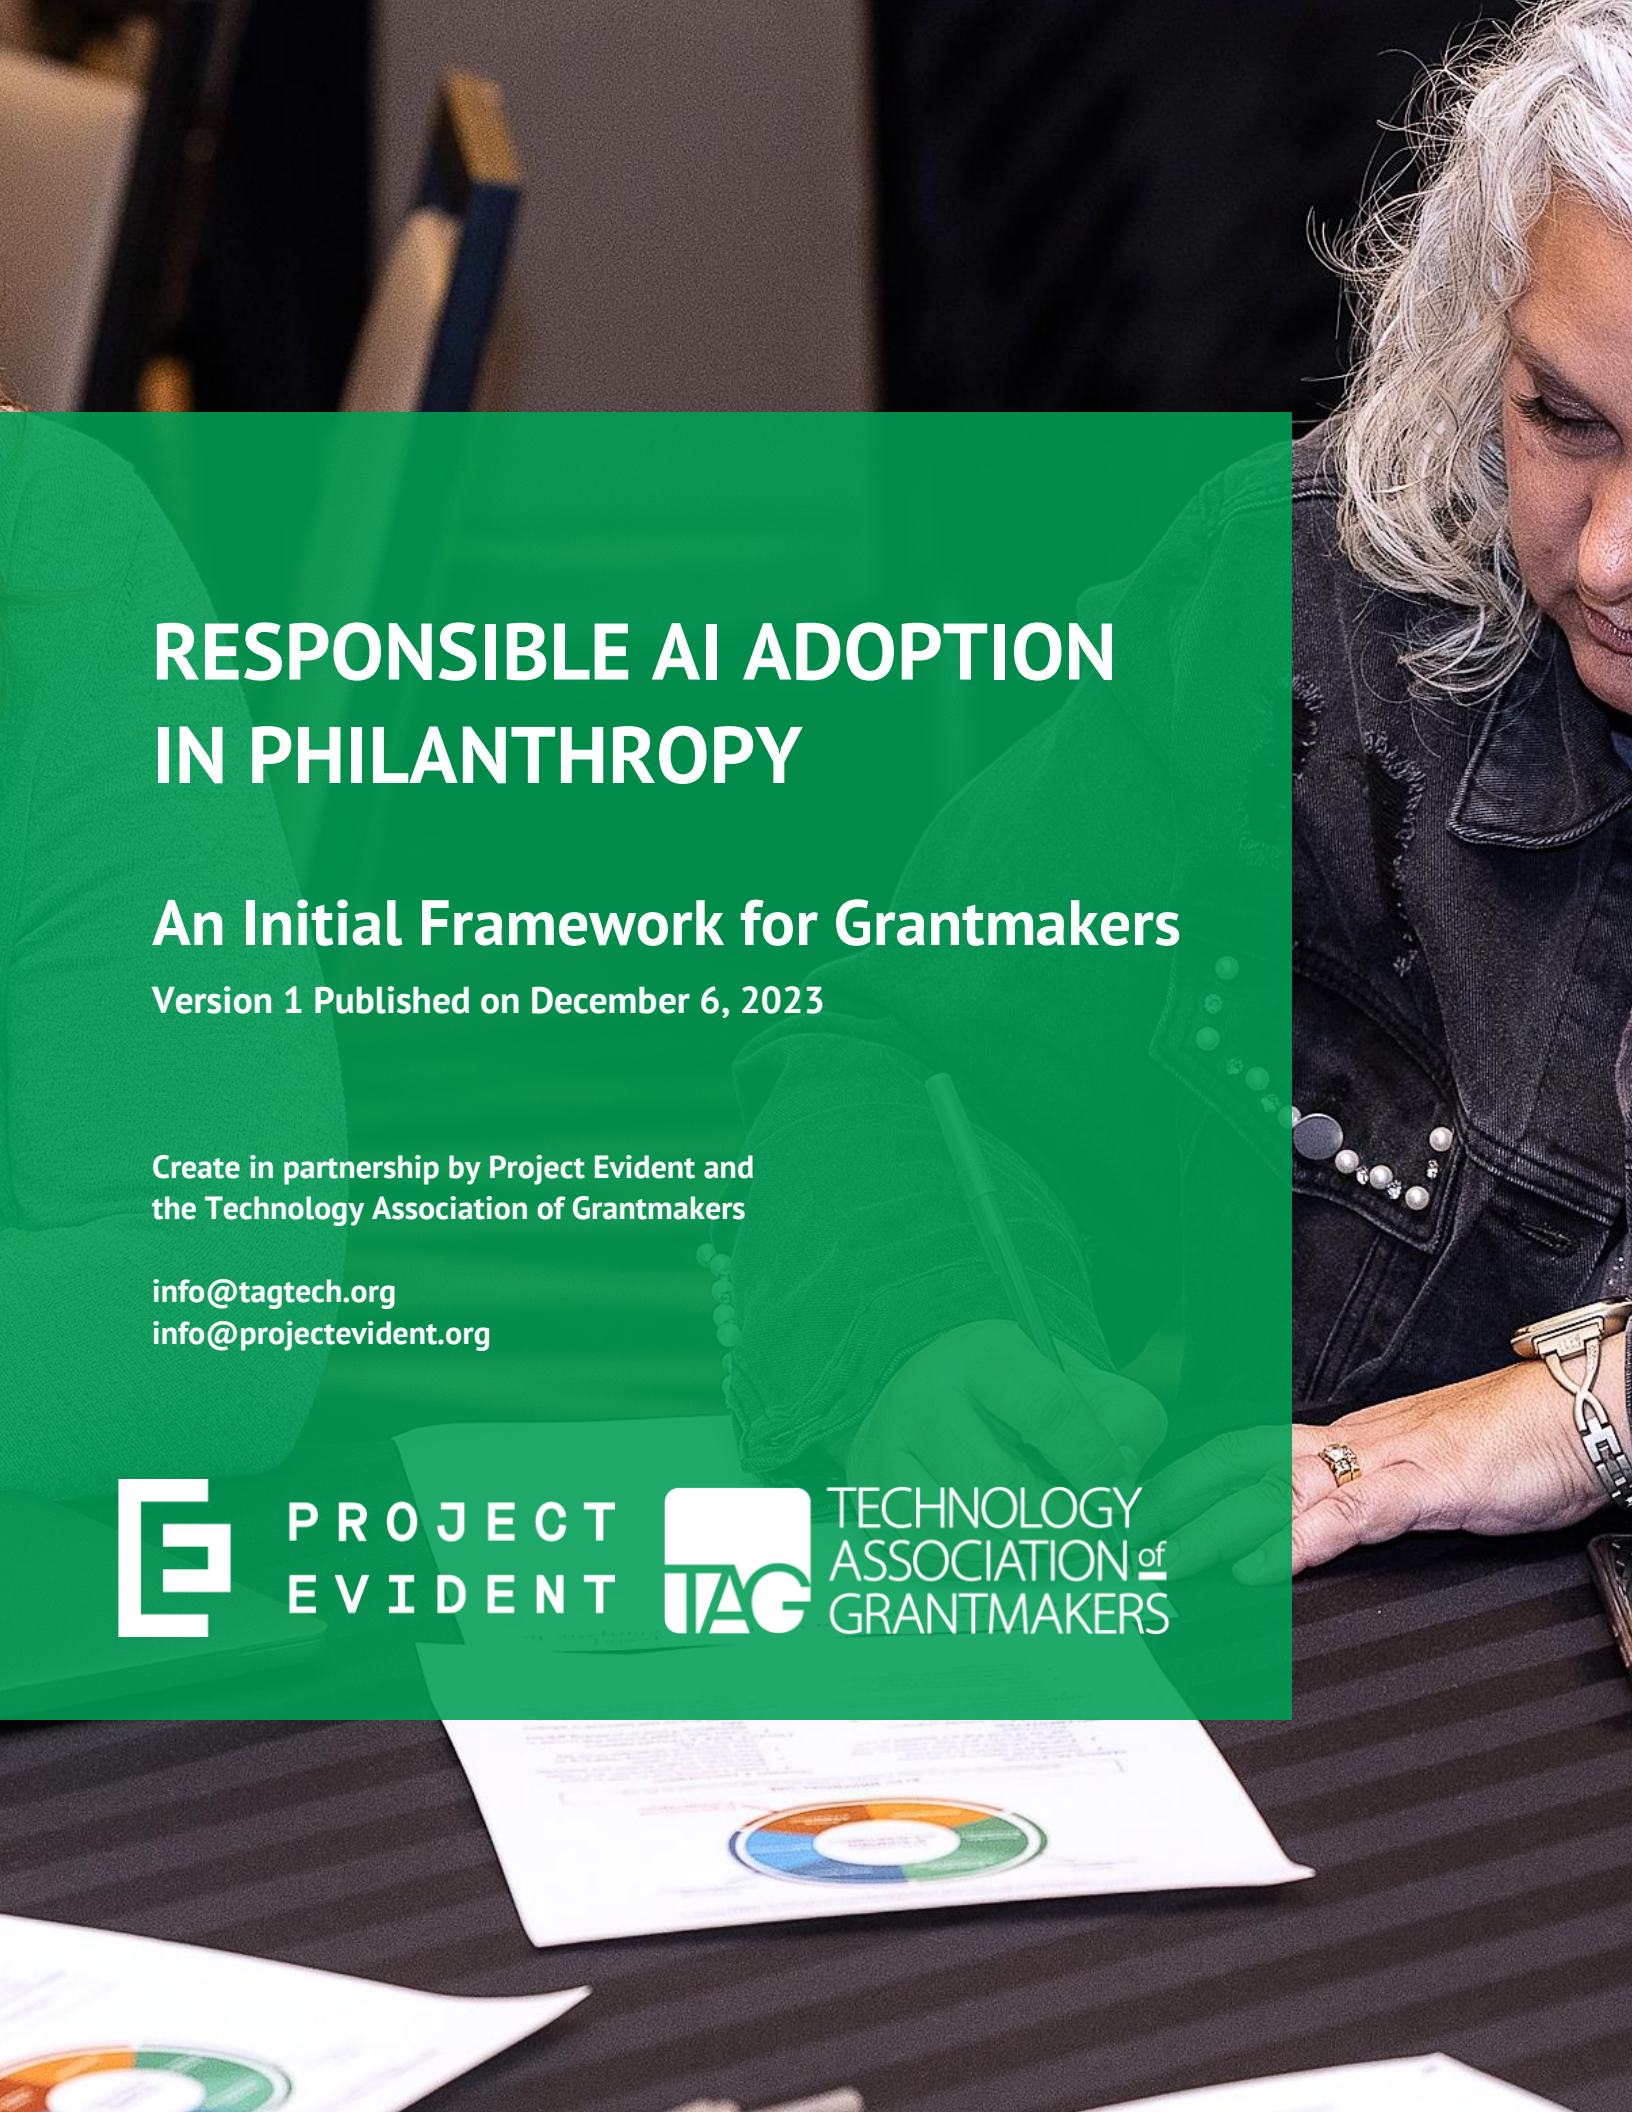

RESPONSIBLE AI ADOPTION IN PHILANTHROPY

An Initial Framework for Grantmakers

Version 1 Published on December 6, 2023

Create in partnership by Project Evident and the Technology Association of Grantmakers

info@tagtech.org
info@projectevident.org

PROJECT EVIDENT

TECHNOLOGY ASSOCIATION of GRANTMAKERS

CONTENTS

OVERVIEW

Responsible adoption of artificial intelligence (AI) is ***human-centered*** adoption.

As society wrestles with the increasing prevalence of AI tools, grantmakers are no outlier in this evolving landscape. The role of charitable organizations is central to determining a future that balances the potential of AI to mitigate inequity and open opportunity with the potential harms to humanity. Grantmakers have a responsibility to ensure that the future of AI aligns with their core values.

While recent advances in the technology landscape may feel swift, AI has been in development for decades and broadly deployed commercially for more than ten years. Grantmakers might think they can take their time in learning how AI can enhance efficiencies or equitable outcomes; however, the reality is that AI is already available to their staff and stakeholders and increasingly integrated into their tech stacks. Because of these societal changes, grantmakers cannot afford to ignore or wait to determine how they will engage with AI in ways consistent with their values.

