WHITE PAPER

Equitable AI Adoption Resource Guide

Understanding Equitable AI Adoption in the Social and Education Sectors

Sarah Di Troia, Sarah Kelley, Laura Schoonmaker

MARCH 2026

PROJECT EVIDENT

About Project Evident

Project Evident harnesses the power of data, evidence, and technology to achieve greater impact. We believe that by empowering practitioners to drive their own data and evidence building while also strengthening the surrounding ecosystem, we can increase the number of effective solutions in the social and education sectors and scale them faster, ultimately producing stronger, more meaningful, and more equitable outcomes for students and communities. Project Evident's OutcomesAI practice provides capacity building, technical assistance, resources, and tools to support practitioners – nonprofits, school districts, and funders. We achieve this by strengthening their ability to use AI to enhance their understanding, improve their impact, support informed decision making, advance R&D, and allocate resources toward achieving better and more equitable outcomes. We recognize the potential for misuse of data, evidence, and technology and seek to limit harmful practices. We serve on the EDSAFE AI Steering Committee and strongly recommend the S.A.F.E. Benchmarks Framework for K-12 AI efforts. Project Evident's differentiator is its use of AI to drive outcomes. We support processes to detect and avoid technology overriding our evaluative work in delivering equitable outcomes.

Acknowledgements

We are incredibly grateful to the teams at Project Evident (Sarah Di Troia, Sarah Kelley, Kelly Fitzsimmons, Peter York, Meaghan Birnie, Mariko Tada, Rhea Tibrewala, and Laura Schoonmaker), Aurora Public Schools (Dusty Burton, Antonio Vigil, Shawn Roloff, Traci Thielen, Johnnie Lacey, Natasha Kellogg), Beyond 12 (Alexandra Bernadotte, Eve Shapiro, Stefan Nilsson, John Lithio, Tiana Sims-Miller), Braven (Moon Lee, Chun Ying Wang), Catalyst Exchange (Rakesh Verma, Jeremy Tullis), Chicago Public Schools (Lorne Rodriguez, Jessica Morris, Casey Fuess, Kelli Johns, Alex Fishman), College Possible (Scott Del Rossi, Derrick Hutchinson), Communities in Schools (Jason Smith, Heather Clauson), Crisis Text Line (Matthew Vanderzee, Margaret Meagher, Vanessa Showalter, Dena Trujillo) CT RISE Network (Julia King, Kristen Negrón), High Tech High (David Hua, Mike Vasquez) New Tech Network, Alix Horton, Sarah Johnson) OneGoal (Daniel López, Kait Sweetman), Per Scholas (Eduardo Hernandez, Tyrone Washington), Quill.org (Peter Gault, Maheen Sahoo, Rebekah Bergman, Daniel Drabik), TalkingPoints (Laila Brenner, Stephanie Kaufman, Emma Berry, Bethany Kreider, Sergio Sanchez), The Calculus Project (Spencer Blasdale, Justin Desai) and members of our Equitable AI Adoption Design Committee (Peter Gault, Sasha Rabkin, Sarah Radcliffe, and Lorne Rodriguez), our subject matter expert reviewers (Farhana Hossien, Erica Van Buren, and Jake Porway) and especially to the Gates Foundation for making this work possible.

WHITE PAPER

Equitable AI Adoption Resource Guide

Understanding Equitable AI Adoption in the Social and Education Sector

Artificial intelligence (AI) and generative AI hold great promise for helping nonprofits expand their services and achieve more equitable outcomes for the people and communities they serve. AI offers a unique opportunity to accelerate progress in addressing persistent social problems by scaling the impact nonprofits can have. Nonprofits and funders share the goal of scaling impact, and AI offers new tools to help achieve it.

A February 2024 working paper by Project Evident and Stanford's Institute for Human-Centered Artificial Intelligence found that approximately 80% of funders and nonprofits believe their organizations would benefit from using more AI, specifically for mission-related work. But knowing how to use AI effectively is another matter. After concerns about bias, the most frequently cited barriers for practitioners and funders were a lack of clarity about how AI will benefit their organizations, and a shortage of sector-specific expertise and resources.

The Equitable AI Adoption (EAIA) Project

The Equitable AI Adoption (EAIA) project aims to fill that gap. To this end, Project Evident is surfacing, creating, and disseminating practical stories of early adopters to understand their AI implementation and distill broadly applicable insights through the development of an EAIA framework and related case studies. This project sought to deeply understand what the core of equitable and effective AI implementation in the social and education sectors looked like by drawing on the experiences of nonprofit practitioners.

