Funding the Future

Grantmaker Strategies in AI Investment

Sarah Di Troia
March 2025

PROJECT EVIDENT

About Project Evident

Project Evident harnesses the power of data, evidence, and technology to create greater impact. We enable practitioners (nonprofits, districts, and state education agencies) and their funders to make better use of these resources. We are driven by the belief that when practitioners are better supported, we are all more likely to achieve meaningful and equitable outcomes.

We offer a range of services designed to strengthen individual leaders and organizations, elevating these experiences to develop tools and resources for the field. Project Evident is at the forefront of next-generation evidence practice in the social and education sectors. We center practitioners and the communities they serve, advance an inclusive research and development approach, and embrace scalable, ethical, and cost-effective technology, capacity-building, and knowledge sharing.

Project Evident engages practitioners and funders who share a sense of "constructive dissatisfaction"—those motivated to improve and grow, and who seek real, transformative solutions backed by evidence. Our team takes a holistic and actionable approach to our work, with a strong commitment to delivering shared outcomes. For more on Next Generation Evidence, please see our e-book. (It is free through a Creative Commons license.)

About OutcomesAI

Project Evident's OutcomesAI practice provides consulting, technical assistance, resources, and tools to support practitioners, nonprofits, school districts, and funders. We do this by strengthening their ability to use AI to understand and improve their impact, support decision-making, advance R&D, and allocate resources toward better and more equitable outcomes. We recognize the potential for misuse of data, evidence, and technology and seek to limit harmful practices. Project Evident's differentiator is in the use of AI for Outcomes. We support processes to detect and avoid any technology overriding our evaluative work in delivering equitable outcomes.

AI offers incredible potential to enhance equitable outcomes for communities, so it must be implemented carefully. Our work will ensure that organizations understand the opportunities AI presents to free staff to focus on mission and enhance outcomes and how to integrate it in ways that prioritize equity and transparency. Project Evident serves on the EDSAFE AI Steering Committee and strongly recommends the S.A.F.E. Benchmarks Framework for K-12 AI efforts. Our society needs the social and education sectors' collective voice in shaping policies and laws on AI integration into our economy. Increasing these sectors' knowledge about AI and encouraging safe experimentation will inform and strengthen that voice.

Acknowledgments

We are extremely grateful to the Design Committee, whose guidance, insights, and leadership helped inform this research and report:

- Haifa Badi Uz Zaman (The Rockefeller Foundation)
- Kevin Bolduc (The Center for Effective Philanthropy)
- Mike Belinsky (Schmidt Sciences)
- Brigitte Hoyer Gosselink (Google.org)
- Sophie Owens (Fast Forward)
- Matt Zieger (GitLab Foundation)
- Whitney Williams (Charles and Lynn Schusterman Family Philanthropies)

Several members of the Project Evident team were essential to executing this work, including Sarah Di Troia, Kelly Fitzsimmons, Greg Keough, Rich Leimsider, Annette Raveneau, Stephen Rockwell, Mariko Tada, Lucklita Theng, and Meaghan Birnie.

This report was made possible by Schmidt Sciences, Google.org, and Charles and Lynn Schusterman Family Philanthropies. Thank you for recognizing the importance of this work and giving your insights and resources.

Google.org Schmidt Sciences CHARLES AND LYNN Schusterman FAMILY PHILANTHROPIES

Cover image: Google DeepMind on Unsplash

Research Participants

True to any research on an emerging practice, finding participants for the interview and survey was hard. We know from our research, Inspiring Action: Identifying the Social Sector AI Opportunity Gap, that ~80% of funders in that sample did not have a separate technology priority of funding area, and ~70% did not anticipate creating a technology funding area in 2024. Concern that AI is biased was the most cited challenge experienced by funders in AI adoption; however, uncertainty about who owns AI learning and implementation could also contribute to lagging AI adoption.

To find funders who had made at least one grant to implement AI, we promoted the survey through the networks of Project Evident and the members of our design committee, organizations that serve grantmakers (e.g., GEO, Peak Grantmaking), and funder collaboratives (e.g., Data Funder Collaborative, Grantmakers for Thriving Youth). We collected data about the AI grantmaking practices of 38 U.S. philanthropies through a survey with 34 respondents and 21 interviews with staff members.

Our sample includes organizations with varying budget sizes and focus areas and represents corporate foundations, family foundations, general foundations, and foundations that pool funding from multiple donors. Notably absent from our sample are community foundations. The majority of our sample (88%) comes from the largest 10% of foundations, defined as those with annual grantmaking budgets over $10M.[1] Compared to the foundation community, the relative wealth of the foundations in this research may limit the practicality or feasibility of adopting some documented practices. This insight underpins our recommendations for investment in shared third-party resources to support AI professional development and technical due diligence.

Six funders more heavily engaged in funding AI implementation shared their internal grantmaking rubrics or scoring criteria. We reviewed these artifacts to look for patterns and created a blended investment rubric that includes areas for inquiry and responses they consider satisfying. Although our research findings show varying engagement, confidence, and funding levels, we see a growing recognition of AI's importance.

[1] FoundationSource. (2023). (rep.). *2023 Report on Private Philanthropy*.

Table of Contents

Funding the Future

Grantmaker Strategies in AI Investment

Introduction

Why this Matters: Innovation and the Role of Philanthropy

We are experiencing unprecedented global innovation; worldwide spending on artificial intelligence (AI) is expected to double by 2028, reaching $632 billion.[2] Unsurprisingly, innovation is a core investment in the for-profit sector. However, investment in innovation, especially technology-enabled innovation, is where the social and education sectors have lagged. At this moment, Philanthropy has a unique and crucial role in fueling new AI-powered approaches to enhancing safe and fair outcomes. Philanthropy is often characterized as flexible capital; funders are unique across all sectors in their positioning to take risks to fund promising pilots and experiments. In our February 2024 work, Inspiring Action: Identifying the Social Sector AI Opportunity Gap, we reported that more practitioners than funders (by over a third) claimed their organization utilized AI. Given that nonprofits rely on funders for capital, we hypothesized that this difference in AI use could impede funding for AI experimentation. If funders are not AI-engaged, they risk falling behind and not capitalizing on their distinctive value proposition of being flexible capital and funding social innovation that is not about driving revenue through data collection. In this new research, we sought to learn how the philanthropic sector is grappling with and funding AI to address social challenges. We found some funders are directly engaging with AI initiatives and restructuring themselves to have AI at the core of their grantmaking strategy. Others are proceeding more cautiously, prioritizing research and careful consideration before making significant commitments. There is tension between "doing" (taking action) and "mapping/understanding" (research and reflection), reflecting a broader question in the sector: how do we balance innovation with responsibility?

> "You can always find a reason to not fund something, to not give organizations the space and capacity to be able to do exploratory work because it can feel precarious or dangerous. Our role as the funder is not for us to mitigate every single possible scenario that can happen as a piece of these projects. You need to empower organizations to do that work as part of a trust-based philanthropy approach."
>
> *– Kevin Bromer, Ballmer Group*

What is not in question among funders in our research is that AI will fundamentally change their work and society. Grantmakers and practitioners have an essential voice in discussing how we should incorporate AI into civil society. While the U.S. does not yet have a federal AI policy,

[2] Worldwide Spending on Artificial Intelligence Forecast to Reach $632 Billion in 2028, According to a New IDC Spending Guide. (2024, August 19). *IDC*

states have already enacted over 300 pieces of AI legislation in 2023-2024.[3] However, the only way to develop a point of view on something new like AI is to use it. The pace of philanthropy's involvement in funding AI innovation will directly affect how quickly grantmakers and practitioners develop a point of view to share about AI and its role in our society. Our theory of change is simple: if grantmakers fund practitioners to use AI, then both will be better equipped to support shaping AI's role in civil society.

