Stanford University
Human-Centered
Artificial Intelligence

PROJECT
EVIDENT

Working Paper
February 2024

Inspiring Action:

Identifying the Social Sector AI Opportunity Gap

Sarah Di Troia
Vanessa Parli
Juan N. Pava
Haifa Badi Uz Zaman
Kelly Fitzsimmons

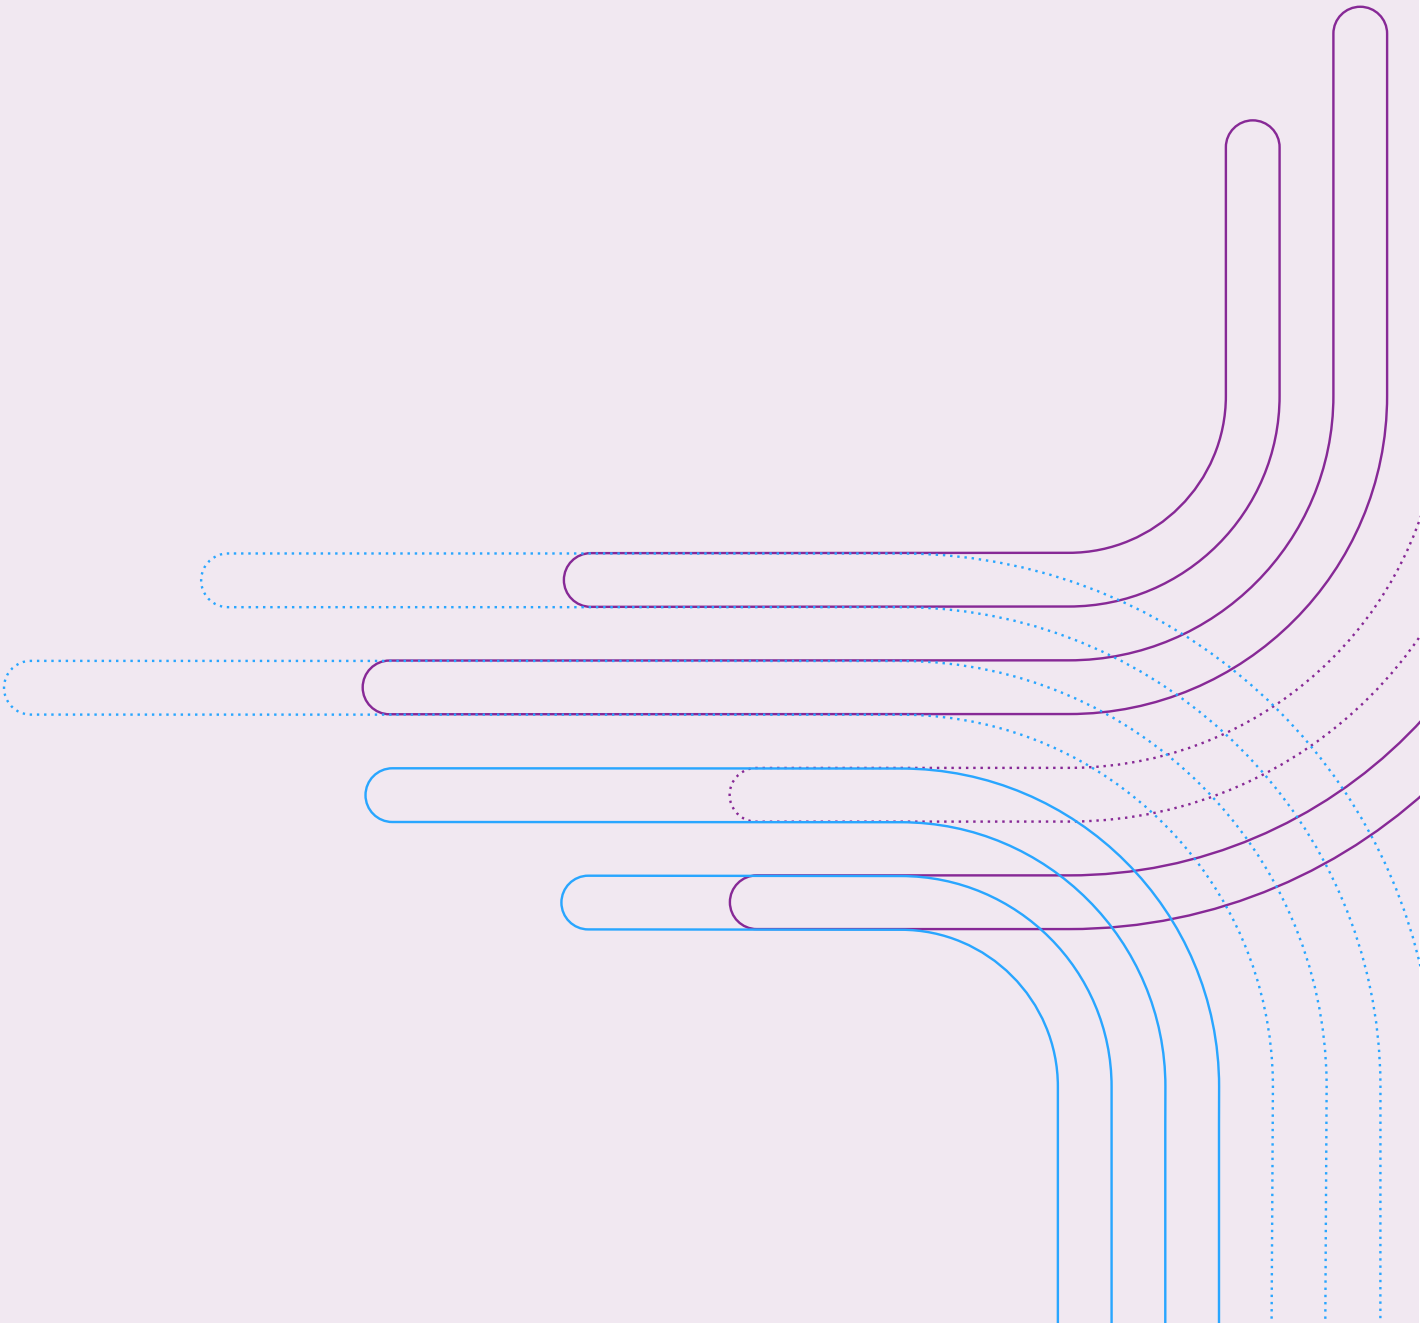

Table of Contents

Disclaimer

The Stanford Institute for Human-Centered Artificial Intelligence (HAI) is a nonpartisan research institute, representing a range of voices. The views expressed in this Working Paper reflect the views of the authors.

About Project Evident

The mission of Project Evident is to harness the power of evidence for greater impact. Project Evident believes that by empowering practitioners to drive their own data use and evidence building while also strengthening the surrounding ecosystem, we can increase the number of effective solutions in the social and education sectors and scale them faster—ultimately producing stronger, more meaningful, and more equitable outcomes for communities. Project Evident is at the forefront of the next generation of understanding and acting on practitioners' use of data and evidence for equitable outcomes in the social and education sectors, advancing an inclusive, R&D approach with appropriate technology, capacity, and knowledge.

About the Stanford Institute for Human-Centered AI

The mission of the Stanford Institute for Human-Centered Artificial Intelligence (HAI) is to advance AI research, education, policy, and practice to improve the human condition. Led by faculty from multiple departments across Stanford University, research focuses on developing AI technologies inspired by human intelligence, studying, forecasting, and guiding AI's human and societal impact, and designing and creating AI applications that augment human capabilities. Through the education work of the institute, students and leaders from a range of disciplines gain a range of AI fundamentals and perspectives. At the same time, the policy work of HAI fosters regional and national discussions that lead to direct legislative impact.

Introduction

This national survey is a collaboration between Stanford's Institute for Human-Centered Artificial Intelligence and Project Evident and was conceived as a project to shed light on the current use of, interest in, and opportunity for AI in the social and education sectors. Over the last decade, AI has reshaped the commercial sector and consumer habits, resulting in significant value creation and profitability—think value created by recommendation systems in e-commerce or streaming services. As it becomes easier to include AI applications (Microsoft Copilot, Google Workspace, OpenAI GPTs) as part of the nonprofit technology stack, the social and education sectors have the same opportunity to deploy AI to create value through enhanced mission-related outcomes.

There are already examples of the social and education sectors using AI to advance their missions. GeoMatch, a machine learning tool, can assist placement officers in finding suitable communities where refugees can prosper. Quill.org deploys nonfiction texts paired with AI-powered writing prompts to provide customized feedback to students and build their reading comprehension and writing skills.

However, in order for the social and education sectors to use AI to advance their missions, they need a coordinated voice in the critical work of defining how AI tools are developed, how they are integrated, and what infrastructure is required to deploy AI for equitable outcomes. The growing consensus is that participation from a broader set of voices is essential to ensure this technology is inclusive and equitable and represents the perspectives of all communities. As a result, approaches such as participatory AI are gaining momentum. Still, many in the social sector feel they lack sufficient expertise to participate in these discussions, risking the exclusion of deep field knowledge in future AI systems and missed opportunities afforded by AI.

We hope this working paper inspires the philanthropic, industry, academic, social, and education sectors to mobilize resources, shape a common learning agenda, and upskill the field so they can practically inform the AI discourse.

