Quill.org**, a nonprofit that provides students with free tools that help them build reading comprehension, writing, and language skills. Quill.org uses an AI-powered chatbot that asks students to respond to open-ended questions based on a piece of text. It then reviews student responses and offers suggestions for improvement, such as writing with clarity and using evidence to support claims. This technology makes high-quality critical thinking and writing support available to students and schools that might not otherwise have access to them. As Peter Gault, Quill.org's founder and executive director, **recently shared**, “There are 27 million low-income students in the United States who struggle with basic writing and find themselves disadvantaged in school and in the workforce. ... By using AI to provide

students with immediate feedback on their writing, we can help teachers support millions of students on the path to becoming stronger writers, critical thinkers, and active members of our democracy.”

2. Reducing bias. There is an abundance of stories illustrating how AI can perpetuate biases, including [police departments using algorithms](#) that flag Black defendants at higher risk of future crime than their white counterparts and [companies using hiring algorithms](#) that disadvantage female candidates. But AI also has the potential to help reduce bias by supporting equitable decision-making. Thoughtfully designed algorithms can disregard variables unrelated to outcomes (such as race, gender, or age) that too frequently color human decision-making, helping nonprofit staff surface patterns and make decisions based on evidence, rather than falling back on human biases and blind spots.

One example of an organization using AI to support evidence-based decision-making is [First Place for Youth](#), an organization that helps foster youth make a successful transition to self-sufficiency and responsible adulthood. First Place for Youth built a recommendation engine that uses precision analytics—a technology that predicts trends and behavioral patterns by discovering cause-and-effect relationships in data—to analyze program administration and case assessment data, and learn from differences in outcomes among youth. This helps staff better understand what has worked for specific populations in the past and what customized supports are most likely to lead to success. Designed with an equity lens, the AI algorithm can make clear whether different demographic groups have equal access to program components. In addition, to avoid replicating any existing biases, it doesn’t match cases based on sociocultural factors such as race that shouldn’t factor into why a child gets selected for different program options.

3. Increasing efficiency. Stories abound of AI applications making mistakes ranging from comic to terrifying, including [Bing’s chatbot sharing dark fantasies and professing its love to a New York Times columnist](#), [an AI-generated Seinfeld parody getting banned from Twitch](#) for making transphobic jokes, and [a Microsoft chatbot being shut down](#) for making racist remarks. But AI-powered applications such as recommendation engines, precision analytics, and natural language processing can help organizations increase their output while reducing human error. Passing off rote and tedious tasks to AI can allow capacity-constrained nonprofit staff to focus more of their time on the strategic and human-to-human work that computers can’t do.

[Crisis Text Line](#), a nonprofit that provides free, text-based mental health support and crisis intervention, uses AI to support efficiency and scale while preserving high-quality personal services. Specifically, the organization trained AI on past texts to recognize high-risk keywords and word combinations, [allowing it to more efficiently triage texters by severity](#). Crisis Text Line also [uses AI as part of its volunteer training](#) to build strong crisis response skills. Its natural language processing algorithm, trained on a set of fictional but realistic cases, allows the algorithm to mimic live conversations between volunteers and clients on topics like anxiety and self-harm. The technology helps Crisis Text Line train its network of

volunteers efficiently and flexibly, in part because volunteers can complete the training at their convenience. Most importantly, it allows staff and trained volunteers to focus more of their time on ensuring that clients receive high-quality, live support.

Supporting Nonprofit AI Development

While examples like the ones above offer exciting illustrations of what's possible, they are unfortunately the exception rather than the norm. The social sector must do more to seize the moment; it must demonstrate what's possible and build the tools and infrastructure required to promote equitable AI.

Our organization, **Project Evident**, helps organizations harness the power of data and evidence for greater impact, and works to build a stronger and more-equitable evidence ecosystem. As part of a recent 18-month **initiative** focused on AI, we convened a cohort of nonprofits eager to use AI for stronger program outcomes and developed **case studies** illustrating what their AI piloting and adoption processes looked like in practice. And to inform the development of effective and equitable AI tools, we are currently partnering with the **Stanford Institute for Human-Centered AI** on a **national survey** to better understand AI usage and learning needs among nonprofits and funders. These projects illuminated several ways the social sector can strengthen the ecosystem for equitable AI.

1. Increase investment in AI tools. Many nonprofits and school districts are eager to understand how AI might support their work but lack the means. “When you look at the funding structures for nonprofits, it typically does not allow for this kind of work to actually happen,” said Jean-Claude Brizard, who leads the education-focused nonprofit Digital Promise and **participated** in our cohort. “Most nonprofits don’t have the time or the resources.” One straightforward solution is for funders to offer grants to both AI “natives” (organizations that already deploy AI to create equitable outcomes) and AI “explorers” (organizations interested in piloting AI but that lack the capital and support they need).