Project Purpose and Goals: Objectives of the Research

Our research examined how funders approach grants for AI implementation among nonprofits and does not address broader issues such as climate implications, regulatory frameworks, and workforce concerns. We defined AI implementation as funding that could support a technical build of an algorithm, purchase of an AI application or tool, support for an AI technical consultant, purchase of compute, purchase of AI credits (e.g., to run ChatGPT, Claude, etc.), or other investments to implement AI. We focused on AI implementation, distinct from AI research or framework development, to understand how grantmakers gauge the risk and ethics of AI use by practitioners and the reward of potentially enhanced outcomes. Analyzing grantmaking practices, processes, and technical due diligence criteria allows us to share emerging practices and accelerate grantmaking to support AI experimentation across the social and education sectors.

> "While technological advancement often overlooks the needs of the most vulnerable, we see a need to prepare nonprofits to leverage this moment of AI transformation."
>
> *– Shannon Farley, Fast Forward*

[3] Curry, H. (2024, October 22). *2024 State Summary on AI*. BSA TechPost.

Key Insights from Funding the Future

- **Practice Makes Perfect:** Funders who have an AI strategy in place or report that more than 50% of the proposals they review include an AI implementation request are more confident in their ability to assess proposals' technical feasibility and ethical impacts.

- **Evolution of Program Officer Roles:** A significant number of respondents use internal or external support to evaluate the technical feasibility or ethical impact of AI implementation proposals. This potentially points to an evolution of funder staff roles and capabilities.

- **Think Sustainability at the Beginning:** Philanthropy has a poor history of supporting ongoing technology investment, yet AI is not a one-time static investment. Funders are placing a greater emphasis on scalability as it relates to the cost and maintenance of systems.

- **Get Clear on Safe and Fair AI:** When we asked for the top three considerations for assessing ethical impacts, the responses were much more diffuse, as opposed to the same question for technical feasibility, where there was greater consensus on what considerations were most important. Different perspectives and definitions of ethical impacts could send a blurry signal to grantees and AI developers about what matters most to guard against negative ethical impacts.

- **No Dumb Questions – But Some Right Answers:** In our interviews about proposal assessment and analysis of investment rubrics, we saw a coalescence of the types of questions funders think are important and the answers they believe would constitute success.

- **A Big, Hazy Future:** There is broad agreement that AI will significantly change how philanthropic funders undertake grantmaking and how grantees pursue their mission. However, there are many different perspectives on what the future could look like, from optimism about sector-wide capacity enhancement to caution about the need for standards for AI accuracy and quality.

Recommendations

AI adoption represents a massive shift for the philanthropic, social, and education sectors. To promote more coherent grantmaking in support of emerging, broader-scale adoption, funders might consider the following recommendations for strategic investing along the adoption curve to promote confidence building, capacity strengthening, implementation, financial planning and support, and progress monitoring. Based on the survey and interview findings, we recommend six action areas for the social and education sectors to consider:

- **Add Capability to Assess Grants:** Most respondents seek assistance evaluating technical feasibility in grant proposals from staff or trusted third parties. However, more attention to building shared assessment capabilities for all funders could accelerate more confident AI grantmaking across the sector.
- **Strengthen Ethical Technical Plans:** Respondents often frame AI ethical assessments within their existing equity and social justice frameworks. Although this is a good start, AI creates more specific concerns about fairness and safety, requiring stronger collaboration and integration of ethics in technical feasibility. Again, this is an opportunity for funders to collaborate and create a shared training resource to support program officers.
- **Build Trust Together:** Half of the respondents identified community involvement in the design process as one of the top three considerations when assessing the ethical impacts of technology. Prioritizing enhanced community involvement in AI project design will help mitigate risk, build trust, and help reduce bias. (Further details can be found in the AI investment rubric in Appendix).
- **Promote Flexibility with General Operating Grants:** Respondents loosely indicate more general operating and fewer project-related grants to support AI adoption. While the data was difficult to parse, prioritizing unrestricted flexible funding is critical given the emergent nature of AI adoption among nonprofits.
- **Accelerate Learning:** ~60% of funders rely on informal networks or do not engage with peers to improve their capacity to evaluate AI proposals. Given the enormity of the AI opportunity, more formal networks and learning groups, aligned and shared grantmaking practices, and common frameworks and field tools could accelerate philanthropic practice and learning about AI grantmaking. Intentional learning is strategic and necessary as the field advances.
- **Enhance Due Diligence through AI-specific Rubrics:** We recommend that funders create an AI-specific rubric that includes technical and ethical queries and lays the groundwork for measuring outcomes in AI proposals, not just outputs that indicate AI use. Given the newness of AI in the social sector, the tenets of trust-based philanthropy, transparency, and collaboration remain even more important for shared learning.

Funders' AI Grantmaking Strategies

Given the emerging nature of funding AI, does it make sense for grantmakers to have an AI funding strategy? This type of grantmaking is an emerging practice even among grantmakers with experience funding AI implementation. Respondents indicated that a relatively small proportion of the proposals they review contain a request to implement AI to enhance outcomes or productivity. Almost half the sample (47%) reviewed 10 or fewer AI proposals in the past 12 months. However, more than 25% have reviewed over 50, suggesting some institutions are already engaged in AI funding while others are just beginning to explore its potential. Among the six funders who host "open calls" for funding proposals, four have seen the number of proposals for AI implementation increase. One interviewee shared, "So we're in FY 25. In FY 24, (we had) three pure AI grants. Since that time, the number has ticked up quite a bit – we're at 20." It seems reasonable that funders should expect to see an increase in proposals requesting AI implementation funding.

Survey respondents were nearly evenly divided between those whose organizations have an explicit AI strategy (47%) and those who have yet to develop an AI strategy (53%). Of the 18 funders who do not have a strategy in place, seven have them under development. In this context, AI strategy refers to having formalized, intentional approaches to AI-related grantmaking rather than responding to AI proposals on an ad-hoc basis. Given the newness of AI for outcomes, it is not surprising that many philanthropic funders are in learning mode. Pete Lavorini of Overdeck Family Foundation shared, "For 2024, our approach was to pay attention and learn. We specifically focused on better understanding whether generative artificial intelligence could meaningfully enhance teaching and learning in an effort to see what investments and strategies made sense for us to pursue going forward."

Those in our survey with an existing strategy expressed greater confidence in their ability to assess proposals in technical and ethical assessments compared to the complete survey sample baseline. These organizations, perhaps unsurprisingly, have a higher share of proposals with a request to implement AI to enhance outcomes or productivity and review a higher percentage of AI implementation-related proposals than those without an AI strategy. Thirteen organizations shared a synopsis of their AI strategy in the survey. The synopsizes had varying degrees of specificity, from the very high level of "advancing AI solutions" to more detailed sharing of strategic goals.

Within these strategies, there seems to be consensus on:

- **AI to Achieve Preexisting Strategic Goals:** This appears to be a logical expansion of a funder's current grantmaking strategy or mission and recognition that AI is already changing the context in which they and their grantees operate (e.g., education and learning, workforce development, etc.).

> "When we invest in and support AI-powered solutions, we're prioritizing accessibility and positive impact for students with learning differences and multilingual learners."

- **Responsible Development and Advancement of AI:** Most funders who provided a more detailed synopsis of their strategy included what domain goals they wanted to achieve with AI and how AI should be designed and developed to prevent potential negative ethical impacts.

