PRACTITIONER GUIDE

NEXT GENERATION
IMPACT ECONOMY EXPERTISE

From Fragmentation to Intelligence:

AI and the Future of Africa's Social and Impact Economies

AUGUST 2026

Reana Rossouw · Next Generation Consultants

CONTENTS

FROM FRAGMENTATION TO INTELLIGENCE

AI and the Future of Africa's Social and Impact Economy — in one page

The argument. Africa's social economy does not suffer primarily from an effort deficit. It suffers from an **intelligence deficit** - a failure of coordination, synthesis, learning and adaptive capability. The sector struggles to know itself, to connect its knowledge, and to act on what it collectively understands. AI, deployed with intentionality, equity and African agency, could be one of the most powerful tools available to address it - not by replacing human judgement, but by augmenting the analytical, connective and synthesising work that under-resourced organisations cannot sustain alone.

WHY IT MATTERS NOW

South Africa's social economy is large and under strain: roughly R13 billion in annual corporate social investment, an estimated 235,000 registered NPOs, and chronic underfunding of the institutional capacity that makes organisations effective. AI does not solve these structural problems. But it dramatically changes what is possible within them - and, deployed without intention, it could just as easily deepen the very inequalities the sector exists to address.

WHAT THE EVIDENCE SHOWS

Across the continent, the most effective applications are already visible: language models built for isiZulu and isiXhosa (Lelapa AI, Masakhane); maternal-health messaging at national scale (MomConnect); AI-assisted youth-employment matching (Harambee, SAYouth.mobi); biometric grant integrity (SASSA). The pattern is consistent: the applications that work solve a concrete daily problem, work in the languages communities actually speak, and embed in infrastructure people already trust.

The four conditions for getting it right

- **African agency** — AI built by and with Africans, governed by African data principles, not imposed from outside.
- **Community consent and data dignity** — the people whose data trains these systems must benefit from them.
- **Governance** — human-in-the-loop on every decision that affects access to services. AI recommends; humans decide.
- **Institutional investment** — funders must pay for the overhead that makes intelligent organisations possible. Not as a concession. As a strategy.

What to do first. If you read nothing else in the full guide, read the Executive Summary, complete the AI Readiness Self-Assessment, and look at the cross-cutting insights from the case studies. If you fund the sector, ask what institutional capacity your grants are building - not just what programmes they pay for. The choices made now, about governance, ethics, inclusion and design, will shape this transition for a generation. The moment to make them intentionally is not tomorrow. It is now.

Reana Rossouw · Next Generation Consultants · The full practitioner guide runs to fourteen sections and eight standalone tools.

THE GUIDE, IN BRIEF

Executive Summary

Reana Rossouw · Next Generation Consultants

R13bn
in annual corporate social investment

235,000
registered NPOs in South Africa

4
conditions for getting it right

14
sections in the full guide

THE PROBLEM

Africa's social and impact economy is at a decisive crossroads. On one side lies the familiar landscape of fragmentation: knowledge silos, disconnected data, duplicated effort, institutional inequality and a sector perpetually stretched between growing social need and constrained resources.

THE POSSIBILITY

On the other lies a genuinely new possibility: the emergence of artificial intelligence as a connective intelligence infrastructure capable of reshaping how social ecosystems coordinate, learn, adapt and make decisions.

> *"**Africa's social economy does not suffer primarily from an effort deficit. It suffers from an intelligence deficit** — a failure of coordination, synthesis, learning and adaptive capability. AI, deployed with intentionality, equity and African agency, could be one of the most powerful tools available to address it."*
>
> — THE CENTRAL ARGUMENT OF THIS GUIDE

A note on language

This guide uses the word "**intelligence**" in a specific and deliberate sense. When it speaks of an "**intelligence deficit**," it does not mean a deficit of artificial intelligence - the technology. It means a systemic failure of coordination, synthesis, learning and adaptive capability: the failure of a sector to know itself, to connect its knowledge, and to act on what it collectively understands.

Artificial intelligence - AI - is one potential tool for addressing this deficit. But the deficit is older than the technology, and the technology alone will not resolve it.

Throughout this guide, "**ecosystem intelligence**" refers to the systemic capacity of social actors to coordinate and learn together; "AI" refers to the technology that augments that capacity.

The distinction matters because conflating the two leads to a category error: treating the adoption of a tool as a substitute for the development of a system.

> *This guide argues that the most profound opportunity of AI within the social economy is not automation.* ***It is transformation: of ecosystem intelligence, institutional capability, collaborative governance and adaptive learning.***

Yet this transformation is neither guaranteed nor inherently equitable. Without intentional design, AI could deepen the very inequalities that Africa's social sector exists to address.

For South Africa specifically, this moment carries particular weight. The country hosts a social economy of significant scale: Approximately R13 billion in annual corporate social investment; an estimated 235,000 registered NPOs; a complex web of government, philanthropy, civil society and social enterprise; yet one characterised by structural fragmentation, underfunding of institutional capacity, and growing demand for measurable impact. AI does not solve these structural problems. But it dramatically changes what is possible within them.

Grounded in current African AI developments, South African sector data, established impact measurement frameworks and emerging case studies from across the continent, this guide maps the AI-enabled opportunity across the full social impact lifecycle: from design and diagnosis through to exit and sustainability. It addresses the sector's diverse stakeholders, its major development portfolios, the ethical risks that must be actively governed, and the practical pathway organisations can follow regardless of their current readiness level.

A note on how this guide was developed

This guide was developed with the substantive assistance of artificial intelligence - specifically, Claude, developed by Anthropic. That fact is worth naming clearly, for reasons that go beyond transparency.

The social economy in South Africa is characterised by what this guide and the author calls an intelligence deficit: a failure of coordination, synthesis, learning and adaptive capability that limits the sector's collective ability to know itself, connect its knowledge, and act on what it understands. One of the guide's central arguments is that AI can help address this deficit; not by replacing human judgement, but by augmenting the analytical, connective and synthesising functions that under-resourced practitioners and organisations struggle to sustain alone.

This guide is an example of exactly that proposition in practice.

What the collaboration involved

The research, conceptual framing, practitioner knowledge, sector relationships and authorial judgement that shape this guide are human. Two decades of work at the intersection of social

innovation, impact investment, ecosystem development and MEL practice in South Africa cannot be replicated by any AI system, and they are the foundation on which everything in these pages rests.

What AI contributed was different in kind, and worth being specific about.

Research synthesis and analytical depth. It is important to note that the author had to first identify the information sources, organise the literature, and verify its relevance to the context of the research and final product. AI was used to synthesise large bodies of research, identify patterns across literature and case study evidence, and surface connections between domains; between MEL practice and AI application, between generational workforce theory and organisational readiness, between data governance frameworks and the lived realities of South African NPOs, that would have taken far longer to develop through manual research processes alone.

Structural thinking and flow. Although the final document is presented as a cohesive research and substantial guide for implementation, it is important to note that it went through five iterations and its development was supported by peer reviews (contributors who supported the author through the various iterations) to protect the accuracy, relevance and completeness of the document. AI then assisted to identify where the document's argument was internally inconsistent, where transitions between sections were abrupt, where a concept was doing too much work, and where the reader needed orientation before being asked to engage with complex material. For example: The identification of the tension between the guide's systems argument and its linear value chain framework, emerged from AI-assisted structural analysis, but the author had to personally make specific decisions to ensure consistency throughout the chapters. The author had to rework significant sections to ensure the support the reader would require through each phase of AI adoption and integration is included. The resolution of that tension strengthened the guide's intellectual integrity.

Accuracy review and risk identification. The case study section of this guide received particular human scrutiny and peer review throughout its various iterations. The most challenging part was to find evidence of AI usage and application that is relevant to the social and impact economy. Once the case studies were identified AI identified mislabelling errors - an ECD platform incorrectly categorised as a GBV support tool, for instance, flagged claims that were promotional rather than evidenced, distinguished South African organisations from non-South African ones that had been inadvertently included. AI also identified where the framing of corporate AI applications overstated the evidence base. In this case, the author had to verify the accuracy of the data, conduct additional research, remove some case studies and confirm the data sources for correctness as well as relevance and materiality. These were not minor editorial corrections. In a document that names specific organisations, accuracy is a matter of credibility and reputation, for the organisations described and for the guide itself.

Intellectual challenge and pushback. The author realised from the beginning of the research process that readers' own perceptions and fears about AI adoption and integration will influence their understanding of the content. For this reason, the author was very specific to have each argument and position challenged (by peer reviewers) to ensure that personal views were balanced and well informed. In this regard, perhaps the most valuable contribution of using AI to challenge broad statements and assumptions was not analytical but dialogic. AI functioned as a thinking partner that was willing to push back: on the overuse of a conceptually overloaded word, on the gap between the guide's stated commitments to community agency and its absence of tools for communities themselves, on the underdeveloped treatment of funder accountability, on the places where the argument was right but the evidence was thin.

> ***A thinking partner who only affirms is not a useful thinking partner.*** *The challenge was generative.*

What this means for the reader

If you work in South Africa's social economy: as a funder, an NPO leader, an MEL practitioner, an intermediary, a social enterprise founder - and you have wondered whether AI is genuinely useful for the kind of work you do, this guide is a partial answer to that question.

Not a complete answer. The work of community engagement, relational trust, programme design and ethical accountability that defines the social economy cannot be delegated to any technology. Those remain irreducibly human.

But the work of making sense of complex, fragmented information; of synthesising research, stress-testing arguments, identifying what is missing, structuring thinking so that others can follow it; is work that AI can genuinely support. Not by doing it for you. By doing it *with* you.

The guide you are reading took shape through that kind of collaboration. Its weaknesses are where the human judgement was insufficient, or where the AI was asked the wrong questions, or where the synthesis failed to capture something that only lived experience could see. Its strengths are, in part, a product of what becomes possible when a practitioner with deep sector knowledge and an AI system with broad analytical capacity work together toward a shared purpose.

That is not a disclaimer. It is the argument of the guide, demonstrated.

The AI collaboration in the development of this guide was conducted in accordance with the data governance principles described in Section Nine. No confidential beneficiary data, client information or proprietary organisational data was shared with AI systems in the development of this guide. All factual claims, case study descriptions and sector-specific interpretations were reviewed and validated by the author before publication.

A Note on how to use this guide

This guide is long. It is deliberately so: the social economy's AI landscape is complex, the stakeholder ecosystem is diverse, and the practical, governance and human dimensions of AI adoption each deserve substantive treatment rather than a checklist. But length is only useful if the reader can find what they need. This section is a map.

SHORT ON TIME

If you read nothing else

Read the Executive Summary, complete the AI Readiness Self-Assessment (Section Ten), and skim the cross-cutting insights from the case studies (Section Six). Under an hour, and it stands alone.

THE FULL ARGUMENT

If you are reading linearly

A progressive argument: structural challenge (1–3), AI reframed (4), application across the ecosystem (5–8), governance and readiness (9–11), people (12), and recommendations (13). Each section builds on the last.

FIND YOUR LANE

If you are navigating by role

Funder, NPO lead, MEL practitioner, board member, social enterprise founder? The table below maps the sections and tools most relevant to each role.

Your role	Start here	Don't miss	Tools - Appendices
CSI Manager or Corporate Funder	• Section Six (stakeholder matrix, your row specifically) • Section Seven (thematic portfolio mapping)	• Section Nine (AI governance for funders) • Recommendation Cluster A (Section Thirteen)	• Tool 1 (Funder's AI Diligence Checklist). • Tool 4 is also useful for understanding what grantees are budgeting for.
NPO Executive Director or Senior Manager	• Section Three (the structural challenge named) • Section Ten (readiness assessment) and Eleven (tool selection)	• The 90-day implementation roadmap (Section Eleven) • Section Twelve (workforce and the AI transition)	• Tool 4 (How to Budget for AI Adoption). • Tool 7 if MEL is part of the role.
MEL Practitioner or Impact Analyst	• Section Eight (AI and impact measurement frameworks) • Section Five (AI across the programme value chain)	• Section Nine (data governance and POPIA) • The case study synthesis (Section Six)	• Tool 7 (A Worked Example for MEL Practitioners). • Tool 4 for budgeting the work.
Intermediary or Ecosystem Builder	• Section Four (the ecosystem intelligence argument) • Section Six (the intermediary row of the stakeholder matrix)	• Recommendation Cluster B (Section Thirteen) • Section Twelve (talent strategy for the sector)	• Tool 5 (The Disintermediation Risk). • Tool 4 for sector capacity-building conversations.
Social Enterprise Founder or Social Entrepreneur	• Section Five (AI across the programme lifecycle)	• Section Six (social enterprise case studies)	• Tool 8 (AI for Social Enterprises). • Tool 4 for budgeting the operational adoption.

	• Section Eleven (practical tool selection matrix)	• Section Nine (minimum viable governance checklist)	
Board Member or Governance Lead	• Section Nine (ethics, risks and governance framework) • The POPIA compliance section (Section Ten)	• Section Twelve (generational workforce - the governance implications) • Recommendation Cluster D (Section Thirteen)	• Tool 2 (Board's Ten Questions on AI). • Tool 1 if the board oversees grant-making.
Community Organisation or CBO Leader	• Section Eleven (where to start - the 90-day roadmap) • The community questions appendix (Section Nine)	• Section Six (case studies of community-facing AI) • Section Three (why the structural challenge is not your fault)	• Tool 3 (CBO Companion). Tool 4 for honest cost conversations with funders.
Policymaker or Government Official	• Section Four (the ecosystem intelligence argument) • Section Seven (thematic portfolio - cross-sector applications)	• Section Thirteen (all recommendations) • The SASSA and spatial mapping cases (Section Six)	• Tool 6 (Government as AI Adopter). • Tool 1 if the role includes grant-making oversight.

A note on how to read the case studies

The case study section in Section Six is the longest single section in the guide, and it is designed to be used as a reference rather than read from start to finish. Each case is self-contained. Readers with a specific thematic interest: health, ECD, youth employment, environmental sustainability - will find the cases grouped by actor type, with the thematic index in Section Seven providing an alternative entry point by focus area.

The cross-cutting insights that follow the case study table are the analytical core of that section and are worth reading in full, regardless of which individual cases are most relevant to your context. They are where the evidence base yields its most important lessons.

A note on language and register

This guide is written for practitioners, not for academics. Where technical concepts are introduced: POPIA, IRIS+, SROI, IMP frameworks, large language models - they are explained in context. Readers who encounter unfamiliar terms will find working definitions in the text rather than in a glossary; an embedded glossary of practitioner terms is provided at the back of this document.

The guide does not assume that its readers are AI enthusiasts, AI sceptics or AI experts. It assumes that they are people doing serious, important work in difficult conditions, for whom AI is a potential tool worth understanding - and that they deserve a guide that is honest about both its potential and its limitations. That is what this guide attempts to be.

READER'S GUIDE

Finding What You Need

If you are looking for a specific tool or framework

The guide contains a number of standalone frameworks, matrices and checklists that can be used independently of the surrounding text. These are the most frequently referenced:

SECTION TEN

AI Readiness Self-Assessment

A diagnostic tool for organisations assessing their current AI readiness across eleven dimensions, 60+ questions. Can be run as a facilitated team exercise over one or two sessions.

SECTION ELEVEN

The Seven-Decision Framework for AI Tool Adoption

A structured decision process for evaluating any AI tool before adoption, with particular attention to POPIA compliance and data governance.

SECTION ELEVEN

AI Tool Usability Matrix

A comparison of commonly encountered AI platforms mapped against data residency, training practices, POPIA risk level and social sector fit.

SECTION ELEVEN

Minimum Viable AI Governance Checklist

A ten-step governance framework for organisations at Readiness Levels 1 and 2.

SECTION ELEVEN

The 90-Day Implementation Roadmap

A phased first-year guide for organisations beginning their AI journey with limited resources and no dedicated technology staff.

SECTION FIVE

Impact Value Chain AI Application Map

A reference table mapping AI tools and applications to each stage of the programme lifecycle, from design through to exit and sustainability.

SECTION SEVEN

Thematic Portfolio Matrix

AI applications mapped to eleven social investment themes including health, education, food security, gender, livelihoods and climate.

SECTION SIX

Stakeholder AI Application Matrix and Case Studies

A compilation of AI applications in South Africa's social economy, organised by actor type with cross-cutting analytical insights.

SECTION TWELVE

AI-Era Competency Framework

A seven-domain competency map for social economy practitioners navigating the AI transition, with role-specific guidance.

SECTION TWELVE

Emerging Career Roles Table

Ten new and hybrid roles emerging at the intersection of AI and social economy practice, with South African entry points.

SECTION NINE

Community Questions Appendix

A guide for community members, written in plain language, covering the questions they have the right to ask any organisation using AI with their data.

If you are looking for supportive implementation tools

The Appendices: Supplementary Tools Pack at the back of this guide contains eight audience-specific tools developed to serve readers whose needs the main argument does not fully address. These are designed for direct, standalone use:

TOOL 1	**Funder's AI Diligence Checklist** A three-part checklist covering pre-grant diligence, active portfolio management and portfolio-level AI strategy. For corporate and foundation funders.
TOOL 2	**Board's Ten Questions on AI** Ten questions every board should be able to answer about its organisation's AI use, in plain language. For directors and trustees.
TOOL 3	**CBO Companion: AI in Plain Language** A plain-language companion designed for community-based organisations with limited resources. For CBO leaders and grassroots practitioners.
TOOL 4	**How to Budget for AI Adoption** Indicative cost categories, budget scales by readiness level, the cost of inaction, and how to ask funders for AI budget. For NPO Executive Directors.

TOOL 5	**The Disintermediation Risk for Intermediaries** The strategic risk map and repositioning argument for intermediary organisations facing AI-enabled disintermediation. For intermediary leadership.
TOOL 6	**Government as AI Adopter** AI applications for the four roles the government plays in the social economy. For policymakers and public sector leadership.
TOOL 7	**A Worked Example for MEL Practitioners** A composite NPO walks through the Seven Decision Framework and a manual-versus-AI-assisted SROI on a flagship programme. For MEL practitioners.
TOOL 8	**AI for Social Enterprises** Revenue, sustainability and the mission tension. For commercially-trading organisations with a social mission.

Table 20 — Consolidated Recommendations Summary (Section Thirteen)
All sixteen recommendations mapped to primary owner, time horizon and cluster, for funders, intermediaries, NPOs, sector bodies and government to locate what applies to them.

ABOUT THE AUTHOR

Reana Rossouw
Founder, Next Generation Consultants

Reana Rossouw is the founder of Next Generation Consultants, a leading impact advisory firm working at the intersection of social innovation, sustainable development, ecosystem change and impact management and measurement. For more than two decades she has worked alongside funders, non-profits, social enterprises, intermediaries and governments across the social, solidarity and impact economies - helping them design impact strategies, measure what matters, and improve impact performance.

Her practice sits where strategy meets evidence. She is known for translating complex, systems-level thinking into tools practitioners can actually use, and for insisting that measurement serve learning rather than compliance. That orientation runs through her published work, and technology platform - the Investment Impact Index™, in addition to the numerous frameworks, guidelines, practice notes and impact measurement methodologies she has developed for the sector. Reana is well known for the annual research she produces on trends and insights shaping the social, solidarity and impact economies in South Africa and across the continent.

This guide reflects the throughline of that career: a conviction that the social economy's deepest constraint is not effort or even funding alone, but its capacity to coordinate, synthesise and learn: and that this capacity, thoughtfully supported, is where the sector's next decade will be won or lost. Reana writes as a practitioner for practitioners, grounded in the South African context and accountable to the communities the sector exists to serve.

To learn more about her work, visit the Next Generation Consultants website, where her latest research reports and practitioner guides are available.

ABOUT THE CONTRIBUTORS

Hilton Theunissen: Director | International Partnerships at GrowthWheel International: Hilton's contribution to this guide was threefold: He reviewed the first draft of the document and his biggest criticism was the lack of the human capital aspect. His contribution was fundamental not only in shaping section 12, but specifically in regards to the skills and competencies the sector will require as it enters the AI era. He also challenged the usability and applicability of the guide – but essentially highlighted the tension between providing career opportunities for future generations, and also how to bridge the AI literacy gap across the social economy ecosystem. My sincerest gratitude to Hilton for ensuring that the human element served as grounding for the entire guide.

Jacqueline Cutten: Founder & Principal Consultant | Cutten Advisers (CuttenAdvisers.com): Jacqueline's contribution was instrumental in the development specifically of the self-assessment tools, structuring of the document, the logic, flow and the tension between knowledge and evidence that I battled with from the start. Her guidance and feedback on drafts 2 and 3 (particularly Sections 10 and 11) was invaluable as it challenged my thinking, stretched my imagination and kept me grounded throughout the writing process. My absolute and most heartfelt appreciation for your patience, your time and your contribution to the content creation process.

Suzaan Rossouw: Commercial Growth Specialist and AI Enthusiast: Suzaan's contribution was critical in the final quality control stages of this guide. From editorial and content review, to structure and design. Above all else - Suzaan's knowledge of AI platforms, her AI structuring and prompting capabilities and AI research skills added dimensions that not only contributed to the overall quality of this guide, but proved that AI native generations are a critical component in future research, strategy and performance management activities. Above all – what I (as a Baby Boomer) am most grateful for is not just the tokens she gifted when my AI time was blocked, but the insight, guidance and support from a Millennial that validated that the AI future is multigenerational. My most heartfelt appreciation for your patience, guidance and support in the finalisation of this guide.