Responsible AI adoption is human-centered AI adoption.

Many grantmakers are already asking important questions about AI. A survey conducted in August 2023 by the Technology Association of Grantmakers (TAG) indicates that the following questions are surfacing with increasing regularity:

- *How can AI be infused into our organization's operations?*
- *What does an ethical AI policy look like in philanthropy?*
- *How do we conduct due diligence on AI tools under consideration?*
- *How do we integrate AI into grant application processes to reduce burden and increase efficiency?*
- *Can AI be used without compromising data security?*
- *How can AI be used to assess program impact while ensuring accuracy?*
- *How can we ensure AI adoption does not worsen existing inequities?*
- *In what ways can AI be strategically employed to augment our overall impact?*

To begin addressing these questions holistically and to guide grantmakers in the tactical and strategic adoption of AI, TAG and Project Evident have partnered to produce this framework. Despite the numerous ethical AI frameworks from government, industry, and membership groups, there is a lack of practical adoption support for how to address the ethical, technical, and organizational issues that attend the different types of AI usage facing grantmakers. The goal of the *Responsible AI Adoption Framework* is to provide tactical guidance of how to address AI usage and experimentation.

This framework is an initial version based on input from nearly 300 practitioners, leaders, and partners in the social sector. It draws on the lived expertise of the community to ensure that grantmakers of all sizes establish safe and equitable AI policies and practices. Learn more about the advisors and community involved in this work on page 14.

The AI landscape is evolving rapidly, and we expect this v1 framework will also evolve as well. We encourage grantmakers to use this framework to start their journey and build a foundation for ongoing AI thinking and learning.

HOW GRANTMAKERS ARE USING AI

Currently, grantmakers are exploring, evaluating, and in some cases, already implementing AI in three key usage areas:

1. Individual
2. Organizational
3. Mission-Related

Very few philanthropic organizations are yet exploring nonprofit enablement via AI; however, we believe that the mission of philanthropy calls for this investment.