Since the release of ChatGPT in November 2022, there has been active debate about whether the social and education sectors should use AI. However, in October 2023, following the White House's Executive Order on AI, Vice President Kamala Harris announced a new initiative committing $200 million in funding to enable AI innovations that are grounded in public interest and democratic rights. Based on a first-of-its-kind survey on AI in the social and education sectors, this working paper aims to document how AI is already being used, understand where these sectors see opportunities with AI, and identify the challenges to continued experimentation and adoption. We hope this working paper inspires the philanthropic, industry, academic, social, and education sectors to mobilize resources, shape a common learning agenda, and upskill the field so they can practically inform the AI discourse.

What Do We Mean When We Say AI?

In the survey, we ask questions about both traditional AI and generative AI and use this same language in the report. Traditional AI refers to systems designed to make specific predictions or decisions based on a particular set of inputs. Generative AI refers to models trained on large amounts of data that learn the underlying patterns to generate new data mirroring the training data. To help orient survey takers, we provide examples of different types of AI systems versus providing technical definitions. Traditional AI examples are recommendation engines or predictive analytics that make suggestions (e.g., Netflix suggesting a movie) or chatbots/virtual assistants that answer specific questions (e.g., Siri/Alexa or an online retailer asking if you need help). *Examples of generative AI are ChatGPT, Claude, Bard, Bing Chat, Stable Diffusion, and DALL-E.

Additionally, we want to understand the types of work for which survey respondents use AI. For the purposes of this survey, work is divided into three categories based on Project Evident's experience of working with nonprofits and is defined as supportive (finance, human resources, technology, communications, etc.), revenue generation (fundraising, business development, sales, etc.), and mission-related (working with clients, implementing programs, or making grants, etc.). We use this same language and definitions throughout the working paper.

What Do We Mean When We Say Social and Education Sectors?

For the purposes of this survey, we define the social sector and education sectors according to the National Taxonomy of Exempt Entities (NTEE) system for classifying organizations. We group the social and education sector respondents into two groups: (1) "Grantmakers," which are respondents that classified their organizations under the NTEE code "Philanthropy, Voluntarism & Grantmaking Foundations," and (2) "Nonprofits," which are respondents that classified their organization under any other NTEE code (with some exclusions). The nonprofit category within our sample includes education-focused nonprofits and public school districts, as they have program functions and staff who may have an AI mandate for the K-12 schools they manage more broadly, but it does not include individual schools. For more information on the type of respondents considered for the survey, see the "Methodology" section below.

**Note that since the release of the survey, generative AI tools are starting to be incorporated in chatbots.*

Working Paper Highlights

1. AI already has a considerable presence in the social and education sectors. 48% of funders and 66% of nonprofit respondents claim their organization utilizes some type of AI. Given that nonprofits rely on funders for capital, differences in levels of use could impede AI experimentation in the social and education sectors.

2. 78% of funders and 77% of nonprofits believe their organization would benefit from using more AI (specifically in mission-related work). This creates an opportunity gap for nonprofits of 14% for traditional AI and 22% for generative AI; for grantmakers, the gap is even larger—26% and 39%, respectively.

3. Education nonprofits use AI significantly more than other nonprofits. This could be due to the substantial investment in education technology companies that has shaped the field of education and the availability of data.

4. While most respondents state they use AI in their work, many do not have an organizational policy guiding AI usage (78% of nonprofits and 72% of funders), which introduces risks such as exposing sensitive data or limiting use and experimentation within the organization. This is especially concerning as both nonprofits and funders have access to community-level data collected for management and outcome tracking and reporting.

5. About 80% of respondents who use AI deploy it for supportive work (finance, human resources, technology, communications, etc.), but only about **60% deploy AI for mission-related work** (working with clients, implementing programs, or making grants).

6. Bias in AI systems is the most cited barrier to AI adoption, followed by challenges in envisioning how AI can be used and a lack of expertise inside the organization. **Nonprofits have a particular concern about the cost of AI technology.**

7. Most grantmaker respondents do not have a specific technology grantmaking priority and do not plan to create one in the next year. Instead, funding that goes toward technology is channeled through other priority funding areas.

Inspiring Action

AI adoption is a massive shift for the social and education sectors; pooled resources, aligned collective action, and shared learning will ensure that this technology does not benefit only the early adopters or most networked or best resourced organizations. A first step in shifting the status quo is to create a shared learning agenda and mobilize funds that fuel research, education activities, and convenings. Based on the survey findings, we recommend six action areas for the social and education sectors to consider:

1. Invest in the development of unified, cost-effective, and scalable upskilling resources for grantees and grantmakers with a focus on AI to support mission attainment. Academic institutions, job training platforms, and technical assistance providers could be powerful partners to create these materials.

2. Enable deep collaboration and experimentation between nonprofits, AI researchers, and AI developers. This will surface the needs and wants of the social and education sectors early in the AI research and design process, ensuring equity is a core principle.

3. Surface and invest in the creation and dissemination of case studies and stories of early adopters to study progress and share insights and findings. These resources will help nonprofits prioritize where to start their AI journey, accelerating learning and use. The education sector, which is ahead in its adoption of AI, is a fertile ground for best practices.

4. When funding AI adoption, think about the systemic barriers that affect the field (e.g., infrastructure, compute, data, upskilling) and might not be apparent within program funding areas.

5. Become engaged buyers of AI tools by learning how bias manifests in models and working collaboratively across grantmakers and nonprofits to define the minimum threshold required of technology companies to address bias in their products. Collective action would lessen the burden on any single entity having to define what is good enough.

6. Review and amend grantmaker and nonprofit operating policies to address AI-related risks and provide staff and stakeholders with training as a continuous activity; topics could include equitable AI use, data security audits, policy reviews, and sharing of recent technological developments.

Current Use of AI

Despite active debate on whether the social and education sectors should use AI, survey responses suggest that AI is already being used, with 48% (26/54) of funders and 66% (118/179) of nonprofits utilizing some type of AI at work [Figure 1]. The exigent conversation for these sectors is not whether to use AI but how best to deploy AI to enhance equitable outcomes.

For those currently using traditional and generative AI, most funders (65% or 17/26) and nonprofits (78% or 92/118) are deploying these tools for supportive work such as finance, human resources, technology, and communications [Figure 2]. Using AI in supportive functions saves staff time by automating routine activities or freeing up staff to focus on higher-value areas of their work.

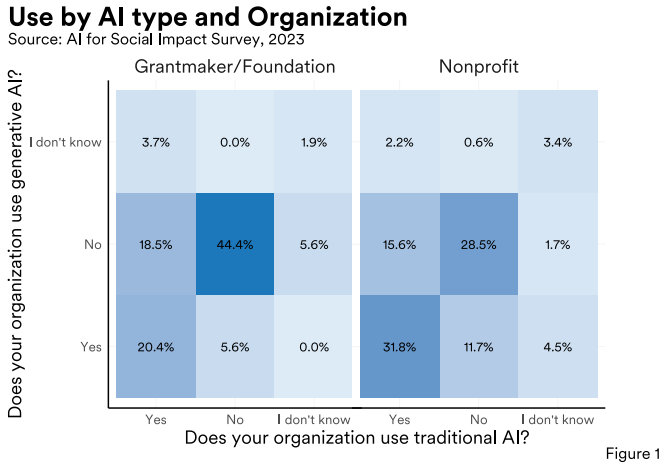

Figure 1

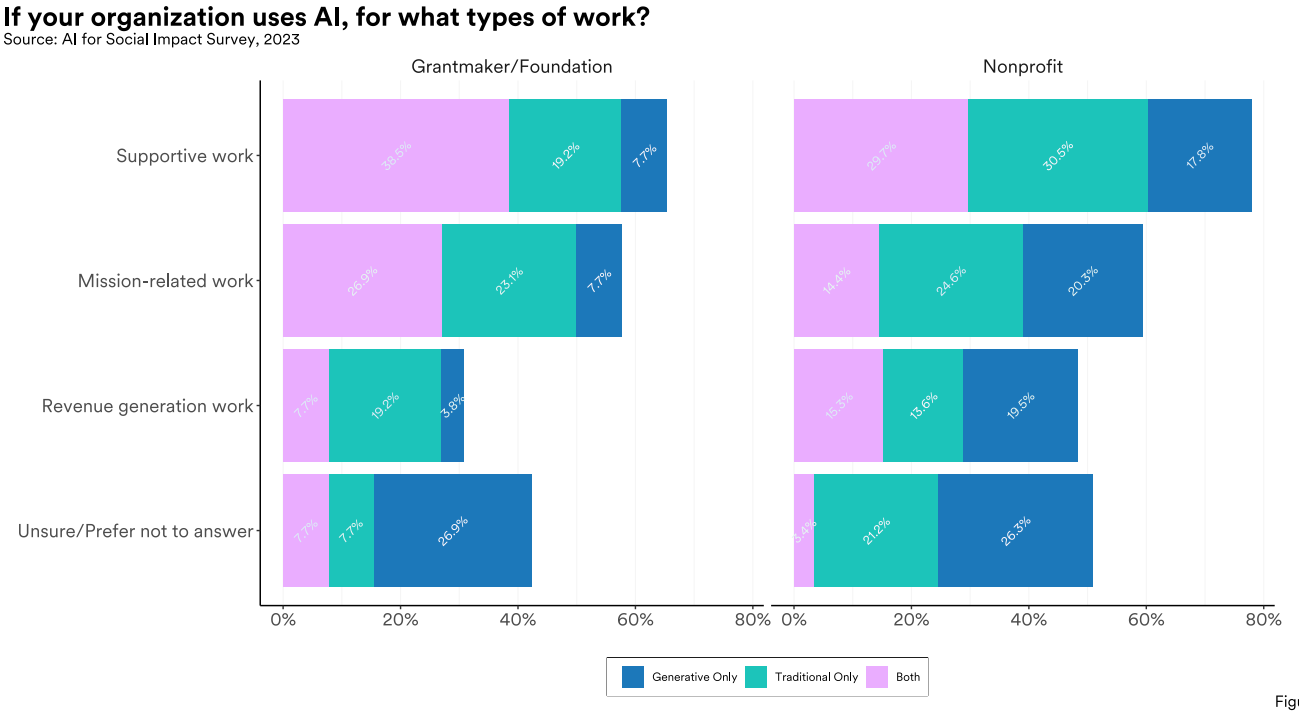

Figure 2

While AI can lead to operational efficiencies, the promise is to use the tools to enhance mission-related outcomes, significantly increasing meaningful community impact. After supportive work, mission-related activities (working with clients, implementing programs, or making grants) are where respondents are using AI.