A final note before you dive into this guide

This guide is the culmination of more than four months of conceptualisation, development, writing and research. It is written from the perspective of a consultant that works with a wide range of stakeholders across the social economy. For this reason it was important to capture not only different applications but also to highlight the potential of what might be possible for the future of the social economy.

This practitioners guide is for organisations working at the intersection of the social, solidarity and impact economies and is time sensitive. Tools, applications, platforms, and case studies referenced were researched and reviewed during March, April, May and June 2026. This guide acknowledges

that technology, platforms and applications will change, indicative of a fast paced environment. For this reason the guide is not about which AI tool or application to use: it is a holistic perspective of how to go about a) thinking about AI for your specific organisation b) consider potential applications within the development sectors your organisation operates c) understand the guardrails within which you will have to operate throughout your organisation's AI journey d) provide guidance on an organisational and strategic level to start integrating AI.

All the best for your AI journey!

PART ONE

The Argument and the Landscape

IN THIS PART

> *In many ways, Africa's social economy suffers from an intelligence deficit rather than an effort deficit alone.*

SECTION ONE

The Social Economy Is Entering an Intelligence Era

Who this section is for: All readers. This section establishes the case for why AI matters to the social economy now, not as a technology story, but as a systemic response to the sector's most persistent structural challenge. Read this before anything else.

Africa's social and impact economy stands at a critical inflection point. Across the continent - and particularly within South Africa - social systems are under immense and compounding pressure.

- **Governments** face mounting fiscal constraints and a shrinking ability to deliver on their social mandates.
- **Civil society organisations** are expected to deliver more with fewer resources, shorter grant cycles and heavier accountability demands.
- **Funders** increasingly require measurable impact, systems-level outcomes and evidence-based decision-making - yet frequently do not invest in the very data and learning infrastructure that would make this possible.
- **Communities** are navigating overlapping crises that include unemployment, food insecurity, democratic instability, educational inequality, public health vulnerability, climate risk and growing social fragmentation.

At the same time, the architecture of social change itself is evolving.

The social economy is becoming more interconnected, more collaborative and more systems-oriented.	**Funders** are moving beyond isolated project funding toward ecosystem investment.	**Intermediaries** are becoming increasingly important as coordinators, aggregators and ecosystem stewards.
Knowledge generation is becoming more distributed.	**Multi-stakeholder partnerships** are replacing siloed interventions.	**Systems change, collective impact and adaptive learning** are no longer peripheral concepts — they are becoming operational necessities.

Yet despite these advances, one of ***the greatest structural challenges facing the social and impact economy remains fragmentation***. Knowledge stays siloed. Data remains inaccessible. Learning stays disconnected. Organisations duplicate effort. Community insight is underutilised. Institutional capability is unevenly distributed. Administrative burden consumes scarce capacity. Funders

struggle to see ecosystems holistically. Smaller organisations remain excluded from strategic infrastructure.

The challenge is not merely one of resource scarcity - though chronic underfunding of the NPO sector remains a structural reality that no technology can solve on its own. Increasingly, the challenge is also one of coordination, intelligence and institutional capability. This is why artificial intelligence represents one of the most consequential shifts to emerge within the social economy in decades.

> ***The funding architecture that AI cannot fix - but funders can:*** *South Africa's social sector is chronically underfunded at the institutional level. The bulk of CSI and philanthropic capital flows to programme costs: the visible, reportable, compliance-friendly expenditure. The data systems, staff development, MEL infrastructure, governance capacity and technology investment that would make programmes more effective and sustainable are routinely excluded from grant budgets or funded at token levels.*
>
> *AI does not solve this. In fact, AI adoption requires exactly the kind of institutional investment that current funding norms actively discourage. The single most powerful intervention available to funders who want AI to deliver value in the social economy is simple: fund the institutional overhead that makes intelligent organisations possible. Not as a concession. As a strategy.*

Much of the global discourse surrounding AI has focused on automation, productivity and technological disruption. Yet within the social and impact economy, the more profound question is not whether AI automates work, but whether it transforms how ecosystems coordinate, learn, adapt and make decisions.

AI has the potential to reshape ecosystem coordination, institutional strengthening, impact measurement, collaborative governance, knowledge accessibility, programme design, philanthropic decision-making, systems learning, community engagement and the very architecture of social infrastructure. Importantly, ***AI may create disproportionate value within social systems precisely because social systems are information-rich, relationship-dependent and coordination-intensive environments.***

The social economy fundamentally operates through knowledge, relationships, learning, trust, systems navigation, institutional capability and adaptive coordination - all areas where AI can significantly augment human capability. However, this transformation is neither inevitable nor inherently equitable. Without intentional governance, AI could deepen existing inequalities, reinforce institutional exclusion, centralise power and accelerate forms of digital dependency that undermine local autonomy.

The question therefore is not whether AI will influence Africa's social economy. It already is. The more important question is: ***what kind of social economy will AI help build?*** Will AI become another extractive technology layer imposed onto African systems? Or can it become a tool for democratising institutional capability, strengthening ecosystem intelligence, accelerating collaborative learning and enabling more adaptive, equitable and locally-rooted social systems?

This guide argues that: *AI could fundamentally reshape the architecture of Africa's social and impact economy, not simply at the organisational level, but at the ecosystem level. And that the choices made now, about governance, ethics, inclusion, localisation, infrastructure and institutional design, will shape the trajectory of the continent's AI-enabled social future.*

SECTION TWO

Understanding the Social and Impact Economy

Who this section is for: All readers, with particular relevance for funders, intermediaries and policymakers who need to understand the South African ecosystem before engaging with AI's implications for it. If you are already deeply familiar with the sector's landscape, you may move quickly through this section - but the POPIA context in the legislative subsection is essential for everyone.

The social and impact economy is often described through the language of organisations, projects and funding streams. Yet increasingly, this framing is insufficient to capture either the complexity of how social change actually happens or the systems-level opportunities that AI presents.

Social change does not occur through isolated institutions operating independently. It emerges through interconnected ecosystems composed of multiple actors, relationships, infrastructures and systems. Within Africa's social economy, these actors include philanthropic foundations, bilateral and multilateral donors, corporate social investment programmes, civil society organisations, social enterprises, intermediaries, academic institutions, consultants and advisors, media organisations, civic movements, grassroots networks, governments at all levels, and communities themselves.

These actors do not function in isolation. They exchange knowledge, capital, legitimacy, data, influence, relationships, capability and trust. **This means the effectiveness of the social economy depends not only on the strength of individual organisations, but on the quality of the connections between them.** In practice, however, many social ecosystems remain deeply fragmented. Different sectors operate with different languages, priorities, incentives and reporting structures. Academic research remains disconnected from implementation practice. Community insight is frequently excluded from formal decision-making systems. Smaller organisations struggle to access funding ecosystems. Learning remains trapped in reports that are seldom synthesised or shared effectively.

The South African Social Economy: A System Under Pressure

South Africa's social economy has distinctive characteristics that shape both its challenges and its AI-enabled opportunities. The country's corporate social investment landscape - driven substantially by B-BBEE legislation and the social and ethics committee requirements of the Companies Act - represents one of the most significant institutionalised social investment ecosystems on the continent. By 2025, CSI expenditures had risen to approximately R13 billion, with 93% directed to NPOs, primarily in education, community development and food security.

> ***Key context:*** *Despite R13 billion in annual CSI expenditure, the bulk of this funding goes to programme costs rather than the institutional capacity, data systems and learning infrastructure that would make programmes more effective and sustainable.* ***AI could reshape both how this capital is allocated and what organisations are able to do with it.*** *Geographic concentration is also stark - Gauteng receives a disproportionate share while Limpopo, the Northern Cape and North West each receive only 1% of average CSI spend. (Source: Trialogue Business in Society Handbook, 2025)*

South Africa's approximately 235,000 registered NPOs operate under the NPO Act (71 of 1997), with NPC and Section 21 company structures adding further legal complexity. The sector is characterised by enormous institutional diversity: from large, well-resourced organisations with professional staff and international donor relationships, to community-based organisations operating with minimal infrastructure, volunteer capacity and hand-to-mouth funding. ***This inequality of institutional capability is one of the social economy's most persistent structural challenges, and one where AI could play a genuinely democratising role.***

The South African Social Economy: A Legislative Perspective

The Protection of Personal Information Act (POPIA, 2020) creates important data governance obligations for any AI deployment in the sector, requiring careful design of consent processes, data storage, and the use of beneficiary information. ***Any AI strategy for the South African social economy must engage seriously with POPIA compliance, not as a compliance burden, but as an opportunity to build community-centred data governance that restores trust and agency to those most affected.***

> ***In many ways, Africa's social economy suffers from an intelligence deficit rather than an effort deficit alone.*** *The challenge is not a lack of commitment. The challenge is that ecosystems struggle to coordinate information, synthesise learning, share intelligence, identify patterns and adapt collectively. This is where AI becomes particularly relevant - not as a productivity tool, but* ***as a potential intelligence infrastructure.***

SECTION THREE

The Core Problem: Fragmentation, Institutional Inequality and Coordination Failure

Who this section is for: All readers. This section names the problems that AI is being asked to address. If you work in the sector, you will recognise them. If you are a funder, this is the argument for why institutional investment in AI capability is a strategic priority, not an overhead cost.

Across South Africa's social economy, fragmentation manifests at every level of the system and carries enormous social costs.

01
ORGANISATIONAL LEVEL

Many NGOs operate with limited strategic, operational and analytical capacity. Smaller community-based organisations often lack access to research support, institutional development, fundraising expertise, data systems, monitoring and evaluation capability, and strategic planning infrastructure. At an ecosystem level, organisations frequently work in parallel rather than in alignment: programmes are duplicated, knowledge is siloed, learning remains fragmented, partnerships are difficult to coordinate, community intelligence is underutilised and data systems are incompatible.

02
FUNDING LEVEL

Philanthropic capital is often distributed through fragmented reporting structures and disconnected grantmaking systems that prioritise short-term measurable outputs over ecosystem learning and long-term institutional resilience. Bridgespan's 2025 research — Delivering on the Promise of Corporate Social Investment in South Africa — based on interviews with 36 corporate, NPO and sector experts — found that most corporates still exhibit a compliance mindset rather than a transformational one, with CSI manager turnover, short grant cycles and transactional relationships all undermining the potential for lasting social change.

03
KNOWLEDGE LEVEL

Enormous volumes of research, evaluations and implementation learning remain inaccessible to practitioners and communities. Reports sit unread. Research remains disconnected from implementation. Community insights are poorly synthesised. This problem is especially acute in Africa, where many social systems operate within environments characterised by institutional inequality, uneven state capability, digital divides, donor dependence, linguistic diversity and infrastructural constraints.

Yet paradoxically, these conditions also create fertile conditions for innovation. Because African systems are often more adaptive, more network-based, more hybrid and less institutionally rigid than many Global North systems, they may be better positioned to integrate AI in ways that genuinely strengthen local agency rather than simply replicating the extractive dynamics of earlier technology waves.

Africa's leapfrog potential in AI is real - but the analogy requires careful, sector-specific qualification.

The mobile money comparison is the most commonly cited evidence for African AI leapfrogging, and it is instructive - but only up to a point. M-Pesa succeeded because it solved a concrete, daily problem (transferring money without a bank account) using infrastructure most Africans already had (a basic mobile phone and a SIM card). It required minimal literacy, minimal data, minimal device capability, and no stable electricity. It was designed for the infrastructure realities of its users from the start.

AI in the social economy faces a fundamentally different set of preconditions. Effective AI use ***requires reliable internet connectivity*** *- which reaches only 43% of Africa's population and is significantly lower in rural areas. It* ***requires capable devices*** *- smartphones or computers that many community-level practitioners and beneficiaries do not own. It* ***requires electricity stability*** *- a serious constraint in a country where load-shedding has cost the economy hundreds of billions of rands and disrupted organisational operations for years. It* ***requires data literacy*** *- the ability to understand, critically evaluate and appropriately apply AI outputs - which cannot be assumed across the sector's diverse workforce. And it* ***requires training data in African languages*** *- which remains critically scarce for most of the continent's languages, creating AI systems that work well in English and poorly in the languages most South African communities actually speak.*

This does not mean leapfrogging is impossible. It means it is sector-specific and infrastructure-dependent. *In areas where preconditions are partially met, leapfrogging is already happening: in urban health systems using AI-assisted diagnostic tools, in youth employment platforms reaching digitally connected job-seekers, in ECD data systems serving practitioners with basic smartphones. In rural agricultural and community development contexts, the preconditions for most current AI tools do not yet exist at scale.*

Microsoft's 2025 AI Diffusion Report confirms that ***no African country has yet reached 20% AI adoption.*** *The Stanford AI Index 2024 notes that 27% of Kenyans use ChatGPT daily - a striking data point, but one that reflects urban, connected, educated Kenyans, not rural subsistence farmers or community health workers in under-resourced clinics.*

The practical implication for social economy actors is this: ***leapfrogging is a possibility in specific sectors and geographies with the right infrastructure investments — not a general principle that justifies skipping the foundational work.***

The sector needs to invest in connectivity (models like Zenzeleni's community networks), in devices and power access, in African-language AI (Masakhane, Lelapa AI), and in the data literacy that makes tools usable. These are not preliminary steps before the real AI work begins. They are the AI work, for much of the sector, right now.

For AI to fulfil its potential in addressing this fragmentation, it must first be reframed. ***The sector needs to move beyond thinking about AI as an automation tool or a productivity hack, and begin thinking about it as a new kind of social infrastructure - one that can help ecosystems coordinate, learn, adapt and make better decisions at the speed and complexity that modern social challenges demand.***

SECTION FOUR

Reframing AI: From Automation to Ecosystem Intelligence

Who this section is for: All readers, and especially board members, senior leaders and CSI managers who are shaping organisational or portfolio-level AI strategy. This section makes the conceptual argument. It is deliberately brief - but the distinction it draws between AI as automation and AI as ecosystem intelligence is the lens through which everything that follows should be read.

*This section carries the guide's central conceptual argument, and it is worth pausing on the language before proceeding. When this guide speaks of "**ecosystem intelligence**," it is not using the word "intelligence" as a synonym for AI. It means something older and more fundamental: **the capacity of a system to know, learn, adapt and coordinate.***

A sector can have sophisticated AI tools and still lack ecosystem intelligence - if those tools are deployed in silos, governed without community accountability, or used to optimise programmes that were never well-designed in the first place. Conversely, the goal of building ecosystem intelligence - genuine systemic learning and coordination - is what should guide AI adoption decisions. The question is not "what AI tools can we use?" but "what would it take for this ecosystem to learn and adapt more effectively - and can AI help with that?"

Most AI discourse remains dominated by narratives of automation. AI is frequently discussed as a productivity tool, a content generator, a workflow optimiser, and - with growing urgency - as a threat to employment. While these discussions are important, they are insufficient within the context of the social and impact economy. The more transformative role of AI lies elsewhere entirely.

BEYOND AUTOMATION

AI fundamentally alters how ecosystems coordinate, learn, adapt, make decisions, process complexity and distribute institutional capability. This is a much deeper transformation than productivity improvement — it is a potential **reshaping of the intelligence architecture of the entire social sector.**

CONNECTIVE INFRASTRUCTURE

AI should not be understood primarily as a tool for replacing people. It should be understood as **connective intelligence infrastructure** — a new layer of capacity that helps systems see patterns, synthesise knowledge and coordinate action they could not previously sustain.

This distinction is critical. The greatest constraints within the social economy are often not labour shortages alone. They are information overload and fragmentation, weak ecosystem coordination, inaccessible institutional capability, and painfully slow learning cycles. AI is particularly powerful in precisely these environments - those characterised by complexity and information fragmentation. It can synthesise vast amounts of information, identify patterns across systems, accelerate knowledge translation, support strategic planning, strengthen institutional capability and improve ecosystem visibility.

What AI Can and Cannot Do

Importantly, intelligence does not equal wisdom. AI can process information at scale. It cannot replace the trust, contextual understanding, facilitation, ethical judgement, community legitimacy, political navigation and relational stewardship that are the true foundations of social change. In fact, the rise of AI makes these human capabilities even more important - and more distinctively valuable.

> ***The future social economy may increasingly become a hybrid system where AI augments institutional intelligence, while humans steward relationships, ethics, legitimacy and contextual judgement.*** *The most effective organisations may not be those with the most advanced technology - they will be those best able to hold both capabilities in productive tension.*

It is also essential to be clear about what AI does not address. Structural underfunding of the NPO sector, short-term grant cycles, donor dependence, the political economy of social investment, and the persistent marginalisation of community voice in decision-making - these are not technology problems. **AI cannot fix a broken funding architecture or a governance deficit. Any narrative that presents AI as a silver bullet for systemic inequality should be treated with deep scepticism.**

African AI Agency: Building Intelligence from Within

Any credible discussion of AI in Africa must centre African-built and African-governed AI solutions. The continent is not only a consumer of AI - it is increasingly a producer and shaper of it. This distinction matters enormously for the social economy.

LELAPA AI

A South African startup co-founded from the Masakhane research collective, has developed InkubaLM - a small language model supporting isiZulu, isiXhosa, Hausa, Swahili and Yoruba. Its product Vulavula provides voice interfaces, translation tools and local-language access to digital services. Co-founder Jade Abbott frames language as “the enabler”: if Africa solves its language barriers, it unlocks every other application of AI. Masakhane itself has released more than 350 open-source models and over 25 publicly available African-language datasets on Hugging Face since its founding in 2018.

THE DEEP LEARNING INDABA

A grassroots pan-African AI community - and related organisations like Data Science Africa, AI Saturdays Lagos and Ghana NLP are building the human infrastructure that AI adoption in the social sector will eventually depend on. Researchers funded by the Gates Foundation have recently launched African Next Voices, the largest AI-ready dataset for African languages, featuring 9,000 hours of recorded speech across 18 languages. These are not peripheral projects. They are the foundation on which a genuinely African AI ecosystem will be built.

*In July 2024, the African Union Executive Council endorsed the Continental AI Strategy at its session in Accra, Ghana - establishing **Africa's commitment to an Africa-centric, development-focused approach to AI aligned with Agenda 2063 and the SDGs.** The strategy calls for unified national AI approaches, regional cooperation, and ethical, responsible AI deployment. Sixteen African countries (of the African Union's 55 member states) currently have formal national AI strategies, with several more in development. South Africa's national AI policy framework is in development.*

SECTION FIVE

The AI-Enabled Social Impact Value Chain

Who this section is for: Programme staff, MEL practitioners, implementers and funders who want to understand where AI adds value across the full programme lifecycle. The value chain below is a practitioner's reference - a way of asking 'where in my work does AI matter?' Use it as a map, not a model. The deeper insight is that AI's most transformative value lies not within each stage but in connecting across them.

One of the most significant opportunities presented by AI lies in its potential application across the entire social impact lifecycle: from initial design and diagnosis through implementation, learning, reporting and eventual exit or sustainability planning. Historically, AI discussions within the social sector have focused narrowly on specific functions such as fundraising, education technology or impact measurement. **The more transformative insight is that AI has implications across every stage of the impact value chain.**

The table below maps AI application areas, specific tools and stakeholder relevance across each stage of the social impact lifecycle. It is intended as a practical reference for organisations seeking to understand where AI can add genuine value in their specific context. It is important to note however, that tools are timebound. For the purpose of this paper therefore, the categories of tools with current examples, ***represent a snapshot in time.***

The value chain that follows is a linear map of something that is not, in practice, linear. It is offered as a practitioner's reference: a way of asking 'where in my work does AI add value?' - not as a model of how change actually happens. Real social systems are iterative, feedback-dependent and non-sequential. The value of mapping AI applications to a linear lifecycle is practical, not theoretical: it helps organisations identify entry points. The deeper insight - that AI's ecosystem-level value lies precisely in connecting across these stages, not just within each one - sits underneath the table, not inside it.