EAIA Resources Overview

This project co-created a variety of resources for practitioners and educators to help guide the sector in AI adoption. This white paper summarizes the eight actions they should take to support equitable AI adoption. It also provides a directory of relevant resources that offer inspiration and guidance for each action. The tools and resources created are:

- **The Equitable AI Adoption Framework** defines eight actions that practitioners and educators should take to support equitable AI adoption. The framework was co-created with twenty-six technology and program practitioners representing thirteen nonprofits. Leading subject-matter experts who served as reviewers, and a Design Committee, comprising practitioners and educators, further refined this resource.
- **Five Case Studies & Associated Webinars** featuring organizations at the cutting edge of AI Adoption, highlighting recommended actions from the EAIA framework.

- **Podcast Series** featuring the lived experience of nonprofit practitioners who developed the EAIA across the eight actions that practitioners and educators should take to support equitable AI adoption.
- **A Virtual Convening** held on March 25, 2026, with 13 recorded sessions featuring leading practitioners and educators at the cutting edge of AI Adoption, as well as subject matter experts.

EAIA Framework Elements

Our framework broke equitable AI adoption down into 8 core topic areas.

- Safe and Fair Practice
- Data Governance & Risk
- Design for Outcomes
- Culture of Learning and Managing Change
- Stakeholder Engagement
- Technology Implementation
- Data
- Monitor and Assess

Safe And Fair Practice

AI is already available for free to many nonprofit practitioners and educators. Even if an organization is not ready to implement AI, it can establish expectations for AI use (whether in the form of guidelines or legally enforceable policies) that align with its mission. If an organization intends to implement AI, establishing expectations for its use is the first step in AI enablement.

Resources:

- Chicago Public Schools Case Study
- Chicago Public Schools Webinar
- Episode 7: Safe and Fair Practice: Per Scholas, Catalyst Exchange, PERTS
- Virtual Convening session featuring Kevin Barenblat, Anna Lerner, Rakesh Verma, Afua Bruce

Design for Outcomes

Designing for outcomes begins with pausing to assess whether AI can truly address an identified challenge or opportunity. Some challenges and opportunities lend themselves to AI-enabled solutions; others do not. AI adoption with program activities is most effective when grounded in a clear program logic model that defines the intended outcomes and pathways to achieve them. Within that structure, AI functions as an enhancement to existing work. It can streamline workflows, strengthen program delivery, and expand the quality or reach of services.

Resources:

- Quill Case Study
- Quill Webinar
- Episode 5: Design for Outcomes with College Possible, TalkingPoints, Calculus Project
- Virtual Convening session featuring Tamar Bauer, Scott Del Rossi, Sam Azar, Matt Hillard

AI can also support evaluation efforts, offering insights into how a current model is performing and where changes are needed. When aligned to program logic and outcome goals, AI becomes a practical tool for advancing impact, rather than an end in itself.

Stakeholder Engagement

When integrating AI into a new or existing program offering, organizations should engage program participants or those responsible for delivering programs early in the design process to map problems and potential AI solutions, with a focus on program outcomes. Because AI tools, workflows, and expectations evolve, feedback mechanisms for program participants must remain active. Ongoing listening and structured input are essential to sustaining alignment, building trust, and increasing the likelihood that AI adoption advances mission-driven impact.

Resources:
- Aurora Public Schools and TalkingPoints Case Study
- APS & TP Webinar
- Episode 6: Stakeholder Engagement featuring
- Virtual Convening session featuring Ariam Mogos, Sarah Gripshover, Stephanie Kaufman, Meme Style

Data

Data underpins any effective technology-enabled program model. The quality of available data directly influences whether technological tools, including AI systems, can meaningfully improve outcomes. Importantly, "good data" extends beyond traditional structured data. Organizations often possess valuable insights embedded in materials they already collect, such as transcripts, case notes, and other qualitative documentation.

Resources:
- B12 Case Study
- B12 Webinar
- Episode 1: Data
- Virtual Convening session featuring Farhana Hossain, Pete York, Erika Van Buren

Regardless of format, the integrity of the data is paramount. Data should be as current, accurate, and complete as possible to ensure that any resulting analysis or AI-enabled output supports sound decision-making and advances mission-aligned impact. When thoughtfully organized and analyzed, these assets can meaningfully inform program improvement.

Data Governance and Risk

Regardless of which data sources an organization elects to use—whether they include personally identifiable information (PII) or are fully de-identified—successful AI adoption requires deliberate attention to privacy, security, and bias. These considerations should not be treated as new obligations introduced by AI, but as extensions of existing data governance responsibilities.