How can we measure the lost opportunity in current AI usage? The opportunity gap is the difference between the current usage of AI and the belief that your organization would benefit from using more AI. The survey revealed that 78% (42/54) of funders and 77% (138/179) of nonprofits believe their organization would benefit from using more AI [Figure 3], creating an opportunity gap between current and desired usage.

While AI can lead to operational efficiencies, the promise is to use the tools to enhance mission related outcomes, significantly increasing meaningful community impact.

If distinguished by AI type, the opportunity gap for nonprofits is 14% (25/179) for traditional AI and 22% (39/179) for generative AI; for grantmakers, it is even larger: 26% (14/54) and 39% (21/54), respectively [Figure 4].

Benefit by AI type and Organization
Source: AI for Social Impact Survey, 2023

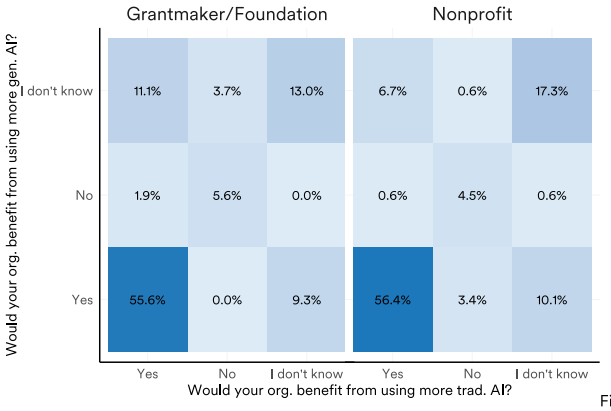

Figure 3

Opportunity Gap
Source: AI for Social Impact Survey, 2023

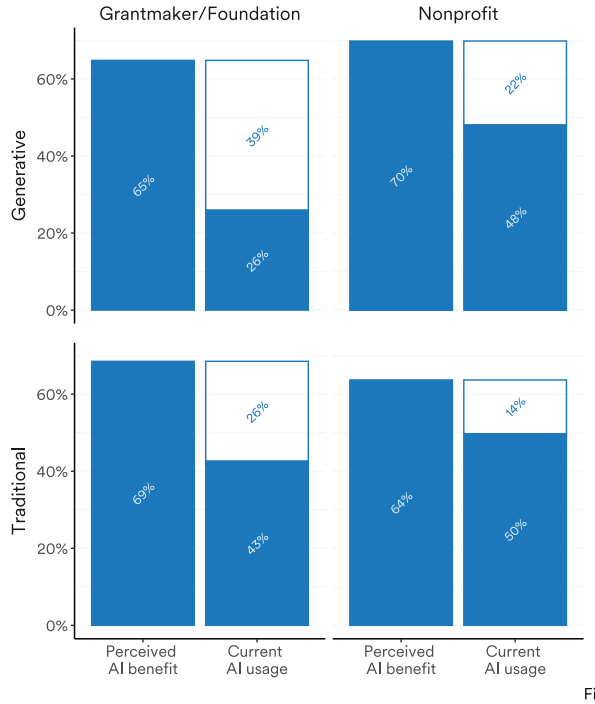

Figure 4

What value could be realized if that 14%-39% opportunity gap was closed? When asked what type of work their organization would benefit from by using more AI, mission-related activities is the leading response, closely followed by supportive work [Figure 5].

For what type of work would your organization benefit from using more AI?
Source: AI for Social Impact Survey, 2023

Figure 5

Nonprofits have ideas about how they envision deploying AI to support their missions:

- "Predicting what educational tools would support specific individuals with disabilities."
- "How these [tools] can be used by child welfare and community orgs to improve their services to families."
- "I'm most interested in the potential of ChatGPT for generating customized content and scenarios. This can make training modules more engaging and relevant to individual learning paths. In addition, the ability of generative AI to produce diverse and inclusive content resonates well with our commitment to creating equitable learning environments."

Grantmakers also see ways to create value using AI beyond finding ways to be more efficient.

- "Could traditional AI improve grantee experience with our foundation?"
- "Using AI to make sense of the mess of climate/ESG analytics that exist on companies . . . so this info can be used by financial institutions, NGOs, etc., to make clearer judgments on who is transitioning fast enough."

Closing the opportunity gap requires leadership from grantmakers and nonprofits: 39% (91/233) of survey respondents identified themselves as executive leadership within a nonprofit organization, and they tend to believe their organization would benefit from using more generative AI compared to employees in other roles [Figure 6]. For grantmakers and for traditional AI, the difference between executive leadership and employees in other roles was not significant.

However, nonprofit leadership interest in experimenting with AI will only go so far without supportive investment from grantmakers. The difference between funders' and nonprofits' AI usage is striking: 66% (118/179) of nonprofits use AI compared with only 48% (26/54) of funders. Because funders provide financial resources to nonprofits, often critical innovation capital for experimentation, the lag in funder AI usage could dampen social and education sector research and development on using AI to enhance equitable outcomes.

AI and fundraising

Fundraising is an always-on activity for most nonprofits and many education organizations. It is not surprising that this is an area nonprofits are eager to learn more about. Among nonprofit survey respondents, 48% (57/118) reported using AI for revenue generation [Figure 2], but 76% (105/138) said they would benefit from more AI in this area [Figure 5]. Nonprofits envision AI as a powerful tool to support these efforts; according to one nonprofit, "I'm particularly interested in learning about how traditional AI can enhance donor engagement and improve fundraising efforts. Understanding how AI can help us analyze donor data, identify potential major donors, personalize communication, and optimize fundraising campaigns would be incredibly valuable." Investment in AI for fundraising may help level the playing field between large organizations with well-staffed development departments and smaller organizations where fundraising may be one of many staff responsibilities. In the long term, AI-powered fundraising tools may enhance overall nonprofit sustainability and resolve the discrepancy between the amount of assets held by organizations whose leaders come from marginalized communities versus those that do not.

Leadership vs. Others | Would your nonprofit benefit from using more generative AI in its work?
Source: AI for Social Impact Survey, 2023

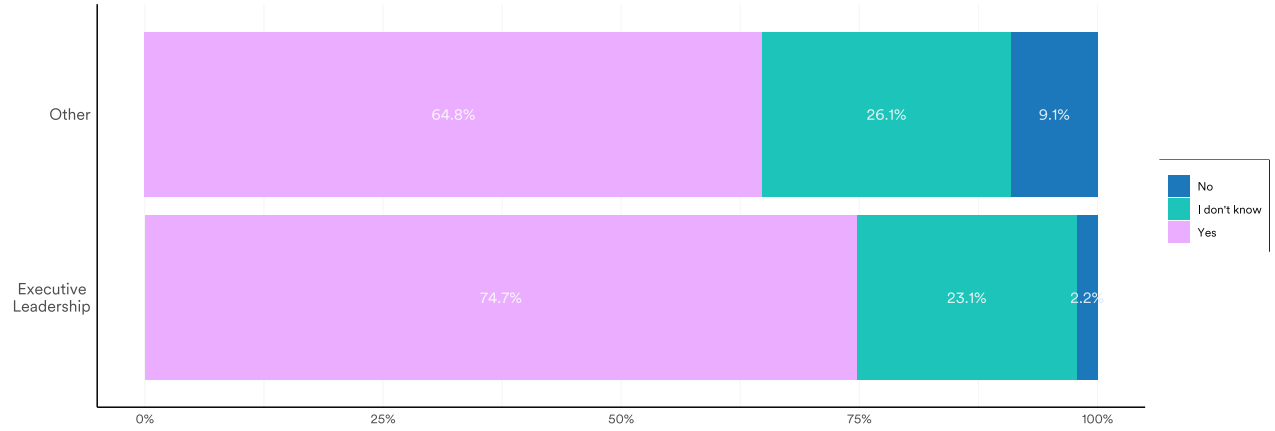

Figure 6

Education vs. rest of field

Over 28% (50/179) of nonprofits' respondents are from the education sector. These respondents use AI significantly more than the other sectors, underscoring a distinct interest and need within the education community to engage with AI-related issues. This forward momentum in the education sector may be a result of readily accessible data or the $26.5 billion in U.S. private edtech venture funding from the start of 2010 through Q3 of 2023.