> ***A note on this framework:*** *the ten-stage lifecycle below organises AI's applications in a linear sequence. This is a deliberate simplification. The more important insight - that AI can connect across these stages, enabling real-time feedback between implementation and design, between exit planning and programme learning - is precisely what makes it transformative at the ecosystem level. Use this table as a reference, not a model.*

Stage	AI Application Areas	Key Tools / Examples	Stakeholders Involved	Risk Level
Design & Diagnosis	Ecosystem mapping, needs analysis, community sentiment	ESRI ArcGIS AI, Ushahidi, Kobo Toolbox + AI analysis, Google Earth Engine	Funders, NGOs, Government, Communities	Medium

Stage	AI Application Areas	Key Tools / Examples	Stakeholders Involved	Risk Level
	analysis, geographic gap identification			
Strategy & Planning	Scenario modelling, trend synthesis, portfolio gap analysis, environmental scanning	ChatGPT, Claude, Perplexity, AI-enabled logic model tools, strategic synthesis platforms	Foundations, Intermediaries, Boards	Low
Fundraising	Proposal drafting support, donor matching, equity analysis, funding gap identification	Instrumentl, Submittable, custom GPT tools, AI-assisted grant writing	NGOs, Intermediaries, Social Enterprises	Low
Programme Design	Predictive modelling, theory of change synthesis, community intelligence, localisation	b.world, Canopy, AI-assisted TOC platforms, participatory design tools	NGOs, Implementers, Communities	Medium
Implementation	Workflow automation, multilingual communication, volunteer coordination, beneficiary support	Vulavula (Lelapa AI), WhatsApp AI bots, Microsoft Copilot, multilingual chatbots	CBOs, NGOs, Beneficiaries	High
Monitoring & Evaluation	Qualitative data synthesis, real-time dashboards, pattern detection, outcome analysis	DHIS2 + AI, b.world, 60 Decibels, SPAR4D, DEval AI tools	MEL specialists, Funders, Implementers	High
Reporting	Automated narrative generation, data visualisation, funder reporting, impact storytelling	Power BI + AI, Tableau AI, Claude/GPT report drafting, GRI-aligned tools	Funders, Boards, Regulators, Public	Medium
Knowledge Management	Repository creation, cross-sector synthesis, translation, institutional memory	Notion AI, Confluence AI, Masakhane translation tools, knowledge base platforms	Academia, Intermediaries, Practitioners	Low
Scaling & Replication	Pattern identification, comparative analysis, replication modelling, ecosystem benchmarking	Network analysis tools, AI comparative research, diffusion mapping platforms	Funders, Policy makers, Intermediaries	Medium
Exit & Sustainability	Institutional sustainability analysis, dependency mapping, scenario forecasting, risk modelling	Scenario planning AI, financial sustainability tools, transition planning frameworks	Boards, Funders, Government, DFIs	High

Table 1: AI-Enabled Social Impact Value Chain: Applications, Tools and Stakeholders Across the Programme Lifecycle

Across this value chain, a number of stages deserve particular attention because of the depth of transformation they may enable.

Ecosystem Intelligence and Diagnostics

One of the most important challenges within the social economy is **limited ecosystem visibility**. Funders, NGOs and policymakers often struggle to understand who is operating where, which communities remain underserved, where duplication exists, what systemic gaps remain and how different actors interact. ***AI can significantly strengthen this ecosystem intelligence through stakeholder mapping, geographic gap analysis, community sentiment analysis, systems mapping, network analysis, risk detection and pattern identification across datasets.***

An intermediary coordinating youth employment initiatives across provinces, for example, could use AI to identify geographic service gaps, map overlapping interventions and detect areas where youth engagement remains persistently weak.

Monitoring, Evaluation and Learning

This may become one of the **most transformative AI application areas** within the social economy. Traditional monitoring and evaluation systems often struggle with the complexity, systems-level dynamics, qualitative richness and adaptive learning demands of contemporary social programmes. ***AI is particularly useful in complexity-rich environments: it can synthesise qualitative data at scale, support real-time learning systems, enable developmental evaluation, recognise outcome patterns and map systems feedback loops.***

An intermediary coordinating multiple implementation partners could use AI to synthesise learning across reports, interviews and community feedback data in ways that would be practically impossible through manual analysis.

Knowledge Translation and Ecosystem Learning

One of the greatest failures of the social economy is that learning often remains trapped within institutions. Programme evaluations sit in report archives. Academic research never reaches practitioners. Community insights are captured in consultations but never systematically analysed. ***AI can help create more connected learning ecosystems through searchable knowledge repositories, automated cross-sector synthesis, multilingual translation of research into local languages, and ecosystem learning dashboards that give funders, implementers and communities a shared view of what the evidence shows.***

Exit, Sustainability and Systems Resilience

Few social organisations systematically analyse long-term sustainability risks, ecosystem dependencies or transition pathways. Yet this may be one of the highest-stakes application areas for AI in the South African social economy: precisely because the financial precarity facing the NPO sector is acute and worsening.

Government contracting has become less reliable as fiscal constraints tighten. CSI cycles are shortening. The cost of operations continues to rise. Many NPOs that have served communities for decades are one funding cycle away from closure - not because their work lacks value, but because they have never built the analytical infrastructure to model their own sustainability, map their dependencies, or plan credible transitions.

AI may support more sophisticated resilience planning through four specific capabilities:

01 **Institutional sustainability modelling** analyses an organisation's income streams, cost structure, funder concentration risk and scenario-based projections to identify vulnerability before it becomes a crisis. Tools that combine financial data with funder relationship intelligence can flag when an organisation is dangerously dependent on a single source of funding, and model what diversification pathways might look like.

02 **Ecosystem dependency mapping** asks a harder question: what would happen to the communities this organisation serves if it closed tomorrow? Which services would go unmet? Which referral pathways would collapse? Which community relationships would be lost? AI can map these dependencies across a portfolio of organisations, giving funders and intermediaries the systemic picture they need to make exit decisions that minimise community harm.

03 **Scenario forecasting** applies AI-enabled modelling to the range of futures an organisation might face: a 30% reduction in CSI funding, the loss of a major government contract, the entry of a new actor into their geography, a shift in funder priorities. Organisations that have stress-tested their model against realistic scenarios are better positioned to respond when those scenarios materialise — as they inevitably do.

04 **Transition and handover planning** addresses the often-neglected question of what happens when a programme ends or an organisation closes. AI can help organisations document institutional memory, identify which community relationships and capabilities can be transferred to other actors, and design handover processes that protect community continuity. In a sector where institutional memory is frequently lost when organisations close or leadership changes, this is not a minor capability.

The practical implication for funders is significant: ***exit planning should not be an afterthought at the end of a grant cycle. It should be a funded, AI-supported component of programme design from the outset.*** An organisation that cannot model its own sustainability is an organisation whose social impact is permanently fragile, regardless of how well its programmes perform.

South African context: *The NPC's State of the Nonprofit Sector reports consistently identify financial sustainability as the sector's most acute structural challenge. An estimated 40% of registered NPOs report being unable to meet their financial obligations at some point in any given year. AI-enabled sustainability modelling is not a luxury for well-resourced organisations, it is a risk management tool that the sector's most financially vulnerable organisations need most. Funders who want their grant investments to create lasting impact must be willing to fund the analytical infrastructure that makes sustainability planning possible.*

PART TWO

AI Across the Ecosystem

IN THIS PART

> *AI's most transformative social value is not efficiency — it is equity.*

SECTION SIX

AI Within the Social Economy: A Stakeholder Matrix across the Ecosystem

Who this section is for: All readers - each stakeholder group will find their own context in the matrix. CSI managers, read your row carefully. Intermediaries pay particular attention to the repositioning argument. Grassroots CBO leaders, the equity gap section is written for and about you. The case study table that follows is the most detailed compilation of South African AI practice currently available.

AI will not transform only individual organisations. It will transform ecosystems. Different actors within the social economy will experience this transformation differently, with different opportunities, different risks and very different starting points in terms of readiness and access.

The following matrix maps these dimensions for the sector's key stakeholder groups.

Stakeholder	Primary AI Use Cases	Key Opportunities	Key Risks	Readiness Level
Corporate CSI Managers	Portfolio analytics, impact reporting, grant compliance, grantee capacity analysis	Shift from compliance to transformation; smarter, evidence-led capital allocation	Algorithmic bias in grant selection; data privacy of beneficiaries	**Medium:** tools exist, adoption remains nascent
Philanthropic Foundations	Ecosystem intelligence, funding gap analysis, grantee support, strategic philanthropy	Democratise access; smarter thematic investment; ecosystem-level visibility	Over-reliance on quantitative data over relationships	**Medium:** some early adopters emerging
Intermediaries	Ecosystem coordination, MEL synthesis, knowledge platforms, institutional strengthening	Evolve into 'ecosystem intelligence hubs'; shared infrastructure role	Disintermediation risk if smaller actors bypass intermediary layer	**Low-Medium:** significant opportunity, low current adoption
NPOs / NGOs (large)	Programme delivery AI, M&E systems, communications, reporting automation	Operational efficiency; access to strategic capability previously unaffordable	Governance risk, data quality challenges, staff capacity	**Low-Medium:** individual use growing
CBOs / Grassroots Orgs	Multilingual chatbots, community feedback tools,	Democratisation of institutional capability; reduced administrative burden	Digital divide, connectivity, AI literacy gap, infrastructure absence	**Low:** major equity gap; most at-risk of exclusion

Stakeholder	Primary AI Use Cases	Key Opportunities	Key Risks	Readiness Level
	basic administrative AI			
Government / DFIs	Policy analytics, service delivery optimisation, SDG tracking, budget modelling	Better evidence-based policymaking; improved social investment returns	Surveillance risk; algorithmic discrimination; procurement constraints	**Low:** regulatory readiness poor; policy vacuum
Academia	Research synthesis, knowledge translation, participatory AI, rapid evidence reviews	Bridge the practice-research gap; faster knowledge use in communities	Publication bias; commercialisation tensions; ethical oversight gaps	**Medium:** tools widely available, governance frameworks developing
Social Enterprises	Product/service optimisation, impact measurement, investor-readiness reporting	Competitive edge; faster impact evidence generation for investment readiness	Overemphasis on measurable scale vs. depth and quality of impact	**Medium:** commercially motivated, faster adoption likely
Beneficiaries / Communities	Voice tools, feedback systems, co-design AI, community-led data collection	Agency and genuine participation in programme design and evaluation	Data extraction without benefit; consent and privacy violations; exclusion	**Very Low:** most under-served; participation remains largely extractive

Table 2: Stakeholder AI Application Matrix: Opportunities, Risks and Readiness Across the Social Economy

Corporate CSI Managers: From Compliance to Intelligence

Corporate social investment in South Africa has long been shaped by the compliance imperative of B-BBEE legislation. The result, as Bridgespan's 2025 research confirms, is that most corporates exhibit a regulatory compliance mindset in their CSI work - a transactional approach that limits social impact. ***AI has the potential to accelerate the shift from compliance to transformation by giving CSI managers ecosystem-level visibility:*** understanding where duplication exists, identifying underfunded geographies and themes (only Gauteng receives meaningful CSI concentration; Limpopo, Northern Cape and North West are persistently under-served), and enabling smarter, more evidence-led capital allocation.

Intermediaries: From Coordinators to Intelligence Platforms

Intermediaries may become some of the most important AI-enabled institutions within the social economy, precisely because they already operate at the intersection of coordination, learning, systems navigation, institutional strengthening and ecosystem stewardship. ***AI could enable intermediaries to evolve from administrative coordinators into ecosystem intelligence platforms - providing AI-enabled institutional support, collaborative learning infrastructure, shared knowledge platforms and adaptive coordination capability to the organisations within their networks.*** This is a profound repositioning, but one that requires intentional investment in mandate, governance,

data infrastructure and financial sustainability. Many South African intermediaries are themselves financially precarious, and the transition to an intelligence-platform model will require funder commitment to the institutional capacity of the intermediary layer itself.

NGOs and Social Impact Organisations: Democratising Capability

For many NGOs, AI may democratise access to institutional capability in ways that were previously unimaginable. Strategic planning support, evaluation expertise, research synthesis, advanced communications, governance assistance and learning systems have historically been accessible primarily to well-resourced organisations with strong funder relationships. ***AI significantly lowers these barriers - enabling smaller organisations to access quality strategic support, proposal assistance, research capability, and learning systems without the overhead of expensive specialist staff or consultants.***

Grassroots CBOs: The Equity Gap That Must Be Addressed

Community-based organisations, the closest to the communities they serve and often the most trusted by them, face the greatest risk of being left behind in an AI-enabled social economy. Digital divides, connectivity gaps, AI literacy deficits and the absence of appropriate tools in African languages mean that without deliberate equity-centred design, ***AI could widen rather than narrow the institutional capability gap between large and small organisations.*** This is not an inevitable outcome. Offline-capable tools, shared infrastructure models, multilingual design (Vulavula and Masakhane's work is directly relevant here), and intermediary-supported AI adoption pathways can all reduce this risk - but only if funders and ecosystem builders prioritise equity as a design principle from the outset.

Beneficiaries and Communities: From Subjects to Participants

Perhaps the most important and most neglected dimension of AI in the social economy is the role of communities themselves - not merely as subjects of AI-enabled programmes, but as participants in their design, governance and benefit. ***AI voice tools, community feedback systems and co-design platforms offer genuine possibilities for deepening community agency.*** But without explicit participatory governance mechanisms and consent-first data design, AI could easily reproduce the extractive dynamics of conventional development: taking community data in, generating institutional intelligence out, and delivering little benefit back to the people whose lives and experiences created the data in the first place.

Strategic Stakeholder Differences: Non-Profit vs. Private-Sector vs Foundation and Intermediary AI

To understand application of AI across the ecosystem, it is helpful to look at how these AI programmes and applications vary across different **institutional arrangements**:

Feature	NGO-Driven AI	Corporate CSR/CSI AI	Foundations and Intermediary AI
Primary Funding	Independent grants, philanthropy, public fundraising.	Fixed corporate budgets tied to B-BBEE scorecards and ESG goals.	Endowments, pooled corporate trust funds, and third-party capital management.
Strategic Intent	Direct community service delivery, crisis response, and hyper-local aid.	Commercial alignment, brand equity, and B-BBEE skills compliance.	Systemic transformation, sector-wide sustainability, and policy advocacy.
Tech Infrastructure	Built on free, open-source, or third-party APIs (WhatsApp/SMS).	Custom proprietary software, enterprise cloud suites, and big data.	Multi-tenant grant platforms, cross-sector data lakes, and localised civic-tech portals.
Data Utilisation	Focuses on qualitative feedback and lived-experience storytelling.	Focuses on quantitative, predictive analytics and operational efficiency.	Standardising M&E across hundreds of NGOs to eliminate reporting fragmentation.
Scalability	Hyper-localised; dependent on community-by-community rollout.	National scaling from day one, leveraging existing commercial logistics and footprints.	Ecosystem scaling; democratising tools so multiple small NGOs can adopt AI simultaneously.
Risk Management	High operational vulnerability due to unpredictable funding cycles.	Commercial risk oversight and strictly defined corporate parameters.	Proactively researching social impacts, such as addressing algorithmic bias or POPIA compliance.

Table 3: Stakeholder AI Application Matrix: Different Institutional Arrangements Across the Social Economy

Case Studies and Application of AI Across the Social Economy Ecosystem (and across the stakeholder universe) in South Africa

South Africa's social economy features robust and rapidly expanding experimentation with artificial intelligence. Across the ecosystem, in non-governmental organisations, community-based organisations, social enterprises, research institutions, corporate CSI programmes, government agencies and intermediaries, AI is being tested, adapted and deployed to address some of the country's most persistent development challenges.

What distinguishes South Africa's AI landscape from many other African countries is both its breadth and its intentionality. Applications range from community health workers using SMS-based advisory tools to national government agencies deploying biometric identity verification systems at scale. Corporate supply chains are being repurposed to redirect food to soup kitchens. Grassroots research collectives are building language models for isiZulu and isiXhosa. Social enterprises are mapping the hidden township economy through gamified mobile platforms. These are not isolated experiments, they are early evidence of a social economy actively reshaping AI to fit local contexts, local languages and local needs.

The case studies below are organised by ecosystem actor type to reflect a central argument of this guide: **that AI's transformative potential in the social economy lies not only in what individual**

organisations can do, but in how different actors: NPOs, corporates, intermediaries, government, grassroots movements and research communities, bring complementary capabilities that, when connected, create ecosystem-level intelligence.

Each case has been assessed for accuracy, materiality and social economy relevance. Where AI applications are in pilot or experimental stages rather than at full scale, this is noted.

> ***A Note on scope:*** *The cases below focus on South African applications and South African organisations. Where an international example is referenced because of its direct relevance to South African practice, ITIKI, for instance, this is clearly indicated.*

Actor Type	Organisation	AI Application and Social Economy Relevance
GRASSROOTS RESEARCH ORGANISATIONS AND ECOSYSTEM BUILDERS		
Grassroots Research & Ecosystem Catalyst	**Deep Learning Indaba**	Founded in 2017 in Johannesburg, the Deep Learning Indaba has grown from a single academic gathering into the most influential grassroots AI movement on the African continent. Its mission is explicit: to transform Africans from consumers of global AI into active builders, researchers and governors of their own AI ecosystem. For the social economy, its significance operates on four levels: • **Social AI incubation:** Incubation of social AI applications through its annual Community Challenge, which has seeded projects spanning agricultural resilience, maternal health, multilingual text triage and climate modelling. • **Brain drain prevention:** Talent retention by creating professional pathways for African data scientists that keep expertise on the continent and connected to local development challenges, rather than absorbed into corporate or international roles. • **Policy and governance:** Ethical AI advocacy by equipping civil society and academic researchers to audit foreign AI models deployed in South African schools, clinics and financial systems - ensuring alignment with local legislation including POPIA. • **Ecosystem founding:** As the birthplace of both Lelapa AI and the Masakhane Research Foundation, the Indaba has demonstrated that grassroots community building is a legitimate and powerful route to institutional AI infrastructure. ***The Indaba is the connective tissue of South Africa's social AI ecosystem. Its contribution is structural, not transactional - it builds the people and relationships that make everything else possible.***
Pan-African Research Foundation	**Masakhane Research Foundation**	Masakhane is a pan-African grassroots research organisation anchoring a distributed network of scientists,

Actor Type	Organisation	AI Application and Social Economy Relevance
		linguists and community researchers committed to building Natural Language Processing (NLP) capabilities for indigenous African languages. Its core mandate is to prevent digital colonisation - ensuring that African communities are not excluded from the AI era because the technology does not speak their languages. Masakhane has published more than 350 open-source models and over 25 African-language datasets on Hugging Face since 2018. The Masakhane Web platform enables domain experts and community speakers to interact with models and provide real-time dialect feedback, allowing continuous retraining to reduce historical linguistic bias. Through its LINGUA Africa programme, developed in partnership with global technology companies and philanthropic funders, Masakhane provides computing resources and grants to African developers building language AI for real-world applications including: • Clinical documentation and health advice accessible in patients' first languages • Localised agricultural climate alerts for smallholder farmers • Legal information translation for rural communities navigating civic and legal systems ***Masakhane's open-source approach is a deliberate strategic choice: by making language infrastructure freely available, it democratises AI capability across the entire social sector - any NPO, CBO or community health programme can build on its foundations without proprietary cost.***
Social Enterprise — Language AI	**Lelapa AI**	Lelapa AI is a South African AI company co-founded from the Masakhane research community, building AI infrastructure specifically designed for African languages and African contexts. It represents one of the most significant examples of African AI agency - AI built by Africans, for African communities, governed by African data principles. Its flagship products have direct social economy relevance: • **InkubaLM:** InkubaLM is an open-access foundational language model supporting isiZulu, isiXhosa, Hausa, Swahili and Yoruba, enabling township businesses, community clinics and NPOs to build voice-activated and text-based tools in the languages their communities actually speak. • **Vulavula:** Vulavula is a resource-efficient language processing API trained to accurately transcribe, analyse and translate code-switched multilingual dialogue, handling the fluid, real-world mixing of isiZulu and English, or Sesotho and Afrikaans, that characterises how most South Africans actually communicate. • **Esethu Framework:** The Esethu Framework is a data governance model co-created with communities to

Actor Type	Organisation	AI Application and Social Economy Relevance
		ensure that the people who contribute linguistic data are ethically credited and compensated. It is a direct response to data colonialism, the extraction of community knowledge for commercial AI without benefit flowing back to the community. ***Co-founder Jade Abbott frames language as "the enabler": solve the language barrier and every other AI application in health, education, livelihoods and governance becomes accessible to millions of South Africans currently excluded from digital services.***
NON-GOVERNMENTAL ORGANISATIONS AND CIVIL SOCIETY		
NPO — Maternal Health	**MomConnect (National Department of Health)**	MomConnect is a national maternal health platform operated through the Department of Health, using SMS and WhatsApp-based messaging to deliver personalised health guidance to pregnant women and new mothers across South Africa. It is one of the largest and most evidence-supported mHealth programmes on the continent. The platform uses automated message sequencing and AI-assisted personalisation to calibrate guidance according to gestational age, health risk indicators and the mother's preferred language. At national scale, it has improved maternal health tracking and contributed to improved infant health outcomes in participating districts. For the social economy, MomConnect illustrates several critical design principles: it operates through infrastructure most South African women already have (a basic mobile phone); it works in multiple South African languages; it is embedded in an existing public health system rather than building parallel infrastructure; and it maintains a human referral pathway for high-risk situations - a nurse or midwife is always accessible when automated guidance is insufficient. ***MomConnect is the paradigmatic example of the 'embedded AI' principle: it works because it solves a daily problem using existing infrastructure, in the language of the user, within a system they already trust.***
NPO — Youth Livelihoods	**Harambee Youth Employment Accelerator and SAYouth.mobi**	Harambee is South Africa's largest youth employment intermediary, connecting marginalised job-seekers to work and learning opportunities at scale. SAYouth.mobi is the digital platform Harambee built to extend this reach - it is now one of the largest datasets on youth unemployment and workforce readiness in Africa, with millions of registered young people. AI is embedded throughout its operations: • **Algorithmic matching:** A dynamic demand-supply matching algorithm evaluates candidates based on location, verified skills and potential, deliberately moving away from CV-based screening, which