> "Leaders must be intentional about designing for equity – engaging diverse community members throughout the design process, and intentionally mitigating potential risks and biases."

- **Capacity Building**: Many funders recognized the current uneven state of AI adoption in the social sector and called out the need for investment in education, capacity building, and knowledge sharing across practitioners as part of their AI strategy. (See Beyond the Check).

> "Our programs and resources give nonprofits the know-how to use AI ethically and effectively."

- **Policy Development:** A subset of four funders named advancing policy with an eye toward fairness as an explicit part of their AI strategy. We are encouraged by this recognition that we are at the beginning of society's AI journey and the need for the social and education sector's collective voices to be included in this critical discussion.

> "Strengthen emergent coalitions to advance public and private policy frameworks and sector strategies for the equitable deployment of AI."

Confidence

Practice Makes Perfect: Funders who have an AI strategy in place or report that more than 50% of the proposals they review include an AI implementation request are more confident in their ability to assess proposals' technical feasibility and ethical impacts.

Confidence in Assessing the Technical Feasibility of AI Implementation

Confidence helps people trust their abilities and instincts, which is essential for taking calculated risks. Assessing AI implementation proposals from grantees presents a new challenge for some funders. It brings them into the realm of not only having to be domain

experts but also being able to assess the technical merits and feasibility of a proposal. Some funders have planned and hired for this. James (Jim) Savage of Schmidt Sciences said, "I am an econometrician. I studied machine learning before they called it AI and led data science, applying tools that you would now call AI." Others draw on technical expertise from different areas within their organizations: "In terms of general due diligence, we are generalists, so very few of us besides our actual technical folks on different teams probably have any sort of deep technical expertise." However, those who feel confident are in the minority; in our sample, only 36% of those surveyed expressed confidence about their ability to assess the technical feasibility of AI implementation in grant proposals. This is notable as our sample reflects those funders in the top 10% of annual giving who have already begun funding AI implementation. We hypothesize that confidence levels would be even lower for the thousands of program officers not in our sample and have not started funding AI implementations.

CONFIDENCE IN ASSESSING AI IMPLEMENTATION IN GRANT PROPOSALS (n=34)

Confidence Level	In Assessing Technical Feasibility	In Assessing Ethical Impacts
Extremely Confident	**12%**	**12%**
Very Confident	**24%**	**38%**
Somewhat Confident	59%	41%
Not Very Confident	6%	9%

We see positive gains in confidence for those who have more time on task in reviewing proposals. Of those who reported that half of all the proposals they review contain a request to implement AI, 67% reported confidence in their technical assessment. However, if less than half of a program officer's proposals contain AI implementation, their confidence drops to 24%. The greater the investment of time, the greater the confidence, which cuts against funding one-off proposals and encourages thoughtful, scaled investment.

We are mindful that hiring for technical expertise may not be financially feasible for many funders. The confidence levels in our sample reveal a potential impediment to spreading the practice of grantmaking for AI implementations. This indicates an opportunity to create third-party external resources with technical assessment expertise accessible to all grantmakers to support due diligence. This could be a grantmaking resource that multiple funders invest in standing up together or a market opportunity for an enterprising entrepreneur or existing intermediary that already serves foundations.

Confidence in Assessing Ethical Impacts

Survey respondents indicated higher confidence in assessing ethical implications versus technical considerations (14% higher for ethical). Many funders express higher confidence in ethical assessment perhaps because they view it as an extension of existing social impact

evaluation frameworks. Jeffrey Jiménez-Kurlander of Surdna Foundation said, “We feel confident leading our due diligence efforts because the socio-technical aspects of our work align closely with the social justice criteria we use in every grant proposal. While we continuously learn and adapt our approach—especially in AI—our confidence in evaluating these dimensions remains strong.” In several interviews, program officers shared how they leverage established methodologies for evaluating social impact. There's a natural extension of current due diligence practices to AI-specific concerns. However, although respondents expressed confidence about assessing ethical impacts, there was less consensus across our sample about what specific ethical considerations matter most. (See Considerations).

> “Over the last year, we focused on the advocacy side that included collective federal and state level recommendations with our expert grant partners, specifically on the different levers of AI from education to regulation to investments – specifically how AI design, development, deployment impacts communities of color.”
>
> *– Lili Gangas, Kapor Center*

Capabilities

Evolution of Program Officer Roles: A significant number of respondents use internal or external support to evaluate the technical feasibility or ethical impact of AI implementation proposals. This potentially points to an evolution of funder staff roles and capabilities.

Most program officers surveyed seek support for evaluating technical feasibility, 59% receive support from colleagues inside their organization, and 68% gain support from subject matter experts external to their organization. Internally, program officers sought guidance on technical feasibility from technically knowledgeable program officers, dedicated AI subject matter experts, engineering teams, and internal IT staff. This internal technical support ranges from casual “what we have right now is purely informal” to structured. Tyler Sussman of Chan Zuckerberg Initiative said, “Our diligence on grants to enable AI involves input from in-house software engineers and data scientists who work on our education team, and also from a central technology team that supports data privacy security and trust infrastructure questions across our initiative. I would call that our technology diligence process.”

Externally, philanthropic funders have a range of ways they work with technical subject matter experts, from paid to pro bono relationships. According to Evan Trout of the Siegel Family Endowment, “We bring in third parties as part of our inquiry; not as a consultant or in a paid relationship necessarily, but more that when we see somebody who has a keen level of insight and who's willing to take an hour to come and talk to us, we really value that.” Several interviewees underscored the importance of traditional grantmaker skill sets and the opportunity to upskill or teach how to do technical due diligence. “We're just looking for

strategic thinkers. Again, we have a pretty strong belief that the skills required to evaluate a technology proposal are very teachable and actually somewhat generalizable across topic areas as well," said Nick Cain of the Patrick J. McGovern Foundation. Regardless of whether a foundation is accessing third-party or internal expertise to support technical feasibility due diligence, this moment is ripe for the professional development of all program officers.

> "If the AI use is a fit and merits technical review beyond what (our) staff can provide, we are building a panel with legal, technical, and industry expertise to advise staff on retainer."
>
> *– from the Annie E. Casey Foundation's investment rubric*

RECOMMENDATION

Add Capability to Assess Grants: Most respondents seek assistance evaluating technical feasibility in grant proposals from staff or trusted third parties. However, more attention to building shared assessment capabilities for all funders could accelerate more confident AI grantmaking across the sector.

Fewer program officers reported using internal (44%) or external collaborators (50%) to support the assessment of AI's ethical impacts. Those who did seek internal guidance worked with AI ethical knowledgeable program officers, engineering teams with AI ethical expertise, and dedicated AI ethicists. Program officers' increased confidence in assessing ethical impact and the fact that bias and fairness are a core part of their work likely reduces the impulse to seek internal or external support. According to Nick Cain of the Patrick J. McGovern Foundation, "We are constantly reading, discussing, and thinking about that issue internally."

Regardless of whether grantmakers are upskilling, hiring for technical skills, or seeking paid or pro bono support from internal and external experts, technical feasibility and ethical impacts are needed in due diligence for AI implementation grants. The relatively high degrees of internal and external collaboration point to an evolution of program staff capabilities and the opportunity for a third-party resource to support technical due diligence.

Considerations

Think Sustainability at the Beginning: Philanthropy has a poor history of supporting ongoing technology investment, yet AI is not a one-time static investment. Funders are placing a greater emphasis on scalability as it relates to the cost and maintenance of systems.