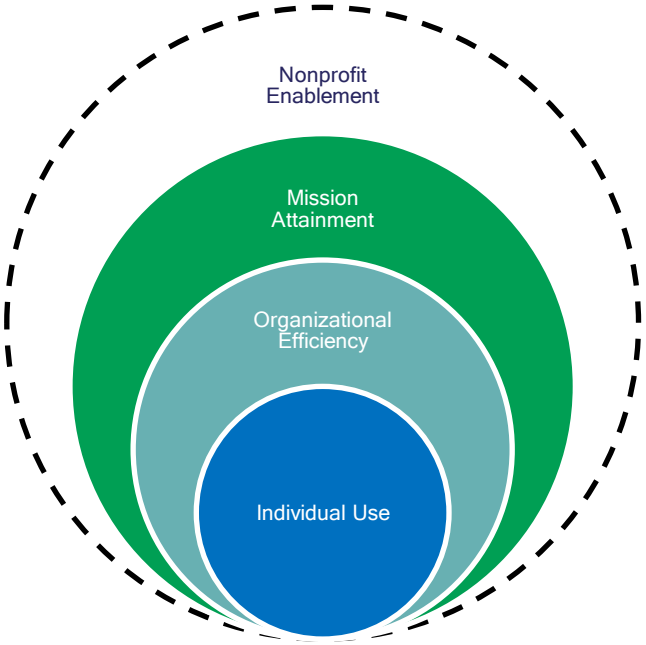

Figure 1: Areas where grantmakers are exploring the use of artificial intelligence tools and platforms

Individual Use of AI refers to staff/stakeholder use of publicly available AI tools that are not provisioned by the organization and may not be easily controlled. These tools are broadly available to the public and available for individual use by staff, contractors, the Board, nonprofit partners, and grantees.

Examples include:

- Inviting Otter.ai or other transcription services to a meeting to take notes
- Using ChatGPT or other public large language models to craft emails, memos, and external reports or to summarize grant proposals

AI for Organizational Efficiency refers to tools provisioned by the organization to save time, predict accurately, and improve operational efficiency. Often, these take the form of AI upgrades to your existing technology stack.

Examples include:

- Turning on transcription and summary services in a Zoom Enterprise account
- Using AI features in grants management software to aid in coding, classifying and summarizing grants
- Deploying Microsoft Co-pilot to analyze and synthesize grantee data

AI for Mission Attainment refers to tools provisioned by the organization for discovery, insight, fundraising, impact assessment, predictive analysis. These are less likely to "appear" in a tech stack and will more likely be upgrades or custom-built tools. However, more off-the-shelf tools are coming…

Examples include:

- Support grantmaking decisions or surface hidden grantmaking bias
- Enable measurement, evaluation and learning across program areas
- Analyze a theory of change for gaps and errors
- Landscape analysis of issue areas, geographies, and demographics
- Universal internal search on prior grants and grantee reports

AI for Nonprofit Enablement refers to the investment in or provisioning of AI tools and solutions for nonprofit partners and grantees.

Examples include:

- Using AI to support fundraising from creating custom outreach, prospecting new donors, and analyzing giving data.
- Deploying chatbots or virtual agents to provide immediate access for clients to routine information or support.
- Using precision analytics to customize program offerings for clients and enhance outcomes
- Support front line staff with recommendation engines to simplify decision-making

While most grantmakers are not yet investing in or exploring the use of AI for nonprofit enablement, we strongly encourage program officers and IT staff to work together in envisioning and investing in this future. As a starting point, grantmakers should prioritize listening and exploration with nonprofits regarding the highest value usage areas for AI in their organizations. Once grounded in this awareness, grantmakers may then choose to invest in AI capacity leveraging strategies such as the following:

- Make grants to support AI professional development and AI use to gain efficiencies and enhance equitable outcomes.
- Fund systematically rather than investing in one-off systems. Fund learning and technology in cohorts with overlapping needs. Share policies, use cases, and knowledge about vendors. Learn more about emerging best practices in funding nonprofit tech at www.tagtech.org/fundingtech.

In all four usage areas, the *Responsible AI Adoption Framework in Philanthropy* described on the following pages offers a starting point for ensuring that your explorations mitigate risk and maximize benefit.

INTRODUCING THE AI FRAMEWORK

Regardless of the usage area for AI noted above – individual, organizational, mission-related, or nonprofit enablement – grantmakers exploring AI adoption must address important considerations as part of their decision-making process. The *Responsible AI Adoption Framework for Philanthropy* groups these considerations into three areas: **Organizational**, **Ethical,** and **Technical**. Each of these must be addressed for AI usage to responsibly center the needs of people.