Actor Type	Organisation	AI Application and Social Economy Relevance
		systematically disadvantages young people without formal work experience or quality schooling. • **Conversational AI:** AI-powered chatbots handle tens of thousands of interactions monthly, guiding work-seekers through registration, surfacing learning resources and conducting audio-based communication assessments. • **Wellbeing monitoring:** Speech and text analytics in Harambee's call centres automatically flag distress signals - words and patterns associated with depression or crisis - enabling immediate human intervention and referral to support. • **Data-driven advocacy:** The platform's aggregated data is used to identify structural barriers to youth employment in real time, informing advocacy to government and private-sector employers about what is actually preventing young people from accessing work. ***Harambee's model is notable for the explicit governance principle it applies: AI recommends, humans decide. No placement or eligibility decision is made by AI alone - the algorithm narrows the field, and a human counsellor makes the call. This is the correct model for high-stakes decisions affecting vulnerable people.***
NPO — Capacity Building	**Khulisa Social Solutions**	Khulisa Social Solutions is a South African NPO working in crime prevention, youth development and substance abuse. In its efforts to understand community dynamics and programme effectiveness, Khulisa has integrated AI-assisted analysis of community-generated qualitative data, including community storytelling, participatory art and lived-experience narratives, to identify patterns of systemic disadvantage and programme impact that conventional quantitative data collection tends to miss. The approach combines AI's capacity to process large volumes of unstructured qualitative data with Khulisa's deep community relationships and contextual knowledge - allowing the organisation to build a richer and more community-grounded evidence base than either approach could produce alone. ***Khulisa's application illustrates an important principle: AI's value in community-facing NPOs is not in replacing human relational intelligence, but in helping organisations make sense of the qualitative richness their community engagement already generates. The insight is in the combination, not the technology alone.***
NPO — Digital Access	**Silulo Ulutho Technologies**	Silulo Ulutho Technologies is a social enterprise with a network of digital access hubs across township and rural communities in the Western Cape and Eastern Cape. Its mission is to bridge the technological divide, providing affordable internet access, digital skills training and

Actor Type	Organisation	AI Application and Social Economy Relevance
		employment preparation to communities historically excluded from the digital economy. Within its training hubs, Silulo uses automated skills assessment frameworks and AI-assisted course guidance to evaluate learners' digital aptitudes and personalise their training pathways. Rather than a one-size-fits-all curriculum, the system identifies individual starting points and accelerates learners toward the specific digital competencies most relevant to available employment pathways in their context. ***Silulo's significance for the social economy goes beyond its direct service: its network of physical hubs in under-connected communities provides the access infrastructure without which AI literacy programmes remain urban and elite phenomena. The hubs are, in effect, AI readiness infrastructure for communities at Readiness Level 1.***
EARLY CHILDHOOD DEVELOPMENT — AN ECOSYSTEM IN MOTION		
NPO / Data Systems — ECD	**DataDrive2030**	DataDrive2030 is a South African non-profit applying machine learning and large-scale data systems to improve early childhood development outcomes. It is one of the most technically sophisticated data-for-impact organisations operating in the South African social sector. Its AI applications are grounded in nationally representative datasets, including the Thrive by Five Index - one of the most comprehensive surveys of ECD quality and child development outcomes in South Africa. Machine learning models applied to this data identify the most significant predictors of early learning outcomes: which combinations of household conditions, structural factors and programme quality have the highest predictive power for child development trajectories. This analysis enables the generation of decision-relevant profiles - specific risk configurations that allow funders, programme designers and government departments to target support to the children and communities at greatest developmental risk before they enter the formal school system. Rather than allocating resources based on broad geographic or income categories, the data enables hyper-local, evidence-based prioritisation. ***DataDrive2030 exemplifies the role of specialist data organisations in the social ecosystem: they do not deliver services, but they produce the intelligence infrastructure that makes other organisations' service delivery more effective and more equitable***.
Social Innovation — ECD Screening	**Innovation Edge and Thrive by Five (with DataDrive2030)**	Innovation Edge is a South African social innovation organisation catalysing systems change in early childhood development. In collaboration with the Thrive by Five Index and DataDrive2030's data systems, it has supported the deployment of automated developmental screening built on the Early Learning Outcomes Measure (ELOM) - a

Actor Type	Organisation	AI Application and Social Economy Relevance
		standardised, validated tool for assessing developmental readiness in children aged 4–5 years. By digitising assessment data collected by community practitioners and processing it through predictive analytics, the system can identify developmental delays at population scale, flagging geographic clusters where developmental stunting is most acute and allowing targeted resource allocation before children enter Grade R. The system integrates cognitive, language and fine-motor development indicators alongside nutritional and socioeconomic data to build a multi-dimensional risk profile for each catchment area. ***The population-level flagging capability is the key innovation here: it shifts early intervention from a reactive, individual-child model to a proactive, community-level risk management model - enabling social investment to reach the right places before the window for effective intervention closes.***
Trust Networks — ECD Health	**DGMT, Zero2Five Trust and One-to-One Africa**	This cluster of ECD-focused trusts and foundations uses AI-assisted data collection and analysis to address childhood stunting and malnutrition, one of the most significant and preventable causes of long-term cognitive impairment in South African children. Community health workers and ECD practitioners collect anthropometric data - mid-upper arm circumference, weight-for-age, height-for-age - using mobile tools integrated with platforms like CommCare. Backend algorithms evaluate these measurements against growth standards in real time, flagging anomalies and automating the data verification process across thousands of clinics and community sites. The result is a continuously updated, quality-verified digital health dashboard that allows programme managers and funders to track nutritional status at population scale, identify deteriorating conditions before they become medical crises, and direct referrals to appropriate care with speed that manual systems cannot match. ***The significance of this application is not the sophistication of the AI - the algorithms involved are relatively straightforward. The significance is that the combination of mobile data collection and automated verification makes the data reliable enough to act on at scale, across settings where data quality has historically been a persistent challenge.***
Social Enterprise — ECD Admin	**GROW ECD**	GROW ECD equips women running informal township crèches, many of whom are unqualified but deeply committed practitioners serving communities with no alternative childcare, with enterprise-grade operational tools that would otherwise be inaccessible to organisations of their scale and resource level. The GROW Giraffe and GROW Meerkat app suites automate the administrative burden of running a registered

Actor Type	Organisation	AI Application and Social Economy Relevance
		ECD centre: attendance records, fee tracking, dietary planning and compliance documentation for Department of Basic Education audits. The system is piloting machine learning modules to automate fee projections and generate grocery budgeting based on child attendance data. The social economy significance of GROW ECD is structural: informal crèches that cannot produce compliant documentation cannot access government subsidies, meaning the children they serve are doubly disadvantaged - attending an under-resourced facility that cannot access the funding designed to support it. By automating compliance, GROW ECD unlocks state subsidies for crèches that were previously excluded from formal funding simply because the administrative burden was beyond their capacity. ***This is AI as an equity tool: not solving a technical problem, but removing an administrative barrier that has systematically excluded the smallest, most community-embedded ECD providers from the funding they are entitled to.***
Digital Platform — ECD Knowledge	**ECD Connect (Incubated by DGMT)**	ECD Connect is a conversational AI platform delivering contextualised early childhood development guidance to practitioners and parents in South African languages through WhatsApp and basic mobile interfaces. It uses Natural Language Processing to respond dynamically to queries about child stimulation, age-appropriate learning activities, behavioural management and nutrition, turning a low-cost smartphone into a 24-hour educational support resource. For isolated home-based crèches with no access to formal trainer support, the platform provides the kind of professional guidance that has historically been available only to ECD practitioners in well-resourced, formally registered centres. The WhatsApp delivery channel is deliberate: it requires no app download, no data-intensive interface and no digital literacy beyond basic messaging. ***ECD Connect and the Zuzi platform (see below) represent a convergence of language AI and community service delivery - both demonstrate that the most significant accessibility innovation in digital health and education tools is often the decision to work through infrastructure communities already use, in languages communities already speak.***
Not for Profit — ECD Learning and GBV rights	**Zuzi**	**Zuzi** - created by the **NGO GRIT** - is a South African AI-assisted digital learning and content platform for ECD practitioners, parents and caregivers. Zuzi also provides immediate, safe access to information on gender-based violence (GBV), bodily rights, and wellbeing. Zuzi targets the approximately 40,000 ECD centres serving an estimated 850,000 children in the 0–4 age range, a system chronically under-resourced in both qualified practitioners and quality learning materials.

Actor Type	Organisation	AI Application and Social Economy Relevance
		The platform uses AI-assisted content personalisation to adapt learning resources to each practitioner's language preference, literacy level and geographic context - recognising that a one-size-fits-all resource designed for a qualified teacher in Johannesburg is often inaccessible to a home-based carer in rural Limpopo. Content is delivered through a low-data mobile interface specifically designed for practitioners with limited digital infrastructure. ***Zuzi's design philosophy - start with the practitioner's actual context, not an assumed context - and is a model for responsible AI design in under-resourced community settings. The personalisation is not a luxury feature; it is the mechanism that makes the content usable.***
HEALTH AND COMMUNITY HEALTH SYSTEMS		
Digital Health Organisation	**HealthEnabled**	HealthEnabled is a South African digital health organisation operating at the intersection of public health, technology and policy. Its primary focus is evidence synthesis and knowledge translation, taking the large volumes of global and African health research and making them practically usable for frontline workers, NPO programme teams and government policymakers. Its AI-enabled applications include the synthesis of clinical research into actionable, multilingual digital health frameworks delivered through WhatsApp, SMS and app-based channels - the same infrastructure community health workers and rural clinics already use. HealthEnabled also supports government departments and NPOs to navigate digital health ecosystems, helping organisations scale evidence-based health programming without building redundant digital infrastructure from scratch. A distinctive feature of HealthEnabled's work is its explicit commitment to gender-intentional and context-aware AI design - ensuring that digital health tools do not default to assumptions that exclude women, rural communities or non-English speakers. ***HealthEnabled occupies a critical but often invisible role in the social health ecosystem: it bridges the gap between what the research says and what practitioners can actually do. In an AI context, this knowledge translation function becomes more important as the volume of available evidence grows faster than any practitioner can process it manually.***
Health Research Organisation	**The Aurum Institute**	The Aurum Institute is a South African health research and implementation NPO specialising in TB and HIV. It combines large-scale clinical trial research with operational programme delivery, generating evidence and immediately translating it into practice within the same organisation. Its AI applications centre on the analysis of large, locally generated health datasets from operational research programmes including the WHIP3TB trial. Machine learning models applied to this data support:

Actor Type	Organisation	AI Application and Social Economy Relevance
		• Predictive modelling of infection hotspots, allowing public health resource allocation to be proactive rather than reactive • Differentiated service delivery design, identifying which patient populations benefit most from which intervention configurations • Optimisation of community health worker deployment, matching CHW availability to predicted demand patterns ***Aurum's evidence-to-policy pipeline is a model for the sector: it generates evidence through rigorous research, applies AI analytics to make that evidence actionable at the programme level, and advocates for policy change based on the combined weight of data and operational experience. This closed loop, from community to data to policy and back, is what intelligent social systems need to function.***
Corporate CSI — Preventive Health	**Discovery**	Discovery is a South African financial services group whose commercial model is built on behavioural incentive AI, its Vitality programme uses machine learning to encourage healthy behaviours through personalised rewards. Through its 'Force for Good' CSI initiatives and employee volunteering programmes, Discovery has partnered with community clinics to apply elements of this behavioural AI framework to community health interventions in under-served areas. The approach applies machine learning to anonymised, aggregated health metrics to identify geographic hotspots where chronic illness burden - particularly diabetes and hypertension - is elevated but unmanaged. Community health workers are then deployed proactively to these areas for preventive care visits, rather than waiting for symptoms to reach a crisis level requiring costly clinical intervention. ***Discovery's approach represents an important but underexplored model: applying commercially developed behavioural AI frameworks to social health challenges, where the commercial product's sophistication subsidises the social application's capability. The governance requirement is strict: anonymised, aggregated data only, with no individual profiling of community members outside their informed consent.***
Corporate CSI — Digital Safety	**Vodacom (Bright Sky SA)**	Vodacom has partnered with the UK-based Hestia charity to localise the Bright Sky application for South Africa, adapting an internationally developed GBV support tool for the South African context, including local emergency pathways, service directories and safety planning resources. The AI layer within Bright Sky classifies risk levels based on user inputs - responses to safety assessment questions - to surface contextually appropriate resources and emergency referral pathways. Text-based triage uses machine learning to route users to the most relevant support based on their

Actor Type	Organisation	AI Application and Social Economy Relevance
		described situation, operating as a first-response tool that connects survivors to human support services rather than replacing them. ***The Vodacom/Bright Sky case illustrates both the potential and the governance imperative of AI in GBV response. The tool is designed as a bridge to human services, not a substitute for them - and its adaptation for South Africa included deliberate work on local language access, local service directories and trauma-informed interface design. This localisation work is the part that is rarely funded and most often skipped - and most consequential for whether the tool actually works for the communities it is intended to serve.***
EDUCATION AND SKILLS DEVELOPMENT		
Educational Intermediary	**Sci-Bono Discovery Centre**	Sci-Bono is Gauteng's largest science and technology centre, functioning as an intermediary between the Department of Basic Education and public schools, particularly those in under-resourced communities facing severe teacher shortages in mathematics and science. Sci-Bono is piloting AI-driven adaptive learning systems and automated assessment tools designed for Grade 12 learners in under-resourced public schools. These systems dynamically adjust the difficulty and format of mathematics and science exercises based on a learner's real-time responses - functioning as an automated tutor that accommodates different learning speeds and compensates, in part, for the absence of qualified educators. It is important to characterise these as active pilots rather than established programmes: Sci-Bono is testing and refining the tools, not deploying them at full scale. This is appropriate institutional behaviour, the complexity of adaptive learning in under-resourced, multilingual, high-pressure exam environments demands careful piloting before scale. ***Sci-Bono's intermediary position - bridging government, schools and technology - is precisely where AI adoption needs institutional champions. Intermediaries who understand both the pedagogical requirements and the operational constraints of under-resourced schools are essential to responsible AI deployment in basic education.***
Foundation — Systemic Education	**DG Murray Trust (DGMT)**	The DG Murray Trust is a South African foundation focused on systemic change in early childhood development, literacy and youth opportunity. Through initiatives including the Accelerated Learning Project, DGMT applies predictive data models to identify educational deficits across rural and township schools with sufficient granularity to guide targeted intervention. Working with the Teaching at the Right Level (TaRL) framework - an evidence-based pedagogical approach that groups learners by current ability rather than age or grade - DGMT uses data analytics to identify where literacy and numeracy gaps are most severe, directing educator

Actor Type	Organisation	AI Application and Social Economy Relevance
		development resources to schools and districts exhibiting the sharpest performance declines. ***DGMT's data-driven approach to systemic education is significant not because of the sophistication of its AI tools, but because of the discipline with which it connects data to decisions. The lesson for the sector is less about technology than about the institutional commitment to evidence-based allocation - a commitment that AI can support but cannot substitute for.***
Corporate CSI — AI Skills	**MTN SA Foundation**	The MTN SA Foundation targets structural youth unemployment through its CSI programmes, with a deliberate focus on building AI and digital competencies among young South Africans in under-resourced communities. Rather than generic IT funding, the Foundation invests in dedicated digital innovation hubs and AI incubation spaces within selected public schools and TVET colleges in underserved areas. Programmes provide machine learning curricula, affordable devices and career guidance tools - systems that analyse learners' aptitudes and suggest relevant digital career pathways aligned to labour market demand. The Foundation's approach is explicitly supply-side: addressing the structural skills gap that prevents township youth from accessing the formal technology economy. ***The MTN Foundation's alignment with South Africa's G20 priorities around digital inclusion is intentional. For the social economy, the significance is in the model: corporate CSI is positioned not as charity but as investment in the talent pipeline that the technology sector - including the social AI ecosystem - will depend on.***
AGRICULTURE, FOOD SECURITY AND ENVIRONMENT		
AgriTech Social Enterprise	**Aerobotics**	Aerobotics is a South African precision agriculture company using AI and drone-based imagery to support crop management for fruit and wine growers. Its Aeroview platform uses computer vision applied to high-resolution aerial imagery to detect tree stress, disease, pest damage and fruit development - giving growers granular, field-level intelligence that was previously available only through expensive manual scouting. For the social economy, Aerobotics' relevance extends beyond its primary commercial market. Accurate yield forecasting and early disease detection reduce food waste in export supply chains - with direct downstream effects on food availability and pricing. Smallholder cooperative programmes and development finance institutions are beginning to explore how precision agriculture tools of this kind can be adapted for smallholder contexts with lower infrastructure requirements. ***Aerobotics is primarily a commercial AgriTech company. Its social economy relevance is indirect but real: reducing***

Actor Type	Organisation	AI Application and Social Economy Relevance
		post-harvest losses and improving supply chain predictability have direct food security implications, and the technology models it has developed may inform future smallholder-oriented adaptations.
African Case Study (Kenya / East Africa)	**ITIKI**	ITIKI is an East African AI project that integrates indigenous weather knowledge with satellite climate data to generate highly localised agricultural forecasts for smallholder farmers. The approach feeds AI models with both formal meteorological data and traditional ecological indicators - bird behaviour, plant flowering patterns, insect activity - that have proven predictive value within specific ecosystems but are typically excluded from formal forecasting models. The result is a forecast calibrated to local conditions, delivered in accessible formats (SMS, audio) directly to farmers, enabling better planting timing, reduced crop failure risk and improved food security at the household level. ***ITIKI is included as an East African example, not a South African one, because it represents a design philosophy with direct relevance to South Africa's smallholder and rural agriculture context: the integration of indigenous knowledge into AI systems, rather than the displacement of it. South African agriculture AI has largely been developed for commercial-scale farming; the ITIKI model suggests a different path for rural food security applications.***
Corporate CSI — Food Security	**Shoprite**	Shoprite, South Africa's largest food retailer, deploys advanced AI-powered replenishment and demand forecasting systems across its national store network - commercial operational technology that has been repurposed to serve a significant social function. The system predicts, with considerable accuracy, when and where surplus fresh food will accumulate across its stores before it reaches its sell-by date. Rather than discarding this surplus, an automated workflow redirects it to pre-verified community food hubs, soup kitchens and NPO food programmes. The AI does not merely identify surplus, it manages the logistics of distribution, matching predicted volumes with registered community partners and generating the documentation required for tax-efficient donation. ***Shoprite's case is one of the strongest examples of commercial AI creating social value as a by-product of operational efficiency. The governance insight is important: the system works because Shoprite invested in verifying its community partners before connecting them to the automated system - the data governance behind the scenes is what makes the social impact sustainable and scalable.***
Corporate CSI — Environment	**Nedbank**	Nedbank's corporate sustainability framework directs CSI investment toward climate action and environmental resilience. Through CSI-supported pilots, Nedbank has

Actor Type	Organisation	AI Application and Social Economy Relevance
		funded predictive analytics work modelling water and energy demand patterns across vulnerable municipalities - combining utility data, weather forecasts and socioeconomic variables to anticipate resource stress before it becomes a crisis. In the South African context, where load-shedding has caused significant economic and social harm over an extended period, predictive modelling of energy demand and grid resilience has direct community impact. The data produced by these models is intended to support community infrastructure planning and clean energy distribution optimisation in under-served municipalities. ***It is important to note that Nedbank's environmental AI work is described as a CSI-supported pilot programme rather than an established, verified system. The framing here is appropriately measured: the potential is significant, but the evidence base for impact at scale is still being developed.***
Green Economy — Intermediary	**GreenCape**	GreenCape is a sector development agency driving the green economy in the Western Cape, operating at the intersection of business, government and academic institutions to remove barriers to the adoption of clean technology. In the AI context, GreenCape functions as an ecosystem builder: identifying and amplifying early-stage AI and IoT innovations with environmental social impact through mechanisms including its annual GreenPitch Challenge, hosted in partnership with the Friedrich Naumann Foundation. The GreenPitch process has supported companies like Dragonfly - a Cape Town SMME developing AI and IoT solutions for waste management - by connecting them with investors and removing regulatory obstacles to scaling. GreenCape's Market Intelligence Reports provide freely available analysis of high-growth green technology sectors, directing both social investors and commercial investors toward areas where AI applications can address resource management and environmental resilience challenges. ***GreenCape's contribution to the AI social economy is catalytic rather than direct: it builds the conditions - market intelligence, investor connections, regulatory navigation, peer networks - that allow AI-enabled environmental social enterprises to scale. This intermediary function is underrepresented in most accounts of AI in the social sector, but it is structurally essential.***
DATA SYSTEMS, M&E, GOVERNMENT AND INFRASTRUCTURE		
Intermediary — M&E and Grant Management	**Tshikululu Social Investments**	Tshikululu Social Investments is one of South Africa's largest social investment intermediaries, administering multi-billion-rand portfolios across multiple corporate social investment trusts. Managing grant portfolios at this scale requires robust data verification systems - and Tshikululu has integrated AI and digitisation into its M&E