Does your nonprofit use AI?
Source: AI for Social Impact Survey, 2023

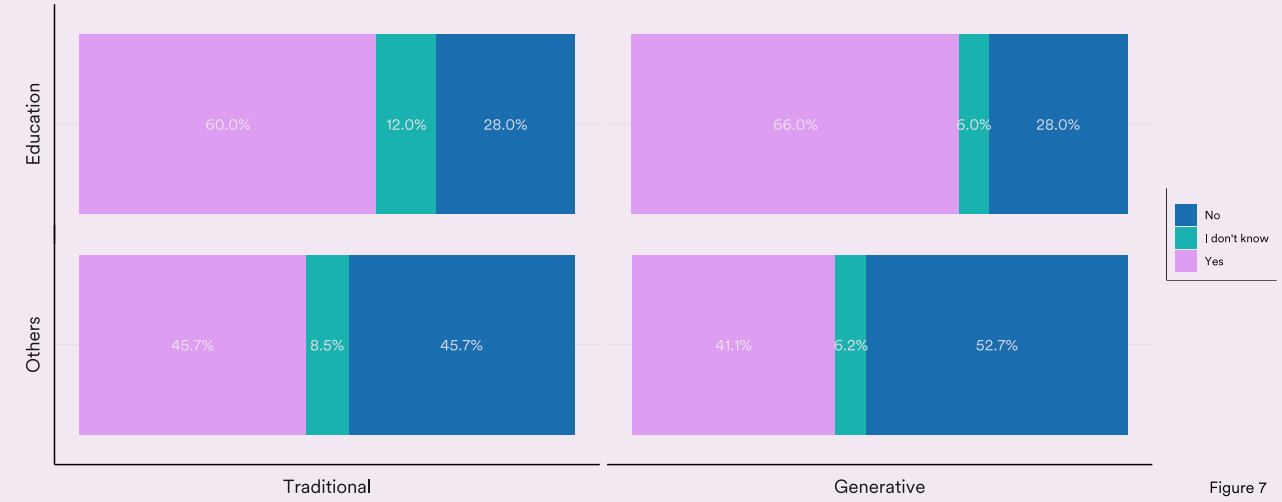

Figure 7

One of the most notable findings is that most (84% or 123/147) nonprofit respondents use generative AI for work at least occasionally, with 1 in 7 individuals using it daily [Figure 8].

If you have used generative AI before, how often do you currently use it for your job?
Source: AI for Social Impact Survey, 2023

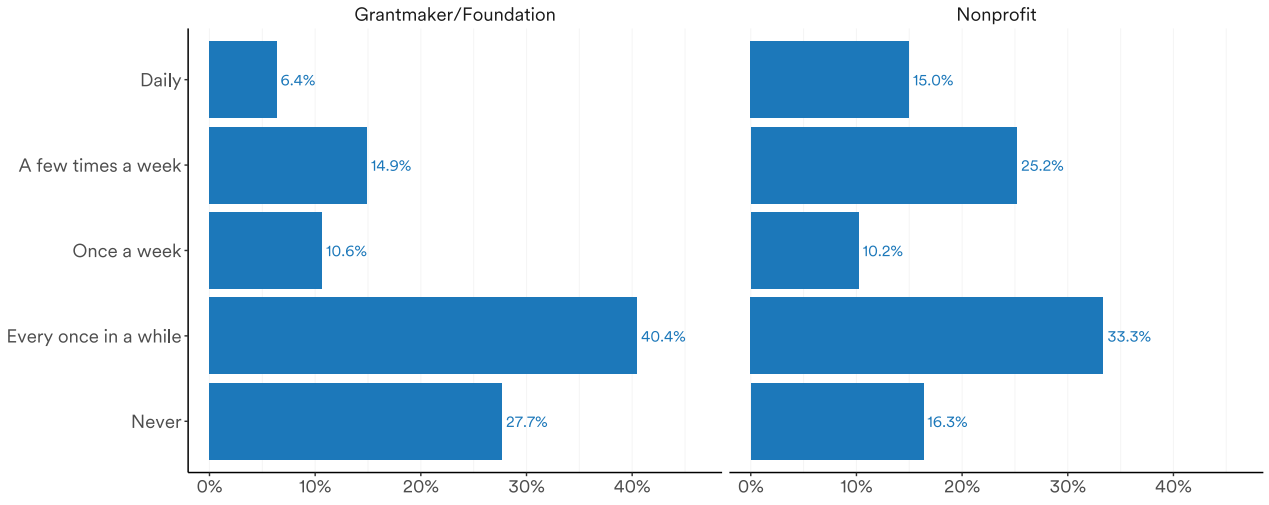

Figure 8

Yet, 78% (139/179) of all nonprofit respondents do not have any policy regulating the usage of generative AI in the workplace [Figure 9].

Does your organization have a policy outlining how generative AI can be used for your work?
Source: AI for Social Impact Survey, 2023

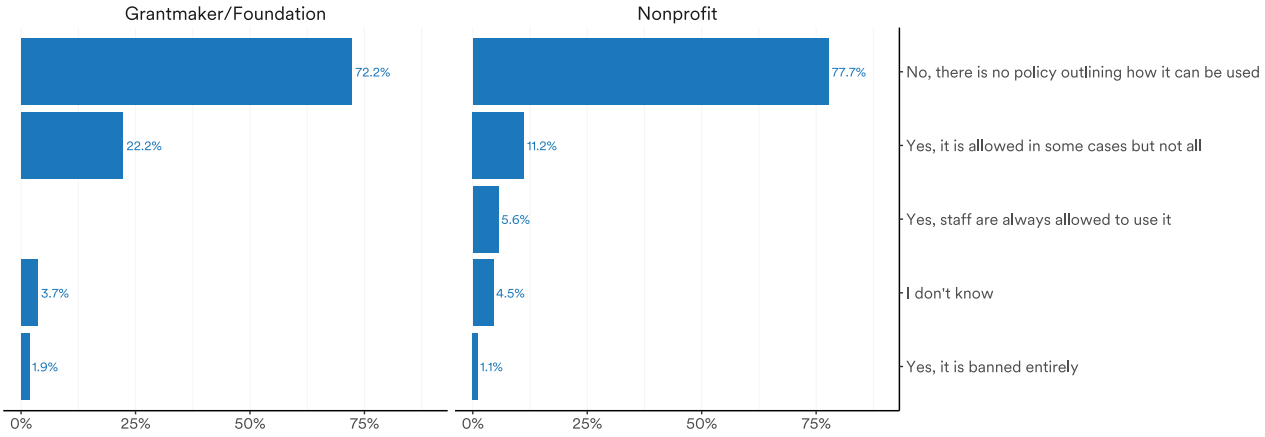

Figure 9

Similarly, 72% (34/47) of grantmaker respondents use generative AI for work at least occasionally, and 72% (39/54) report not having a policy. The lack of clear AI policies or understanding of how current data privacy and governance policies apply to AI introduces risks such as exposing sensitive data or limiting use and experimentation within the organization. This is especially concerning as both nonprofits and funders have access to community-level data collected for management and outcome tracking and reporting.

The majority of respondents report that AI has a very positive or somewhat positive impact on their organizations [Figure 10]. While this result could be because participants who agree to take a survey about AI are inclined to feel favorably toward the technology, it also likely reflects grantmakers' and nonprofits' generally positive AI experiences and/or outlook so far.

What kind of impact has AI had on your organization's work?
Source: AI for Social Impact Survey, 2023

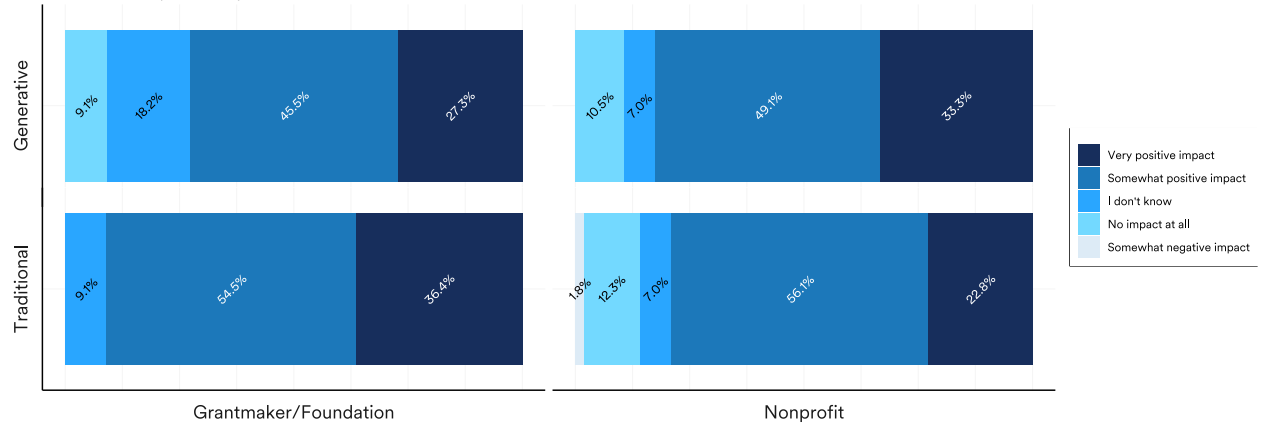

Figure 10

Challenges to AI Adoption

Survey respondents felt that the most significant barrier to AI adoption is concerns about bias [Figure 11]. Equity is at the heart of the mission pursued by most in the social and education sectors. Since 2015, most AI tools have been developed by the commercial sector. While research teams within the commercial sector work to ensure AI tools are fair, bias still shows up, and determining whether the model is trustworthy can be difficult. Nonprofits and grantmakers need to be equipped as savvy buyers of AI tools. They need to understand how biases manifest in these tools so they can experiment while maintaining their organization's ethical standards. One way to address bias in AI systems is through buyers taking a proactive stance in identifying instances of potential bias. As nonprofit and grantmakers enter the market for AI, concerns about equity will need to be addressed. The social and education sectors could work together to define the minimal viable anti-bias thresholds for AI tools to address their shared concerns about bias.

