



Pursuing the Promise of AI: A Guide for African Nonprofits to Advance Their Missions

By Madhuri Mukherjee, Aditi Sharma, Titi Kabi, Ynis Isimbi,
Louis Moussi, Courtney Gehle, and Sohini Pal

OCTOBER 2026

Table of Contents

Introduction	3
About the Research	6
Four Stages of Mission-Driven AI Adoption.....	7
Stage 1: Explore & Prepare.	9
What is the problem we are trying to solve?	10
Is AI the right tool to solve the problem?	11
Do we have the right data, in the right language, for the AI problem we want to solve?	12
Are our teams and leadership confident and comfortable with using AI?.....	14
Stage 2: Prove & Learn	16
Should we build, buy, adapt, or partner?	17
Does our use case work in real-world conditions?	18
How are we defining success, and is our pilot measuring up?	19
Stage 3: Embed & Operate.	21
Can we integrate the use case with our existing programmes and delivery systems?	22
Do we have a working AI governance policy?	22
Do we have the talent and resources to sustain this in the long term?.....	24
Stage 4: Sustain or Scale.	26
What does scale look like for this use case?	27
Are we ready to scale through public systems?.....	27
Missing Links: The Support That Nonprofits Are Asking For	28
Looking to the Future.....	31
Appendix 1. Methodology.....	32
Appendix 2. Select Resources	32
Appendix 3. List of interviewees	33
Endnotes	34

Introduction

When Kenyan nonprofit [Jacaranda Health](#) launched its health service, PROMPTS, its mission was to solve a simple but urgent problem: providing new and expectant mothers with the information they needed to get the right care at the right time.

Launched in 2018, PROMPTS began as a broadcasting service, sending women health information via SMS. When mothers texted back with their own questions, the organisation staffed a helpdesk to respond to them. As more mothers used the service, however, the number of questions grew. Soon, their nurse-staffed helpdesk was receiving as many as 100 questions a day, making it difficult to manage the work manually. There was also the risk that questions needing immediate attention could be lost amongst the routine enquiries that Jay Patel, director of technology at Jacaranda Health, calls “avocado questions” (such as, “can I eat avocado while pregnant?”), which are important, but less urgent.

To address this bottleneck, Jacaranda Health began to experiment with artificial intelligence (AI)ⁱ. Initially, the main purpose was triage – using existing AI tools to prioritise enquiries, answer the “avocado questions,” and escalate to human agents as needed. But Jacaranda Health soon encountered another challenge: language. Mothers were using “code-mixed” language (a blend of Swahili, English, abbreviations, and slang) that was difficult for off-the-shelf natural language processing systems to interpret. In 2022, the team set out to develop a more sophisticated multilingual model that could better understand code-mixed questions and respond more quickly.

Over time, Jacaranda Health began to consider whether PROMPTS could move beyond simply answering questions to proactively managing risk. This led to testing newer features that could provide more nuanced care. For example, its latest model can assess an expectant mother’s risk level based on her clinical history, past conversations, and socioeconomic factors such as household income and education level. In recent years, as generative AI has become more advanced and accessible, Jacaranda Health has developed its own large language models, trained in multiple African languages, enabling PROMPTS to expand beyond Kenya into countries such as Ghana, Tanzania, and Nigeria.

Jacaranda Health’s AI journey was shaped by how it approached two simple questions.

The first is “why?”

ⁱ AI is an umbrella term for computer systems that can perform tasks that typically require human reasoning, problem solving, learning, and decision making. In this report, we use the term to refer to machine learning approaches, including deep learning and generative AI tools (e.g. large language models), and their applications in areas such as natural language processing, speech, computer vision, and predictive modelling).

Jacaranda Health didn't begin with the intent to "do AI"; the technology followed the problem. Sometimes, the answer is that AI isn't the right fit. Many other nonprofits face pressure to adopt AI and not get "left behind." As a result, they often start with "what can we do with AI?" instead of "why?" Debbie Rogers, chief executive officer at [Reach Digital Health](#) in South Africa, says, "I think a lot of people have a little bit of AI blindness at the moment. It's going to solve all the problems, it's cool, there's money out there for it ... and we seem to be forcing the solution into proposals rather than looking for the problem first."

“I think a lot of people have a little bit of AI blindness at the moment. It's going to solve all the problems, it's cool, there's money out there for it ... and we seem to be forcing the solution into proposals rather than looking for the problem first.”

DEBBIE ROGERS, CHIEF EXECUTIVE OFFICER
REACH DIGITAL HEALTH

The second question is "how?"

When organisations begin using AI to advance their missions, they quickly realise the path from identifying problems to seeing results can be winding. For Jacaranda Health, as Patel says, it required repeated experimentation, testing, and retracing of steps. Along the way, the organisation had to grapple with difficult questions: What can AI do for us and our mission? How are people actually using our product? When should staff intervene? Can we afford it? In Africa, these questions become even more challenging, given infrastructure and capacity constraints, limited access to electricity and the internet, and gaps in digital and AI literacy – all of which can make or break AI solutions.

These constraints explain why, in our conversations with nonprofits across Africa, we found that while almost all believe in the potential of AI, few are using it to advance their organisations' missions. This aligns with global trends. Our recent [survey of nonprofits](#), conducted in partnership with [Nonprofit Technology Enterprise Network \(NTEN\)](#), found that 70 percent of nonprofit leaders and staff globally agree they aren't taking advantage of meaningful opportunities to use AI in their organisations. Only 8 percent report having a formal AI roadmap in place.¹

Many organisations are using AI tools for day-to-day tasks that increase efficiency. (See "[Choosing Your AI Path](#).") While that can be a strong starting point for organisations to become more comfortable with AI tools, if AI adoption stops at productivity, much of its potential to advance the organisation's mission – especially in resource-constrained settings – is left on the table.

Choosing Your AI Path

The Bridgespan Group report, "[Choosing Your AI Path: A Framework for Nonprofit Leaders to Make Strategic Choices](#)," identifies three dimensions of AI use:

Augment: How can AI enhance quality, reduce staff burden, increase speed, inform decision making, and enable new ways of working?

Advance: How can AI strengthen, scale, or create programmes and services – and how might organisations need to adapt as AI changes their fields?

Advocate: Where might organisations have a role in shaping responsible AI development and adoption through governance, policy, partnerships, research, and engagement with technology builders?

When we talk about “mission-driven” AI use or adoption in this research, we refer to the Advance dimension, where AI directly strengthens programmes and services. This doesn’t always mean community-facing tools; it can also include decision-making or insight-gathering tools that help improve programme delivery and design. We recognise, however, that some Augment functions can indirectly contribute to advancing a mission – for example, by freeing up staff time – but we haven’t included those in examples of mission-driven AI use.

Impact is possible if early signs are anything to go by. PROMPTS, for example, is integrated with Kenya’s county health systems and has reached four million mothers with support throughout their pregnancies and postpartum journeys. A randomised controlled trial found that it has tangibly improved care-seeking behaviour amongst mothers.ⁱⁱ This, along with several other examples throughout our report, illustrates that when used systematically, AI can help nonprofits innovate faster, reach more people, and strengthen their impact. In fact, some of this is already underway: AI is helping farmers with advice on weather and crop diseases; making diagnostics easier and more accurate for community health workers; and helping teachers create better learning aids for students.