Actor Type	Organisation	AI Application and Social Economy Relevance
		architecture to address the chronic problem of data fragmentation and reporting inconsistency across its grantee base. The AI-assisted system automates the real-time collection, verification and analysis of community-reported impact data submitted by NPO grantees. It validates that metrics are accurate, consistent and compliant before subsequent funding tranches are released - replacing the manual checking processes that previously consumed significant staff capacity and still allowed errors to pass undetected. ***Tshikululu's case is significant for two reasons. First, it demonstrates that AI-enabled data governance can improve both the efficiency and the integrity of intermediary grant management, creating a more reliable evidence base for both funders and implementing organisations. Second, it illustrates the leverage potential of AI in the intermediary layer: when an intermediary with portfolio-wide reach improves its data systems, the quality improvement flows through to every organisation in its grantee network.***
Community Spatial Mapping	**Community Asset and Poverty Mapping Initiative**	South Africa's spatial geography remains deeply shaped by apartheid-era planning. Marginalised townships and informal settlements are geographically isolated, under-mapped and poorly served by the spatial planning data that drives public and social investment decisions. The Community Asset and Poverty Mapping (CAPM) initiative directly addresses this by using AI-powered computer vision to generate high-resolution, continuously updated spatial data about under-mapped communities. The technical approach combines multiple data sources and methods: • Convolutional Neural Network models analyse satellite imagery to identify and classify land use patterns - distinguishing townships, informal settlements, suburbs, industrial areas and villages with high accuracy across large geographic areas. • Night-time light emission data is used as a proxy for economic activity and relative wealth, enabling dynamic mapping of development change over time. • The Asset-Based Community Development framework is embedded in the data model, mapping community strengths (health clinics, communal farming spaces, skill clusters) alongside poverty indicators, countering the deficit-only framing that shapes most poverty mapping exercises. By comparing satellite imagery over time, the models track how communities grow, densify and change - providing policymakers and social investors with the kind of dynamic, longitudinal spatial intelligence that was previously available only for wealthy, well-documented urban areas. ***This initiative is significant not only for its technical sophistication but for its explicit politics: it names spatial apartheid as the problem it is designed to address, and it builds***

Actor Type	Organisation	AI Application and Social Economy Relevance
		community asset mapping, not just deficit mapping, into its methodology from the outset. This is what responsible social AI design looks like.
Government — Social Security	**South African Social Security Agency (SASSA)**	SASSA distributes social grants to over 26.5 million South Africans - the largest social grant system in sub-Saharan Africa. The integrity of this system depends on accurate beneficiary verification, fraud prevention and the prevention of payments to deceased individuals. AI and biometric technology are deployed across several functions: • **Biometric Beneficiary Enrolment:** Compulsory fingerprint scanning and facial recognition at physical offices, linked in real time to the Department of Home Affairs database, verify identity and confirm proof of life for existing beneficiaries. • **Automated Means Testing:** Algorithms cross-reference SASSA registries with credit bureau, banking and government databases to evaluate applicants' financial circumstances and determine eligibility, replacing labour-intensive manual means testing. • **Digital Life Certification (in development):** Digital 'e-Life' certification, using facial recognition and liveness detection through self-service kiosks and mobile applications to verify that recipients are alive without requiring a physical office visit - is currently in piloting and rollout phases, not yet universally deployed. ***SASSA's AI applications occupy a different position in the social economy to most other cases in this table: this is government AI at population scale, with direct bearing on the income security of the country's most vulnerable households. Its governance implications are therefore among the most significant: any error, bias or system failure affects millions of people with no alternative income. The balance between fraud prevention and beneficiary access is a live and contested governance challenge - AI improves detection, but must not create barriers for legitimate recipients. Appropriate human review pathways and accessible appeals mechanisms are non-negotiable in this context.***
Infrastructure — Rural Connectivity	**UCT ICT4D Research Group and Zenzeleni Networks**	The University of Cape Town's ICT for Development research group, in partnership with the Zenzeleni Networks community network project, has pioneered a model of community-owned wireless infrastructure that directly addresses the connectivity barrier to AI adoption in rural South Africa. Zenzeleni, meaning 'do it yourself' in isiXhosa, enables rural communities to own and operate their own last-mile wireless networks, providing affordable internet access to areas that commercial telecommunications operators have no commercial incentive to serve. The model has been replicated in over a dozen communities across South Africa

Actor Type	Organisation	AI Application and Social Economy Relevance
		and studied internationally as a template for community-owned digital infrastructure. For the AI social economy, the significance is structural and foundational: AI tools require connectivity. A community health worker without internet access cannot use a diagnostic support tool. A smallholder farmer without a data connection cannot receive AI-generated crop advisory alerts. A CBO without reliable internet cannot use cloud-based M&E or reporting tools. Community network models like Zenzeleni are not a peripheral detail, they are the infrastructure layer that makes AI adoption by rural CBOs, community health workers and smallholder farmers physically possible. ***This is the most important infrastructure case in the table, and the one most often absent from AI in development discussions. It is included here as a deliberate reminder: before the question 'which AI tool should we use?' can be answered, the question 'does the community we serve have the connectivity to use any AI tool?' must be answered first.***
Social Enterprise — Economic Mapping	**Timbuktu**	Timbuktu is a South African social enterprise developing gamified mobile applications that engage township residents in mapping the economic activity, assets and needs of their own neighbourhoods. Residents use the app to log local businesses, services, skills and community needs, generating a community-sourced dataset of the informal and township economy. AI aggregates and analyses this community-generated data to produce dynamic digital records of economic ecosystems that formal market research and government datasets routinely miss. South Africa's informal economy is substantial - estimates suggest it accounts for a significant share of economic activity in township areas - but it is largely invisible to the data systems that drive social investment and policy decisions. ***Timbuktu's application is still relatively early-stage and the specific figures attributed to the 'hidden economy' should be treated as estimates rather than verified data. What is significant is the design philosophy: community-generated data, gamified participation, and AI synthesis to produce intelligence that no external researcher or government survey could replicate. This is community data sovereignty in practice.***

Table 4: AI Application Case Studies Across the South African Social Economy

What the Evidence Tells Us: Cross-Cutting Insights from South African AI Practice

The case studies above represent a comprehensive compilation of AI applications in South Africa's social economy. Taken individually, each tells a story about a specific organisation, a specific problem and a specific technology response. Taken together, they reveal something more important: the conditions under which AI actually creates value in the South African social sector, and the patterns that explain why some applications gain traction where others stall.

Eight insights emerge from the evidence.

INSIGHT 1

The most effective applications solve a concrete problem, not an abstract strategic one

The cases that demonstrate clearest impact share a defining characteristic: they address a specific, recurring friction point in the lives of practitioners or community members. MomConnect delivers health guidance to expectant mothers at the moment they need it. GROW ECD removes the administrative burden that prevents township crèches from accessing government subsidies. Harambee matches job-seekers to opportunities at the point of application. Shoprite's AI redirects surplus food to soup kitchens automatically - solving a logistical problem that was previously solved manually or not at all.

In each case, AI is addressing a problem that a practitioner or beneficiary encounters every day, not a problem that a strategist identifies from a distance. This mirrors the logic that made mobile money successful: M-Pesa worked because it solved a concrete daily problem using infrastructure people already had, without requiring new behaviours or new systems. The AI applications gaining durable traction in South Africa share this design discipline.

The implication for programme designers is direct: before asking 'where can AI add value?', ask 'what does this practitioner, health worker or community member experience as their most persistent, most frustrating daily problem?' Start there.

INSIGHT 2

Language is the first infrastructure barrier — and organisations that solve it first unlock everything else

The concentration of language-focused organisations in these case studies - Lelapa AI, Masakhane, Vulavula, ECD Connect, Zuzi - is not coincidental. It reflects a structural reality: **AI tools designed for English-speaking users are not merely less convenient for isiZulu, isiXhosa or Sesotho speakers. They are effectively inaccessible.** A health advisory chatbot that cannot communicate in the language of the user it is meant to serve is not an imperfect tool - it is a non-tool for that user.

Jade Abbott's framing – “language is the enabler” - is borne out consistently across the evidence base. Every application that achieves meaningful community-level reach in South Africa has

answered the language question before most other design questions. The organisations building the most durable community-facing AI applications have treated language not as a localisation step at the end of product development, but as the primary design constraint from the outset.

For funders and programme designers, this has a direct procurement implication: preference South African AI tools with African language capability - Vulavula, Masakhane's open models, InkubaLM - over English-dominant alternatives, particularly for any community-facing or practitioner-facing application.

INSIGHT 3

The most durable applications are embedded in existing community infrastructure — not built alongside it

The applications that have achieved meaningful scale share a structural feature: they did not build new community infrastructure to deliver their AI applications. They embedded AI into the infrastructure communities already trusted and already used. MomConnect runs through mobile numbers South African women already have. Harambee reaches youth through WhatsApp flows they already navigate. Shoprite's food redistribution runs through a supply chain and community partner network already in operation. ECD Connect and Zuzi deliver through WhatsApp - no app download, no account creation, no new platform to learn.

This is not merely a design preference - it is a survival requirement. Social economy AI applications that require communities to adopt new platforms, new devices or new digital behaviours face an adoption barrier that most will not clear. The design question is not 'what is the best possible technology for this use case?' but 'what is the best technology for this use case that can be delivered through systems this community already uses and trusts?'

INSIGHT 4

The infrastructure layer is a precondition, not a parallel track

The UCT/Zenzeleni case stands apart from every other case in this compilation because it addresses a dependency that the other cases tend to assume rather than examine: **AI tools require connectivity. Without reliable internet access, the most carefully designed AI tool is inaccessible to the communities that need it most.**

In rural South Africa, connectivity is not a given. Load-shedding disrupts even where connectivity theoretically exists. Device access is uneven. Data costs remain high relative to household income. The SASSA case adds a further dimension: even government AI systems at scale face connectivity-related equity challenges, as biometric verification systems require physical office visits that are burdensome for elderly, disabled or geographically remote beneficiaries.

The practical implication: the social economy's AI strategy must include deliberate investment in the infrastructure layer - connectivity (Zenzeleni and similar community network models), devices (MTN Foundation's hardware provision model), power stability and data affordability. These are

not preliminary steps before the 'real' AI work begins. For a significant portion of the social sector's intended beneficiaries, they are the AI work.

INSIGHT 5

Corporate and community AI carry different DNA — and both are necessary

The case study evidence makes visible a structural complementarity that is rarely explicitly designed for: corporate AI and community AI have fundamentally different strengths, and the highest-impact ecosystem configurations are those where both are present and connected.

Corporate AI: Shoprite's supply chain intelligence, Discovery's behavioural analytics, Nedbank's environmental modelling, MTN's career guidance systems, brings operational scale, commercial data infrastructure, financial sustainability and the ability to deploy at national scale from day one. These are real, non-trivial capabilities that community organisations cannot replicate.

Community and NPO AI: Khulisa's qualitative data analysis, Zuzi's contextualised learning content, GROW ECD's crèche administrative tools, Harambee's safeguarding sentiment analysis, brings community embeddedness, lived-experience proximity, cultural competency and the trust that makes community engagement possible. These are also real, non-trivial capabilities that commercial AI systems struggle to replicate.

The cases that approach genuine systemic change are those where these capabilities connect: Tshikululu's data governance creating a shared evidence standard across corporate and NPO grantees; DGMT's data analytics shaping both its own programming and the programmes of its implementation partners; the Harambee/SAYouth.mobi ecosystem connecting corporate employer data with community youth intelligence. The design imperative is to create the conditions for these connections - not to position corporate and community AI as alternatives.

INSIGHT 6

Governance is the variable that separates pilots from programmes

A striking pattern in the South African evidence base is **the correlation between explicit data governance investment and sustainable AI application.** Organisations that have built governance frameworks into their AI work: Tshikululu's data verification architecture, Lelapa AI's Esethu Framework for community data compensation, Harambee's explicit 'AI recommends, humans decide' principle, SASSA's biometric verification audit trails - have built applications that can be scrutinised, defended and sustained.

Organisations that have deployed AI opportunistically, without data governance, community consent frameworks or accountability mechanisms, tend not to appear in verifiable evidence bases: not because they do not exist, but because ungoverned AI applications rarely generate the kind of attributable, defensible impact that makes them documentable or replicable. **The absence of ungoverned cases in this compilation is itself a finding: the organisations whose AI work can be described with specificity and evidence are, almost without exception, organisations that invested in governance as a design requirement, not as an afterthought.**

For the social economy, this has a direct implication: governance is not bureaucracy. It is the mechanism that transforms an interesting pilot into a credible programme.

INSIGHT 7

The African AI ecosystem is being built from the grassroots up — and this matters

The Deep Learning Indaba, Masakhane and Lelapa AI are not peripheral figures in South Africa's AI story. They are its intellectual and organisational foundation. The fact that multiple community AI tools, language models and social enterprise applications across this case study compilation trace their origins to these grassroots research communities is evidence of something structurally important: **the African AI ecosystem is being built by practitioners motivated by the specific challenges of African communities, not by technology companies optimising for global scale.**

This matters for the social economy for two reasons. First, African-built AI is more likely to be trained on African data, to function in African languages, to embed African governance principles and to remain responsive to African community needs. Second, an ecosystem built on grassroots research community infrastructure - relationships, open-source norms, knowledge sharing, talent retention - is more robust and more equitable than one built on proprietary commercial platforms.

The investment implication is direct: funding the Deep Learning Indaba, Masakhane and organisations of their kind is not peripheral social good. It is an investment in the foundational infrastructure of South Africa's social AI ecosystem.

INSIGHT 8

AI's most transformative social value is not efficiency — it is equity

Across the case study evidence, the applications that generate the most significant social value are not those that make existing systems marginally more efficient. **They are those that give previously excluded actors access to capabilities they could not otherwise afford or access.**

GROW ECD gives women running township crèches the compliance infrastructure to unlock government subsidies previously inaccessible to them. ECD Connect gives isolated home-based childcare providers the kind of professional guidance that has historically been available only to formally trained educators in well-resourced centres. Lelapa AI gives NPOs, clinics and community businesses the ability to build digital tools in the languages their communities actually speak. Tshikululu's data governance gives the smallest grantees in its portfolio the data infrastructure to demonstrate impact in ways they could not sustain independently.

In each case, the value of AI is not that it makes a capable organisation marginally more capable. It is that it makes a structurally disadvantaged actor capable of things the existing system prevented them from doing. **This is AI as an equity infrastructure - and it is the framing that should guide social economy AI investment decisions.**

A note on this evidence base

The case studies above represent a snapshot of AI application in South Africa's social economy as documented in mid-2026. They have been selected for materiality, verifiability and social economy relevance. Where AI applications are in pilot or experimental stages, this is noted. Where descriptions rely on organisational self-reporting or secondary sources rather than independent evaluation, readers are encouraged to seek updated information directly from the organisations concerned. The social economy's AI landscape is evolving rapidly - the specific tools, platforms and programme configurations described here will change; the design principles they illustrate will not.

SECTION SEVEN

AI Across Development Portfolios: Thematic Applications

Who this section is for: Funders, programme designers and thematic specialists who think in development portfolios. This section translates AI's potential into the specific language of your portfolio - education, health, food security, gender, livelihoods, climate and more. Use it to identify the most relevant entry points for your specific focus area.

Social economy funders, government departments, intermediaries and implementing organisations think and work in development portfolios: education, health, food security, livelihoods, gender, climate and more. **Understanding AI's potential through a portfolio lens is essential for making the opportunity concrete and actionable for sector practitioners.** The following table maps AI-enabled opportunities, real African case studies and South African relevance to each major development portfolio.

Development Portfolio	AI-Enabled Opportunities	African Case Study / Example	South African Context & Relevance
Education & Skills Development	**Adaptive learning AI,** tutors, attendance tracking, teacher support, TVET optimisation	Kenya: personalised learning pathways reducing dropout (McKinsey 2024); Eneza Education SMS tutoring across East Africa	High dropout rates; NEET crisis; **AI can support township schools, ECD centres and TVET colleges - especially for under-resourced educators**
Health & Wellbeing	**Diagnostic AI**, TB/HIV screening, maternal health chatbots, telemedicine, community health worker support	minoHealth AI Labs (Ghana): 14 chest condition diagnostics; Intron Health (Nigeria): NLP for clinical settings; SA: radiology and TB screening pilots	Deep rural-urban health inequality; overburdened public health system; **AI to support CHWs in under-served communities**
Food Security & Agriculture	**Crop advisory AI**, food insecurity forecasting, smallholder support, nutrition programme optimisation	WFP Nigeria: reservoir computing for sub-national food insecurity forecasting; Crop2Cash AI agricultural advisory	1 in 5 South Africans food insecure; **AI for community food gardens, school nutrition programmes and early warning systems**
Livelihoods & Employment	**Skills matching platforms, job coaching AI**, informal sector tools, SMME support	iCog Labs Ethiopia: Amharic-language AI for employment services; WhatsApp farming advisory in Senegal	NEET crisis: South African youth (15–34); **AI-enabled SETA, NPO and DSD skilling programmes for unemployed youth**
Gender Equity & GBV Response	**Safe digital reporting tools**, legal aid chatbots, GBV pattern analysis, survivor support systems	ImpactHER 2024: 86% of African women lack basic AI proficiency - the gender digital divide is acute and must be designed against	GBV epidemic in SA; **AI-enabled safe reporting, legal navigation, and survivor case management tools - with trauma-informed design**

Development Portfolio	AI-Enabled Opportunities	African Case Study / Example	South African Context & Relevance
Climate & Environmental Resilience	**Climate risk modelling**, disaster early warning systems, green economy planning, community resilience tools	Digital Earth Africa; Gates Foundation African Next Voices dataset; AI-enabled climate advisory services	Coastal, agricultural and urban climate risk; **AI for community climate resilience planning in informal settlements**
Governance & Civic Participation	**Public sentiment analysis**, misinformation monitoring, participatory budgeting tools, service delivery tracking	Ghana, Rwanda, Tunisia: national AI strategies with civic dimensions; Ushahidi platform for crisis mapping and civic intelligence	GNU coalition environment; **AI for community accountability, participatory governance and local government monitoring**
Child & Youth Development	Learning tools, child protection pattern detection, **mentorship AI**, ECD programme support	UNESCO AI for Africa Initiative: Education AI across 18 African languages; child-safe AI design principles	Children in care and child protection system; **AI to support ECD centres, youth-serving NPOs and mentorship programmes**
Social & Community Development	**Community needs diagnostics**, volunteer coordination, place-based planning, social cohesion tools	Deep Learning Indaba grassroots AI communities building local capacity; community-led AI design initiatives	Township and informal settlement work; **AI for street-level service mapping, community safety and social cohesion**
Financial Inclusion	**AI credit scoring** for marginalised communities, micro-insurance design, SMME support tools	M-Pesa + AI analytics; mobile money AI tools across East and West Africa; DSN Nigeria training AI talent	Section 18A, SMME-supporting NPOs; **AI to unlock financial access and economic participation for excluded communities**

Table 5: AI Applications Across Development Portfolios: African Case Studies and South African Context

Priority Themes for South Africa

Within the South African context, several portfolio areas deserve particular emphasis given the country's specific social and political economy.

01
EDUCATION & SKILLS

The education and skills development portfolio carries enormous strategic importance given South Africa's NEET crisis: approximately 45.6% (9.6 million) of South African youth are not in employment, education or training (aged 15–24: 3.9 million individuals, 37.6% of this age group; aged 25–34: 5.7 million individuals, 53.3% of this age group). AI-enabled adaptive learning, teacher support tools and TVET programme optimisation could make a meaningful difference in the country's most under-resourced schools and training institutions, provided connectivity and device access challenges are addressed.

02
GENDER EQUITY & GBV RESPONSE

The gender equity and GBV response portfolio is both urgent and complex. South Africa has one of the highest rates of gender-based violence in the world. AI-enabled safe reporting tools, legal navigation assistants and survivor case management systems carry genuine life-saving potential, but must be designed with extreme care around trauma-informed principles, digital safety and the protection of survivor identity and data.

03
LIVELIHOODS & EMPLOYMENT

The livelihoods and employment portfolio is perhaps the most contested terrain for AI in the South African context, given legitimate concerns about AI-driven job displacement in an economy already characterised by 32–35% unemployment. Social economy organisations working in this space need to engage deliberately with the dual reality: AI may displace certain forms of work while creating opportunities in others, and the sector's role is to ensure that communities — particularly youth — are equipped for both dimensions of this transition.

SECTION EIGHT

AI and Impact Measurement: Strengthening the Standards Infrastructure

Who this section is for: MEL practitioners, impact investors, intermediaries and funders who work with or require structured impact measurement. This section maps AI onto the standards infrastructure you already use - IRIS+, SROI, GRI, SDGs, the IMP's Five Dimensions. AI does not replace these frameworks. It makes them more powerful and more accessible.