TOP CONSIDERATIONS FOR TECHNICAL FEASIBILITY

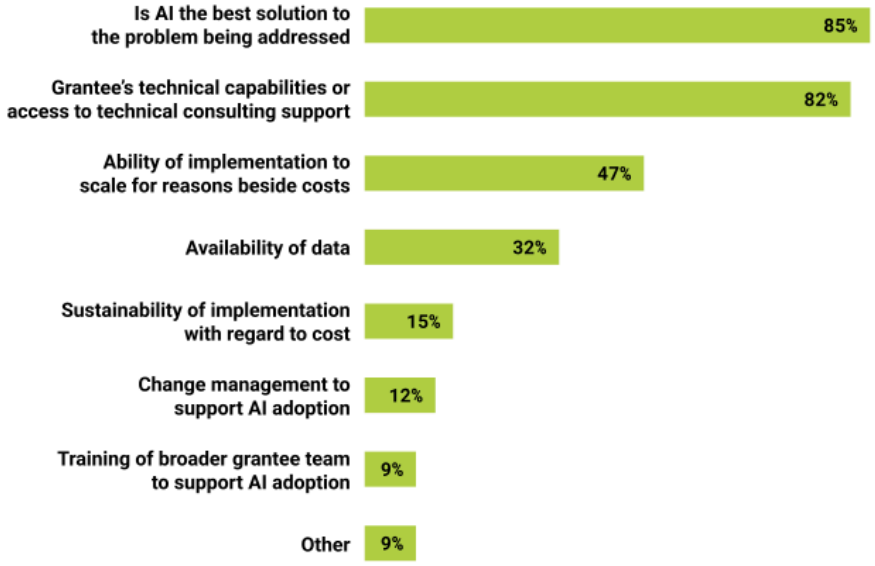

85% = 29 of the 34 respondents selected this.

(These questions allow for multi-selection, up to 3.)

The survey asked respondents about their top three considerations in assessing AI implementation proposals for technical feasibility and ethical impacts. There was consensus on two top considerations for technical feasibility (problem definition and technical talent), and both appear prominently in the draft AI investment rubric (see Appendix), which we developed from the investment artifacts shared by respondents.

> "I feel like sometimes we get so quickly caught up in 'Is it responsible?' 'Is it good?' 'Is it technically feasible?' And that's all important, but none of it matters if the solution itself doesn't seem as though it will enable a very promising application for social impact outcomes."
>
> *– Brigitte Hoyer Gosselink, Google.org*

However, there was less consensus around other considerations. Sustainability is present in the investment rubric, and many of our interviewees discussed it. Still, only 15% of survey respondents selected this as one of the top three considerations in assessing technical feasibility. Ongoing investment in technology has been a historical weakness in philanthropic funding; in 2024, only 25% of nonprofits received foundation grants that included budgets for technology.[4] This is particularly challenging because AI is not an inexpensive or one-time

[4] NTEN and Heller Consulting. (n.d.). (rep.). *2024 Nonprofit Digital Investments Report*.

investment; licenses and compute are ongoing expenses, as is the maintenance of models to contain bias or drift, the decline in a machine learning model's predictive accuracy over time due to changes in the underlying data distribution. Amber Oliver of Robin Hood Learning + Technology Fund said, “I was concerned that there were not very many tools that we had faith were rigorous. And I had even less faith in the defensibility of most of the solutions I was seeing. I was also concerned that nonprofits didn’t have access to the ongoing financial or human resources to be creating those kinds of tech solutions.” Thinking about sustainability from the beginning can help right-size investment and expectations between funders and grantees.

Similarly, the people's side of technology adoption is another area where more attention may be needed, given the potential staff impacts of AI’s increased use within organizations. Funders should look at their organization’s AI adoption and reflect on any staff resistance and their investment in supporting staff adoption and training to guide understanding of the importance of change management and professional development. Matt Zieger of GitLab Foundation noted that, ”Successful AI adoption often involves overcoming significant governance and change management hurdles. Projects have shown that even technically sound solutions can struggle without strong governance frameworks and effective change management strategies.”

Get Clear on Safe and Fair AI: When we asked for the top three considerations for assessing ethical impacts, the responses were much more diffuse, as opposed to the same question for technical feasibility, where there was greater consensus on what considerations were most important. Different perspectives and definitions of ethical impacts could send a blurry signal to grantees and AI developers about what matters most to guard against negative ethical impacts.

TOP CONSIDERATIONS FOR ETHICAL IMPACTS

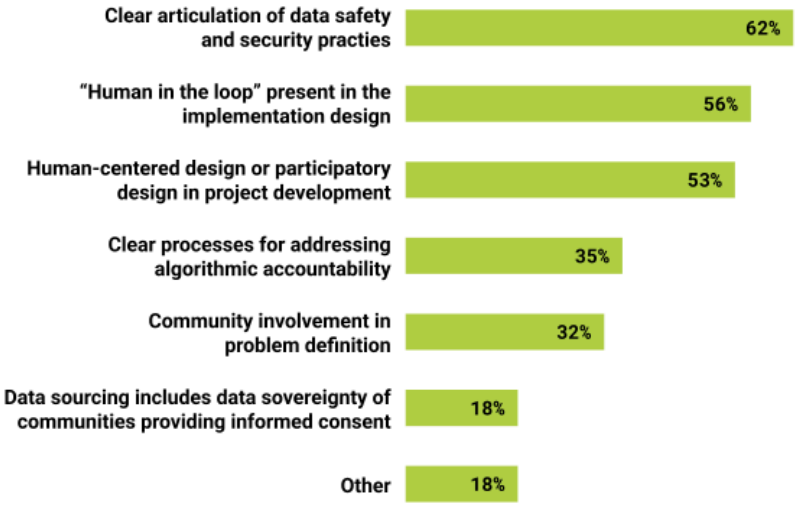

(These questions allow for multi-selection, up to 3.)

Although respondents expressed greater confidence in assessing ethical impacts versus technical feasibility, there was less consensus about what specific ethical considerations matter most. Data safety ranks as the top concern, but beyond this, perspectives diverge. This lack of alignment on ethical priorities may reflect uncertainty within organizations about the ethical implications of AI or a lack of clarity on how to manage risks. The dispersion across funders about what is essential in considering the ethical impacts of AI could send a confusing message to practitioners about what matters most when guarding against negative ethical impacts in AI design and development.

RECOMMENDATION

Strengthen Ethical Technical Plans: Respondents often frame AI ethical assessments within their existing equity and social justice frameworks. Although this is a good start, AI creates more specific concerns about fairness and safety, requiring stronger collaboration and integration of ethics in technical feasibility. Again, this is an opportunity for funders to collaborate and create a shared training resource to support program officers.

Less than 35% of respondents consider community involvement in the problem definition as a top three consideration when assessing ethical impacts, and just over 50% selected human-centered or participatory design. Leaving the community out of the process could fuel mistrust of AI and confuse grantees and developers about what matters most in responsible AI development. If philanthropists believe their voices and their grantees' voices are needed to help shape AI in society, communities that could be affected by AI implementation grants must be brought into the process from the beginning. In the words of one interviewee, "What role do community members play in the development of whatever product you're trying to build?"

"Having people with the lived experience that we're trying to impact at the design table is very important to us."

– *Whitney Williams, Charles and Lynn Schusterman Family Philanthropies*

RECOMMENDATION

Build Trust Together: Half of the respondents identified community involvement in the design process as one of the top three considerations when assessing the ethical impacts of technology. Prioritizing enhanced community involvement in AI project design will help mitigate risk, build trust, and help reduce bias. (Further details can be found in the AI investment rubric in Appendix).