Organisations don’t necessarily need to be “tech forward” or invest millions of dollars in training models to realise this. As the World Bank’s latest *World Development Report* puts it, low-cost “small AI” tools can “for the first time, put scarce knowledge about what matters most for economic progress within the reach of millions – through a text message, a voice call, and a variety of other low-cost but sharp instruments for clinics, courts, schools, farms, and small businesses.”²

ii A randomised controlled trial showed PROMPTS has helped improve care-seeking – 91 percent of PROMPTS mothers report four-plus prenatal visits, compared to 66 percent nationally (Kenya DHS 2022).

In Africa, where many countries face shortages of skilled workers and resources, the impact of AI could be transformative. Even if AI fails to live up to its hype, worst-case estimates from the report suggest that “potential growth rates in developing economies would [still] rise above the dismal average of the first half of the 2020s. At a time when development progress has dipped to a 75-year low, nothing else on the horizon offers as much upside.”³

About the Research


Our research aims to bridge the gap between AI’s potential to advance a nonprofit’s mission and its ability to realise that potential in practice. It does so by helping nonprofits carefully consider the “why” (and whether), and the “how” of adopting AI to advance their missions, so that they can use AI more responsibly and equitably to deliver on their organisations’ missions.

Our research draws on interviews and a focus group discussion with more than 40 nonprofits and social enterprises, funders, experts, and intermediaries, working in sub-Saharan Africa across sectors like health, education, and livelihoods, amongst others, to understand how organisations have used AI to advance their missions. We explored how they started, how their approaches evolved, the strategic and organisational choices they faced along the way, and the wider ecosystem conditions that shaped these journeys.

Of course, this doesn’t mean that every nonprofit can or should adopt AI for mission-driven use. We came across many instances where organisations have chosen not to adopt AI for mission-driven use or to do so in limited ways. One throughline that emerged from our research is that there is no “ideal” way to adopt AI to advance a mission. Our research doesn’t advocate for mission-driven AI adoption as the “end goal” that all nonprofits must strive for. Instead, we aim to distil and share insights from organisations doing this work to show what it can look like in practice, so nonprofits at the start of their AI journeys can make informed decisions about using AI to reach their goals. We hope this research can serve as a starting point for African nonprofits to consider what AI can make possible for their missions and help funders and ecosystem actors identify where investments can make mission-driven AI adoption easier and more equitable.

Four Stages of Mission-Driven AI Adoption

Across our interviews, we found that the journey towards mission-driven AI use typically moves through four stages.



Throughout our conversations, organisations emphasised that mission-driven AI is shaped by its environment. Several external factors overlap with organisational ones to determine when and how AI can be effective. These include existing digital infrastructure (access to internet and smartphones, for example), policies and regulations on AI use, and the availability of local, context-specific resources and mechanisms for reaching communities (such as frontline workers, schools, government platforms, and community networks). The opportunities and constraints of the local context in which nonprofits operate can significantly influence how AI is implemented, sustained, or scaled.

While adoption paths vary, each stage of the journey presents its own set of strategic choices. In the sections that follow, we unpack considerations at each stage as key questions for leaders to reflect on, and describe how these considerations interact with enabling conditions and how nonprofits can navigate these choices.

Explore & Prepare

In this stage, organisations move from general interest in AI to considering whether, and how, to use it to advance their missions. Typically, this involves identifying a specific mission-driven use case – understanding how a user will interact with an AI tool to achieve an objective – that is worth testing and laying the groundwork for its pilot.



What is the problem we are trying to solve?

“One of the things that’s [currently] missing in this space is ... asking, ‘What problem are we solving [with AI]?’”, says Sarah Ndinya, global director of education systems and technology products at [Educate!](#), a nonprofit focused on skills development for African youth. “Instead, it’s often, ‘let’s build it, and then try to fit it in.’”

Many nonprofits begin with a “tool-first” approach to AI, sometimes driven by funder enthusiasm. This can lead to a rush to build an AI solution before they are clear about the problem they are trying to solve. But as organisations across our interviews emphasised, AI needs to follow the problem, not the other way round.

Take, for instance, [VillageReach](#), which works to improve primary healthcare across Africa. Its flagship AI tool in Mozambique began with a problem: the erosion of community trust in public health services, which led to declining coverage of services such as immunisations. To address this, the organisation developed an AI tool that maps community concerns from the national health hotline against existing public health data to identify community sentiment, track changes in behaviour, detect outbreaks, and develop targeted interventions. As Dércio Duvane, senior manager for health information systems at VillageReach, explains: “We did not start with AI; we landed on AI after doing the required analysis around the problem.”

“We did not start with AI; we landed on AI after doing the required analysis around the problem.”

DÉRCIO DUVANE, SENIOR MANAGER,
HEALTH INFORMATION SYSTEMS,
VILLAGEREACH

For [Maisha Meds](#), a Kenya-based developer of software and data for healthcare providers, the technology followed the problem. Its malaria claims reimbursement programme began with a single spreadsheet, but manual auditing processes couldn’t keep up as it expanded. As a result, pharmacies had to wait longer for reimbursements, which, in turn, affected affordable access to quality-assured malaria testing and treatment.

To solve this, Maisha Meds built Madai (Swahili for “claims”), a back-end auditing platform that analyses pharmacy records, automates certain approvals or flags, and reduces the average review time from six days to less than one day. It also partnered with [Audere](#), whose AI computer-vision model verifies images of malaria test results submitted by pharmacies as part of each claim, which was one of the biggest bottlenecks in manual review.

Where to turn for inspiration

For nonprofits that may not be proficient with technology, it can be hard to decide whether AI is the right tool for the job. But at this stage of the adoption journey, it is more of a judgement call based on a hypothesis. In Stage 2, real-world testing and evidence from your pilot might change a “yes” to a “no”. That is a legitimate outcome, not an indication of failure.

Many nonprofits recommend using field experts or advisory board members as sounding boards. See Bridgespan’s [“Choosing Your AI Path: Applying the Framework for Nonprofits”](#) and [compiled social-sector AI resources](#) for additional guidance.

Is AI the right tool to solve the problem?

Taking a problem-first approach also goes hand in hand with asking this fundamental question before rushing in. We heard several variations on that theme in our interviews. Many organisations noted that while it can be tempting to view AI as the hammer for every nail, that might not always be the case. Instead, organisations need to consider whether AI can help deepen their impact or reach more people at lower costs. Nelson Nogueira, chief technology officer at South Africa-based [Tiko](#), which works to improve healthcare for adolescent girls, adds: “Ask how you’d solve the problem if you had no tech at all, then how tech would help, then – only then – how AI would help.”

For an organisation like [Village Enterprise](#), the answer was a clear “yes” for its programme to boost the entrepreneurial skills of rural African households living in extreme poverty.

“Our entrepreneurs rely heavily on mentors. As we scaled our programme, we found individual mentors supporting dozens, and sometimes hundreds, of businesses ranging from retail to animal rearing to crop farming and so on. No single mentor can possess the expertise needed to advise every type of enterprise. That’s where we saw AI adding value, serving as a first point of contact where entrepreneurs can access guidance and answers, while enabling mentors to focus on higher-value coaching and support,” says Meshack Mbinda, senior director of technology solutions at Village Enterprise.