Any credible discussion of AI in the social economy must engage with the sector's existing impact measurement and reporting infrastructure - not to replace it, but to **understand how AI can deepen and accelerate its application.** The social economy already has a sophisticated (even if unevenly adopted) set of measurement frameworks, investment standards and reporting requirements. ***AI does not make these redundant; it makes them significantly more powerful.***

Framework	What It Measures	AI Integration Opportunity	SA Relevance
IRIS+ (GIIN)	Standardised impact metrics for investors; 5 Dimensions of Impact (What, Who, How Much, Contribution, Risk)	AI to automate indicator tracking, cross-portfolio benchmarking, narrative generation from structured data	Growing SA impact investing market; DBSA, IDC, NEF increasingly require IRIS+ or equivalent alignment
SROI (Social Value UK)	Monetises social value created relative to investment; stakeholder-led; ratio of impact to cost	AI to accelerate stakeholder interviews, outcome mapping, financial proxy identification and model recalculation	NPOs and intermediaries can use SROI to demonstrate value to CSI managers and government funders
GRI Standards	Sustainability reporting on environmental, social and governance performance; internationally recognised	AI to generate GRI-aligned narratives from operational data; conduct gap analysis; benchmark against peers	JSE-listed companies increasingly require GRI; CSI funders and their NPO partners must align for reporting
SDG Alignment	Tracks contribution to the 17 Sustainable Development Goals and 169 sub-targets	AI to map programme activities to SDG targets; generate automated SDG contribution dashboards	Government and bilateral donors require SDG alignment; AI significantly reduces the manual mapping burden
Impact Management Project (5D) (Impact Frontiers)	Structured evidence across What, Who, How Much, Contribution, and Risk dimensions	AI to synthesise evidence across all 5 dimensions from mixed quantitative and qualitative data sources	Adopted by leading SA foundations and impact investors; increasingly referenced in grantee reporting requirements
b.world (AI-Native IMM)	AI-powered impact design, data collection, analysis and	Native AI integration: build logic models, collect data, analyse results,	Emerging platform; intermediaries could deploy

Framework	What It Measures	AI Integration Opportunity	SA Relevance
	storytelling platform built for social organisations	communicate outcomes - end to end	as shared MEL infrastructure for grantee portfolios
60 Decibels / Lean Data	Rapid beneficiary feedback system; captures community voice at scale through standardised mobile surveys	AI to analyse open-ended survey responses; identify patterns in beneficiary feedback across large portfolios	Currently underused in SA social sector; AI makes community voice collection and synthesis significantly more feasible

Table 6: AI Integration Across Impact Measurement Frameworks and Standards

Several dimensions of this integration deserve particular attention

01 The IRIS+ and Impact Investing Intersection

The Global Impact Investing Network's IRIS+ system provides a comprehensive set of standardised metrics across five dimensions of impact - What, Who, How Much, Contribution and Risk. As South Africa's impact investing market matures - with development finance institutions like the DBSA, IDC and NEF increasingly requiring structured impact reporting - the ability to automate IRIS+ indicator tracking, cross-portfolio benchmarking and narrative generation from structured data becomes increasingly valuable. ***AI makes this not only more efficient but more comprehensive: drawing insights from qualitative programme data that traditional reporting systems would miss.***

02 Social Return on Investment and the Case for Community Value

The SROI methodology - which monetises social value relative to investment - offers NPOs and intermediaries one of the most powerful tools for demonstrating value to CSI managers and government funders who think in financial terms. The methodology's stakeholder-led approach, in which communities define what outcomes matter to them, is philosophically aligned with the participatory AI principles that should govern social economy technology design. ***AI can accelerate the most time-consuming elements of SROI - stakeholder interview synthesis, outcome mapping, financial proxy identification and model recalculation - making it accessible to organisations that currently cannot afford the consultancy costs of a full SROI analysis.***

03 Real-Time Learning and the MEL Revolution

The emergence of AI-native impact measurement platforms such as b.world - which uses AI to support logic model design, data collection, outcome analysis and impact storytelling in a single integrated workflow - points toward a significant shift in how social organisations manage evidence. Combined with rapid beneficiary feedback tools like 60 Decibels, which capture community voice at scale through standardised mobile surveys, these platforms enable a form of real-time adaptive learning that has historically been available only to large, well-resourced organisations. ***AI has the potential to democratise this capability across the sector.***

SECTION NINE

Risks, Ethics and the Governance Imperative

Who this section is for: All readers - but especially governance leads, board members, data protection officers and anyone responsible for an organisation's ethical accountability. This is not optional reading. The risks here are real, the mitigations are practical, and the governance framework is designed to be applied, not archived.

Despite its transformative potential, AI also introduces significant risks that the social economy must actively govern. The following framework addresses eight categories of risk with structured mitigation strategies and references to relevant standards - because identifying risks without providing actionable governance responses is insufficient.

Risk Category	Description	Mitigation Strategy	Relevant Framework / Standard
Data Colonialism	AI systems trained predominantly on non-African data; extracts community value without local benefit or ownership	Prioritise African-built models (Lelapa AI, Masakhane); establish data sovereignty protocols; community data trusts	AU Continental AI Strategy (2024); AU Data Policy Framework; POPIA (SA)
Algorithmic Bias	Urban, donor-centric, and English-language bias embedded in AI outputs; reinforces existing institutional privilege	Audit AI outputs regularly for bias; establish community review panels; require multilingual tools	UNESCO Recommendation on AI Ethics (2021); NIST AI Risk Management Framework
Digital Inequality	CBOs and rural organisations structurally excluded from AI benefits due to connectivity and infrastructure gaps	Offline-capable tools; shared infrastructure models; intermediary AI support hubs for under-resourced organisations	GSMA AI for Development; IDRC AI4D Africa Programme; Bridgespan SA capacity-building research
Over-Quantification	AI privileges measurable data over relational and community-based knowledge; undermines qualitative understanding	Combine AI with participatory M&E; protect qualitative knowing; use Most Significant Change alongside AI analytics	SPAR4D; Developmental Evaluation (Patton); Most Significant Change methodology
Data Privacy & Consent	Beneficiary data exposure; absence of informed consent processes; inadequate data security in resource-constrained orgs	Adopt POPIA compliance frameworks; design consent-first data systems; develop community data governance protocols	POPIA (SA, 2020); GDPR-aligned principles; UNHCR Data Protection Guidelines
Loss of Human-Centred Practice	Risk that automation replaces facilitation, empathy, social work and community trust-building	Position AI explicitly as augmentation, not replacement; invest in human relational capability alongside AI tools	Social Work Act (SA); NPO Act compliance; professional ethics codes; Ubuntu principles

Risk Category	Description	Mitigation Strategy	Relevant Framework / Standard
Vendor Dependency	Over-reliance on foreign AI platforms creates long-term financial and strategic dependency; sustainability risk	Preference for open-source tools; build shared infrastructure; invest in local African AI capacity	AU Agenda 2063; SA National AI Policy Framework (in development); POPIA data localisation provisions
Governance Gaps	Absence of AI governance policy in the majority of NPOs, funders and government departments	Develop sector-wide AI governance guidelines; create NPO AI charters; include AI in board governance training	King V (governance); SAICA professional ethics; South African Institute for Advancement - Inyathelo NPO governance standards; ISO/IEC 42001 AI Management

Table 7: AI Ethics and Governance Framework for the Social Economy: Risks, Mitigations and Standards

Data Colonialism: The Structural Risk

Perhaps the most fundamental risk is one that is rarely named in mainstream AI ethics discussions: data colonialism. Many of the AI systems being deployed in African contexts were built on datasets that poorly represent African realities, African languages, African social dynamics and African community priorities. When these systems are used to analyse African communities, make recommendations about African programmes or inform decisions about African funding allocation, they embed the biases and assumptions of the contexts in which they were built.

The answer is not to reject AI, but to insist on African-built, African-governed and African-language AI solutions wherever possible - and to support the ecosystem of organisations (Lelapa AI, Masakhane, Deep Learning Indaba and their peers) that are building this alternative. The Gates Foundation's investment in the African Next Voices dataset, providing 9,000 hours of multilingual speech data across 18 African languages, is an example of the kind of foundational infrastructure investment that makes African AI possible.

The Gender Dimension of AI Risk

The gender dimensions of AI risk in Africa are stark and demand explicit attention. An ImpactHER survey in 2024 found that 86% of women across 52 African countries lack basic AI proficiency, 60% have not had digital skills training, and 50.2% do not have adequate internet access - 37 percentage points lower than African men. In rural areas, 84.3% of women respondents reported no digital device ownership. Any AI strategy for the African social economy that does not explicitly address the gender digital divide will deepen rather than reduce gender inequality within the sector.

Over-Quantification and the Protection of Relational Knowing

The social economy depends heavily on relationships, trust, lived experience and contextual understanding. These are not peripheral concerns - they are the substrate through which social change actually happens. **AI risks privileging measurable, digital, text-based data over the relational and community-based knowledge that experienced social sector practitioners have spent careers developing.** This is not merely a philosophical concern. It has practical consequences: for how communities are listened to, for how programmes are designed, and for what counts as evidence of impact. The sector must actively protect participatory and qualitative methodologies: Most Significant Change, developmental evaluation, community-led assessment - not as alternatives to AI-enabled analysis but as complementary and sometimes primary sources of insight. The most powerful applications of AI in the social economy will combine computational intelligence with deep human wisdom.

A Community Member's Guide to AI in Programmes That Serve You

This section is for community members and community organisations, not the funders and programme managers who are the guide's primary audience. If you are a beneficiary, a parent, or a community leader whose data is being collected, these are the questions you have the right to ask and have honestly answered. Tool 3 in the Appendices covers similar ground for organisations; ask these directly.

Questions communities should ask organisations using AI with their data

THE BASICS

What data about me or my community is being collected, by whom, why, and who else will see it? How long will it be kept, and what happens when the programme ends?

Is AI being used to analyse it, or to make decisions about my eligibility, risk level or support?

If an AI system makes a recommendation about me that I think is wrong, who do I speak to and how do I challenge it? Has it been tested to make sure it doesn't disadvantage people like me?

ABOUT YOUR RIGHTS

Do I have to share this data, and what happens if I say no? Can I see, correct or delete it?

Who is responsible for protecting it, and how do I contact them if something goes wrong?

ABOUT BENEFIT

Will I or my community see the results of this data analysis? How and when?

How will the insights from this data be used to improve services for my community?

Will my community have a say in how the findings are used?

Under South Africa's Protection of Personal Information Act (POPIA), these are not aspirational questions — they are legal rights, regardless of whether the organisation is an NPO, a corporate, a government department or a research institution. An organisation that cannot answer them clearly and honestly has not yet earned the right to collect your data.

PART THREE

Readiness, Adoption and People

IN THIS PART

> *The social economy's AI journey will ultimately succeed or fail not at the level of technology but at the level of people.*

SECTION TEN

An AI Readiness Pathway for Social Economy Organisations

Who this section is for: Operational managers, IT leads, programme directors and anyone making concrete tool selection decisions. This section moves from principles to practical rulings. The seven decision gates, the tool usability matrix, and the governance checklist are designed to answer the questions practitioners ask most frequently, and most urgently, but that existing guidance rarely addresses.

A common barrier to AI adoption in the social sector is the sense that the technology is beyond reach - relevant only to well-resourced organisations with sophisticated data infrastructure and specialist staff. This is a misconception. **AI tools are increasingly accessible at every level of organisational capacity, from free consumer-grade applications to enterprise platforms, and the sector's AI journey does not need to start at the most sophisticated end.**

A pathway readiness tool

The following readiness pathway provides a practical framework for organisations to **understand their current position and identify appropriate AI entry points at each stage of development.**

Readiness Level	Characteristics	AI Entry Points	Support Needed
Level 1: **Aware**	Organisation has awareness of AI; no structured use; predominantly paper-based or basic digital systems	AI literacy training; basic prompt engineering; use of free tools (ChatGPT, Claude.ai); attend sector AI events	Digital literacy support; device and connectivity access; peer learning networks; no-cost introductory resources
Level 2: **Experimenting**	Individuals using AI tools informally; no organisational policy, governance or coordinated approach	Formalise use; develop an organisational AI policy; introduce AI into report writing, communications and basic M&E	AI governance policy template; facilitated experimentation workshops; peer networks; funder support for capacity
Level 3: **Integrating**	AI embedded in specific functions (M&E, communications, fundraising, reporting); some governance structures in place	Extend AI across the programme lifecycle; connect to impact measurement frameworks (IRIS+, SROI, GRI); staff training	MEL capacity building; data management systems; change management support; technology partner relationships
Level 4: **Transforming**	AI is core to organisational strategy and operations; ecosystem intelligence functions beginning; data governance strong	Ecosystem-level AI platforms; shared intelligence infrastructure; sector thought leadership; collaborative tools	Strategic AI partnerships; shared data standards; ecosystem governance frameworks; funder investment in infrastructure
Level 5: **Leading**	AI-native organisation; contributes to shaping sector	Develop sector-specific tools; share open-source solutions;	Policy advocacy platforms; sector convening and

Readiness Level	Characteristics	AI Entry Points	Support Needed
	AI norms; co-creates AI solutions with communities and peers	advocate for AI equity; build African AI ecosystem	knowledge sharing; African AI ecosystem investment and partnerships

Table 8: AI Readiness Pathway for Social Economy Organisations: A Five-Level Framework

The pathway is not a linear progression that all organisations must follow in sequence. Some organisations will begin at Level 2 or 3, having already experimented with AI tools informally. Others may remain appropriately at Level 1 for an extended period while building foundational digital literacy and infrastructure. The key insight is that AI adoption should be matched to organisational context, not driven by technology enthusiasm or funder pressure.

For funders and intermediaries, the readiness framework has a direct implication: supporting AI adoption in the social sector is not primarily about providing access to tools. It is about investing in the foundational conditions: digital literacy, data governance, change management capacity and peer learning networks that make effective tool adoption possible.

A self-assessment readiness tool

The following self-assessment tool is provided to help boards, management, and practitioners in social economy organisations including nonprofits, NGOs, social enterprises, funders, and intermediaries understand their readiness to adopt artificial intelligence responsibly and effectively.

It covers 11 dimensions with 60+ questions. It can be completed individually or used as a facilitated team exercise. The richest insights often come from comparing how different roles score the same dimension - the gaps between scores are frequently more revealing than the scores themselves.

How to Use This Assessment

Each question is rated on a 1–4 scale:

Score	Rating	What it means
1	Not yet	This is not addressed in our organisation. No meaningful awareness, policy, or practice exists.
2	Emerging	There is some awareness or early conversation, but no consistent practice, policy, or accountability.
3	Developing	Work is actively underway. Progress is visible but incomplete or inconsistently applied.
4	Ready	This is well established, consistently practised, documented, and accountable.

A note on the 1–4 scale: *The even-numbered scale is intentional. A 5-point scale tends to cluster responses at 3 ("we're doing okay"). Forcing a choice between 2 and 3 requires a more honest assessment of whether practice is genuinely established.*

Role Guidance

Each question indicates the role(s) it is primarily intended for. You do not need to answer questions outside your role but you may find value in reading them to understand the full picture. The assessment is designed to be completed by:

BOARD MEMBERS & TRUSTEES	Focus on strategic, governance, and future readiness dimensions.
EXECUTIVE DIRECTORS & CEOS	All dimensions are relevant.
MANAGEMENT & OPERATIONS LEADS	Focus on operational, financial, people, and data dimensions.
PRACTITIONERS & PROGRAMME STAFF	Focus on programmatic, people, and ethics dimensions.
FUNDERS & INTERMEDIARIES	All dimensions, with particular attention to ecosystem and strategic readiness.

Dimension 1: Strategic Readiness

Primary audience: Board / Leadership

To assess whether leadership understands the strategic implications of AI and whether AI adoption is aligned to mission, sustainability, and future relevance.

Question	Role(s)	1	2	3	4	N/A
Our board and senior leadership have a shared, documented understanding of what AI means for our mission and operating context.	*Board & Executive*	○	○	○	○	○
We have explicitly discussed how AI could either advance or undermine our core values and commitments.	*Board & Executive*	○	○	○	○	○
AI adoption is referenced or integrated into our current strategic plan or theory of change.	*Board & Executive*	○	○	○	○	○
Our leadership can articulate specific opportunities AI presents for our sustainability and long-term relevance.	*Board & Executive*	○	○	○	○	○
We have considered how AI adoption positions us relative to peers, funders, and the communities we serve.	*Board & Executive*	○	○	○	○	○
Our strategy explicitly addresses how we will remain accountable to our beneficiaries and stakeholders as we adopt AI.	*Board & Executive*	○	○	○	○	○
Leadership actively engages with external thinking (research, peer networks, sector guidance) on AI in the social economy.	*Board & Executive*	○	○	○	○	○

Strategic Readiness — Scoring summary

Total score (add all ratings)	/ 28	21–28: Ready | 15–20: Developing | 8–14: Emerging | <8: Not ready
Key observations / notes		

Strategic Readiness — What your score means

21–28 Ready	Strong foundation. Ensure strategy is reviewed annually and that AI alignment is tested against mission in practice.
15–20 Developing	Formalise AI in your next strategic planning cycle and assign clear accountability for delivery.
8–14 Emerging	Consider a board-level AI briefing and facilitated strategy session to move from awareness to intent.
1-7 Not ready	Critical gap. Leadership needs a structured process to develop shared understanding before any AI adoption.

Dimension 2: Governance, Risk & Ethics

Primary audience: Board / Governance

To assess whether governance structures, accountability mechanisms, and ethical frameworks are in place to guide responsible AI adoption.

Question	Role(s)	1	2	3	4	N/A
Our board has a named governance lead or committee responsible for overseeing AI-related decisions.	*Board*	○	○	○	○	○
We have a written AI policy or are actively developing one that sets ethical boundaries and use case criteria.	*Board & Executive*	○	○	○	○	○
We have a process for assessing the ethical risks of specific AI tools before adoption.	*Management*	○	○	○	○	○
Our organisation understands its obligations under relevant data protection and AI-related legislation (e.g. POPIA, GDPR).	*All roles*	○	○	○	○	○
We have mechanisms for staff, clients, and communities to raise concerns about AI use and for those concerns to reach decision-makers.	*Management*	○	○	○	○	○
We actively consider how AI tools could introduce or amplify bias against the communities we serve.	*All roles*	○	○	○	○	○
Our governance processes include external or independent review of significant AI decisions.	*Board*	○	○	○	○	○

Governance, Risk & Ethics — Scoring summary

Total score (add all ratings)	/ 28	21–28: Ready | 15–20: Developing | 8–14: Emerging | <8: Not ready
Key observations / notes		

Governance, Risk & Ethics — What your score means

21–28 Ready	Leading practice. Share your governance framework with the sector and help build collective standards.
15–20 Developing	Strengthen external review and community accountability mechanisms to close remaining gaps.
8–14 Emerging	Assign a governance lead and begin drafting an AI policy - even a lightweight one signals accountability.
1-7 Not ready	Urgent. Governance gaps expose your organisation to legal, reputational, and ethical risk. Prioritise a basic AI policy and accountability structure.

Dimension 3: Operational Readiness

Primary audience: Management

To assess whether AI can improve organisational efficiency, workflows, decision-making, and internal systems.

Question	Role(s)	1	2	3	4	N/A
We have mapped our core internal workflows and identified where inefficiencies, bottlenecks, or manual processes exist.	*Management*	○	○	○	○	○
We have piloted or are actively exploring AI tools for administrative or operational tasks (e.g. reporting, scheduling, communications).	*Management*	○	○	○	○	○
Our management team uses data and evidence to make operational decisions. AI would enhance rather than replace this practice.	*Management*	○	○	○	○	○
We have internal capacity (staff time or budget) to implement, test, and iterate on operational AI tools.	*Management*	○	○	○	○	○
Our IT infrastructure (cloud storage, device access, connectivity) is sufficient to support new digital tools.	*Management*	○	○	○	○	○
We have clear procurement and vendor management processes that would apply to AI tool selection.	*Management*	○	○	○	○	○
Staff at all levels have access to and are able to use our core digital systems reliably.	*Management & Practitioners*	○	○	○	○	○

Operational Readiness — Scoring summary

Total score (add all ratings)	/ 28	21–28: Ready | 15–20: Developing | 8–14: Emerging | <8: Not ready
Key observations / notes		

Operational Readiness — What your score means

21–28 Ready	Optimised. Focus on integration, staff feedback loops, and sharing learnings with the broader sector.
15–20 Developing	Systematise your approach. A digital road map and procurement criteria will help scale what is working.
8–14 Emerging	Pilot one operational AI tool with clear success criteria and a structured review process.
1-7 Not ready	Start here. Map your core workflows and identify 1–2 low-risk operational tasks where AI could add value.

Dimension 4: Programmatic Readiness

Primary audience: Practitioners / Management

To assess whether AI can strengthen service delivery, impact outcomes, inclusion, learning, and scalability.

Question	Role(s)	1	2	3	4	N/A
We have a clear, measurable theory of change that could guide how AI is integrated into our programmes.	*Practitioners & Management*	○	○	○	○	○
We have explored specific AI applications that could directly improve service quality or beneficiary outcomes.	*Practitioners*	○	○	○	○	○
We actively collect data on programme activities, outputs, and outcomes that could be analysed with AI.	*Practitioners & Management*	○	○	○	○	○
We have considered whether AI tools we might adopt are accessible and appropriate for the communities we serve (language, literacy, connectivity).	*Practitioners*	○	○	○	○	○
Our programmes have the flexibility and adaptive learning culture needed to incorporate new tools without disrupting service delivery.	*Practitioners*	○	○	○	○	○
We have engaged beneficiaries or community members in thinking about how AI might affect their experience of our services.	*Practitioners*	○	○	○	○	○
We have a framework for monitoring unintended consequences or harms to beneficiaries from AI-assisted programmes.	*All roles*	○	○	○	○	○

Programmatic Readiness — Scoring summary

Total score (add all ratings)	/ 28	21–28: Ready | 15–20: Developing | 8–14: Emerging | <8: Not ready
Key observations / notes		

Programmatic Readiness — What your score means

21–28 Ready	Leading. Document and share your approach. The sector needs models for responsible AI in service delivery.
15–20 Developing	Centre beneficiary voice and equity considerations more explicitly in your AI programme design.
8–14 Emerging	Identify one programme area with strong data and a willing team to carefully test AI enhancement.
1-7 Not ready	Pause before proceeding. Ensure your theory of change and data collection are solid first.