Beyond the Check

21% of our survey respondents provide the majority of general operating support for AI implementation grants. The median foundation in 2022 awarded 32% of total grant dollars as general operating support, and only 5% awarded all of their grants in this way.[5] The data is difficult to parse due to shifting definitions of general operating support and a lack of overall tracking across the philanthropic sector. However, given the emergent nature of AI in nonprofits, it is unlikely that the plans outlined in a grant proposal will come to fruition as expected. According to Matt Zieger of GitLab Foundation, "There is only so much you can assess for viability in early-stage projects without relying on the fact that leadership is going to have to navigate some really important pivots along the way." Program-related support could create an undue burden on practitioners.

> "That's one of the nice things with these general funds and the more flexible funding grants. We are looking at your organization's KPIs, so we are not getting in the weeds of evaluating your use of money."
>
> *– Grantmaker*

General operating support, a form of flexible funding, benefits AI implementation. AI tools are immature, and there are not yet off-the-shelf products that the social sector can easily use. The social and education sectors need entrepreneurial organizations to tinker with many approaches, and general operating support is uniquely good at enabling that space for innovation.

RECOMMENDATION

Promote Flexibility with General Operating Grants: Respondents loosely indicate more general operating and fewer project-related grants to support AI adoption. While the data was difficult to parse, prioritizing unrestricted flexible funding is critical given the emergent nature of AI adoption among nonprofits.

Many funders offer assistance beyond the grant dollars. 53% of survey respondents indicated that they provided other types of capacity building to grantees working on AI implementation projects. Broadly, 65% of U.S. foundations provide capacity-building support to grantees[6]. Given that AI is a new practice for most in the social and education sectors, more, not less than the average, would likely support grantee readiness and the adoption of AI.

[5] Sato, G., & Dayal, S. (2023, September 13). 4 Things We Learned About Foundations and General Operating Support. *Candid.*

[6] Scheid, P., & Helé, K. (2022). (rep.). *How Funders are Strengthening Nonprofit Capacity: Findings from a Field Scan*. The William and Flora Hewlett Foundation.

Capacity-building support reported by our survey respondents included:

TYPES OF CAPACITY BUILDING Response (n=18)	% of sample	n
AI educational content	**61%**	**11**
Access to engineering/developer talent	**50%**	**9**
Access to product development/UX talent	39%	7
AI technical training sessions (e.g., prompt engineering, building AI agents, etc.)	39%	7
Cloud infrastructure usage: compute, storage, AI token usage, etc.	33%	6
General AI mentors	22%	4
Data engineering (e.g., cleaning, infrastructure management, etc.)	22%	4
Other	6%	1

"You get communication support, and for technical projects, we do an assessment of what each project needs and get specialists specifically for those projects."

– J. Bob Alotta, Mozilla Foundation

Seeking Support

Grantmakers are seeking ways to improve their capacity to evaluate AI proposals. Some respondents have found peer learning communities, with 59% reporting being part of informal or formal learning networks. (See Appendix). Given the "all teach all learn" moment, more formal networks could accelerate philanthropic practices and learning about AI grantmaking. Pooled funds with an AI focus allow many grantmakers to learn together through an investment vehicle that de-risks AI grantmaking. Fast Forward brings together 39 funders to support technology-powered practitioners, and the Robin Hood Learning + Technology Fund includes six. Respondents offered up a range of external resources they have found helpful, including articles, frameworks, and tools created by other organizations. (Full list in Appendix).

RECOMMENDATION

Accelerate Learning: ~60% of funders rely on informal networks or do not engage with peers to improve their capacity to evaluate AI proposals. Given the enormity of the AI opportunity, more formal networks and learning groups, aligned and shared grantmaking practices, and common frameworks and field tools could accelerate philanthropic practice and learning about AI grantmaking. Intentional learning is strategic and necessary as the field advances.

Funders have a clear opportunity to collaborate and create shared resources to support program officers and the entire philanthropic sector. This includes creating common evaluation frameworks and building pools of technical expertise. Given the broad-based need for AI education content and that 61% of respondents report already providing such support (see above), this could be a robust area for collaborative investment to upskill the sector.

Investment Rubrics and Assessment

No Dumb Questions - But Some Right Answers: In our interviews about proposal assessment and analysis of investment rubrics, we saw a coalescence of the types of questions funders think are important and the answers they believe would constitute success.

Folks in the philanthropic sector are fond of saying, "If you've seen one foundation, you've seen one foundation." Perhaps it is unsurprising that the unique market in which funders operate, absent competition or traditional customer feedback loops, has given rise to design dispersion vs. concentration of grantmaking practices. Within our sample, we saw grantmakers are developing nuanced approaches to evaluating AI projects. We analyzed six investment artifacts (rubrics and scorecards) shared by respondents. We sought areas of overlap in their topics for assessment, questions to probe, and what they deemed a strong grantee response. Key shared focus areas in their assessments include problem definition, data considerations, and ethical impacts. However, just because it is an AI investment, it does not mean the basics of a grantmaking strategy go away.

"The framework of neglectedness, importance, and tractability is a really helpful one to keep in the back of mind."

– James (Jim) Savage, Schmidt Sciences

Based on those artifacts and our research findings, we have developed a draft grantmaking rubric that addresses technical and ethical considerations in assessing investment in AI implementation. The rubric includes questions for program officers to consider, such as: How was AI determined the most appropriate and effective solution compared to non-AI alternatives? Does the organization clearly understand the technical roles and expertise required to achieve its goals, and does it have a capable team with the necessary skills to design, implement, and maintain the solution? Who is the tool designed to empower, and how does it address rights and equity considerations? Importantly, the rubric also includes what these problem officers deem as strong grantee responses.

The entire rubric can be found in the Appendix.

RECOMMENDATION

Enhance Due Diligence through AI-specific Rubrics: We recommend that funders create an AI-specific rubric that includes technical and ethical queries and lays the groundwork for measuring outcomes in AI proposals, not just outputs that indicate AI use. Given the newness of AI in the social sector, the tenets of trust-based philanthropy, transparency, and collaboration remain even more important for shared learning.

Practitioner View

"Trust-based philanthropy and unrestricted funding – if I could underline two things, it's that."

While funding AI proposals may be an emerging practice for grantmakers, numerous nonprofits have been using AI for years and are seeking funding to support their work. Since grantmaking is, at its heart, a relationship between a funder and a practitioner, we spoke to five AI-powered nonprofits to illuminate their perspectives on raising money for AI.

All the practitioners we interviewed feel enormous pressure to move quickly and innovate with AI. Program interventions in the social and education sector are traditionally human intensive and, therefore, high cost and hard to scale. AI can lower the cost of operations and program interventions, making some areas of the sector more attractive to for-profits. One of our interviewees noted, "AI is moving so fast. If the good guys don't do something about it, the bad guys are, and they will do it." As discussed in the introduction of this report, innovation requires resources, and right now, practitioners are caught between the rapid pace of technology evolution and the slow adoption of grantmakers funding AI proposals. "If we don't have the resources to seize on this moment to innovate, then we're going to be left behind. And then for-profit entities with profit as their main driving factor will take over the market, and we will lose this opportunity." While funders are balancing the tension between taking action and understanding AI, AI-powered practitioners understand the importance of the social and education sectors' support in shaping AI's role in civil society.

"I would strongly argue that nonprofits are some of the best folks to tackle this new technology and think about how to use it for the good of humanity."