Enabling conditions factor into every answer to this question. Even when AI appears to be a good fit for a problem, it might not be the right answer, depending on the systemic realities in which the nonprofit operates. We found that many organisations chose not to pursue certain AI use cases, or limited mission-driven AI use to only part of their work because the use case did not align with the conditions in which they operated. (See also [Stages 2](#) and [3](#).)

“Ask how you’d solve the problem if you had no tech at all, then how tech would help, then – only then – how AI would help.”

NELSON NOGUEIRA, CHIEF TECHNOLOGY OFFICER, TIKO

Resources to assess organisational readiness to use AI

- [NTEN's tech effectiveness framework](#)
- [Project Evident's AI readiness diagnostic](#)
- [Nethope's Digital Nonprofit AI Readiness assessment](#)
- [Techsoup's Digital Assessment tool](#)

[Ubongo](#), a Tanzania-based nonprofit that creates educational content for children, has limited its use of AI. “Our mission is not to use more and more advanced learning tools,” says Christina Bwana, chief operating and product officer. “It’s to make learning fun and accessible. Right now, AI isn’t accessible in [many parts of] Africa. So, for us, the more traditional [channels like television and radio] are better, because you can access more kids [through these channels].”

Do we have the right data, in the right language, for the AI problem we want to solve?

Organisations highlight that simply “having data” isn’t the same as having data that is usable by AI models. For many organisations, programme data exists but is fragmented, siloed, or not digitised, which makes it unusable. “A lot of [nonprofits] have been collecting data for years, but it isn’t structured in the right way to either train a model or to have a dashboard that has sufficient data for them to make decisions,” says Chido Dzinotiwei, co-founder and chief executive officer at [Vambo AI](#), an organisation that trains AI models on African languages.

In these cases, organisations may first need to focus on digitisation, or even an overhaul of their existing data, before they can consider developing an AI use case. This is what [Imagine Worldwide](#) did. “So, our first investments were in building foundational infrastructure – a very robust and secure data collection and transmission pipeline, from student tablets, to local servers or portable servers, and then to the cloud,” says Rapelang Rabana, co-chief executive officer at Imagine Worldwide, which builds the literacy and numeracy skills of children in Africa. “It was critical for the successful implementation of our current, non-AI, learning programme that the data pipeline was fault-tolerant in the context of intermittent connectivity and power. Without this, we realised we also wouldn’t have the data that would be useful for future edtech innovations, including AI.”

It is not only about having the data infrastructure in place; the data quality is also critical. As Julius Mbeya, co-chief executive officer at Kenyan nonprofit [Lwala Community Alliance](#), observes: “Before deployment, organisations need to ask whether their data is good enough to reduce hallucinations [falsifications] and whether it is sufficient to train models that can generate useful and reliable outputs.”

Gaps in data, especially context-specific data, can also amplify AI bias – its tendency to produce prejudiced results – and reinforce historical inequities. Bias can arise in AI tools due to several factors, including the training data used, the design of the model itself, or existing human biases that seep into how AI models are built and used. Thus, when thinking about the

“right data,” organisations need to also consider where those biases might be present in their training data, and how to mitigate this risk. For example, [Generation Kenya](#), an employment-enabling organisation that recruits, trains, and links youth to meaningful employment through short-term skilling programmes, is cautious about using AI tools to screen and shortlist learners because available tools carry the risk of bias and can unfairly affect employment opportunities.

The data question is also closely tied to that of language. Most AI models are trained in English or other “high-resource” languages. Local-language or domain-specific datasets are rarely available for training AI models in Africa. This leaves organisations with a strategic trade-off: either use existing tools based on existing models that often don’t perform well in local languages and contexts, or invest in creating local datasets to train their own AI models or adapt existing ones, which is often an expensive process. (See “[Should we build, buy, adapt, or partner?](#)” in Stage 2.)

For [Viamo](#), the lack of local-language datasets has been an ongoing challenge to scaling its voice-first generative AI model, Ask Viamo Anything, which makes AI accessible through a simple phone call. “Access to language models restricts where and how quickly we can launch,” says Sulakshana Gupta, vice president of philanthropy partnerships. “There was a time when we tried to create our own language [datasets] from scratch. But we quickly realised that [it’s better] to wait for languages to become available and then be able to plug into them, because we simply can never do enough post-training. Only big tech companies have the [resources] for that. The other thing is cost. To give you an example, in Rwanda, we had a local company that could deliver beautiful Kinyarwanda, but it cost 10 times more than accessing [a big-tech LLM model]. So, [we have to ask] at what point is a language good enough and scalable? This is why Viamo recently joined a global funder and other industry partners in a commitment to ensure that AI reaches last-mile users in their own languages.”

Which approach is most suitable depends on what an organisation wants to achieve with the use case it has identified. As Dzinotyiwai says, “Not every nonprofit needs data for generative AI because they’re not going to train their own language model. Having an objective in place [of what you want to do with AI] can feed into the type of data collection framework you’re going to curate. Once that objective is set, it becomes easier to create a data collection tool, or even just update your data collection pipeline.”

“Not every nonprofit needs data for generative AI because they’re not going to train their own language model. Having an objective in place [of what you want to do with AI] can feed into the type of data collection framework you’re going to curate.”

CHIDO DZINOTYIWEI, CO-FOUNDER AND CHIEF
EXECUTIVE OFFICER, VAMBO AI

Enabling conditions: Public datasets in local languages and regulatory frameworks around data

The lack of availability of local-language datasets is a shared ecosystem challenge across Africa, one that is only beginning to be addressed through shared infrastructure.

Data readiness also depends on the availability of public datasets, which are more readily available in some countries than in others. For example, VillageReach's AI-enabled outbreak prediction model in Mozambique was made possible by both community-reported data from the national hotline and data from existing public health systems.

Regional data localisation and protection laws can also be an important consideration. Storing data, especially sensitive data, locally is a common regulatory requirement. On the one hand, this can be critical for reducing the risk of data breaches and ensuring countries have greater agency over how their data is stored and used. On the other hand, it can considerably limit which use cases an organisation is able to test, even when all other factors are in place.

"Many of the countries that we work with still insist that all health systems [be] locally hosted," says Debbie Rogers, CEO of Reach Digital Health, "which just cuts out many of the experiments that we've been able to run with non-locally hosted models, and increases the cost considerably." The regulatory balance between enabling innovation and protecting people's data can be hard to strike.

Are our teams and leadership confident and comfortable with using AI?

The success of a mission-driven use case also depends on how comfortable and confident staff in the organisation feel about AI. AI literacy goes beyond simply training teams to use AI tools; it also involves helping them understand what AI can and can't do, what constitutes responsible use, and where it can pose risks. "You cannot move from zero directly to AI," says Thomas Kongongo, chief technology officer at [The BOMA Project](#), which works to build sustainable livelihoods for those living in extreme poverty across Africa.

For many organisations, AI use might begin with individual "first adopters" experimenting with available AI tools. But as almost all nonprofits note, leadership buy-in is key, especially when exploring mission-driven AI use. Kongongo says, "Leadership alignment determines [both] investments and the speed at which we [organisations] move. [You can't have] a new technology without a programme champion."