Dimension 5: Data & Technology Readiness

Primary audience: All Roles

To assess whether the organisation has the digital infrastructure, systems, data capabilities, and data governance necessary for AI adoption - including community data rights.

Question	Role(s)	1	2	3	4	N/A
Our data is structured, consistently collected, and stored in accessible, secure systems - not primarily in spreadsheets or paper records.	*Management*	○	○	○	○	○
We have a data governance policy that covers how data is collected, stored, used, and shared.	*Management*	○	○	○	○	○
We have considered the rights of the communities and individuals whose data we hold in relation to AI use.	*Management & Board*	○	○	○	○	○
Our digital systems (CRM, case management, reporting tools) are integrated enough to allow data analysis across programmes.	*Management*	○	○	○	○	○
We have processes for ensuring data quality, accuracy, and completeness.	*Management & Practitioners*	○	○	○	○	○
We understand what data AI tools would need to function and whether we can provide it legally and ethically.	*Management*	○	○	○	○	○
We have a cybersecurity policy and have assessed the security implications of introducing AI tools.	*Management*	○	○	○	○	○

Data & Technology Readiness — Scoring summary

Total score (add all ratings)	/ 28	21–28: Ready | 15–20: Developing | 8–14: Emerging | <8: Not ready
Key observations / notes		

Data & Technology Readiness — What your score means

21–28 Ready	Ready. Maintain vigilance on data security and community consent as AI tools evolve rapidly.
15–20 Developing	Address data quality and security gaps - and ensure community data rights are explicitly governed.
8–14 Emerging	Prioritise data governance policy, system integration, and staff data literacy as pre-AI work.
1-7 Not ready	Pre-AI work needed. Invest in data infrastructure and governance before adopting AI tools.

Dimension 6: People, Culture & Capability Readiness

Primary audience: Management

To assess organisational culture, workforce readiness, AI literacy, and change capability.

Question	Role(s)	1	2	3	4	N/A
Staff at all levels understand what AI is and are not fearful of its implications for their roles.	*Management*	○	○	○	○	○
We have a culture of learning and experimentation. Staff feel safe to try new tools and report what does not work.	*Management*	○	○	○	○	○
We have identified AI champions or early adopters within the team who could support wider AI adoption.	*Management*	○	○	○	○	○
We have a plan (or capacity) to provide AI literacy and upskilling across different staff levels.	*Management*	○	○	○	○	○
Our HR policies and employment frameworks are being reviewed or updated to address AI-assisted work.	*Management*	○	○	○	○	○
Leadership actively models a balanced, curious (rather than dismissive or uncritical) approach to AI.	*Board & Management*	○	○	○	○	○
We have engaged staff on how AI might change their roles and have addressed concerns transparently.	*Management*	○	○	○	○	○

People, Culture & Capability Readiness — Scoring summary

Total score (add all ratings)	/ 28	21–28: Ready | 15–20: Developing | 8–14: Emerging | <8: Not ready
Key observations / notes		

People, Culture & Capability Readiness — What your score means

21–28 Ready	Strong. Maintain this learning culture as AI tools change - and avoid complacency about ongoing literacy needs.
15–20 Developing	Identify and resource AI champions, and link upskilling to actual tool adoption plans.
8–14 Emerging	Run a structured AI literacy process and create safe spaces for concerns to surface.
1-7 Not ready	Culture first. Address fear and uncertainty through honest, open communication before introducing AI tools.

Dimension 7: Financial & Resource Readiness

Primary audience: Board / Management

To assess whether the organisation has the resources, investment appetite, and sustainability outlook required for AI adoption.

Question	Role(s)	1	2	3	4	N/A
We have a budget line (or could create one) for technology investment and AI-related costs.	*Management & Board*	○	○	○	○	○
Leadership understands that AI adoption has ongoing costs (maintenance, training, subscriptions) - not just one-off implementation costs.	*Board & Executive*	○	○	○	○	○
We have explored AI-related funding opportunities (grants, partnerships, tech philanthropy) relevant to our context.	*Management*	○	○	○	○	○
Our fundraising and income strategy accounts for how AI might change funder expectations or reporting requirements.	*Management & Board*	○	○	○	○	○
We can make the case to funders and stakeholders for responsible AI investment as part of our long-term sustainability.	*Management & Board*	○	○	○	○	○
We have assessed the cost-benefit case for specific AI tools - not just potential upside, but implementation risk and total cost.	*Management*	○	○	○	○	○

Financial & Resource Readiness — Scoring summary

Total score (add all ratings)	/ 24	19–24: Ready | 13–18: Developing | 7–12: Emerging | <7: Not ready
Key observations / notes		

Financial & Resource Readiness — What your score means

19–24 Ready	Well-positioned. Model responsible investment for the sector and advocate for funder support of AI readiness.
13–18 Developing	Develop a multi-year AI resource plan and link it to your fundraising strategy.
7–12 Emerging	Build a basic AI investment case and begin engaging funders on technology sustainability.
1-6 Not ready	Resource constrained. Explore shared infrastructure, coalitions, and tech philanthropy before individual investment.

Dimension 8: Ethics, Equity & Social Impact Readiness

Primary audience: All Roles

To assess whether the organisation applies a social justice and equity lens to AI. A critical dimension for social economy actors that goes beyond standard governance.

Question	Role(s)	1	2	3	4	N/A
We have considered how AI tools trained on biased or non-representative data could disadvantage our beneficiaries.	*All roles*	○	○	○	○	○
We actively consider whose voices and perspectives are absent from AI systems and how this affects our work.	*All roles*	○	○	○	○	○
We have an explicit commitment to not using AI in ways that could survey profile, or disempower the communities we serve.	*Board & Management*	○	○	○	○	○
Our use of AI is transparent to beneficiaries, funders, and the public. We do not obscure AI-assisted decisions.	*All roles*	○	○	○	○	○
We understand that AI can concentrate power and are taking deliberate steps to prevent this in our programmes and operations.	*Board & Management*	○	○	○	○	○
We participate in or are connected to sector-level conversations about the ethical use of AI in the social economy.	*Management & Board*	○	○	○	○	○

Ethics, Equity & Social Impact Readiness — Scoring summary

Total score (add all ratings)	/ 24	19–24: Ready | 13–18: Developing | 7–12: Emerging | <7: Not ready
Key observations / notes		

Ethics, Equity & Social Impact Readiness — What your score means

19–24 Ready	Leading. Share your equity framework openly and help define sector norms for responsible AI.
13–18 Developing	Deepen community voice and link your ethics practice to sector-wide advocacy.
7–12 Emerging	Build an equity lens into every AI decision. Engage communities and use external ethical review.
1-6 Not ready	Non-negotiable. For social economy organisations, equity must precede adoption. Do not proceed with AI until this dimension is addressed.

Dimension 9: Ecosystem & Systems Change Readiness

Primary audience: Board / Leadership

To assess whether the organisation understands AI from a systems and ecosystem perspective including sector collaboration, advocacy, and collective action.

Question	Role(s)	1	2	3	4	N/A
We understand how AI adoption by our organisation could affect the broader ecosystem of organisations and communities we operate within.	*Board & Management*	○	○	○	○	○
We are in dialogue with funders about how AI is shaping their priorities and expectations of grantees.	*Management & Board*	○	○	○	○	○
We are collaborating (or open to collaborating) with peer organisations on AI tools, learning, or shared infrastructure.	*Management*	○	○	○	○	○
We have considered our role in advocating for responsible AI policy and regulation in the social sector.	*Board*	○	○	○	○	○
We are alert to the risk that AI could commodify, deprofessionalise, or erode trust in our sector.	*Board & Management*	○	○	○	○	○
We track how technology companies and AI platforms are influencing the social economy and plan our response.	*Management & Board*	○	○	○	○	○

Ecosystem & Systems Change Readiness — Scoring summary

Total score (add all ratings)	/ 24	19–24: Ready | 13–18: Developing | 7–12: Emerging | <7: Not ready
Key observations / notes		

Ecosystem & Systems Change Readiness — What your score means

19–24 Ready	Systems-savvy. You are well-placed to lead collective AI advocacy and shape sector-level frameworks.
13–18 Developing	Deepen advocacy engagement and consider how your organisation can influence broader AI governance.
7–12 Emerging	Join existing sector AI working groups and share your learning openly.
1-6 Not ready	Isolated. Connect with sector peers and funder networks on AI, you do not need to solve this alone.

Dimension 10: Partnerships & Vendor Readiness

Primary audience: Management

To assess whether the organisation has the capacity to identify, evaluate, and manage AI vendors and technology partners responsibly.

Question	Role(s)	1	2	3	4	N/A
We have clear criteria for assessing whether an AI vendor aligns with our values, data standards, and risk appetite.	*Management*	○	○	○	○	○
We understand the contractual and data-sharing implications of using third-party AI tools.	*Management*	○	○	○	○	○
We have or are building relationships with trusted technology partners who understand the social economy context.	*Management*	○	○	○	○	○
We have a process to assess whether a vendor's AI system has been independently audited or tested for bias.	*Management*	○	○	○	○	○
We avoid vendor lock-in by ensuring our data and workflows are portable if we change tools.	*Management*	○	○	○	○	○
We engage in co-design or feedback processes with technology providers rather than only as passive consumers.	*Management*	○	○	○	○	○

Partnerships & Vendor Readiness — Scoring summary

Total score (add all ratings)	/ 24	19–24: Ready | 13–18: Developing | 7–12: Emerging | <7: Not ready
Key observations / notes		

Partnerships & Vendor Readiness — What your score means

19–24 Ready	Strong. Model responsible vendor engagement for the sector and advocate for social economy-specific tech standards.
13–18 Developing	Seek partnerships with values-aligned technology providers and invest in co-design relationships.
7–12 Emerging	Establish basic vendor assessment criteria before any AI tool procurement.
1-6 Not ready	Vulnerable. Without vendor criteria and data governance, AI tool adoption carries significant legal and ethical risk.

Dimension 11: Future Readiness & Scenario Planning

Primary audience: Board / Leadership

To assess adaptive capacity, resilience, and preparedness for AI-driven disruption including in the social economy's funding, workforce, and service landscape.

Question	Role(s)	1	2	3	4	N/A
Our leadership has engaged in structured scenario planning or futures thinking about AI's impact on our operating environment.	*Board & Executive*	○	○	○	○	○
We have considered multiple futures including ones where AI disrupts our funding model, workforce, or service relevance.	*Board*	○	○	○	○	○
We have identified early warning signals that would prompt us to accelerate, pause, or change our AI strategy.	*Board & Management*	○	○	○	○	○
Our organisation has a culture of adaptive learning. We regularly review and update strategies in response to changing context.	*All roles*	○	○	○	○	○
We have considered how AI might create new competitors or substitutes for what we currently offer.	*Board*	○	○	○	○	○
We have a horizon-scanning practice that connects AI developments to our specific sector and community context.	*Board & Management*	○	○	○	○	○

Future Readiness & Scenario Planning — Scoring summary

Total score (add all ratings)	/ 24	19–24: Ready | 13–18: Developing | 7–12: Emerging | <7: Not ready
Key observations / notes		

Future Readiness & Scenario Planning — What your score means

19–24 Ready	Adaptive. Ensure scenario insights feed back into strategy and operational plans, keep it live.
13–18 Developing	Institutionalise scenario planning in your strategic cycle and build an external intelligence network.
7–12 Emerging	Begin with a facilitated futures conversation at board level, even a half-day session is a strong start.
1-6 Not ready	Reactive risk. Without future thinking, your organisation is vulnerable to disruption. Commission a structured scenario process.

Overall Readiness Summary

Complete this page after finishing all 11 dimensions. Transfer your dimension scores to build your readiness profile.

#	Dimension	Score	Max	Readiness level
1	Strategic Readiness	/	28	Not ready / Emerging / Developing / Ready
2	Governance, Risk & Ethics	/	28	Not ready / Emerging / Developing / Ready
3	Operational Readiness	/	28	Not ready / Emerging / Developing / Ready
4	Programmatic Readiness	/	28	Not ready / Emerging / Developing / Ready
5	Data & Technology Readiness	/	28	Not ready / Emerging / Developing / Ready
6	People, Culture & Capability Readiness	/	28	Not ready / Emerging / Developing / Ready
7	Financial & Resource Readiness	/	24	Not ready / Emerging / Developing / Ready
8	Ethics, Equity & Social Impact Readiness	/	24	Not ready / Emerging / Developing / Ready
9	Ecosystem & Systems Change Readiness	/	24	Not ready / Emerging / Developing / Ready
10	Partnerships & Vendor Readiness	/	24	Not ready / Emerging / Developing / Ready
11	Future Readiness & Scenario Planning	/	24	Not ready / Emerging / Developing / Ready
	OVERALL TOTAL	/	288	

Overall Readiness Level

Not ready	Emerging	Developing	Ready
Below 72 / 288 Below 25% Foundational preconditions for AI adoption are not yet in place. Focus on digital literacy, basic data governance and AI awareness before proceeding with any AI tool deployment.	**72–143 / 288** Between 25–49% Some dimensions are developing but significant gaps remain. Targeted investment in the lowest-scoring dimensions will build readiness. Selective, low-risk AI adoption possible in areas of higher readiness.	**144–215 / 288** Between 50–74% Solid foundation with meaningful progress across most dimensions. Selective AI adoption is possible with appropriate governance. Prioritise closing remaining gaps in governance, ethics and people capability.	**216–288 / 288** Between 75–100% Strong readiness across the majority of dimensions. Proceed intentionally with ongoing learning, governance review and community accountability. Use this profile to support peer organisations and contribute to sector norms.

SECTION ELEVEN

A Practical AI Tool Selection Guide for the Social Economy

Who this section is for: Practitioners. This section addresses the operational decisions that practitioners most frequently ask about: Which tools are safe to use when sensitive beneficiary data is involved? When does data residency rule out an overseas provider altogether? When does a chatbot fit a problem, and when does a knowledge assistant or workflow automation better serve the need? What are the non-negotiable governance steps before any AI tool touches personal data? These questions remain rare in the published literature on AI in the social sector, yet they are the decisions that practitioners face daily. This section of the guide is structured as a decision framework, a set of practical rulings for the most common situations, rather than a comprehensive technology review, which would be obsolete within months of publication.

> ***A note on tool mentions and ratings.*** *Specific tools and vendors are named throughout this section because practitioners need concrete starting points. Mentions are not endorsements. The POPIA risk ratings reflect an assessment as at this publication date and may change as products evolve, terms are updated and the South African data protection landscape continues to mature. Practitioners are responsible for their own due diligence at the point of adoption.*

The Seven Key Decisions Before Any AI Tool Adoption

Before deploying any AI tool in a social economy context, seven questions must be answered. The answers determine which tools are permissible, what governance steps are required, and what ongoing accountability mechanisms must be in place.

Decision Question	When This Applies	Practical Ruling	SA-Specific Considerations
Does this use case involve identifiable beneficiary data?	Any tool that will process names, ID numbers, health info, location data, case files, financial records, or any data linkable to an individual	**YES → POPIA applies**. You must have a lawful basis for processing (consent, legitimate interest, contractual necessity). The tool's data processing agreement must be reviewed before adoption. You cannot rely on implied consent. **NO** → Standard tool risk assessment applies.	POPIA section 11 requires explicit, specific and informed consent or a recognised lawful basis. 'We told people at intake' is not sufficient - consent must be informed about AI-specific uses.
Where does the tool store and process data?	Any cloud-based AI tool (ChatGPT, Claude, Gemini, Microsoft Copilot, Google Workspace AI, etc.)	**DATA STAYS IN SA OR POPIA-COMPLIANT JURISDICTION** → Proceed with DPA review.	Most major US AI providers (OpenAI, Anthropic, Google) process data in US or EU data centres by default.

Decision Question	When This Applies	Practical Ruling	SA-Specific Considerations
		DATA LEAVES SA (e.g. processed in US, EU, Asia) → Binding cross-border agreement required under POPIA section 72. You need a written transfer mechanism — not just the provider's standard terms. **PROVIDER CANNOT CONFIRM JURISDICTION** → Do not use with beneficiary-identifiable data.	Microsoft Azure South Africa regions are available. Microsoft 365 Copilot with SA data residency is viable for organisations that need local processing. AWS Cape Town region is also available.
Does the tool train on user-submitted data?	Free tiers of consumer AI tools; tools with unclear terms; tools that mention 'improving our models' in terms of service	**TRAINS ON YOUR DATA** → Do not use with any client, beneficiary or sensitive organisational data. Use only for non-sensitive, non-identifiable tasks. **DOES NOT TRAIN ON YOUR DATA** (enterprise agreements) → Confirm in writing via DPA before processing personal data. **UNCLEAR** → Assume it does. Treat as above.	ChatGPT free tier trains on conversations by default (opt-out available). ChatGPT Enterprise, Claude for Enterprise, and Microsoft 365 Copilot have explicit no-training commitments. Always use enterprise tiers when processing any person-identifiable information.
Is the tool interacting directly with beneficiaries?	Chatbots on helplines, WhatsApp AI assistants, voice interfaces, community feedback tools, client-facing portals	**YES** → Additional obligations apply: informed consent to AI interaction (POPIA); mandatory human escalation pathway for sensitive disclosures; trauma-informed design review; multilingual access (language rights); crisis protocol integration. **NO** (staff-facing only) → Standard deployment applies.	For GBV helplines, child protection, mental health or legal aid services: NEVER deploy a fully automated chatbot without a human escalation pathway. SA social work ethics codes and the Children's Act impose duty-of-care obligations that automated systems cannot fulfil.
Which AI tool type fits this use case?	Organisation is deciding between chatbot, knowledge assistant, or workflow automation	**CHATBOT → Fits**: high-volume, structured queries from known populations (FAQ, appointment booking, intake screening). **Does not fit:** emotionally complex, crisis, diagnostic, or legally consequential conversations.	WhatsApp Business + AI (e.g. via Twilio or Vonage) works well for community outreach in SA given ~90% WhatsApp penetration. However, WhatsApp's data residency is outside SA — do not use for case-specific beneficiary data.

Decision Question	When This Applies	Practical Ruling	SA-Specific Considerations
		KNOWLEDGE ASSISTANT → **Fits**: staff-facing research synthesis, policy navigation, donor communication drafting, MEL report drafting. **Best fit** for most NPO knowledge work. **WORKFLOW AUTOMATION** → **Fits:** repetitive structured tasks (data transfer, form population, report generation from structured data, email routing). **Does not fit:** judgement-dependent or exception-heavy processes.	
Does the tool require AI to make or recommend decisions about individuals?	Beneficiary eligibility screening, risk scoring, programme matching, loan/grant recommendations, child protection flagging	**YES, AI DECIDES** → This is high-risk automated processing under POPIA section 71. Individuals have the right to human review of automated decisions that significantly affect them. You must implement: (a) human review pathway, (b) explainability (why was this decision made), (c) appeal mechanism, (d) bias audit. **YES, AI RECOMMENDS** (human decides) → Lower risk but still requires bias audit and transparency to affected individuals. **NO** (analytical only, no individual decisions) → Standard deployment.	Harambee Youth Employment Accelerator's AI-enabled matching system uses AI to recommend - not decide - candidate-employer matches, with human review at every placement decision. This is the right model: AI narrows the field, humans make the call.
What is the sensitivity classification of the data?	All data processing decisions	**SPECIAL PERSONAL INFORMATION** (health, race/ethnicity, religion, sexual orientation, criminal record, biometric, child data) → POPIA chapter 6 applies. Processing requires explicit consent or specific lawful exemption. Heightened security obligations. Do not use non-enterprise AI tools. Conduct a Privacy Impact	Health data, HIV status, GBV case information, and child records are all Special Personal Information under POPIA. Many social sector databases combine these - classify your data before choosing any AI tool.

Decision Question	When This Applies	Practical Ruling	SA-Specific Considerations
		Assessment (PIA) before deployment. **PERSONAL INFORMATION** (name, contact, location, financial) → Standard POPIA applies. DPA required. Enterprise-grade tool required. **ANONYMISED / AGGREGATE DATA** → POPIA does not apply (if properly de-identified). Consumer-grade tools acceptable for analysis.	

Table 9: Seven Decision Gates for AI Tool Adoption — Practical Rulings for Social Economy Practitioners

Tool Usability Matrix: Which AI Tools Are Safe for What?

The following matrix maps **the most commonly encountered AI platforms against their data residency, training practices, POPIA compliance posture and social economy suitability**. This is a snapshot as of mid-2026; tool terms of service, data residency options and pricing structures change frequently. Always verify current terms before deployment and record the verification in your procurement documentation.