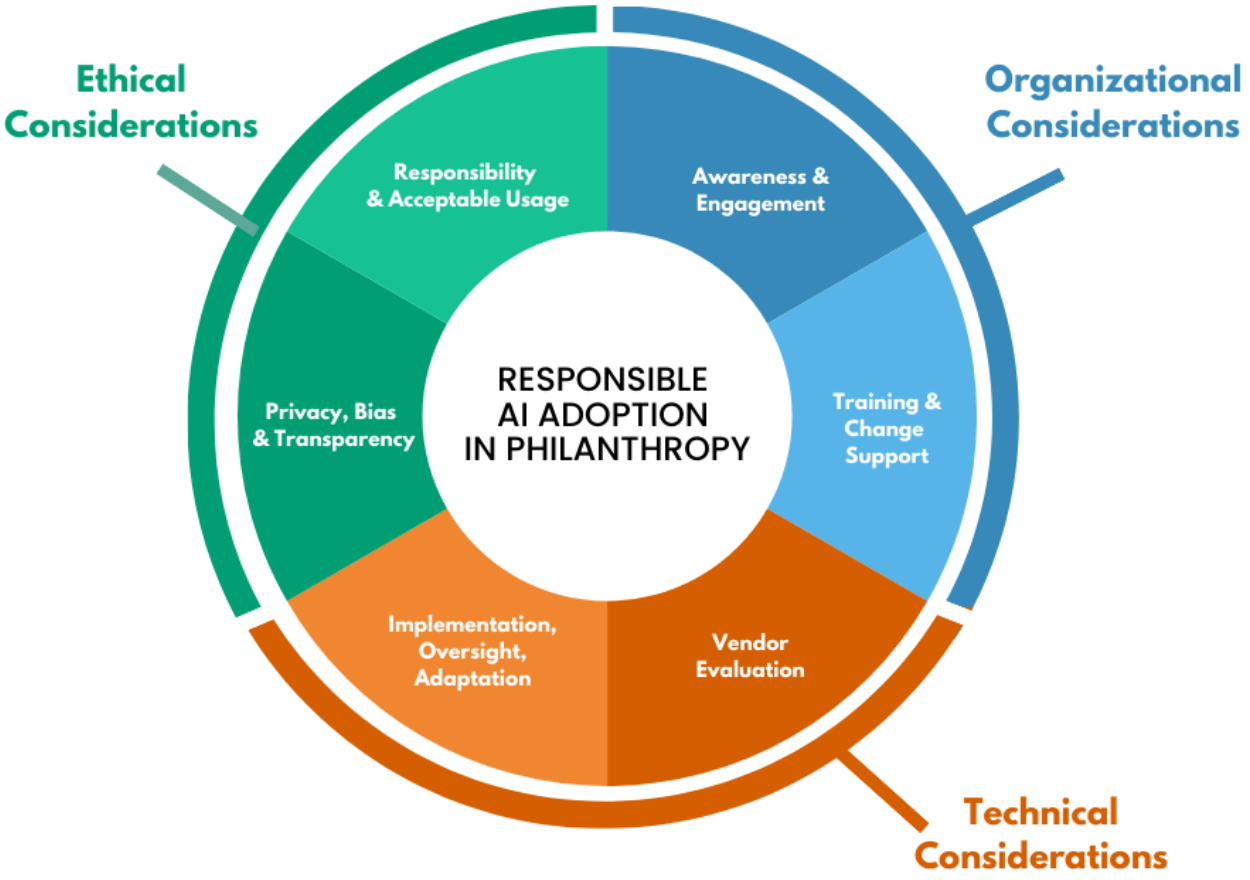

Organizational Considerations: Centering people throughout the AI adoption process begins with building awareness across the organization and engaging a diverse group of users to work alongside the technology staff. Training and change support should be viewed as an ongoing and critical investment to achieve the benefits of an AI tool and avoid potential risks. View a list of organizational consideration on page 9.

Ethical Considerations: Responsibility and Acceptable usage require grantmakers to think about their own values and in what places and cases AI tool usage adheres to their mission and values. This is especially critical for AI tools over which a foundation has control in that it is investing in upgrades or development of AI tools. Privacy, Bias, and Transparency are the ongoing practices a grantmaker will

put in place to ensure the tools are working as planned and to safeguard against risk. View a list of ethical considerations on page 10.

Technical Considerations: Within technical considerations, vendor selection for values alignment and security is the priority consideration before turning to tactical approaches of implementation. Because not all AI tools are specifically provisioned by the grantmaker, vigilance and diligence are required of technology leaders to learn and ask about the security of data and how equity was incorporated in tool development. View a list of technical considerations on page 11.

Each area of consideration is explored more deeply on the following pages.

Exploring and adopting AI is not a linear process, nor is it an isolated one.

Importantly, these considerations are not for technology leaders to address on their own. Human-centered AI requires listening to key stakeholders throughout these considerations including leadership, staff, and grantees who may be affected by the adoption of AI or see fresh potential in its usage. Moreover, AI usage is an iterative process, often beginning with a pilot project that includes open feedback loops, adaptation, and expanded usage only when the application offers real value to the organization and its people. At multiple points in the process, the *Responsible AI Adoption Framework for Philanthropy* may be leveraged to ensure you are including vital considerations in your evolving roadmap. Exploring and adopting AI is not a linear process, nor is it an isolated one.

How to Get Started

If you are just getting started, it's best to begin at the top of the framework with **organizational** and **ethical** considerations. While a technical question such as "Should we use an AI notetaker?" may prompt your journey, it's important to ground your evaluation in identification of the value for your organization, build awareness, and assess risk.

Before applying the framework for AI adoption, first consider the business case. *What is the problem you are trying to solve? How does your organization currently address this problem? What are other approaches to solving this problem that do not include AI? How would AI solve this problem?*

Only if AI is your most practical solution should you move forward.