This also helps organisations put preliminary guidelines around AI use in place.

“As leaders, it’s our responsibility to establish clear policies and guardrails for the responsible use of AI,” says Joseph Mudenge, vice president for monitoring, evidence, and evaluation at Village Enterprise. “Clear guidance helps teams understand how to use these tools appropriately and consistently, recognising that unintended risks and harms can arise even when actions are well-intentioned.”


While leadership support is important, some organisations point out that exploring mission-driven AI solutions need not be driven from the top down. Tiko’s Nogueira says that while AI is embedded in its annual goals so “people know how to focus on it,” the exploration of mission-driven AI use is a bottom-up process in which staff members volunteer to test potential use cases. But for this to work, the organisation has invested in levelling the playing field through periodic structured training for staff. “We have [everyone from] highly technical people to people using workplace tools like Slack for the first time,” says Nogueira. “Expecting everyone to suddenly adopt AI without broader digital literacy is unrealistic.”

Some staff may be resistant to adoption, or fear job loss as AI use expands within the organisation. Navigating this requires transparent change management. For some organisations, taking an innovative approach can make this transition smoother. Generation Kenya, for instance, names an AI champion in each of its country offices and gives champions space to experiment and share learnings. For [Harambee Youth Employment Accelerator](#), which works to solve youth unemployment in South Africa and Rwanda, what worked was gamification. “[We organised] promptathons [for] a gamified prompting competition,” says Kuben Nair, chief technology officer. “And we ended up having pretty much the whole company join. It surprised us how much interest and engagement [there was] even without driving it [through formal change management].”

To be sure, this is not the stage to consider hiring a full-fledged AI team or specialised experts. It is enough for organisations to have a basic understanding of AI, its opportunities and risks, and most importantly, the confidence and curiosity to experiment with it.

Prove & Learn

In this stage, an organisation pilots a specific mission-driven AI use case to gather evidence of its practical usefulness, including costs and scalability. The aim is to determine whether the AI use case performs sufficiently well in real-world settings to justify further investment.



Should we build, buy, adapt, or partner?

While some organisations may have reflected on this question in Stage 1, it becomes critical in Stage 2. This classic organisational “build or buy” decision includes options for training and developing an in-house AI solution, adapting existing solutions to meet its specific needs, or partnering with other organisations to co-develop solutions or use open-source resources.ⁱⁱⁱ

There is no simple answer to which approach is “right.” The choice to build, buy, adapt, or partner depends on the problem an organisation is looking to solve, its comfort or expertise with AI tools, and the context in which it operates.

Organisations emphasise that this decision is dynamic – the approach that proves effective in Stage 1 may not remain appropriate as the use case scales. Moreover, the options are not mutually exclusive. For example, an organisation could invest in building a custom AI tool while also using a general-purpose one.

Take Jacaranda Health. When developing PROMPTS, the organisation spent several years building its own language model because no suitable solutions existed at the time. Patel says, “When GPT came along, we tried it. But at the time, it was very expensive, [and] it wasn’t good enough for our context, which is informal Swahili mixed with English. Over the years, frontier models have gotten really good. If I were to start the AI journey today, I would not be building my own model.”

But Patel is quick to emphasise that piloting with its own model had its benefits. “I’m happy that we did start with our own thing because that helped us figure out all of the scaffolding – evaluation pipelines, safety, human-in-the-loop workflows, through which we can deliver AI at scale.”

Audere, a nonprofit that works to improve healthcare decision making with AI tools, has taken a more practical approach from the beginning, customising existing models to suit the context of its work. “Yes, you might need an innovative customisation for your offering, but doing a customisation is different to having to test, vet, quality control, and reinvent the entire wheel,” says Dino Rech, chief executive officer.

[Digital Green](#) followed a similar approach. It invested significant time and resources to localise an existing model to on-the-ground conditions for farmers in specific regions of India, Ethiopia, Kenya, and Nigeria. Its AI assistant for farmers, FarmerChat, is domain-specific and hyper-localised. “We built FarmerChat on existing models and fine-tuned them with agricultural content that is localised and aligned with local language and context,” says Kebede Ayele, vice president for global programme operations at Digital Green.

Building may be more costly and time-intensive at first, but it can be a good fit for organisations with highly specific work. For others, buying, adapting, or partnering may deliver enough value while allowing organisations to retain accountability. Some organisations also recommend

ⁱⁱⁱ Software, data, or other resources that people can freely use, modify, or distribute because the design/code is publicly available.

putting this decision through a practical test. For example, if the use case is completely dependent on a specific vendor, it might make more sense to build. If it's a question of finding another partner if one goes away, it might be okay to buy or adapt.

Does our use case work in real-world conditions?

Even when a use case seems promising in principle, testing it in real-world situations may reveal unanticipated gaps. For instance, when Reach Digital Health tested a generative AI workflow to collect user data in place of a traditional survey, its hypothesis was that the conversational style would yield better data and fewer users would drop off. To its surprise, users preferred the survey. This sent the organisation back to the drawing board to understand where AI could create real value.

This is where organisations might need to re-evaluate the “is AI the right tool?” question from Stage 1 against the higher bar of “is AI the right tool, here and now?” As with Stage 1, a “no” response doesn't mean failure. Even if a specific use case doesn't move forward, it can help organisations narrow down areas to further test. And, crucially, this stage can help organisations reach that conclusion *before* committing resources to scaling the wrong solution.

Nithya Ramanathan, chief executive officer and co-founder of [Nexleaf Analytics](#), which works on data and technology solutions for health workers, stresses the importance of building “operational pilots” designed from the outset to reflect the context in which the pilot will be embedded if scaled. “It is about building something into somebody's daily workflows from the very beginning,” says Ramanathan. She likens the alternative to lending someone a car. “You can drive it around for a little bit, but that doesn't mean you can [keep it running]. If from day one it is really your car, you have to drive it, [fill] the gas, and change the oil. That is what I mean by an operational pilot.”

Context-specific testing can also reveal gaps in particular parts of the use case. For Tiko, models developed in the United States had to be retrained with locally relevant data. “When we introduced facial recognition ... it did not work in the African context. So, we had to collect images and [have our partners] train their model. The same went for voice biometrics – we had to collect over 30,000 voice samples to train the model on our context,” says Nogueira.

The hidden costs of pilots

Pilots often benefit from funding support that may not be available at scale. Many organisations note that their AI pilots have received subsidies, grants, or credits. This, combined with the small number of initial users, can make use cases appear less expensive than they are. Once these pilots are embedded into programmes, however, costs can rise substantially – both for the tool itself and for the overhead and human resources required to maintain it. This may prompt organisations to re-evaluate the build/buy/adapt/partner decision. In some cases, if costs prove too prohibitive, some organisations might also reconsider adopting a mission-driven AI use case. (See [Stage 3](#) for more on operating costs.)