"How to use generative AI ethically, especially since our nonprofit's work deals with challenging and complex topics, including human behavior, racism, antisemitism, and forms of bigotry?"

Collective action would call attention to their needs in product development and lessen the burden on any single entity to define what is good enough to allay equity concerns. The question, in the words of one respondent, is "how to use generative AI ethically, especially since our nonprofit's work deals with challenging and complex topics, including human behavior, racism, antisemitism, and forms of bigotry?"

After concerns about bias, barriers to AI adoption relate to envisioning how AI can be used and a lack of

What challenges are you experiencing?
Source: AI for Social Impact Survey, 2023

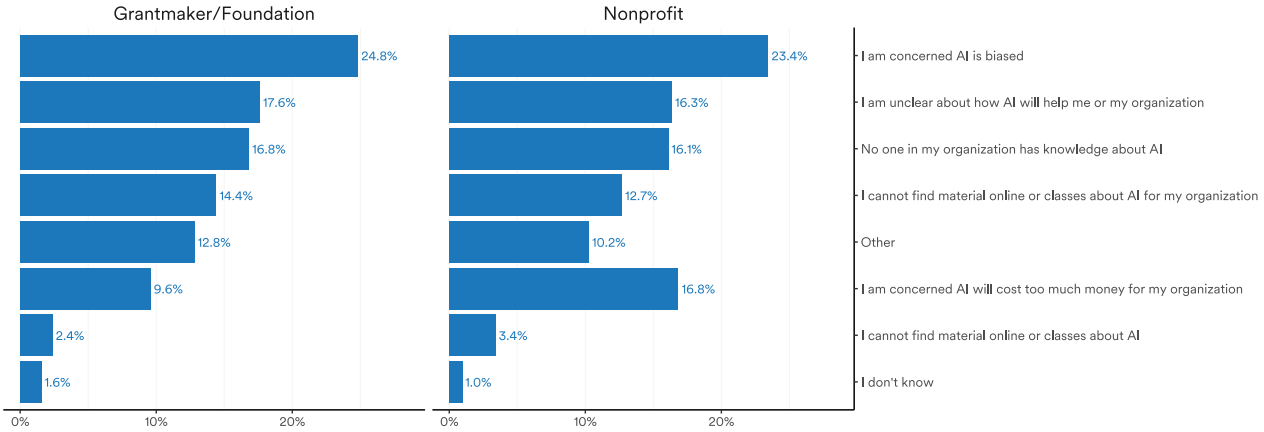

Figure 11

The social and education sectors cannot rely on one-to-one technical assistance models for educational resources as these will only reach the best-resourced or more established organizations.

subject matter experts inside organizations. It is unlikely that each nonprofit will be able to employ AI subject matter experts given the intense technology hiring demands of the commercial sector; however, the social and education sectors have a long history of sharing and documenting what works. Grantmakers investing in AI should surface narrative examples of how to deploy AI, emphasizing designing for equity. Dissemination of these examples will help nonprofits prioritize where to start their AI journey, and accelerate learning and use. Additionally, enabling deep collaboration and experimentation between nonprofits, AI researchers, and AI developers can help all organizations realize impactful social and education sector AI use cases. It will educate nonprofits on "how to use generative AI for mission work in a safe manner," as requested by a survey respondent, and educate AI developers on the dynamics of nonprofit organizations. The social and education sectors cannot rely on one-to-one technical assistance models for educational resources as these will only reach the best-resourced or more established organizations.

In addition, nonprofits have a particular concern about the cost of AI technology. There is a long-standing starvation cycle of grantmakers and governments not fully covering nonprofits' indirect costs, including technology. Especially given the investment required for generative AI, grantmakers and nonprofits need to consider the cost of purchasing, adaption, training, and the ongoing maintenance of models. While low or no-code AI tools make it easier for non-technical staff to access the models, it is still primarily the technology provider that evaluates a model's bias, capabilities, and risks. Providers of the most well-known generative AI tools differ on how transparent they are in these evaluations, meaning buyers need to understand how the evaluations are conducted, trust the evaluation was sufficient, and/or develop processes to perform their own evaluation before purchasing.

Among the 41 nonprofits expressing "other" concerns, roughly 1 in 3 cited worries regarding the technology's best practices, such as privacy, information integrity, copyright, safety, and security. Regarding AI education, figuring out where to begin was a focal point, emphasizing the desire to find resources explicitly tailored to the needs of nonprofits. One of these needs is that for nonprofits, time constraints become just as crucial as other resource constraints by virtue of their small teams. Creating a shared scalable AI education and upskilling resource would be a huge asset for helping the nonprofit and education sectors begin to learn about and experiment with AI. A shared resource would also enable knowledge to flow in the sectors democratically without preferencing organizations with greater financial assets.

One-third of the 16 grantmakers expressed "other" concerns. Like nonprofits, they mention privacy, information integrity, safety, and security concerns. They also show a similar need for guidance on where to start learning about the technology. Again, a shared scalable AI education and upskilling resource would be a huge asset for funders and help unify knowledge across grantmaker staff so grantees have aligned perspectives when interacting with program officers across different funders.

The Learning Agenda

As mentioned earlier in the "Current Use of AI" section, a significant opportunity gap exists between current use and the belief held by respondents that their organization would benefit from more AI. A first step to closing this gap and addressing the discussion points in the "Challenges to AI Adoption" is education. Over 90% (210/233) of survey respondents said they are interested in learning more about how to apply traditional or generative AI to their work [Figure 12].

"I would appreciate a broad and diverse set of practical examples and use cases (actually in practice) to inspire ideas for applications at our org. The more tangible and realistic, the better."

Are you interested in learning more about AI for use at work?
Source: AI for Social Impact Survey, 2023

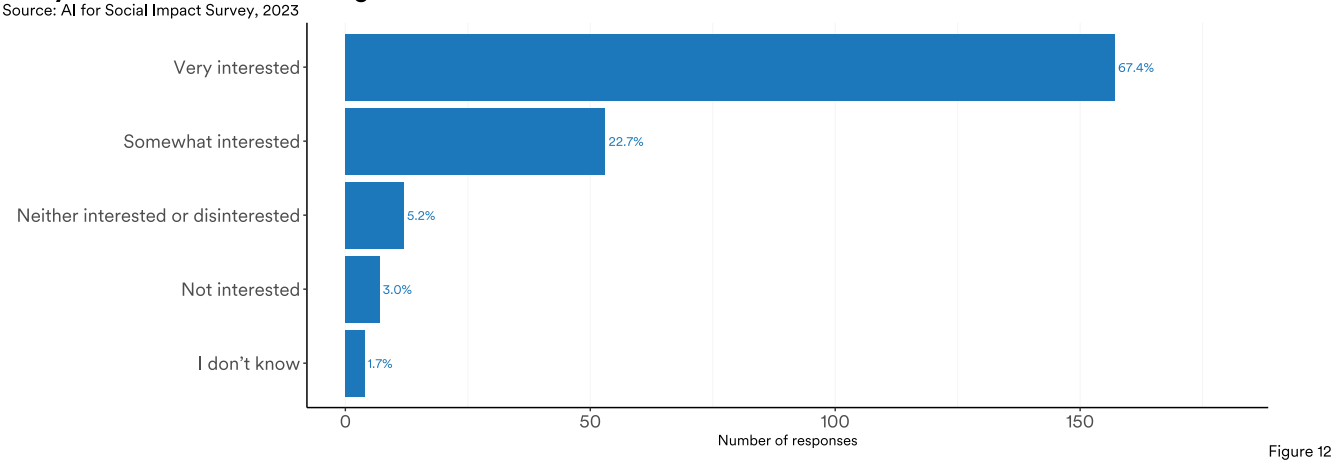

Figure 12

However, knowing where to start a learning journey is hard if you are not able to envision how AI can create value—something survey respondents identified as difficult in "Challenges to AI Adoption." When respondents were asked what they want to learn about AI, one person said, "I don't have the imagination to know the answer to this question." Another stated, "Without having a better sense of what's possible/applicable in a nonprofit context, this is difficult to answer." Several respondents offered not the content they wanted to learn but their need for case studies and best practices to help them prioritize where to start. At a minimum, the field needs stories of early adopters' experience with AI to accelerate learning and use.

> "I would appreciate a broad and diverse set of practical examples and use cases (actually in practice) to inspire ideas for applications at our org. The more tangible and realistic, the better, rather than hypothetical far-off applications we would be unlikely to operationalize successfully."

> "Case studies of good use of traditional AI as well as pitfalls/legal challenges."
>
> "I would like to hear if other Community Foundations have adopted this tool and how."
>
> "Wins and pitfalls for those who are venturing into AI, what tools are being used, and practical how-tos."

Areas respondents identified as those they are most interested in learning about in terms of using AI at work fall into four buckets:

Predictive Analytics: Nonprofit respondents appear keen on using predictive analytics for mission-related work, such as "program design" or to gain "foresight into the effectiveness and reach" of their initiatives. Grantmakers, on the other hand, are interested in using predictive analysis for support-related work, such as "to improve operations (e.g., anticipate workload, etc.)" or "financial planning."