CATALOGUE OF CONSIDERATIONS FOR AI

The *Responsible AI Adoption Framework* is a dynamic wheel that addresses considerations for AI use at grantmaking organizations. However, depending on the size and technology maturity of your organization, you may only engage with some of these considerations.

As noted earlier, responsible AI adoption is centered on people. You are invited, when exploring the checklist of considerations below, to include human-centered design as well as the impacts on staff, stakeholders, partners, grantees and more throughout your evaluation.

Organizational Considerations

Before addressing technical concerns, it's vital to build awareness, engagement, and support within your organization. Below are key considerations.

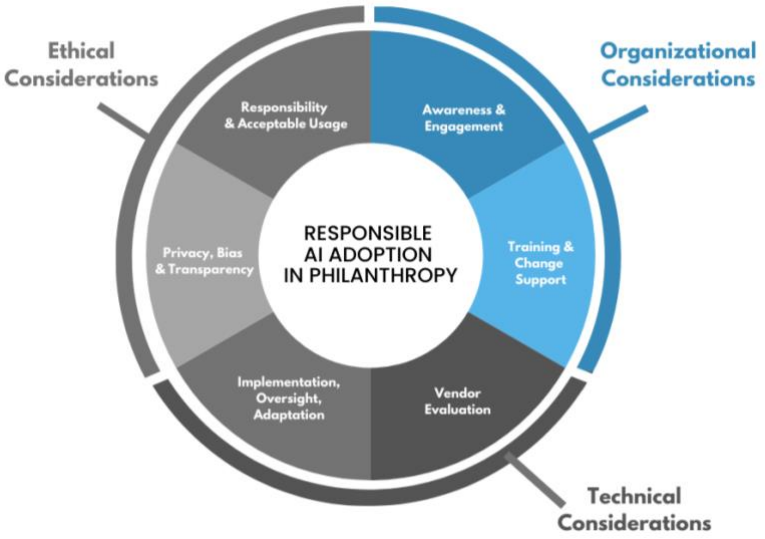

Awareness & Engagement

- ☐ Identify owners/decision makers in the process given your business case. Understanding your business case may help you assign the decision maker and surface other staff and stakeholders who should be a part of your AI adoption process.
- ☐ Find and engage early adopters among the staff. If you do not have early adopters, support a handful of staff to experiment with AI tools.
- ☐ Create a strong two-way learning process between experimenters and IT staff.
- ☐ Launch or work with existing relevant staff taskforce or advisory committee (such as data governance or a super user group) to identify processes that AI could enhance.

As you move into organization-wide usage,

- ☐ Incorporate staff (and where relevant, grantees) in tool selection. Include diversity of roles and relative comfort levels with technology.
- ☐ Have an internal discussion to inform policy and practice regarding anticipated labor outcomes (e.g., staffing impacts, job redesign, job quality effects) and which processes AI cannot replace at an organization.
- ☐ Assess how AI will evolve or sit within mission and culture.
- ☐ Build an internal communication plan to clearly share how and why AI is being used.

Training & Change Support

- ☐ Create a learning hub for staff to share experience, tips, tricks, and losses. Peers could pair to monitor and rate their productivity with and without tool use.
- ☐ Create an ongoing internal resource of case studies/examples of how philanthropies are deploying AI tools/approved examples of AI use in ways that support values.

As you move into organization-wide usage,

- ☐ Host continuous professional development about changing AI tools including different types of tools, how to assess against organization values, tensions regarding ethical/environmental/labor, and highlight good usage. Treat evolving AI tools much like security training.
- ☐ Establish a clear internal communication policy on when AI is being used to generate outputs and when there is a degree of uncertainty associated with outputs.
- ☐ Share case studies (good and bad) on use, ethics, and impacts, and participate in field's efforts to build more responsible tools/understand how to implement more responsibly.

Ethical Considerations

While new AI tools hold an allure, this framework strongly recommends resolving ethical concerns and mitigating risks before moving forward with implementation. Below are key considerations.