Resources:
- Episode 4: Data Governance and Risk
- Virtual Convening session featuring Kelly Fitzsimmons, Margaret Meagher, Erin Mote, Var Shankar

Organizations should conduct a thorough review of the terms of service and data practices associated with any AI solution under consideration.

Ultimately, effective data governance centers on human judgment, particularly in high-risk contexts. AI systems are designed to augment, not replace, human expertise. When AI-generated insights or recommendations have implications for stakeholders' health, safety, or fundamental rights, structured human review is essential. Maintaining this safeguard helps ensure that AI adoption advances mission-driven outcomes while protecting those the organization serves.

Culture of Learning and Managing Change

Organizational change does not succeed without deliberate and sustained attention to the people it affects. Organizations seeking to adopt AI must first assess their readiness for ongoing change and continuous learning. Successful AI implementation depends not only on selecting tools but also on cultivating an environment that can adapt alongside them. Organizations should follow change management best practices and meaningfully engage staff and stakeholders in their AI journey.

Resources:

- Chicago Public Schools Case Study
- Chicago Public Schools Webinar
- Episode 8: Culture of Learning and Managing Change
- Virtual Convening session featuring Tristan Penn, Rich Leimsider, and Albert Chen
- Virtual Convening session featuring DD Di Troia, Lorne Rodriguez, Alina Yang, Summer Jackson

AI adoption requires both ongoing formal training and dedicated time for hands-on practice, experimentation, and reflection. Creating conditions for staff to explore AI's capabilities supports informed use, builds confidence, and increases the likelihood that implementation efforts translate into meaningful improvements in outcomes.

Technology Implementation

AI implementations that meaningfully advance program outcomes often extend beyond the basic functionality of out-of-the-box chatbots. While publicly available tools provide a useful entry point, organizations seeking deeper impact may benefit from customizing AI systems to reflect their specific context and program model. These approaches enable nonprofits to align AI outputs more closely with their mission, client populations, and operational realities.

Resources:

- Quill Case Study
- Quill Webinar
- Crisis Text Line Case Study
- Crisis Text Line Webinar
- Episode 2: Technology Implementation
- Virtual Convening session featuring Stephen Rockwell, Mike Yates, Eduardo Hernandez, Karl Rectanus

At the same time, AI technology continues to evolve rapidly. Organizations must therefore commit to ongoing learning and periodic reassessment of their AI strategy. As capabilities advance, new opportunities for responsible and outcome-oriented implementation will emerge.

Monitor and Assess

Sustained value from AI program implementation depends on structured, ongoing evaluation of model performance to ensure accuracy, validity, and to mitigate bias or unintended harm. For models that operate in high-risk or client-facing contexts, real-time monitoring and clearly defined escalation pathways are essential to protect stakeholders' health, safety, and rights. Human review remains a critical safeguard when AI-generated outputs have meaningful consequences.

Resources:

- Beyond 12 Case Study
- Beyond 12 Webinar
- Episode 3: Monitor and Assess
- Virtual Convening session featuring Sarah Kelley, Laila Brenner, Peter Gault, Pete York

In addition to monitoring output quality, organizations should assess whether AI implementation is advancing intended program outcomes. Evaluating technical performance alone is insufficient. AI systems should be examined for their contribution to impact, ensuring that technological adoption strengthens mission delivery rather than simply introducing new tools.

This framework is not meant to be prescriptive, but rather a starting point for reflection and planning that organizations can adapt to their own context and stage of AI adoption. The resources in this guide were built by and for practitioners navigating these questions in real time. The goal is to go beyond guidance to building shared knowledge and confidence that any organization — regardless of size, resources, or technical expertise — can use AI to deliver better, faster, and more equitable outcomes at scale for the communities it serves.

Appendix

We deeply appreciate the contributions of the practitioners who made this framework a reality. Their knowledge and insights allowed us to root this framework in the needs of the social and education sectors.