Enabling conditions: AI literacy of ecosystem actors

As in their own organisations, many point out that AI literacy across funders, governments, and other partners is also uneven. This makes early decisions around use cases more difficult since other actors may have a limited understanding of how to use AI to advance their missions. As **Rapelang Rabana**, co-CEO of Imagine Worldwide, says, “I can move my teams at a lightning pace in terms of applying AI solutions. But ministries in many countries are still building the digital fluency and expertise around AI. We’re working closely with our government partners to help them become better-equipped to assess new tools, manage risk, and determine adoption. Investing in government capacity is crucial to drive responsible adoption.”

There is also an equity consideration here. Training programmes, particularly those aimed at frontline service-delivery programmes, often do not fit the practical realities of their work. **Nelson Nogueira**, chief technology officer of Tiko, says, “An ecosystem coordinator on the ground is packed [with work] from nine to six. They can’t [spare the time] to go watch all these [training] videos.”

How are we defining success, and is our pilot measuring up?

The next consideration is assessing progress. As with many other technological solutions, organisations often conflate usage with success. While usage indicators – such as the number of prompts or active users – can serve as useful proxies, mission-driven AI use cases can also assess whether the pilot is moving towards the outcomes the organisation hopes to achieve.

For example, when Educate! launched its AI livelihoods coach, its aim was to help young people through the critical survival period of the first three months after starting their own business. Its AI chatbot coached young entrepreneurs through one of the major challenges in this period: profit and expenses, specifically credit management. But instead of simply tracking whether the youth used the tool, the organisation chose to monitor daily income and credit data from users. Over a 10-day period, users saw a 57 percent uplift in their profit, validating that they were on the right track.

Outcomes, however, can take time to manifest, and may be difficult to measure at this stage. In such cases, organisations recommend creating strong community feedback loops to gather qualitative input. Digital Green, for example, built farmer feedback loops into its FarmerChat AI product to refine language, advice, and product quality from the outset. It also tracked whether farmers actually applied the advice, rather than merely tracking whether they used the service.


There is also a return-on-investment consideration to be made around outcomes – that is, whether the results come at a cost that is viable for further investment and scale, and can be sustained over time. (See [“Do we have the talent and resources to sustain this long term?”](#) in Stage 3).

Early guidelines on using AI safely and responsibly

We noted the importance of examining data for bias in Stage 1. In Stage 2, it is helpful for organisations to establish early guidelines around responsible use, even if a comprehensive AI governance policy may be hard to formulate. Organisations highlight that the mission-driven AI use case is likely to interact directly or indirectly with people and communities. So questions around data (what kinds of data should we collect? what is absolutely essential, especially with respect to personally identifiable information?); accuracy (how will we determine whether the AI is generating errors, hallucinations, or biased results?); and humans-in-the-loop (when and where is human intervention needed?) are worth considering in this stage. (See [“Do we have a working AI governance policy?”](#) in Stage 3 for more detail.)

Embed & Operate

If the organisation proceeds with the pilot, this is the stage where it embeds the use case into its regular programme delivery.



There are two parts to this stage. The first is understanding how the use case can be expanded to fit within the organisation's existing programmes, processes, and delivery systems (including public systems, if that is an organisation's primary mode of delivering its services). The second is recognising what must change in the organisation's internal operating structure (governance mechanisms, budgets, change management, talent, etc.) for the use case to be sustained. Across our interviews, organisations agreed that transitioning into this stage was perhaps the hardest, yet the most rewarding.

Can we integrate the use case with our existing programmes and delivery systems?

Enabling conditions matter in Stages 1 and 2, but infrastructure or connectivity bottlenecks become truly tangible once organisations expand their pilots. Where organisations often struggle is building their pilots as “islands of excellence”, designing for where they would like enabling conditions to be, rather than where they actually are.

Nonprofits that have successfully integrated their AI use cases into their core programming emphasise the importance of embedding AI tools within existing systems and workflows, rather than creating parallel ones. For example, Audere's HealthPulse AI, which uses computer vision to interpret rapid diagnostic tests, is integrated into Rwanda's existing primary-care platform, which frontline health workers are familiar with.

For [Shule Direct](#), an ed-tech organisation making quality educational resources accessible, the integration question centred on matching existing user behaviour to infrastructure realities. Its AI tool for teachers was initially built as a web dashboard with the assumption that government-issued tablets in primary schools would support adoption. But in the pilot, the organisation quickly realised that the hardware availability did not equal readiness: the tablets went largely unused in practice, and teachers' comfort with technology remained low. So it pivoted to WhatsApp, a platform already embedded in teachers' daily routines, enabling them to access the tool's full functionality through a more familiar channel.

Do we have a working AI governance policy?

As pilots evolve into programmes, organisations can turn early AI governance guidelines into operating and design principles. While specific guiding principles vary based on context, a few recurring themes emerged in our conversations:

Humans in the loop. Most organisations emphasise that human oversight is irreplaceable, even for highly effective AI use cases, and should be built into the structure of any tool. It is also important for organisations to establish where AI can be used and where there are “no-go” areas (e.g. workflows such as those pertaining to sensitive health data or those prone to AI bias) that require mandatory human oversight. [Dandelion Africa](#), a Kenyan nonprofit that focuses on women's health and economic empowerment, describes this as a “do-no-harm” approach.

"For sensitive data related to gender-based violence or financial information, we do not use AI," says Julia Ayaya, monitoring and evaluation manager at Dandelion Africa.

Community agency and data consent. Mission-driven AI use cases typically rely on data about people and communities. Organisations highlight that communities should have a voice in what data is collected and how it is used, and that they co-create data governance mechanisms to build long-term ownership. "The people who generate the data often carry the most risk," says Angela Oduor Lungati, executive director of [Ushahidi](#), a global nonprofit that helps communities collect and share information to raise voices, inform decisions, and influence change. "Yet they're the last to benefit from it. Data collected from a community must reflect that community's lived reality and stay governed by that community. If the people who provide the data can't access it, we haven't built shared infrastructure; we've just moved ownership somewhere else." However, data governance can be challenging, especially when organisations don't own the platforms or tools they operate, or when they function across multiple geographies with distinct regulatory frameworks. Beyond data collection and governance, embedding robust feedback loops at every stage can also help strengthen community agency.

Inclusion. Organisations underscore the need to embed inclusion into the principles of AI use cases, whether by building the tool on platforms that match the end users' literacy or comfort levels, or considering how those who are already marginalised – for example, through lack of telecommunications or digital infrastructure – will access these solutions. Most organisations emphasise that AI should be regarded as one solution, not the only solution. In many cases, traditional alternatives (SMS, phone calls, or in-person engagement) need to remain in place to avoid excluding those without access to advanced digital tools.

Evidence. In a constantly changing AI landscape, it can be tempting to act quickly or use the latest technology, but organisations we spoke to recommend a deliberate approach. "There is a hype cycle around AI," says Kartik Dhar, chief technology officer at [Project ECHO](#), a global nonprofit headquartered at the University of New Mexico, United States, "and organisations need discipline and leadership maturity not to chase whatever is trending." It's important, therefore, to invest time in disciplined experimentation and rigorously evaluate AI output before adoption. As Dhar says, "Organisations need to first focus on architecting a scalable digital infrastructure layer, on top of which AI can be implemented. Without a solid foundation, AI projects cannot scale and often get stuck in the pilot phase. Our mission is not to build something for the next six months; we build for the next decade and beyond."

“There is a hype cycle around AI, and organisations need discipline and leadership maturity not to chase whatever is trending.”