Virtual Assistants/Chatbots: Most nonprofit and grantmaking respondents expressed interest in using virtual assistants and chatbots for support-related work. The primary goal would be to simplify and reduce workload, particularly administrative tasks, and assist with various duties such as "streamline responses to frequently asked questions," "answer general questions" from communities, and "improve intake process."

Data Collection and Analysis: Nonprofit and grantmaking respondents indicated they would like to use traditional AI for the general collection, organization, and evaluation of data. One nonprofit expressed a desire for "improved case management through automated tracking of communications and other data, and assistance with evaluation of collected data," while a grantmaker saw the benefit of "using AI to support historical analysis of impact."

ChatGPT/Content Creation: Many respondents were interested in generative AI for content generation, especially for support-related tasks such as "developing tailored content to use for marketing and fundraising." More specifically, nonprofits and grantmakers saw the greatest benefit from using generative AI to produce more engaging marketing and communications material such as data visualizations, website content, social media posts, press releases, and "emails and articles."

What is the future of grantmaking?

Not surprisingly, nonprofits are keen to use generative AI to "accelerate grant applications & proposals" and, more generally, "ChatGPT for grant-writing, acknowledgment letters, donor content." At the same time, grantmakers want to learn how to use generative AI technologies to improve their grantmaking efforts, and they see potential benefits in every stage of the process, from discovering potential grantees to writing and evaluating grant proposals. As one grantmaker respondent asked, "How it might be used to generate and review grant proposals, conduct field scans of various sectors?" Grants management systems are beginning to release AI upgrades in their software, in some cases supporting grantee writing and in others helping program officers sort through applicants and summarize applications. We should expect a future scenario where there is AI assistance on both sides of the grant application process. We are already seeing grant process changes with trust based philanthropy and many funders, such as MacKenzie Scott, have already been experimenting with new methods that AI can accelerate and elevate.

How Do Funders Think About Grantmaking for AI?

Because funders play a critical role in providing innovation capital to the AI sector, as part of the survey we asked grantmakers how they plan to invest in AI moving forward. Most respondents do not have a specific technology grantmaking priority and do not plan to create one in the next year [Figure 13]. Instead, the majority of funding that goes toward technology is channeled through other priority funding areas. While the number of respondents is small, the trend to invest in technology through other program areas means that funders must educate more staff about AI to facilitate AI grantmaking. Additionally, making technology grants primarily through priority funding areas could limit grantmakers' ability to invest in the underlying systemic factors that perpetuate the data divide between the commercial and the social and education sectors, including the absence of widely used data platforms designed for the social sector, data refineries, or the structuring of social impact data. Grantmakers should think strategically about the barriers to and opportunities for equitable access to AI for outcomes across the social and education sectors, in addition to specific program areas.

Does your organization currently have a separate technology grantmaking priority or funding area?
Source: AI for Social Impact Survey, 2023

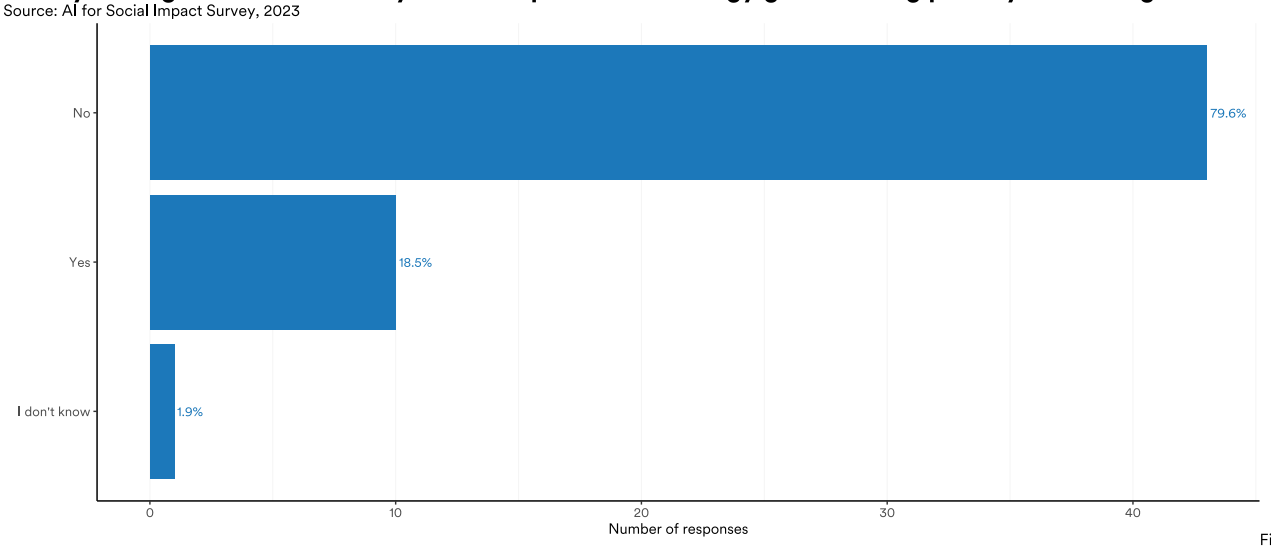

Figure 13.1

If not separate, does your organization currently fund tech projects within other priority grantmaking areas?

Source: AI for Social Impact Survey, 2023

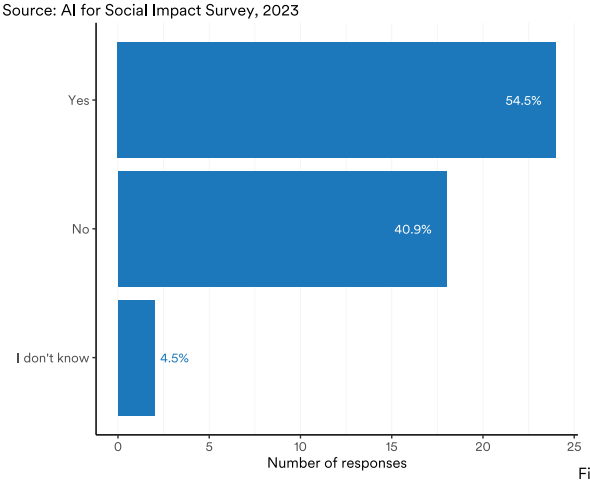

Figure 13.2

If you have not, do you expect to create a separate technology funding priority in 2024?

Source: AI for Social Impact Survey, 2023

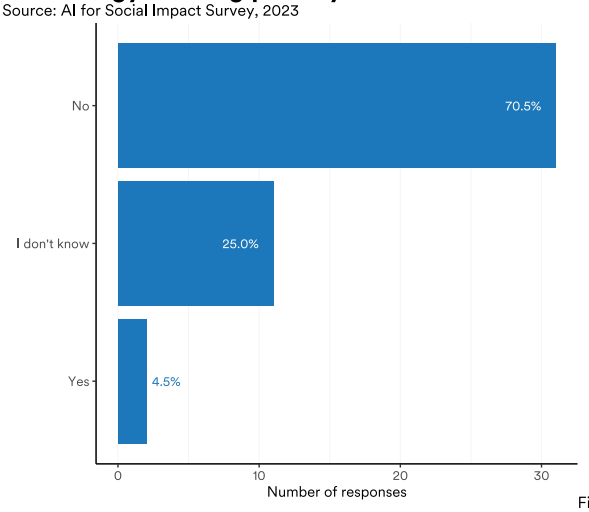

Figure 13.3

Conclusion

The social and education sectors are actively experimenting with AI, but our survey findings show that there is untapped potential in using AI for mission-related impact. While concerns about bias in AI and lack of internal expertise show up as barriers to adoption, social sector leaders can be empowered to leverage AI if provided more resources (e.g., case studies, affordable AI tools), guidance (e.g., organizational policy, training), and peer learning opportunities (e.g., communities of practice, education programs).

The benefits of embracing AI are significant; however, successful AI adoption that is equitable calls for targeted and collective action. Most importantly, the deeper systemic challenges of knowledge and infrastructure gaps that are not atypical to the social sector can be addressed through cross-sector collaboration between philanthropy, academia, and civil society. These players in the AI ecosystem share a broad consensus about the critical role that an empowered civil society can play in ensuring that AI design, application, and governance advances more equitable outcomes. However, it is time to act on this desire by taking concrete steps that tackle the resource and learning barriers uncovered in this working paper.

**We recognize that this survey is just the tip of the iceberg, and there are many critical perspectives that can bolster our understanding and recommendations. As Stanford HAI and Project Evident chart a path forward, we invite social and education sector leaders to share any feedback or additional insights not reflected in this Working Paper. You can reach us at nonprofit-ai-survey@stanford.edu.*

Demographics

This survey was completed by 179 nonprofit organizations, the majority of which are from the education sector [Figure 14], and 54 grantmaking organizations.