The five interviewees have had mixed success in raising money and were able to provide a nuanced perspective of the different types of grantmakers currently funding AI. For one organization, the introduction of AI was a catalyst for significant investment: "Our largest foundation funder before our AI program was around $200,000, and we had two of them. After launching an AI product, we received funding from many different foundations, resulting in

millions of dollars." AI accelerators led by technology companies were the first open calls for AI proposals. Our interviewees frequently cited these vehicles as incredible opportunities to learn from technical experts and their peers. However, there is also recognition that these grants are not entirely altruistic, "they're doing it because they're trying to get your business. Everything isn't free. They're giving you money, but they're also charging you money. It's in their interest to do this." Philanthropy's reliance on corporations to be the first movers in funding AI comes with ongoing costs for practitioners, costs grantmakers may not be prepared to fund sustainably. (See below, Technology is a Program Cost – Not Overhead). This could also contribute to a "leaky pipeline" for funding where accelerators provide start-up capital, but there are no funders willing to invest in growth and sustaining AI applications.

Practitioners understand that funders are generally hesitant to embrace innovation, and pooled funding was seen as a vehicle to bring more funders to the table "because there is hesitancy of being the first; there's a hesitancy of doing something alone, especially when it's a big bet. You really do need the funders to be aligned and making each other feel good and de-risking the investment together." However, while pooled funding can help more grantmakers learn how to fund AI, it also has drawbacks for practitioners in terms of fundraising relationship cultivation and ownership: "The trust level may not be there (with some pooled funders) because they have an incentive to fundraise as well. And in some cases, we're pursuing the same funders." More grantmakers funding AI and more types of grantmakers coming to the table would increase the overall probability of practitioners accessing resources to fund innovation and give practitioners a greater choice in selecting philanthropic partners.

Moving slowly to fund AI or not investing in AI professional development for program officers shifts the responsibility of education from grantmaking organizations to the practitioners they serve at a time when practitioners have constrained resources for their work. Given that funding AI proposals is an emerging practice for grantmakers, it is not surprising that our interviewees have disparate experiences in raising money across funders. This puts the burden on practitioners to calibrate their pitches to meet program officers where they are at in AI knowledge. In addition to the challenge of raising money to fund AI, practitioners "struggle to articulate and explain (their) work at a high enough level that it's comprehended and at a deep enough level that you're trusted." It also might mean they must leave out critical aspects of their approach to designing and developing safe and fair AI to enhance outcomes because "you've kind of got these two ends of the chasm. One where it's just 'tell us about the impact.' And then the other end where it's like, 'tell us about the technology, but not tell us about what is the context of use, what's the scaffolding support, implementation, and human factors" thus limiting full knowledge sharing and partnership between funders and practitioners. And practitioners want full partners in this work, it is new and risky and they need help, "if you are asking me about the dataset and how it was compiled and how I'm taking into account inherent bias towards the global north, you get it. I want to work with you. I'm not teaching you why we need custom datasets."

There are several messages the interviewees wanted funders to hear.

The first is that technology is a program cost – not overhead.

The philanthropic sector has a long history of expecting nonprofit budgets to look magically different from those of for-profits and only require 25% operational costs. Given the high costs of AI, especially generative AI, and the need for ongoing investment to maintain models and keep APIs functional (Application Programming Interfaces are tools that allow two different software systems to communicate with each other), the social and education sectors will not be able to adopt AI broadly unless this reclassification occurs. Right now, "I think that there are a handful, a small contingent of investors who do understand that technology is part of your infrastructure and will fund that," so our interviewees still have to contort themselves to be able to spend the dollars where they are needed.

> "Tech is still seen as overhead rather than program. And so we have to basically not call it technology. Like if you say AI in fancy ways so people pay attention and then don't talk as much about the tech team."

The other critical message is that with tight budgets, AI for efficiency is also important.

One of the essential value propositions of AI is its ability to enhance efficiency across industries by automating tasks, optimizing processes, and providing data-driven insights. Despite the long history of donors and grantmakers scrutinizing practitioner's administrative and fundraising costs, our interviewees were hard-pressed to find funders who saw value in investing in AI to create efficiencies. This stance on not funding AI for efficiencies leaves practitioners once again behind for-profit companies in general technology adoption. "How do we get support from foundations for all these amazing ways that you can use AI and innovative ways to serve your beneficiaries? But then here are all the ways that you could be using AI internally so that you can be more effective and efficient and drive even more impact with the resources that we're giving you." In an ideal world, practitioners would not have to choose to ignore AI efficiencies by only focusing AI for outcomes, and funders would welcome both types of pitches because "it is a win-win for everyone."

> "Get over generative AI, like knock it off. Fund AI, fund technology period, and fund it how it would be best served for the back end or the front end. Because organizations need to become more efficient and more effective to make every dollar count."

Funders should also expect a new line item in AI-powered nonprofit budgets: research and development.

Technology companies spend anywhere from 10-25% of their budgets on R&D.[7] By contrast, federal education R&D spending represents 1.8% of federal per-student funding for K-12 public

[7] How Software Companies Can Get More Bang for Their R&D Buck. (2019, November 22). *BCG*.

schools, and state/local R&D funding is even scarcer, with only one of the five largest U.S. school districts having a dedicated research budget line item.[8] AI-powered nonprofits know they need R&D and are trying to secure investment so they can "experiment in a way that a lot of for-profit companies take for granted having R&D arms and being able to build these hubs within their own organization to experiment quickly and fail fast." R&D is an ongoing and critical investment to leverage evolving technology advances and create new features to enhance outcomes. While R&D investment can be cobbled together through general operating grants, it would be better for funders to embrace the reality that AI-powered practitioners have a product at their core and should have budgets that bear some similarity to technology companies. Because of the experimental nature of R&D, funders need to commit to the tenets of trust-based philanthropy, "I think in innovative spaces like AI, foundations need to have trust in the organization. And they need to let go of traditional notions they have around assessing and evaluating grantees because it's just unrealistic. If they want to fund innovation, they need to do that." That is not to say organizations want to step back from outcomes; they want funders to understand that outcome measurement "is trickier with generative AI than with traditional machine learning to clearly say whether we got it right or got it wrong."

> "Having funders explicitly invest in setting up evaluation benchmarks and testing suites around the specific task or problem as well as capacity building. Setting up those systems early will enable nonprofits long-term to be more successful and produce higher quality impact."

As grantmakers step into learning about AI and funding AI, they have much to learn from the practitioners who are already blazing a path for AI innovation. Funders who truly want to enable and support their grantees will consider including practitioners' experience and perceptions in a shared AI learning agenda for the sector without placing the burden of program officers' learning on their grantees.

Looking Forward

A Big, Hazy Future: There is broad agreement that AI will significantly change how philanthropic funders undertake grantmaking and how grantees pursue their mission. However, there are many different perspectives on what the future could look like, from optimism about sector-wide capacity enhancement to caution about the need for standards for AI accuracy and quality.

Our research indicates a sector grappling with the opportunities and challenges in how it approaches AI. While it offers significant potential to enhance efficiency, optimize resources, and strengthen capacity, there is concern about resource misallocation and quality control. The

[8] Alliance for Learning Innovation. (2024). (rep.). *Why the Federal Investment in Education R&D Must Increase*.

collective sentiment of research participants suggests a maturing outlook that balances optimism with a realistic appreciation of implementation complexities and risks and the reality that AI will not fade as part of the grantmaking landscape.

There is shared optimism for:

Operational Efficiency: "Streamlining resources for constrained nonprofits and increasing internal efficiency that allows founders more bandwidth."

Resource Optimization: "More foundations will begin using AI solutions next year. Which will help unlock resources for the social sector further."

Sector-wide Capacity Enhancement: "AI will make it faster & more affordable for organizations to build scalable technology."