KARTIK DHAR, CHIEF TECHNOLOGY OFFICER, PROJECT ECHO

Sustainability. This can look different for different nonprofits, but most organisations note that it is important to consider what sustainability could mean for the AI use case, whether it is scaling through partnerships with government (see [Stage 4](#) for more) or collaboration with other ecosystem players to build shared infrastructure around AI and reduce replication of work.

Environment. The environmental consequences of using AI are undeniable. Organisations can consider where AI use increases energy and water consumption, emissions, or vendor dependency in ways that conflict with organisational values or capacity, and how mitigation strategies – such as assessing environmental impact or reviewing vendor practices – can help.

Do we have the talent and resources to sustain this in the long term?

As AI use cases move from pilots into regular programme delivery, organisations may need to rethink some of their existing operating models to lay out who owns the work internally, who is accountable for it, and how decision rights are structured. The obvious answer – hiring technical experts – might not be the right one. Several organisations we spoke to pointed out that hiring the “right” talent isn’t just about technical expertise, but rather the ability to make complex judgement calls in a dynamic and high-risk environment.

The nature of technical expertise needed is also evolving. “It’s becoming less and less about [building] machine learning teams [especially if an organisation is working with frontier models]. What’s more important is having someone with opinion and taste, [say] a strong product person,” says Patel of Jacaranda Health. Of course, this does not mean having no in-house technology expertise at all. “An in-house tech team can help you move faster, and [it] has the full context [of your work],” says Patel.

For some organisations, specific concerns around AI – such as the need for data privacy – created demand for talent with specialised expertise. For others, it led to a restructuring of roles or departments. For example, Educate! built a dedicated product function for its AI use case and separated product development from programme delivery.

How much does AI really cost?

Upfront costs: For organisations building their own AI models, preliminary costs required to make the model operational include creating data infrastructure where it is absent, data localisation and integration, and investing in quality control and safety. Costs can often exceed the expense of building the model. Even for those adapting existing models, the upfront costs of localising can be high because most frontier models are neither built for nor trained in African contexts. For those buying or partnering, upfront costs could include data localisation costs, together with licensing and vendor fees.

Usage costs, also called inference or compute costs: These are incurred every time a user “queries” or “prompts” the AI. These tend to get higher depending on the complexity of the query (say, images instead of text) and the number of users.

Operating costs: These include the ancillary costs of running the model, including monitoring and reviewing the quality of AI outputs, updating models, ensuring data privacy and security, addressing user concerns, and hiring and training staff for human escalation, to name a few.

Several nonprofits have also had to deal with difficult decisions around job loss as roles have shrunk or been made redundant. Organisations noted that significant change management is required internally at this stage to ensure alignment with new operating models.


The second major consideration around operations is cost. While pilots can provide indicative estimates, organisations quickly realise that costs are much higher once AI becomes embedded in their programme delivery.

Beyond the upfront costs of accessing or building a model, there are also substantial usage costs. AI often becomes more expensive as usage grows. As organisations scale their use cases, they might find that these costs are too high to sustain. Usage costs also vary depending on modes of delivery and the type of content. For example, WhatsApp delivery and multimedia-based systems are more expensive than text-based alternatives. But for nonprofits working in those specific contexts, these costs might be inevitable to reach their target communities.

In addition to technical costs, operating costs are often overlooked. These can include maintaining data pipelines to keep AI use cases relevant; recruiting talent and developing systems to monitor, curate, and improve outputs; contracting with vendors and partners; training staff and leadership; change management; and covering legal costs to comply with evolving policy and regulatory frameworks.

Sustain or Scale

Once a use case has been embedded into programmes, the pressing question for organisations is “what next?” Not many organisations we spoke to have reached this stage in their AI adoption journeys, and scale is largely aspirational. Those that have, at least to some degree, recommend the following aspects.



What does scale look like for this use case?

Scale can look very different for different use cases.⁴ For some, it is about replicability: expanding the use case to different geographies or user groups. Others are better suited for government adoption, enabling wider reach, or for other ecosystem actors to integrate. Yet others are best left at the status quo or even re-evaluated to identify opportunities for further expansion.

At this stage, organisations need to ask searching questions: Does the use case continue to deliver value once embedded in core programming? Would its impact grow if more users gained access, or if it were adapted to other problems? Where do we see our AI use cases over the next two, five, or 10 years? Can we continue to support this if it is scaled for wider use?

The answers depend not just on the organisation's ambition and internal operating structures, but also on enabling conditions. Infrastructural challenges, such as connectivity, publicly available data systems, and trusted delivery mechanisms that show up in stages 2 and 3, can become more binding at this stage, determining whether a use case is scalable. The challenge, then, is matching the organisation's ambition with the reality on the ground, where populations might not be ready for AI solutions.

Cost comes into play, of course. As with Stage 3, operating costs will continue to grow with usage. Organisations should consider who pays to support long-term sustainability beyond their programmes. Most say the funding they receive rarely factors in the cost of long-term AI use, and when they attempt to scale through government partnerships, costs can quickly become too high for public systems to afford.

"Government agencies that want to launch and sustain AI services are not thinking about the long-term compute costs," says Gupta of Viamo. "At scale, this could easily run into hundreds of thousands of dollars each month, by our calculation. Are [public systems] really prepared to absorb such costs? That is why we are trying a commercial model in partnership with telcos where users pay a nominal fee to access our voice AI services. That in turn creates a way of paying for compute [costs]."

Are we ready to scale through public systems?

When organisations scale, they cite partnering with the government as one of the most common ways of doing so. But integration with public systems can come with its own set of challenges.

Organisations repeatedly note that government systems have not yet caught up with AI use in a way that allows for clear decision-making. A majority of African countries do not have national AI strategies; those that have them face challenges operationalising them.⁵ Very few African countries, for instance, have AI-specific regulatory frameworks.

Existing AI policies neither dovetail perfectly with nor adequately address complex issues around data procurement, sovereignty, privacy, and liability. This means investing time and resources in an uncertain, evolving regulatory landscape.

While acknowledging that regulations around AI are important, organisations such as [Qhala](#), which develops digital platforms and products for Africa, argue that much of this slows down innovation. One of its workarounds is to build what it calls “regulatory sandboxes,” which Wanjiru Mburu, director of research development, says are “a way to support innovators to test their innovation and see whether it meets the regulations of the land.”

Interviewees point out that partnerships with the government require patience and sustained investment of time and resources. With AI, governments are often unaware of what the technology can and cannot do, making it even harder to move the needle. “What made PROMPTS scale was county government relationships, which took two to three years to establish,” says Patel of Jacaranda Health. “The tech was never the hard part. Government relationships and distribution matter more.”

“What made PROMPTS scale was county government relationships, which took two to three years to establish. The tech was never the hard part. Government relationships and distribution matter more.”

JAY PATEL, DIRECTOR OF TECHNOLOGY
JACARANDA HEALTH

Missing Links: The Support That Nonprofits Are Asking For

Across our conversations, a common theme is that progress towards mission-driven AI use doesn't happen in a vacuum. Even when organisations make all the “right” choices, critical gaps need to be addressed before the sector can see tangible, sustainable impact through AI. Nonprofit leaders frequently mentioned five constraints for which funders can provide catalytic support.