Advanced Education Research and Development Fund (AERDF)	Design Committee	Sasha Rabkin, Chief of Program Strategy and Innovation
Altoona Public Schools	Design Committee	Sarah Radcliffe, Director
Braven	Framework	Moon Lee, Chief Technology Officer Chun Ying Wang, Head of Insights and Strategic Planning
The Calculus Project	Framework, Podcast	Spencer Blasdale, Chief Program Officer Justin Desai, Director of School and District Support Dr. Adrian Mims Sr., Founder & CEO
Catalyst Exchange	Framework	Rakesh Verma, Vice President, Product and Technology Jeremy Tullis, Partner
Chicago Public Schools	Design Committee	Lorne Rodriguez, Enterprise Generative AI Manager
College Possible	Framework	Scott Del Rossi, Vice President, College and Career Success Derick Hutchinson, Senior Director, Salesforce
Communities in Schools	Framework	Jason Smith, Vice President, Technology Heather Clawson, Chief Program and Innovation Officer
Crisis Text Line	Framework	Matthew Vanderzee, Chief Technology Officer Margaret Meagher, Chief Impact Officer

CT Rise Network	Framework	Julia King, Monitoring, Evaluation, & Learning Manager Kristen Negrón, On Track Mentor Manager
High Tech High	Framework	David Hua, Chief Technology Officer Mike Vasquez, Principal
Line of Sight Consulting, LLC	Subject Matter Expert	Dr. Erika Van Buren, Chief Executive Officer
New Tech Network	Framework	Alix Horton, Managing Director Sarah Johnson, Coordinator, Leadership Team
OneGoal	Framework	Daniel A. López, Managing Director of Program Kait Sweetman, Vice President Technology & Data
Per Scholas	Framework	Eduardo Hernandez, Executive Vice President, AI Technology Solutions Tyrone Washington, Executive Vice President, Information Technology
Project Evident	Subject Matter Experts	Farhana Hossain, Senior Evidence Advisor Peter York, Chief Data Scientist
Quill.org	Design Committee	Peter Gault, Founder & Executive Director
TalkingPoints	Framework	Laila Brenner, Head of Philanthropy Stephanie Kaufman, Head of Product
Tech:NYC	Subject Matter Expert	Jake Porway, Entrepreneur in Residence, Decoded Futures

Development of the Equitable AI Adoption Framework

The Equitable AI Adoption Framework was built through a highly iterative co-development process, grounded not only in the expertise of subject-matter experts but also in the experience of practitioners on the front lines of AI implementation. The result is a framework that reflects not just best practices, but also the practical realities of nonprofits navigating these issues.

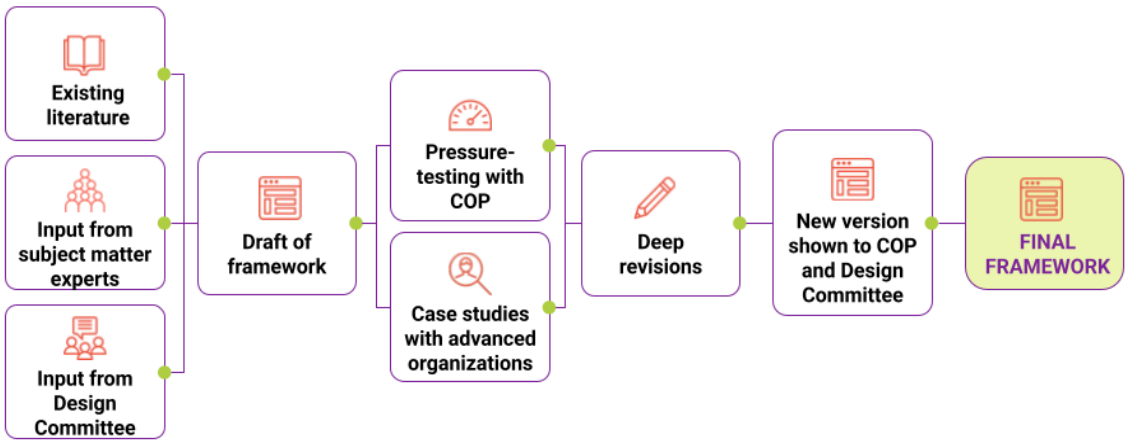

In addition to the content, they had three key pieces of feedback: structure the framework around the different needs that organizations using AI for different use cases (i.e., administrative vs. programmatic) will have; increase clarity on 'what good looks like'; and simplify the language. First, this process revealed that the considerations nonprofits need to make in AI adoption differ substantially depending on how they use AI. Nonprofits that are using AI for administrative, "back office" tasks face very different considerations from those that seek to use AI in their program delivery. For example, having a program logic model is pivotal to effective AI implementation in a program context but not in administrative contexts. And among those using AI for program delivery, there are additional concerns that come about with AI adoption in 'high trust' areas: where AI systems directly face the clients being served, use sensitive data, or make high-stakes decisions. This insight led us to structure the framework around different use cases. Second, we added more clarity on what a 'good' implementation of each category looked like. Finally, we radically simplified the language and sought to avoid excessive jargon, making it interpretable to non-technical audiences.