There seems to be caution around:

Implementation Reality vs. Expectations: "I think we're very much in the trough between 'hype' and meaningful implementation."

Need for Standards: "There's a common need for bench-marking around accuracy and quality. The field might benefit from shared guidelines or expectations that folks investing in AI for social impact can align with or sign onto. This could create more clarity for applicants and grantees about what we're looking to invest in across the sector."

Learning Curve Recognition: "Greater universal consideration of technology whereas it has often been a specialty of some and ignored by others."

There seems to be concern about:

Resource Misallocation: I'm worried people will fund AI projects without awareness of whether they're the best solution to specific problems."

Quality Control Challenges: "RFPs will be (are already) majorly disrupted by LLM-generated proposals."

Conclusion

The philanthropic sector is at a critical moment, and funders have an opportunity to shape the development, use, and impact of AI by practitioners. Success will require a delicate balance of innovation and responsibility, technical expertise, and ethical consideration. But above all else, it will require a change in how grantmakers work and fund.

This is a pivotal turning point for society. It is inconceivable that any practitioner or funder can meet the needs of this moment alone. Although collaborative funding, capital aggregation, and pooled funding have been in practice for decades, we sit at a unique moment where the specific nature of AI-enabled innovation requires scaled, shared, and aligned funding. A move toward embracing collaboration as a routine part of fulfilling their missions is a needed change ahead for the philanthropic community. We have made numerous recommendations for places where joint funding for shared learning resources and supportive services will be a faster and more cost-effective way to enable AI innovation and adoption across the philanthropic, social, and education sectors. We hope this research inspires new dialogue and partnerships to achieve this. Perhaps these types of partnerships could lead the way to even bigger collaboration to invest in a much-needed, new shared infrastructure, including data sets, open source code, and cost-effective and climate-friendly compute that will be needed for all members of the social and education sector to move beyond experimentation and into full adoption of AI to enhance equitable outcomes.

We hope this research will generate conversations inside and across funders. We leave you with the following questions and hope that you will share your thoughts with us in an online survey. This survey will **remain open until September 5, 2025**. We will share what we learn from your responses back to the field in the fall of 2025.

1. What would be the most valuable diligence tools, methodologies, and practices for funders as the pipeline of AI-powered funding proposals grows?
2. What resources should be invested in and shared across funders to accelerate learning and AI innovation for grantmakers and practitioners?
3. How can we partner to invest in infrastructure so that the entirety of the social and education sectors can wholly and permanently benefit from AI?

> "I think the great equalizer for everyone (funders and practitioners) right now is that we're all figuring this out."
>
> *– Practitioner*

Appendix

Interview List

- Alex Nawar (OpenAI)
- Amber Oliver (Robin Hood Learning + Technology Fund)
- Brigitte Hoyer Gosselink (Google.org)
- Cameron White (NewSchools)
- Elyssa Lewis (Skoll Foundation)
- Emily Anthony (Salesforce)
- Evan Trout (Siegel Family Endowment)
- Govind Shivkumar (Omidyar Network)
- J. Bob Alotta (Mozilla Foundation)
- James “Jim” Savage (Schmidt Sciences)
- Jeffrey Jiménez-Kurlander (Surdna Foundation)
- Kevin Bromer (Ballmer Group)
- Kevin O’Neil (The Rockefeller Foundation)
- Lili Gangas (Kapor Center)
- Matt Zieger (GitLab Foundation)
- Michelle Shevin (Ford Foundation)
- Nick Cain (Patrick J. McGovern Foundation)
- Pete Lavorini (Overdeck Family Foundation)
- Shannon Farley (Fast Forward)
- Tyler Sussman (Chan Zuckerberg Initiative)
- Whitney Williams (Charles and Lynn Schusterman Family Philanthropies)

Surveyed Organizations

Organizations self-selected their organization type as follows:

Corporate Foundation (n=4)	Family Foundation (n=14)		General Foundation (n=9)	Intermediary Foundation, Pooled Funds, Other (n=7)
OpenAI	William and Flora Hewlett Foundation	The Robertson Foundation	Cooperative AI Foundation	NewSchools (2)
Google.org (2)	Valhalla Foundation	Bezos Family Foundation	Schmidt Futures Sciences	Robin Hood Learning + Technology Fund
Autodesk Foundation	Samvid Ventures	Charles and Lynn Schusterman Family Philanthropies	Skoll Foundation	Dave Thomas Foundation for Adoption
	Siegel Family Endowment	The Tow Foundation	Alfred P. Sloan Foundation	Fast Forward
	Gates Foundation (4)	Chan Zuckerberg Initiative	Annie E. Casey Foundation	Tipping Point Community
	Unknown Family Foundation		Patrick J. McGovern Foundation	Ballmer Group
			The Pew Charitable Trusts	
			The Ian Potter Foundation	
			The James Irvine Foundation	

Org-Type Definitions Used to Self-Select:

- **Family Foundation:** Family members play an active role in setting the foundation strategy and are on the board of directors.
- **Corporate Foundation:** Grantmaking is affiliated with a for-profit corporate entity, and corporate representatives play an active role in setting the foundation strategy and are on the board of directors.
- **Intermediary Foundation/Pooled Funds:** The Foundation is not endowed but raises money from other foundations/corporations/individuals to create an annual grantmaking budget.
- **General Foundation:** No definition provided.
- **Other**: No definition provided.
- **Community Foundation:** No definition provided. (Not selected by any survey participants.)

Methodology

Project Evident designed an interview protocol consisting of 15 standardized questions to structure and guide interviews to most effectively elicit key insights and trends in this emerging space of AI grantmaking. These questions were focused on three key areas: (1) baseline questions to understand the context, (2) grantmaking process (criteria, due diligence, risk assessment, evaluation), and (3) learning and knowledge sharing. Interviewers were instructed to follow this protocol to ensure that each interviewee had a similar experience and that findings could be aggregated across the 21 individual interviews. Interviews were virtual and Zoom-based, with recordings and transcriptions provided by Zoom. The initial raw transcriptions did not easily support thematic analysis; as such, Project Evident used an AI service (Revoldiv) to re-transcribe the video files to capture the text of the back-and-forth conversations better. An additional human review of these AI-transcribed interviews ensured accuracy.

For thematic and qualitative analysis, Project Evident compiled all interviewee responses within the interview protocol framework – both raw transcripts and an AI-condensed version (using ChatGPT to condense the raw transcriptions to the core elements). A human review process ensured the AI-condensed transcripts did not materially change the interviewee's answers. Findings elevated within this report come from ChatGPT thematic analysis (with a human review to ensure no single finding/output was over-indexed and elevated without supporting data) to identify trends and patterns in interviewee responses. This work was compared to non-technical thematic analysis, including tagging similar responses to identify themes, identify additional areas of inquiry, and ensure the accuracy of any AI output. These findings were further tested by Project Evident staff during team discussions.

The Emerging AI survey design process leveraged the interview protocol to ensure these two processes were complementary. A survey with 25 questions was designed and deployed on November 15, 2024, and was open for responses until December 31, 2024. Over 80 individuals completed the survey; however, based on eligibility and sorting (those actively engaged in AI grantmaking), 34 advanced and completed the entire survey. Results were captured in Qualtrics and downloaded to Microsoft Excel, where descriptive statistics were conducted on the quantitative questions. Qualitative results were analyzed following a similar process to that employed for the interviews.

Grammarly Pro was used to edit this report.

Resources

Our survey asked respondents to write in communities of practice and resources they find helpful, but we did not request links to those resources. We have supplied links to the best of our ability; however, any errors in identifying resources are ours.