1. Different types of funding support across each stage of the journey

A common constraint is the mismatch in how nonprofits' AI journeys progress and how they are funded by philanthropy. Funding for building their expertise in technology has historically been low.⁶ Most available AI funding is concentrated at the beginning of the journey, usually around pilot deployment, when costs are lowest. Even in the early stages, funding is often tied to specific use cases or geographies, restricting nonprofits' ability to experiment, learn, and, most importantly, fail. Many also cite the pressure of including an AI solution in their programmes to access funding before clearly articulating the problem.

Nonprofits also note that although the costs of building and testing AI products have dropped over the past few years, philanthropic funding doesn't always account for the heavy burden of upfront costs – such as building data infrastructure, adapting models to local contexts, hiring specialists – leading to an overreliance on foreign models that may not work well in Africa.

How can this gap be closed?

- **Funders can match funding to the stage of investment.** For Stage 1, nonprofits need unrestricted, risk-tolerant capital that isn't tied to a specific AI model or platform. For Stage 2, they need operational grants that explicitly cover operating costs. To actively deploy and scale in Stages 3 and 4, respectively, nonprofits need multiyear, flexible patient capital that goes beyond one- or two-year grant cycles and can tolerate long-term infrastructure investment and research.
- **Funders can identify portfolio organisations that might be ready to work with AI** and support their experimentation.

2. Shared infrastructure with an eye on responsible use

Another key gap in the African social sector is the lack of shared AI infrastructure, including local-language datasets, sector-specific data for AI to build on, open-source models adapted to African contexts, shared APIs and middleware,^{iv} and common model evaluation protocols. Without investment in these areas, nonprofits struggle and incur heavy costs, trying to do it all alone or by needlessly duplicating others' efforts.

More recently, organisations have been working to close this gap: [AI4D](#), for example, fosters safe, inclusive, and responsible AI ecosystems to accelerate progress on international development challenges, and brings funders and practitioners together to do this; Mozilla's [Common Voice](#) provides a free, open-source platform for community-led data creation; [Amini](#) creates data infrastructure so nations can own, process, and govern their data securely and locally; Ushahidi helps create locally relevant data stacks; [Zindi](#) provides a professional network for data scientists and AI builders in emerging markets; and [Deep Learning Indaba](#) brings together African AI researchers and practitioners to build the future of AI in Africa. More such networks and initiatives are needed.

Of course, shared infrastructure also comes with risk. Building it is not enough; protocols are also needed around privacy, data sovereignty^v and localisation, and cybersecurity. Nonprofits note that investments in governance need to go hand in hand with creating shared infrastructure.

iv An API, or application programming interface, is a set of protocols that allow different applications to communicate with one another. Middleware is software that connects different operating systems, databases, or applications.

v Data sovereignty is the legal principle of data being subjected to the laws and regulations of the country or region where it was generated.

How can this gap be closed?

- **Funders can invest in shared assets for the sector** – datasets, open-source models, and shared evaluation protocols – and strengthen existing African-led collaboratives and open-source communities. Alongside, they can help advocate for more responsible and transparent use of these assets.
- **Funders can engage government and ecosystem actors** to develop common standards around shared public goods to promote interoperability.
- **Funders can encourage better coordination** between nonprofits and infrastructure builders through forums and convenings to encourage the sharing of best practices and reduce duplicative efforts.

3. Access to technical and strategic guidance

Many organisations cannot fill gaps in strategic and technical expertise through hiring alone; there is a need for intermediaries that can help them with mission-driven AI adoption, from building capability to advising on regulatory frameworks to facilitating partnerships and collaboration. Since this sector is relatively new, intermediaries are limited and under-resourced.

How can this gap be closed?

- **Funders can help build intermediary capacity**, supporting organisations that can coach nonprofits on technology, strategy, governance, and implementation. Initiatives such as the Patrick J. McGovern Foundation's Learning Hub or The Agency Fund's learning and development support for grantees are examples.
- **Funders can support communities of practice for AI for social impact** through convenings, resource libraries, peer learning, and dedicated staff to manage and coordinate them. They can also help pair AI-native organisations with nonprofits that have deep implementation expertise but less experience with technology.

4. Collective capacity building

Nonprofits point to gaps in AI literacy, governance, and responsible use, not just within their organisations but across the ecosystem, from governments and regulators to funders themselves. Organisations we spoke to feel that AI expertise is concentrated within a small group of organisations or specialists. Investing in building collective capability specific to Africa will lead to a more resilient social sector, improving decision making, informing regulation and policy, and increasing overall participation.

How can this gap be closed?

- **Funders can build and strengthen peer-learning networks** that deepen their understanding of the risks and opportunities of AI. Many nonprofits note that, like their grantees, funders could also benefit from asking themselves whether AI is the right solution to invest in for the problems they are looking to fund. These networks can also encourage the sharing of knowledge across nonprofits, intermediaries, and governments.

Looking to the Future

The AI landscape in Africa is evolving rapidly. The next few years will be decisive in determining the direction it takes. For African nonprofits, this poses an incredible opportunity to shape the AI agenda on the continent, ensuring that AI is used meaningfully as a powerful force for good, helping reach more people with life-changing programmes and services.

Madhuri Mukherjee is a partner and **Aditi Sharma** is a principal at The Bridgespan Group, based in Nairobi. **Ynis Isimbi** is an advisor at Bridgespan, also based in Nairobi. **Titi Kabi, Louis Moussi, and Courtney Gehle** are a manager, a former consultant, and an associate consultant, respectively, based in Bridgespan's Johannesburg office. **Sohini Pal** is an editorial director at Bridgespan, based in Delhi.

The authors extend special thanks to Angela Oduor Lungati, executive director at Ushahidi, and Avela Gronemeyer, managing director at Inno Yolo, for insights that helped shape this research. We also thank Mark McKeag, Ntefeleng Nene, Jan Schwier, Preeta Nayak, Nate Wong, and Natalia Alarcón of The Bridgespan Group for their insightful reviews and contributions to this report.

We are also grateful for the organisations that took the time to share their insights with us and helped shape the findings of this report.

Appendix 1. Methodology

Primary interviews: We spoke to 33 nonprofits and social enterprises, six funders, and five experts and intermediaries to understand the current AI landscape in Africa and the challenges and opportunities in the sector.

Focus group discussions: We held a focus group discussion with six community-driven nonprofits to understand how AI is showing up in community-centred contexts, what it could mean for communities, and what it would take for AI to be used in ways that are practical, responsible, and grounded in community realities.

Secondary research: We conducted a comprehensive review of secondary literature on AI in Africa and cross-referenced findings with global research from The Bridgespan Group and others.

Appendix 2. Select Resources

From Bridgespan

Brine, Derek, Nate Wong, Jake Porway, and Robyn Porteous, "AI Can't Be Ignored: Exploring the Opportunities for Nonprofits and the Social Sector," The Bridgespan Group, 22 July 2025, www.bridgespan.org/insights/exploring-ai-opportunities-for-nonprofits-and-the-social-sector.

Compiled Social Sector AI Resources, The Bridgespan Group, www.bridgespan.org/insights/compiled-social-sector-ai-resources.