Nonprofit | What type of service does your organization provide?
Source: AI for Social Impact Survey, 2023

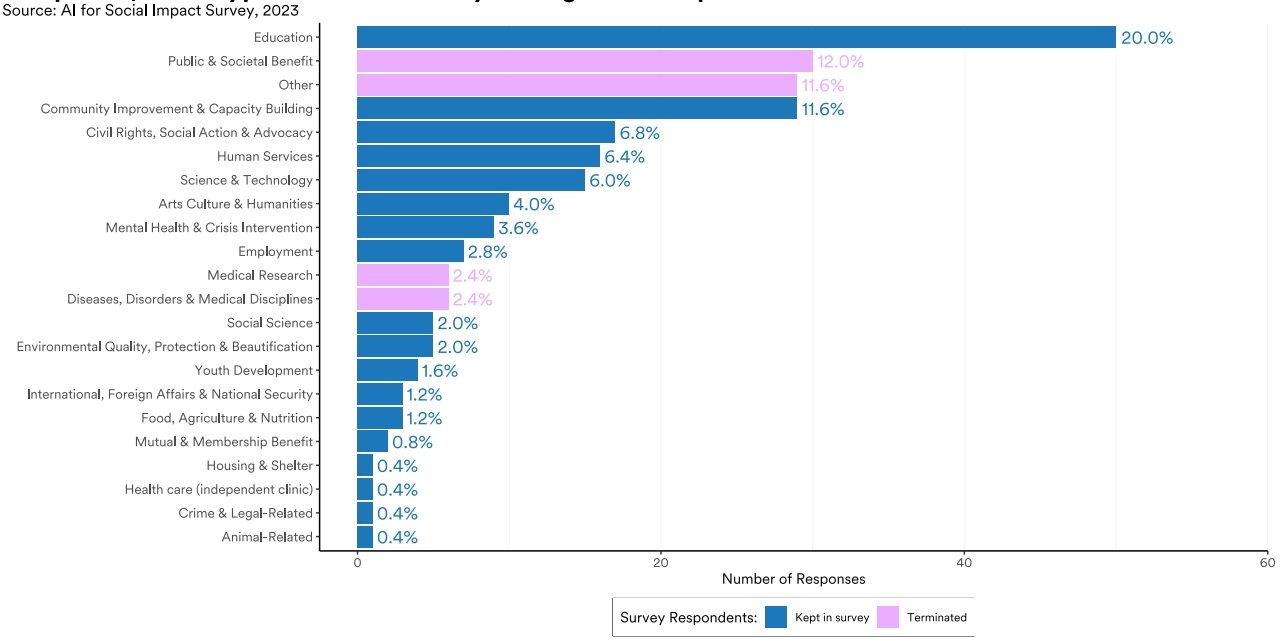

Note: Those organizations identifying as "House of Worship/Religion" and "Technical Assistance Provider/Consultant" had already been terminated at this point in the survey (see Methodology).

Figure 14

Executive leadership is the most common role held by respondents—representing half of those from nonprofit organizations and a third from grantmaking institutions [Figure 15].

Which of the following best describes your role at the grantmaker?
Source: AI for Social Impact Survey, 2023

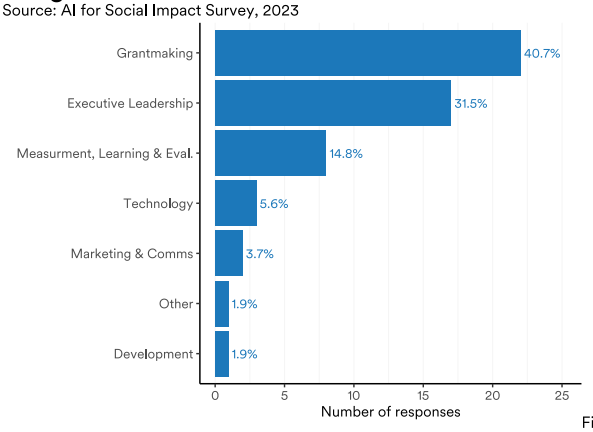

Figure 15.1

Which of the following best describes your role at the nonprofit?
Source: AI for Social Impact Survey, 2023

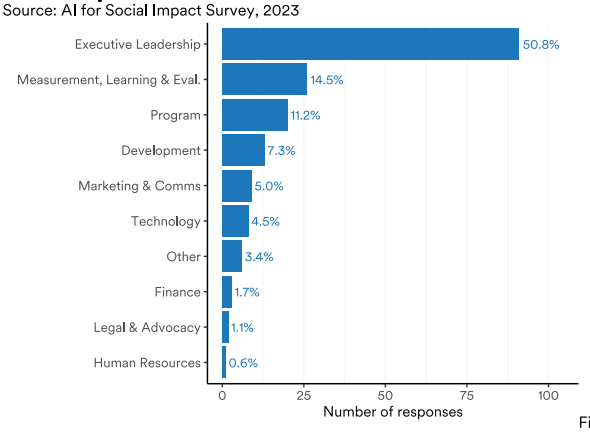

Figure 15.2

Overall, the demographic composition of the respondents (the majority being non-Latinx white women) does not significantly deviate from national averages for nonprofit demographics.

What is your gender?
Source: AI for Social Impact Survey, 2023

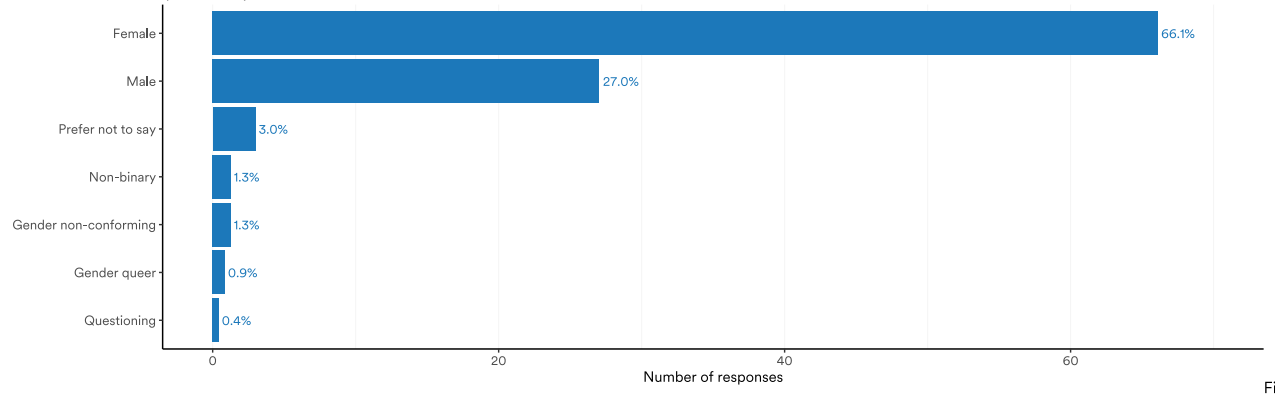

Figure 16

Do you identify as Hispanic or Latinx of any race?
Source: AI for Social Impact Survey, 2023

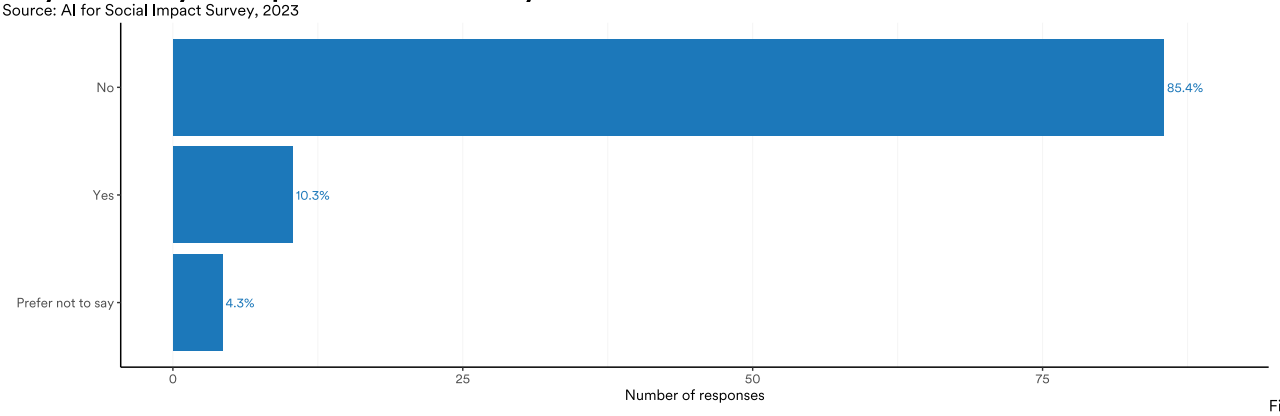

Figure 17

What is your racial or ethnic identification?
Source: AI for Social Impact Survey, 2023

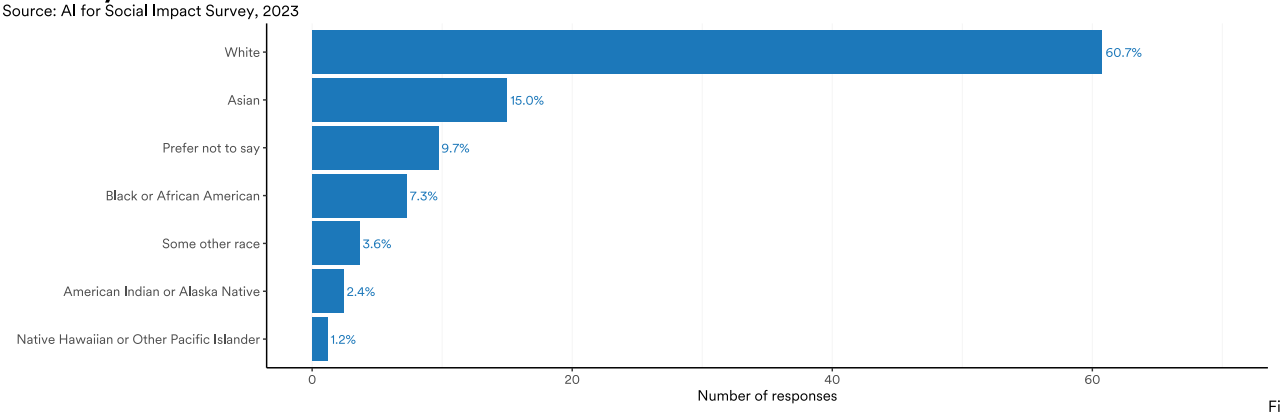

Figure 18

Survey results show a difference in the usage and perceptions of AI based on demographic factors. We found that gender and generational membership impact the frequency of generative AI use. Among our respondents, men tend to use generative AI more frequently than women [Figure 19], and AI use tends to decrease with age [Figure 20]. Though the latter observation is not too surprising and could point to younger adults being more likely to be early adopters of innovations, it may also reflect how generative AI has improved productivity for novice workers while having minimal impact on experienced ones. Additionally, we find that between white and non-white respondents, there is a significant difference in their interest in learning more about AI, with non-white respondents expressing more and stronger interest than their white counterparts, with some white respondents even exhibiting indifference to the technology [Figure 21].