Foundations should also invest in their own learning, and explore whether and how AI can make grantmaking processes more efficient or improve funding strategies. One foundation in our network, for example, uses AI to read and analyze past grant reports to surface patterns around elements of effective funding strategies. The funder engaged a team of evaluators and data scientists to train large language models, with **human-in-the-loop** validation, to transform qualitative text data from thousands of proposals and reports into structured, longitudinal data. The team then combined this data with primary evaluation data (including from surveys, interviews, and focus groups), then cleaned, processed, framed, structured, and labeled it according to the initiative's theory of change. Finally, using machine learning algorithms with human/subject-matter-expert guidance and review, the team **descriptively, predictively, and causally modeled** it. The results provide foundation staff with on-demand data ready for report writing, interactive reports, visualizations, and dashboards, strengthening its ability to understand what's worked and what hasn't. This helps ensure that new grantmaking strategies build on prior knowledge and allows the

foundation to better share what it's learned with grantees and the field. They also create efficiencies that allow program officers to spend more time building strong relationships with grantees, prospective grantees, and communities—the kind of work only a human can do.

2. Collaborate on efforts that address AI and equity. Another issue is that organizations have few financial or regulatory incentives to consider the ethical and equity implications of AI development. The for-profit sector designs most AI tools and processes with profitability and scalability in mind.

Researchers, policy makers, technical assistance providers, funders, nonprofit practitioners, and community members need to work in alignment to advance thoughtful practices and policies around equitable AI. This work could include pushing for legislation and regulation, developing frameworks and standards for the field, engaging with communities to develop tools that address their needs, or sharing knowledge and best practices about testing and monitoring AI algorithms for bias.

Several efforts are currently underway that could help provide the social sector with better support and guidance. US Senate Majority Leader Chuck Schumer **announced** in June that he would be hosting a series of AI forums to help Congress develop comprehensive, bi-partisan AI legislation. The first forum held in September **included** labor and civil rights activists in addition to tech experts. Another effort is the **Distributed AI Research Institute**, which researches how AI can disproportionately affect marginalized groups, and develops frameworks for AI development and practice that center community voice and seek to advance equity. These are both promising examples of forward-thinking, cross-sector efforts to address equity and AI.

3. Stay grounded in the fundamentals. Organizations should have a few components in place before piloting AI for equitable outcomes: a clear strategic purpose, adequate technology systems, the capacity to design for justice and equity, and a strong culture of learning. These are not new capacities, nor are they specific to AI, but they are frequently overlooked and underfunded in the social sector.

Project Evident frequently works with organizations to develop strategic evidence plans—roadmaps for continuous evidence-building and program improvement. When a nonprofit is looking to conduct an evaluation, either because board members or funders require it or leaders think it's “the right thing to do,” we often suggest taking a step back and asking: What are you trying to learn? What outcomes are you seeking to achieve? Having that programmatic clarity is important to developing an evidence strategy, and the same holds true when building recommendation engines or other AI-powered tools. In fact, an AI algorithm is simply a translation of a logic model that lays out a program's inputs and intended outcomes.

Nonprofit leaders should make sure their organizations have the proper foundation in place before engaging with AI, and others in the social sector should support them with the funding and tools they need to build fundamental capacities. The Philadelphia-based social service agency **Gemma Services** is an **example** of an organization that was able to successfully attract funding to develop an AI tool, by

demonstrating to a local funder (the [Scattergood Foundation](#)) and the City of Philadelphia how the potential application of AI would not only benefit Gemma Services' team, but also provide valuable learnings for funders and policymakers concerned with mental health across the Philadelphia region.

The potential of AI in the social sector is immense. Thoughtfully designed AI applications can help close the data gap, reduce bias, and make nonprofits more effective. But to fully harness AI's potential, we must build a robust social sector infrastructure that provides the necessary resources, prioritizes equity concerns, and invests in fundamental organizational capacities. With more resources, better alignment, and greater intentionality, we can help the social sector harness the power of AI for stronger and more equitable outcomes for all.

Read more stories by [Kelly Fitzsimmons](#).



Kelly Fitzsimmons is the founder and CEO of [Project Evident](#), an organization dedicated to building a stronger and more equitable evidence ecosystem. A committed social innovator, Fitzsimmons previously served as vice president/chief program and strategy officer at the Edna McConnell Clark Foundation, co-founded Leadwell Partners and New Profit Inc., and currently serves as a Leap Community Ambassador, a member of Results for America's Invest in What Works Federal Standard of Excellence Advisory Committee.

DOI: 10.48558/7oeI-4447