"Fundamentals of Nonprofit Strategy," The Bridgespan Group, www.bridgespan.org/insights/nonprofit-strategy.

Wong, Nate, Derek Brine, Natalia Alarcón, Paige Boehmcke, Eliza Patterson, and Zach Slobig, "Choosing Your AI Path: A Framework for Nonprofit Leaders to Make Strategic Choices," The Bridgespan Group, 20 August 2026, www.bridgespan.org/insights/choosing-your-ai-path-a-framework-for-nonprofit-leaders-to-make-strategic-choices.

Wong, Nate, Derek Brine, and Robyn Porteous, "Closing the Nonprofit Funding Gap in the Age of AI," The Bridgespan Group, 27 October 2025, www.bridgespan.org/insights/closing-the-nonprofit-funding-gap-in-the-age-of-ai.

Other resources

The 2025 AI for Humanity Report, Fast Forward, 2025, www.ffwd.org/2025-ai-for-humanity-report.

Adams, Rachel et al., *The State of AI in Africa: A Landscape Study*, Global Center on AI Governance, March 2025, www.globalcenter.ai/research/ai-in-africa-landscape-study.

AI as a Catalyst for Social Impact in Kenya's Philanthropic Sector, East Africa Philanthropy Network, 2026, <https://eaphilanthropynetwork.org/eapn-publications/ai-as-a-catalyst-for-social-impact-in-kenyas-philanthropic-sector/>.

Generative AI Evaluation Playbook for the Development Sector, The Agency Fund, August 2025, <https://eval.playbook.org.ai/>.

"AI Evidence Playbook," Abdul Latif Jameel Poverty Action Lab (J-PAL), February 2026, www.povertyactionlab.org/ai-evidence-playbook.

AI for Impact: The Role of Artificial Intelligence in Social Innovation, World Economic Forum, April 2024, www3.weforum.org/docs/WEF_AI_for_Impact_Social_Innovation_2024.pdf.

AI Readiness and Adoption in the Nonprofit Sector in 2024: Results from GivingTuesday's AI Readiness Survey, GivingTuesday, 2024, <https://ai.givingtuesday.org/ai-readiness-report-2024/>.

"Building with GenAI: Girl Effect's Journey to Smarter, Safer health Chatbots", Girl Effect, April 2025, https://ge-prd-eu-west-1-webapp-assets-01.s3.eu-west-1.amazonaws.com/girleffect/media/documents/Girl-Effect_AI_ML-Whitepaper_jFZ61vX.pdf.

Equitable AI Adoption Framework, Project Evident, February 2026, <https://projectevident.org/resource/equitable-ai-framework/>.

Global AI Diffusion: Q2 2026 Trends and Insights, Microsoft AI Economy Institute, January 2026, www.microsoft.com/en-us/corporate-responsibility/topics/ai-economy-institute/reports/global-ai-diffusion-report/.

Humeau, Eugénie and Tanvi Deshpande, "AI for Africa: Use cases delivering impact," GSMA, July 2024, www.gsma.com/solutions-and-impact/connectivity-for-good/mobile-for-development/gsma_resources/ai-for-africa-use-cases-delivering-impact/.

Ndandani, Mbali, "The Role of AI and Digital Tools in Africa's Development, Issue Brief No. 762," Observer Research Foundation, December 2024, www.orfonline.org/research/the-role-of-ai-and-digital-tools-in-africa-s-development.

"PJMF Learning Hub," Patrick J. McGovern Foundation, <https://learn.mcgovern.org/>.

Appendix 3. List of interviewees

Nonprofits and intermediaries

Adopt a Generation Africa (AGA)	mothers2mothers (m2m)
Amini	Nexleaf Analytics
Amp Health	Nguzo Africa
Audere	One Acre Fund
The BOMA Project (BOMA)	Project ECHO
CAMFED	Qhala
Clear Global	Reach Digital Health
Dandelion Africa	Saide
Digital Green	SHOFCO
Educate!	Shule Direct
Farmerline	Teaching at the Right Level (TaRL) Africa
Generation Kenya	Tiko
Global Center on AI Governance (GCG)	Ubongo
Harambee Youth Employment Accelerator	Ubunye Foundation
Imagine Worldwide	Ushahidi
Jacaranda Health	Vambo AI
Kytabu	Viamo
Lwala Community Alliance	Village Enterprise
Maisha Meds	VillageReach
Maono Africa	

Funders

The Agency Fund	Mulago Foundation
Grand Challenges Canada	Patrick J. McGovern Foundation
Luminate	Shortlist Futures
	(Gender & AI Livelihoods (GAIL) Program)

Endnotes

- 1 *State of Nonprofit AI Adoption and Governance*, NTEN and The Bridgespan Group, 2026, www.nten.org/posts/publication/state-of-nonprofit-ai.
- 2 *World Development Report 2026: The Promise of Artificial Intelligence*, World Bank Group, August 2026, www.worldbank.org/en/publication/wdr2026.
- 3 Ibid.
- 4 Marina Fisher, Alyson Zandt, Angie Estevez Prada, Paola Jimenez-Read, and Bradley Seeman, "A New Look at Strategic Pathways for Solving Social Problems at Scale," The Bridgespan Group, 29 April 2025, www.bridgespan.org/insights/a-new-look-at-strategic-pathways-for-solving-social-problems-at-scale.
- 5 "AI Governance in Africa: Insights from a policy dialogue with 12 countries," OECD, 30 April 2026, www.oecd.org/en/publications/oecd-artificial-intelligence-case-studies_c517fcf5-en/ai-governance-in-africa_1ff55135-en.html.
- 6 Nate Wong, Derek Brine, and Robyn Porteous, "Closing the Nonprofit Funding Gap in the Age of AI," The Bridgespan Group, 27 October 2025, www.bridgespan.org/insights/closing-the-nonprofit-funding-gap-in-the-age-of-ai.

THE BRIDGESPAN GROUP

BOSTON 287 Columbus Avenue, 4th Floor, Boston, MA 02116 USA. Tel: +1 617 572 2833

JOHANNESBURG Bridgespan Africa, The MARC, Tower 1, 3rd Floor, Corner Maude and Rivonia Road, Sandown, Johannesburg, South Africa. Tel: +27 11 012 9280

MUMBAI Bridgespan India Private Limited (registered address), 11th Floor, Platina, G Block, Plot C 59, Bandra Kurla Complex, Mumbai 400051, India. Tel: +91 22 6628 9624

NEW YORK 333 Seventh Avenue, 11th Floor, New York, NY 10001 USA. Tel: +1 646 562 8900

SAN FRANCISCO 88 Kearny Street, Ste. 200, San Francisco, CA 94108 USA. Tel: +1 415 627 4500

SINGAPORE Bridgespan Southeast Asia Ltd., 38 Beach Road, South Beach Tower, Level 27, Singapore 189767. Tel: +65 6216 6000

WASHINGTON DC 1201 K Street NW, 4th Floor, Suite 444, Washington, DC 20005



www.bridgespan.org

contact@bridgespan.org

contactindia@bridgespan.org

This work is licensed under the Creative Commons BY-NC-ND License.
To view a copy of this license, visit <http://www.bridgespan.org/about-us/terms-and-conditions>