If you have used generative AI before, how often do you currently use it for your job?
Source: AI for Social Impact Survey, 2023

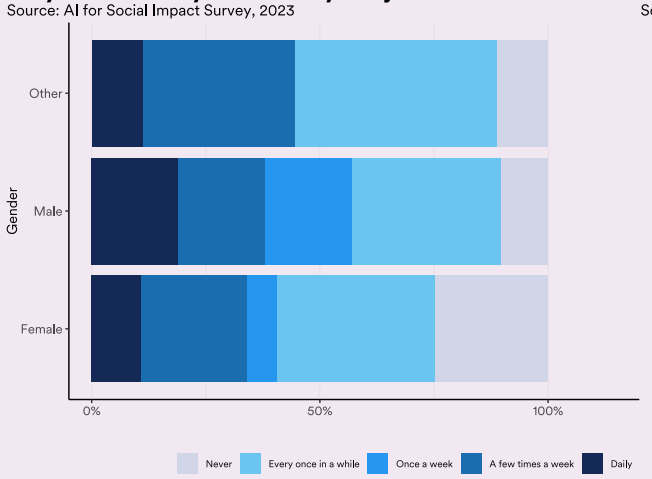

Figure 19

If you have used generative AI before, how often do you currently use it for your job?
Source: AI for Social Impact Survey, 2023

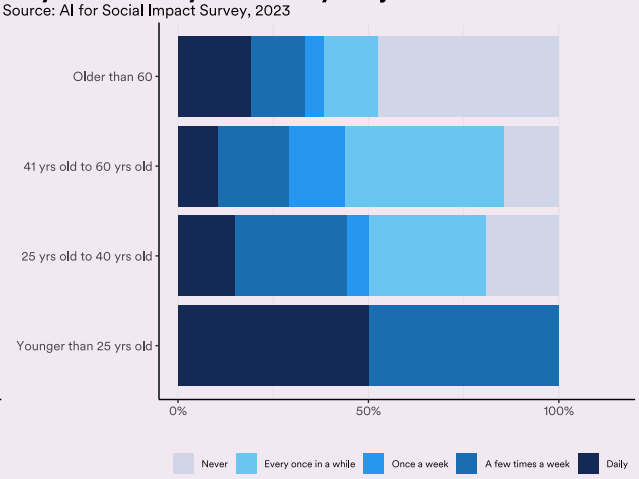

Figure 20

Are you interested in learning more about traditional AI or generative AI?
Source: AI for Social Impact Survey, 2023

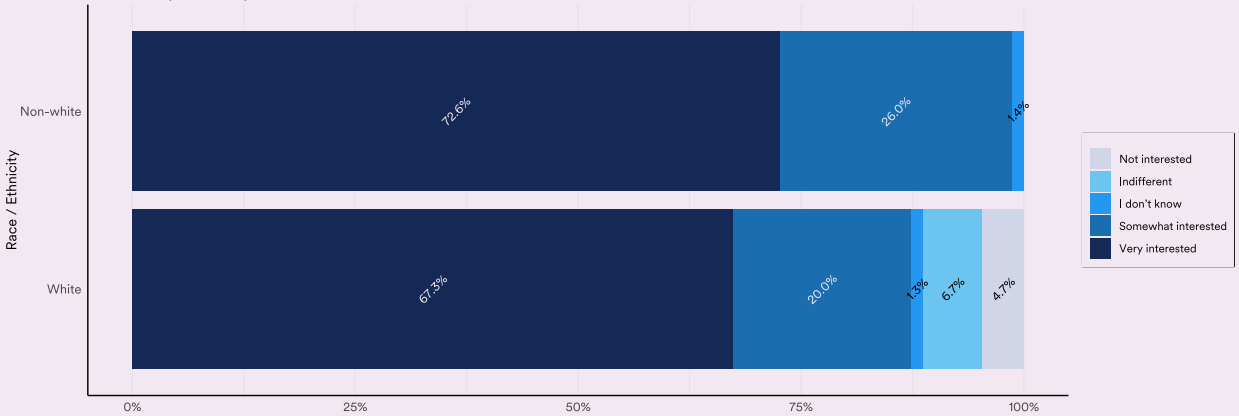

Figure 21

Methodology

Data Collection

The survey was launched in fall 2023 and implemented online using Qualtrics. We exclusively considered respondents employed by nonprofits or grantmakers, constituting 59.1% of those who began the survey. From this group, we narrowed the selection to salaried workers, based on the assumption that full-time staff would have a more informed perspective on current use of, interest in, and opportunity for AI. Salaried workers accounted for 63.2% of those working for nonprofits or grantmakers. Lastly, within nonprofits, we further refined the sample by excluding those in "Medical Research," "Diseases, Disorders & Medical Disciplines," "Public & Societal Benefit," and "Other" services. Approximately 28.7% of interested survey participants met these criteria, resulting in 233 observations.

The survey was shared publicly with newsletter subscribers on the mailing lists of Stanford HAI, Project Evident, the Stanford Center on Philanthropy and Civil Society (PACS), the Stanford Social Innovation Review (SSIR), Leap of Reason Ambassadors community, and the Technology Association of Grantmakers 2023 conference community. In addition to newsletters, the survey was promoted on various social media platforms and at the Nonprofit Management Institute conference organized by SSIR in 2023. Lastly, our wider network of social and education sector contacts played a crucial role in disseminating the survey.

Data Analysis

Open-ended questions on AI interests: Our content analysis encompassed a two-stage process, starting with a quantitative approach that involved tallying word frequencies. Responses were first tokenized—a procedure where they are broken down into individual words—and then the words were standardized to lowercase, with stop words such as "a," "the," and "on" removed to focus on relevant terms. The word "ai" was also eliminated from the responses, given that it was as frequent as it was uninformative for these questions.

The specific bigrams of "predictive analytics," "recommendation engine," "virtual assistant," and "stable diffusion" were intentionally left untokenized. This was because these terms were highly relevant, and frequently used, and allowed us to draw meaningful distinctions, such as between "data analytics" and "predictive analytics." "Virtual assistants" and "chatbots" were treated as the same token. Moreover, where possible, we addressed conjugations, synonyms, and typos, treating variations such as "predictive analysis" and "oredictive analytics [sic]" as the same token as "predictive analytics," just as we treated "chat gpt" as "chatgpt." Ultimately, to enhance clarity in our visual representations, we only considered words that appeared more than five times for nonprofits. For grantmakers, owing to a smaller dataset, words were included if they appeared more than twice.

Following this quantitative phase, we transitioned to a qualitative close reading of the responses to better understand and articulate their underlying motivations.

Cross-Tabulation

Significance test: We employed Fisher's exact test to assess statistical significance. While the chi-square test is customary when evaluating associations between categorical data, the Fisher's test is preferred for smaller sample sizes, such as in this study. Given

the exactness of the test, the Monte Carlo method was used for approximation when the test was too computationally intensive, specifically for the cross-tabulation of Q3 and Q14; and Q27 (generation) in relation to Q12, Q13, and Q14. A significance level of 0.05 was set. In other words, any relationship that meets this criterion indicates that there is less than a 5% probability that the relationship occurred by chance. Relationships with a significance level of 0.1 were also acknowledged and treated as suggesting evidence, albeit weak, of a trend.

Additional notes on classification: Given that not all categories are necessarily relevant, we often selected one or two relevant ones from each response and grouped the remaining ones under the label "Other." For Q1, we compared the responses of "Grantmakers" versus "Nonprofits" to identify significant differences. For Q3, we compared "Education" against non-educational nonprofits, and for grantmakers, we set "Private Grantmaking or Independent Foundation" versus all others. In Q4, we categorized respondents working in nonprofits as either "Executive Leadership" or "Other," while those in grantmakers were classified into three categories: Grantmakers, Executive Leadership, and Others. Demographic data was cross-tabulated by gender (male, female, or other), generation (all categories), and race (white or non-white), with the "prefer not to say" category excluded in each case.

Acknowledgments

We thank the *Stanford Social Innovation Review* and the Stanford Center on Philanthropy and Civil Society for their support in survey distribution. We also thank Rishi Bommasani, Daniel Zhang, Caroline Meinhardt, Christine Raval, Drew Spence, Benjamin Xie, Johannes Eichstaedt, and Kevin Bolduc for their time, insight, and energy reviewing the survey and Working Paper before publication.