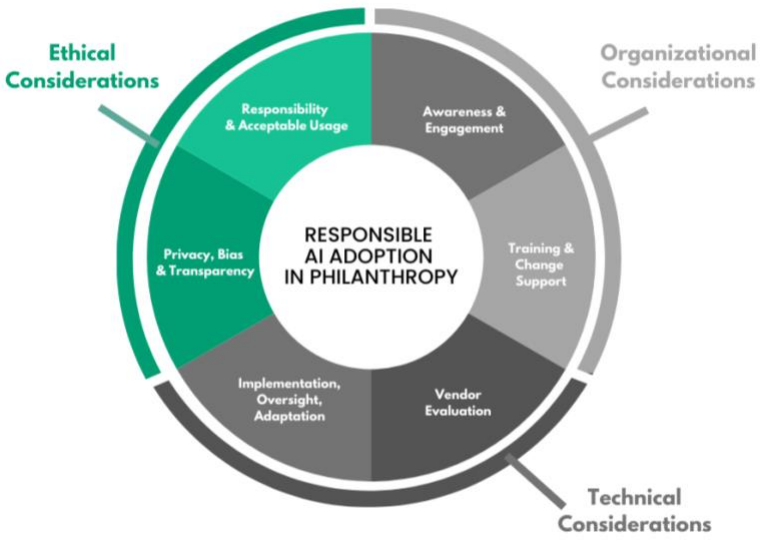

Privacy, Bias, and Transparency

- ☐ Before sanctioning any form of AI usage, even at the individual level, determine how data will be collected, stored, and shared by AI tools.
- ☐ Build or adapt your data governance policy identifying what organizational data can or cannot be entered into non-provisioned and provisioned systems, there can be different rules for different tools. Focus on protecting grantee data and information.
- ☐ Be explicit and make commitments about allowable ethical risks (if any) at your organization.

As you move into organization-wide usage,

- ☐ Launch or work with an existing relevant diverse staff taskforce or advisory committee on AI to monitor privacy, bias, and transparency.
- ☐ Establish a clear external communication policy and practice on notifying when AI is being used to generate outputs and when there is a degree of uncertainty associated with outputs.

- ☐ Create the ability for people/organizations to opt out-in/out of data collection.
- ☐ Establish internal or external resources for human accuracy checks. Create a path for people to request a manual review of AI generated output.
- ☐ Ensure that humans are "in the loop" when using AI technology for relationship management.

Responsibility & Acceptable Usage

- ☐ When evaluating AI usage at the individual level, assume not all staff will use all tools in the same ways; think broadly when creating clear instructions and policies.
- ☐ Ensure that data privacy policies align with data being reported or collected or how AI is being used at the organization.

As you move into organization-wide usage,

- ☐ Ensure that staff are aware that existing guidelines/policies are still applicable.
- ☐ Encourage AI use only where appropriate, by those in roles that can benefit and not be inhibited. Create a list of acceptable/unacceptable use cases for different types of tools.
- ☐ Conduct a technology impact audit that accounts for environmental impacts of AI.

Technical Considerations

After (or concurrent with) your exploration of ethical and organizational consideration, the following technical elements should be a part of your approach to responsible AI adoption.

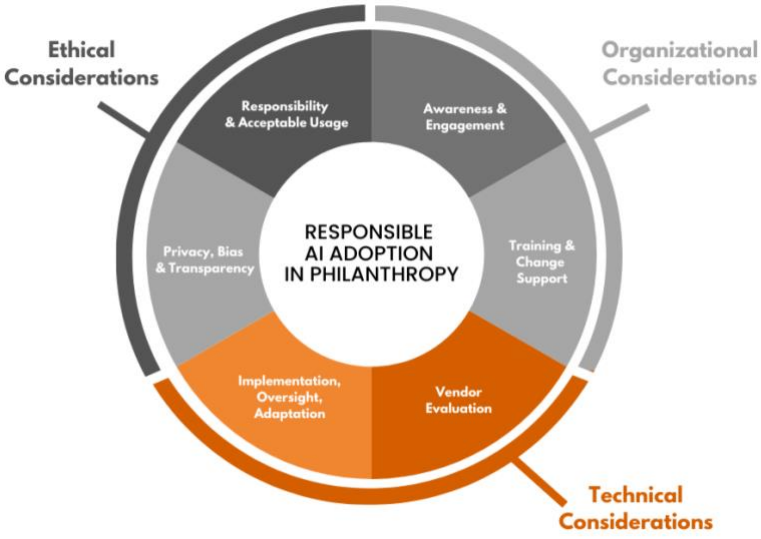

Vendor Selection

- ☐ Review your current data privacy policies. Know the foundation on which the AI adoption framework will build.
- ☐ Work with a staff taskforce or advisory committee to identify AI products potentially in usage at the individual level, organizational level, or for mission attainment.
- ☐ Increase tech fluency for program teams (and program fluency for tech teams).
- ☐ Evaluate vendors based on data privacy, security, and bias. Leverage resources like the Foundation Model Transparency Index (https://crfm.stanford.edu/fmti/) and the Responsible AI Commitments for start-ups (https://www.rilabs.org/responsibleai-commitments).

- ☐ Categorize AI vendors by risk level. Have a clear "do not use" list of tools which are whitelisted on your system. Consider sharing your list with networks such as the Technology Association of Grantmakers to benefit peer organizations.
- ☐ Work with AI vendors who value a diverse workforce, take AI safety seriously, and are transparent about how models work and what data they were trained on.
- ☐ Be clear with vendors on ownership of content and whether your or grantee data will be used for training or product enhancement.
- ☐ Consider a private LLM (your data only) and curate 3rd party trusted data sources for added content.

Implementation, Oversight, Adaptation

- ☐ When addressing individual usage, avoid "full no" policies. Champion controlled experimentation.
- ☐ Select a smaller number of more promising/safe tools to enable experimentation.
- ☐ Create policies and practices that differentiate the types of AI tech, the reasons for use, and the people/communities impacted.
- ☐ Work with staff to select areas to build capacity through automation: edit writing, financial forecasting/ budgeting, or reporting.
- ☐ Invest time in piloting, iterating, red teaming, learning, and human validation/evaluation before full deployment.