Formal Communities Of Practice Cited By Respondents	
• AI Taskforce • Opportunity AI • All Tech Is Human • Data Funders Collaborative • FRIDAI • AI in Education Funders Collaborative	• K-12 AI Education Funders • Partnership on AI • Pathways and Workforce Funders Collaborative • Technology Association of Grantmakers (TAG) • TechEquity Collaborative • P150

Other Resources Cited by Respondents:

- AlphaXIV, The Generative AI Ethics Playbook
- Data.org, 3 Rules to Accelerate AI Inclusion and Impact
- EDSAFE AI Alliance, S.A.F.E. Benchmarks Framework
- Educating All Learners Alliance, Prioritizing Students with Disabilities in AI Policy – Policy Brief
- Effective Institutions Project, A Funder's Guide to AI Governance and Strategy
- Fast Forward, The Philanthropist's Guide to Nonprofit AI Investments
- Fund.ai Conference
- Google AI Principles
- Edmund Korley, Building LLMs for the Social Sector: Emerging Pain Points
- Jakob Mökander, Margi Sheth, David S. Watson, and Luciano Floridi, The Switch, the Ladder, and the Matrix: Models for Classifying AI Systems
- Office of Educational Technology, Designing for Education with Artificial Intelligence: An Essential Guide for Developers
- Partnership on AI
- Project Evident and Stanford's Institute for Human-Centered Artificial Intelligence, Inspiring Action: Identifying the Social Sector AI Opportunity Gap
- Rich Sutton, The Bitter Lesson
- UC Berkeley Labor Center
- Newsletters from A-Street, EdTech Insiders, and GSV

PROJECT EVIDENT

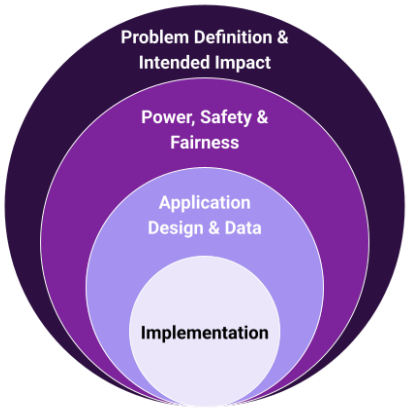

AI Grantmaking Rubric

CATEGORY	SUBCATEGORY	QUESTIONS	ASSESSMENTS
Problem Definition & Intended Impact	Problem Definition	• Is the problem clearly defined, with a well-founded methodology, aligned with the organization's Theory of Change and priorities? • How was AI determined the most appropriate and effective solution compared to non-AI alternatives? • Who are the stakeholders, including the solution's beneficiaries, and how are their needs and definitions of success incorporated into the solution's design and plan?	• Alignment with the organization's Theory of Change • Cost-benefit analysis of other solutions, including non-technological solutions • Thorough landscape and user research to ensure the product addresses an unmet need, incorporating stakeholder and community perspective • Clear quantitative measures of success are determined with stakeholder input and integrated into the project's early stages
	Intended Impact	• What societal value does this application deliver, and how likely would similar work occur without funding? • If successful, how will it impact social outcomes and equity? • Are there liability risks or potentially harmful consequences, such as contributing to the development of more dangerous systems, associated with using this AI?	• Explanation of how the application will foster more equitable outcomes and credible projections of future scale

CATEGORY	SUBCATEGORY	QUESTIONS	ASSESSMENTS
Power, Safety, & Fairness	Power	• What information does the model collect and share, and how does the target audience interact with technology, considering any history of harm or low trust? • Who is the tool designed to empower, and how does it address rights and equity considerations? • Is the use of AI disclosed, can participants opt out, and who creates and monitors the rules governing the tool? • Who owns the innovation, and who can benefit from and build upon it?	• The tool design process included representatives of the community or population it affects, domain experts, and diverse perspectives to ensure its accuracy and relevance • Individuals whose data is used have given consent, and the use of AI is communicated transparently, enabling affected communities to protect their interests and provide oversight
	Safety & Fairness	• What potential negative impacts and risks does the AI solution pose, including the consequences when it is wrong? How thoroughly has the organization explored second and third-order effects? • What mitigation procedures are in place if a red flag is raised during the project's execution? • How will potential harmful consequences or risks be identified, and who is responsible for raising red flags if necessary?	• Active testing to minimize bias for the intended audience, ensuring fairness in the results • Proactive safety and fairness compliance with a robust mitigation plan to address privacy, bias, and other risks throughout the project

CATEGORY	SUBCATEGORY	QUESTIONS	ASSESSMENTS
Application Design & Data	Application Design	• Has the developer/organization considered the tool's deployment in different contexts, and has the model's accuracy and limitations been communicated transparently? • Is there clarity about how the solution works, is it technically feasible, and what evidence supports its ability to deliver on its promises? • How will the tool be tested, rolled out, and monitored, and what methods will ensure its results are clearly communicated to the audience? • What measures has the team taken to optimize the model's accuracy, including precision and recall, and who is responsible for monitoring and reporting its performance?	• Comparison of the AI tool's accuracy to the human process it augments or replaces, and its results are clear and understandable to users and clients • Low-fidelity wireframes, mock-ups, or prototypes to gather feedback from potential users and validate the design with plans to use rapid prototyping and iterative testing • Project plan includes methods to incorporate user feedback throughout the development process, with a clear plan for testing outputs and ensuring the tool meets the needs of its constituents • Consideration of the feasibility of a human-in-the-loop approach, ensuring that it is reasonable to expect human oversight at scale, supported by appropriate training/resources
	Data	**Off-The-Shelf** • Do the tool's data privacy, security policies, and inherent biases align with our definition of success, and how does the organization address potential biases in pre-trained models and test for biased outcomes? **Fine-Tune Model** • How has the organization vetted the data for bias and quality, and do they have the data required to develop a model that accomplishes the proposed task? • How will the organization protect the privacy and security of this data?	**Off-The-Shelf** • Adherence to compliance standards **Fine-Tune Model** • Ownership of necessary data for building a product or training the model • Data has been collected with proper consent, and there is a credible process for regular vetting and cleaning **Custom-Build Model** • Plan to address model drift

CATEGORY	SUBCATEGORY	QUESTIONS	ASSESSMENTS
Implementation	Talent	• Does the organization clearly understand the technical roles and expertise required to achieve its goals, and does it have a capable team with the necessary skills to design, implement, and maintain the solution?	• Leadership support for the project's outcomes • A clear plan for acquiring necessary expertise through current staff, full-time hires, or part-time contractors, with a rationale for their choices
	Cost	• What are the concerns around the solution's long-term maintainability and sustainability, and how will it be maintained financially and technically if it is successful? • What are the costs associated with the tool? Is the funding request appropriate for the scope of work? How will changes or adjustments be handled if needed in the future?	• Vision for a sustainable funding model for the tool's long-term upkeep, including consideration of ongoing costs such as subscription fees or pricing based on usage • Flexible budget, allowing for resource reallocation as necessary • Cost estimates are based on reasonable assumptions, benchmarks, and real quotes from vendors

Definitions

Off-the-shelf AI tools are pre-built, ready-to-use software applications or services designed to perform specific tasks using artificial intelligence without requiring extensive customization.

Fine-tuning a model refers to the process of adjusting a pre-trained machine learning model to perform better on a specific task or dataset by refining its parameters using task-specific data.

A **custom-build** AI model is a tailored artificial intelligence solution designed to address specific business challenges or tasks using proprietary data and specialized algorithms to align with unique organizational needs, ensuring higher accuracy and adaptability to niche requirements.