As you move into org-wide usage,

- ☐ Create a governance or oversight group of nonprofits with visibility into a funder's usage of AI so that problems can be unearthed and addressed.
- ☐ Budget both for AI implementation AND ongoing risk assessment and governance.

ADVANCING A RESPONSIBLE AI FUTURE

Beyond AI adoption ***within*** grantmaking organizations, foundations have a critical obligation to advance the responsible use of AI for nonprofit enablement. Recognizing the transformative potential of AI technologies as well as the justifiable concern, philanthropy plays a pivotal role in shaping the ethical, technical, and organizational dimensions of AI usage within the nonprofit sector.

In the organizational realm, foundations are invited to invest in workforce development for nonprofits, recognizing the need for skill-building in navigating the AI landscape. Grantmakers might consider identifying or funding case studies of AI usage that effectively address inequity and mitigate or counteract bias.

On the ethical front, grantmaking organizations are urged to take a leadership role by setting practical standards for AI usage of social sector data. While avoiding lobbying and campaign activity, foundations might consider advocating for a regulatory environment that ensures AI usage has adequate oversight and that social sector data is handled responsibly and ethically. Consider supporting AI usage policies that promote transparency, fairness, and accountability within the social sector.

In the technical sphere, foundations are encouraged to fund and publish AI tool reviews or audits, share case studies of AI implementations to foster a collective learning environment, and consider issuing AI certifications for vendors catering to the nonprofit space. Grantmakers might also consider funding investigative efforts into vendor algorithms, development processes, and data ownership, with the ultimate goal of publishing the findings to promote transparency and security.

As a key stakeholder in the future of civil society, philanthropy holds a responsibility to invest in and guide the ethical and effective deployment of AI technologies. Doing so in genuine partnership with nonprofits will foster a collaborative ecosystem that maximizes positive impact and ensures the responsible use of AI for the betterment of society.

PUBLICATION ADVISORS

This publication is an initial draft developed with guidance and insight from numerous advisors. The following people were instrumental in providing perspective and counsel in concert with this framework.

Steering Committee

- Ann Puckett, Director of Information Technology, Grand Rapids Community Foundation
- Chantal Forster, Executive Director, Technology Association of Grantmakers
- Hazem Mahmoud, Director, Data Solutions, Patrick McGovern Foundation
- Jake Porway, Co-Founder, DataKind
- Jonathan Goldberg, Vice President, Learning & Impact, Surdna Foundation
- José Aron-Diaz, Director of Information Technology, Bezos Earth Fund
- Kelly Fitzsimmons, Founder and CEO, Project Evident
- Ngozika Egbuonu, Director of Programs, Technology Association of Grantmakers
- Sarah Di Troia, Senior Advisor Product Innovation, Project Evident

Design Committee

- Afua Bruce, Founder & Principal, ANB Advisory Group
- Amy Sample Ward, Chief Executive Officer, NTEN
- Andrew Means, Founder, Good Tech Studio
- Chad Berg, Director, IT, Bill & Melinda Gates Foundation
- Chigozie Okorie, Director, Business Technology, Annie E. Casey Foundation
- Daniel Miller Runfola, Professor of Applied Science and Data Science, William & Mary
- Daniela Weber, Deputy Director of NetHope's Center for the Digital Nonprofit, NetHope
- Eli Sugarman, Senior Fellow, Schmidt Futures
- Graham MacDonald, Vice President, Technology and Data Science and CIO, Urban Institute
- Guli Basu, Chief Technology Officer, Doris Duke Charitable Foundation
- Julio Lopez, Manager of Technology, Gates Archive
- Laura Goff, Vice President of Information Technology, Marin Community Foundation
- Leon Wilson, Chief of Digital Innovation & Chief Information Officer, Cleveland Foundation
- Michelle Shevin, Senior Program Manager, Ford Foundation
- Muneeb Cheema, Deputy Chief Information Officer, Hewlett Foundation
- Siddhartha Jha, AI and Digital Innovation Lead, Foundation Botnar
- Stephanie Bell, Senior Research Scientist, AI and Shared Prosperity, Partnerships on AI
- Zach Goldstein, Chief of Strategic Innovation & Technology, Full Frame Initiative

COMMUNITY PARTICIPATION

Aligned with the input and guidance of the Steering and Design Committees, sincere appreciation is extended to the nearly 300 participants who evaluated the efficacy of this framework during an in-person workshop at the Technology Association of Grantmakers (TAG) conference in November 2023 in Nashville, TN, USA. Facilitated by TAG and Project Evident, the 90-minute working session engaged grantmakers, vendors, and nonprofit staff in an assessment of the framework's utility within real-world scenarios. The invaluable feedback provided during this workshop has directly influenced every aspect of this initial framework.