AI Tool / Platform	Data Residency	Trains on Data?	Safe for Beneficiary-Identifiable Data?	Best Use in Social Economy	POPIA Risk Level
ChatGPT Free (OpenAI)	US servers	YES (default; can opt out)	NO: do not use with personal data	Non-sensitive drafting, research, staff learning	**HIGH**
ChatGPT Enterprise (OpenAI)	US servers (EU option available)	NO (contractual)	Only with cross-border transfer agreement under POPIA s.72	Organisational knowledge work, report drafting, MEL synthesis	**MEDIUM:** transfer mechanism required
Claude for Enterprise (Anthropic)	US / EU servers	NO (contractual)	Only with cross-border transfer agreement	Complex analytical work, policy synthesis, grant writing	**MEDIUM:** transfer mechanism required
Microsoft 365 Copilot (M365 E3/E5)	SA data residency available (Azure SA)	NO (contractual)	YES: if SA data residency configured and DPA in place	Full organisational suite: drafting, data analysis,	**LOW:** best option for SA organisations

AI Tool / Platform	Data Residency	Trains on Data?	Safe for Beneficiary-Identifiable Data?	Best Use in Social Economy	POPIA Risk Level
				meeting notes, reporting	handling personal data
Google Workspace AI (Gemini)	SA servers available	NO (for paid Workspace)	YES: with Workspace Business/Enterprise DPA and SA region	Email drafting, document summary, Sheets data analysis	**LOW-MEDIUM:** verify SA residency in admin console
WhatsApp Business + AI	Meta servers (US/EU)	NO (for Business API)	NO for case-specific data YES for general outreach	Community outreach, appointment reminders, FAQ bots	**HIGH** for case data; **LOW** for anonymous outreach
Kobo Toolbox + AI analysis	AWS (EU/US)	NO	Only with cross-border agreement	Field data collection and basic AI-assisted analysis	**MEDIUM:** use anonymised datasets for AI analysis
b.world	Cloud (EU)	NO	Only with cross-border agreement and DPA	End-to-end impact measurement: design, collect, analyse, report	**MEDIUM:** strong for anonymised programme data
Vulavula / Lelapa AI	SA-based infrastructure	NO	YES: SA-hosted, African language support, sector-aligned DPA	Community-facing voice/language tools; multilingual interfaces	**LOW:** recommended for SA social sector use
Power BI + Microsoft AI	Azure (SA residency available)	NO	YES: with M365 enterprise agreement and SA residency	Impact dashboards, funder reporting, programme analytics	**LOW:** strong choice for data-rich organisations
DHIS2 + AI plugins	Self-hosted or WHO cloud	NO (self-hosted)	YES: if self-hosted or with appropriate DPA	Health programme M&E, community health worker data	**LOW:** (self-hosted) **MEDIUM** (cloud)
Open-source LLMs (local deployment)	On-premise / local server	NO	YES: if deployed on organisation-controlled infrastructure	High-sensitivity use cases; organisations with IT capacity	**VERY LOW:** gold standard for sensitive data, high setup cost

Table 10: AI Tool Usability Matrix for the South African Social Economy: Data Residency, POPIA Risk and Best-Fit Use Cases

Important Note: This matrix reflects platform-level characteristics as of 2026. *Individual deployment configurations can change risk profiles significantly. A Microsoft 365 Copilot deployment without SA data residency configuration carries different risks than one with it. A ChatGPT Enterprise account without a signed Data Processing Agreement (DPA) carries higher risk than one with a DPA in place. Always document your specific configuration, not just the platform name, in your AI governance records.*

Chatbot vs. Knowledge Assistant vs. Workflow Automation: Choosing the Right Tool Type

One of the most common points of confusion in social economy AI adoption is choosing the wrong tool type for a given problem. Deploying a chatbot for a task that needs a knowledge assistant, or building workflow automation for a process that requires human judgement, creates both wasted investment and governance risk. The following framework maps the three most common AI tool types to the problems they are best suited for.

Tool Type	What It Is	Best Fit — Deploy When:	Poor Fit — Avoid When:	SA Social Economy Example
Chatbot (Conversational AI)	An AI system that converses with users in natural language, responding to queries, guiding processes or collecting information through dialogue	• High volume of structured, repetitive queries from a defined population • Questions have known, bounded answers (FAQ, intake screening, appointment booking) • Users are comfortable with digital self-service • Human escalation pathway is in place for sensitive situations • Problem does not require nuanced judgement or emotional attunement	• Emotionally complex or crisis interactions (GBV disclosure, suicidal ideation, child protection) • Legally consequential conversations requiring professional advice • Situations where the query set is unpredictable or open-ended • Populations with low digital literacy or no reliable connectivity • Any interaction where failure to escalate to a human carries life-safety risk	Harambee's candidate intake screening chatbot handles high-volume structured queries from job seekers about programme eligibility and application status - bounded queries, defined population, human counsellor available for complex cases

Tool Type	What It Is	Best Fit — Deploy When:	Poor Fit — Avoid When:	SA Social Economy Example
Knowledge Assistant (AI-Augmented Research and Drafting)	An AI system that helps staff find, synthesise, analyse and communicate information - typically staff-facing rather than beneficiary- facing	• Staff need to process large volumes of documents, reports or research • Drafting tasks are time-consuming and follow known formats (grant proposals, MEL reports, donor updates, policy briefs) • Knowledge is distributed across the organisation and hard to find • Staff AI literacy is sufficient to review and verify AI outputs critically	• Tasks requiring primary data collection or field observation • Decisions with legal or professional accountability (social work assessments, medical triage) • Interactions requiring emotional intelligence or relational trust-building • Situations where AI output will not receive human review before action	A programme officer at a large NPO uses Claude or Microsoft Copilot to synthesise 40 evaluation reports into a cross-programme learning brief - saving 3 days of manual work while applying their own judgement to validate and contextualise the AI's synthesis
Workflow Automation (AI-Enabled Process Automation)	AI systems that handle repetitive, structured tasks by triggering actions, routing information, populating forms or generating standard documents from structured data	• Process is highly repetitive with defined inputs and outputs • Task involves transferring data between systems or populating standard documents • Exception rate is low and exceptions are easily identified • Time and accuracy gains are significant relative to manual processing	• Process involves judgement-dependent exceptions that are frequent or consequential • Task requires contextual interpretation that varies significantly across cases • Errors in automation would cause harm (incorrect grant amounts, missed beneficiary referrals) • Staff do not have capacity to audit automation outputs regularly	Automating the generation of standard grant payment notifications from a financial system to a communication platform - structured inputs, standard output, low exception rate, easily audited

Table 11: Chatbot vs. Knowledge Assistant vs. Workflow Automation — Decision Framework for Social Economy Practitioners

When Does Data Residency Rule Out an Overseas Provider?

POPIA section 72 requires that personal information may only be transferred to a third party in a foreign country if that country has an adequate level of data protection, or if the responsible party (your organisation) has a binding transfer mechanism in place. This is not a technicality - it is a legal obligation that creates direct liability for your organisation and, indirectly, for the individuals whose data is transferred without lawful basis.

The practical ruling is straightforward:

- **If the data is personal information** (names, contact details, location, financial information) → A binding Data Processing Agreement (DPA) with the overseas provider is required before transfer. The DPA must specify what data is transferred, for what purpose, under what security standards, and with what data subject rights protections.
- **If the data is Special Personal Information** (health, race/ethnicity, religion, sexual orientation, criminal record, biometric data, child data) → A DPA alone is not sufficient. You need either: (a) the data subject's explicit consent to the specific cross-border transfer, or (b) a specific regulatory exemption. This is a higher bar, and most overseas AI providers cannot meet it for all use cases.
- **If the data is properly anonymised** (no reasonable possibility of re-identification) → POPIA does not apply. You can use any AI tool without transfer restrictions. The caveat: anonymisation is technically demanding. Replacing names with initials or removing ID numbers is not sufficient for anonymisation if the remaining data fields allow re-identification.

The practical implication: for most social sector organisations handling beneficiary data, Microsoft Azure South Africa or Google Cloud South Africa (with appropriate Workspace configuration) are the lowest-risk options because they provide SA-based data residency. For organisations that cannot yet access enterprise agreements, the safest approach is to anonymise data before using any AI tool, and to maintain a clear internal classification of which datasets may be used with which tools.

> ***Worked example:*** *A GBV counselling organisation wants to use an AI knowledge assistant to help counsellors navigate legal aid options. The AI will be staff-facing only - it will not process individual case files. Counsellors ask general questions: 'What are the procedures for obtaining a protection order in the Western Cape?' This use case does not involve personal data processing. A consumer-grade AI tool is acceptable because no personal information is being submitted. If the counsellor were to paste in a case file summary including the client's name and circumstances, the tool selection decision changes entirely - that would require an enterprise tool with a DPA and appropriate data residency.*

A Minimum Viable AI Governance Checklist for Social Economy Organisations

For organisations at Readiness Levels 1–2 that need a practical governance foundation before adopting any AI tool, the following checklist represents the minimum viable governance posture. It

is not comprehensive, a full governance framework is documented in the AI Ethics and Governance table in Section Nine of the guide - but it prevents the most common and most consequential governance failures.

#	Governance Step	Who Owns It	When Required
1	Classify all datasets the organisation holds by sensitivity: Special Personal Information / Personal Information / Anonymised / Public	Management (with legal input if available)	Before any AI tool adoption
2	For any AI tool that will access personal information: obtain the provider's Data Processing Agreement (DPA) and confirm whether it covers your jurisdiction and use case	Management	Before deployment
3	Confirm data residency: where will your data be stored and processed? For SA-based organisations handling personal data, prefer SA-hosted options (Azure SA, Google Cloud SA)	Management / IT	Before deployment
4	Determine POPIA lawful basis for each AI-enabled data processing activity (consent, legitimate interest, contractual necessity, legal obligation)	Management (with legal input)	Before deployment
5	Develop an organisational AI usage policy - even a one-page document that specifies which tools are approved, what data may be used with each, and what is prohibited	Executive / Board	Within 30 days of first AI tool adoption
6	Brief staff on the AI usage policy, including what to do if they encounter a situation the policy does not cover	Management	At policy adoption and for each new staff member
7	For any beneficiary-facing AI application: design and document the human escalation pathway - what triggers escalation, who receives the escalation, and how quickly	Programme staff / Management	Before any beneficiary-facing deployment
8	For AI tools making or recommending decisions about individuals: document the human review pathway and the appeals mechanism; brief affected individuals that AI is involved in decisions	Management / Programme staff	Before deployment; ongoing
9	Review AI tool terms of service annually - data residency options, training practices and DPA terms change frequently	Management	Annual
10	Appoint a named AI governance contact (this does not need to be a specialist - it needs to be someone with accountability for AI governance questions)	Board / Executive	Immediate

Table 12: Minimum Viable AI Governance Checklist for Social Economy Organisations — Readiness Levels 1–2

A note on proportionality: *These governance steps are not bureaucratic obstacles: they are the minimum standard of due care that responsible AI use requires. An organisation that adopts a chatbot*

without a DPA, uses a free-tier AI tool to process case files, or deploys a beneficiary-facing system without a human escalation pathway is not being agile - it is creating legal exposure for itself and harm risk for the people it serves. Governance is not the enemy of innovation in the social sector. It is the foundation that makes responsible innovation possible.

A Practical Starting Point: The First 90 Days and First Year

For organisations at Readiness Levels 1 and 2, the frameworks and matrices in this section can feel overwhelming. The governance checklist has ten steps. The decision framework has seven gates. The tool matrix has twelve platforms. Where does an organisation with limited capacity, uncertain digital infrastructure and a board that has never discussed AI actually begin?

The answer is not to tackle everything at once. It is to make a small number of high-leverage moves in a deliberate sequence. The following roadmap is designed for organisations beginning their AI journey with limited resources and no dedicated technology staff.

Phase	Timeframe	Focus	Key Actions	Who Leads
Foundation	Days 1–30	Awareness and baseline	Complete the AI Readiness Self-Assessment (Section Ten). Identify your lowest-scoring dimensions. Classify your data: what is Special Personal Information, Personal Information, anonymised, or public? Identify one AI champion in your team - someone already experimenting informally.	Executive Director with Board awareness
First Use	Days 31–60	Safe, low-risk experimentation	Select one non-sensitive, non-identifiable use case for AI - drafting a communication, summarising a public document, generating a report outline. Use a free consumer tool (ChatGPT, Claude.ai) with anonymised or non-personal content only. Document what worked and what didn't. Share with the team openly.	AI champion with management support
Governance	Days 61–90	Minimum viable governance	Draft a one-page AI usage policy using the minimum viable checklist (Table 12). Specify which tools are approved, what data may be used with each, and what is prohibited. Brief all staff. Identify whether any current AI tool use involves personal data - and if so, initiate the DPA and data residency review.	Executive Director with Board approval
Deepening	Months 4–6	Extending to programme functions	Identify one programme function where AI could add genuine value: MEL synthesis, proposal drafting, funder reporting, beneficiary feedback analysis. Design a structured pilot with clear success criteria. Apply the seven-decision framework (Table 9) before adopting any new tool.	Management with programme staff
Learning	Months 7–9	Review and share	Conduct a structured review of the pilot: what did AI improve, what did it miss, what governance gaps emerged? Document the learning. Share it with at least one peer organisation or within your intermediary	AI champion with Executive Director

			network. Revisit the readiness self-assessment - has your score changed?	
Formalising	Months 10–12	Building for the next level	Update your AI usage policy with what you have learned. Identify the next two or three AI applications to explore. Brief your board with a short AI update as a standing agenda item. Consider whether your readiness level has shifted - and what that means for your next year's AI investment.	Board and Executive Director

Three principles for this roadmap

FIRST

Start with staff-facing tools before beneficiary-facing ones

The governance complexity, consent requirements and duty-of-care obligations multiply significantly when AI interacts directly with the communities you serve. Build internal competence and governance discipline first.

SECOND

Prioritise anonymised and aggregate data before personal data

Most of the highest-value AI applications for social sector knowledge work — synthesising evaluations, drafting reports, analysing research — can be done with data that does not identify individuals. Start there. Build the governance infrastructure for personal data use carefully and deliberately.

THIRD

Treat failure as evidence, not embarrassment

The AI tools available in 2026 are powerful but imperfect. They will produce outputs that are wrong, incomplete, biased or inappropriate for your context. This is not a reason to abandon AI adoption, it is a reason to maintain human review, to document what fails as carefully as what succeeds, and to build an organisational culture where AI experimentation is a learning practice rather than a performance.

The governance checklists, decision frameworks and tool selection matrices in this section are only as valuable as the people who apply them. A minimum viable governance checklist means nothing if no-one in the organisation has been briefed on it. A POPIA compliance assessment requires someone with the literacy to conduct it. A human escalation pathway for a beneficiary-facing chatbot requires a practitioner who understands both the technology and the community it serves.This is the bridge to the section that follows. Every practical governance step in this section depends on a workforce that is prepared, capable and supported. And yet the social economy's workforce is navigating an AI transition with minimal investment, unclear career pathways, and almost no sector-specific guidance on what AI means for their roles, their careers and their professional identities. The tools are only as powerful as the people holding them. Section Twelve turns to those people.

SECTION TWELVE

The Human Dimension: People, Skills, Careers and Generational Readiness in the AI-Enabled Social Economy

Who this section is for: HR leads, executive directors, board members responsible for talent strategy, funders thinking about workforce investment, and every practitioner navigating the AI transition personally. This section is about the people who make the sector work - and about whether the sector is setting them up to navigate what is coming.

The Dimension We Mostly Overlook

Every section of this document up to this point has examined what AI can do within the social economy: how it can reshape ecosystems, strengthen measurement, improve programme design, deepen community engagement and accelerate institutional learning. There is, however, a dimension of equal importance that risks being eclipsed by the excitement of technological possibility: **the people who must actually navigate this transformation**.

The social economy is fundamentally a people's economy. Its practitioners are its most important resource - not its technology, its data or its platforms. Social workers, programme officers, community health workers, MEL specialists, fundraisers, ECD practitioners, gender-based violence counsellors, community facilitators, youth mentors and dozens of other roles represent the irreplaceable relational core of the sector. Any serious conversation about AI in the social economy must therefore also be a conversation about who these people are, what the AI transition means for them, how different generations will experience it differently, what new capabilities the sector urgently needs, and what future careers are being created that could attract the next generation of purpose-driven talent.

This section confronts a challenge that is simultaneously a crisis and an opportunity. The social economy in South Africa and across Africa already faces a significant and compounding skills deficit: ***a gap between the capabilities the sector needs and those it currently possesses.*** AI does not make this deficit smaller in the short term. If anything, it sharpens it. **But it also creates a pathway for the sector to become a more compelling, more capable and more professionally rewarding destination for precisely the talent it most needs.**

> *The social economy cannot afford to let the AI skills conversation happen only in the corporate and technology sectors. If it does, the talent gap will widen into a chasm - and the sector will find itself operating increasingly sophisticated AI tools with a workforce that was prepared for a very different world. The time to invest in people is now, not after the technology arrives.*

This section of the guide is structured around five interconnected themes: the current skills deficit and what it means in an AI context; how different generations are experiencing the AI transition

and what this means for workforce strategy; the new roles and careers the AI era is creating within the social economy; the competency framework required for AI-era practitioners; and a talent strategy for attracting, retaining and developing the people the sector will need.

The Skills Deficit: What We Have, What We Need

South Africa's social economy already faces a workforce crisis that predates AI. A PNP Staffing Group report found that 59% of NPOs reported it was significantly harder to fill staff positions in 2024 than in previous years, with 55% citing the inability to offer competitive salaries as a primary barrier. This is not merely a recruitment problem - it reflects structural underpayment, limited career pathways, high burnout and the sector's difficulty in competing for skilled professionals against the corporate and public sectors.

Against this backdrop, **the AI transition introduces a new and urgent layer of skills demand that the sector is currently ill-equipped to meet.** The following table maps the sector's existing capacity gaps across key skill dimensions and assesses the urgency of each in an AI context.

Skills / Capacity Dimension	Current State in SA Social Economy	AI Era Urgency	Consequence of Inaction
Digital and data literacy	Low to medium across most NPOs; critical gaps in CBOs and rural organisations	**High:** AI tools require baseline digital fluency to operate, interpret and govern	Deepening two-tier sector: AI-enabled large orgs vs. permanently excluded small orgs
M&E and impact measurement	Uneven; compliance-driven reporting dominates; few organisations practice adaptive MEL	**Very high:** AI-enabled MEL requires data thinking, not just data collection	Continued inability to demonstrate impact; loss of credibility with funders
Strategic and systems thinking	Strong in well-resourced intermediaries; weak in implementing NPOs and CBOs	**High:** AI tools surface patterns but humans must interpret them systemically	AI output misread or misapplied; poor strategic decisions reinforced by data
Fundraising and resource mobilisation	Chronic shortage; over-reliance on known funders; limited prospect development capability	**Medium:** AI tools can assist but cannot replace relationship intelligence	Continued funding precarity; inability to diversify income streams
Leadership and governance	Founder-dependent; weak board capability in smaller organisations; high executive turnover	**High:** AI governance requires board-level understanding of risk, ethics and oversight	Ungoverned AI use; data breaches; reputational and legal risk
Programme design and theory of change	Often donor-driven; limited community co-design; weak theory of change discipline	**High:** AI can strengthen design but only with human clarity on change logic	Poorly designed AI-enabled programmes produce poor outcomes efficiently
Community facilitation and social work	Strength of the sector; deeply relational; culturally competent	**Medium:** Must be explicitly protected as AI adoption grows	Loss of the sector's core differentiator if AI

Skills / Capacity Dimension	Current State in SA Social Economy	AI Era Urgency	Consequence of Inaction
			replaces rather than augments
Research synthesis and knowledge use	Very weak; research-practice gap is wide; academic knowledge rarely reaches implementers	**High:** AI can bridge this gap but only with human capacity to commission and interpret	Perpetual reinvention; evidence fails to accumulate across the system

Table 13: Skills Deficit Landscape: Current State and AI Era Urgency in the South African Social Economy

What the Skills Gap Actually Requires

What Table 13 reveals is that the social economy's AI-era skills challenge is not primarily a technical challenge. The sector does not need to produce data scientists or machine learning engineers in large numbers. What it needs is something more specific and more difficult: a broad elevation of digital and data literacy across the entire workforce, combined with focused investment in a small number of specialist roles - and a deliberate, sustained effort to protect and deepen the relational, facilitative and community-embedded capabilities that are the sector's most distinctive strength.

This analysis has three practical implications that are frequently missed in AI skills discussions.

The distinction between AI literacy and AI expertise matters enormously for planning. The sector cannot afford - financially or strategically - to make everyone an AI expert. Nor does it need to. What every practitioner needs is sufficient AI literacy to use tools responsibly, recognise their limitations, maintain critical evaluation of outputs, and understand the governance principles that govern their use. This is a different investment from technical AI training, and it is achievable across the workforce at a fraction of the cost. The competency framework in this section specifies what "core" AI literacy looks like at practitioner level - it is primarily about judgment, not technical skill.

The relational and facilitative capabilities in the bottom rows of Table 13 are not a comfort to practitioners worried about displacement. They are a strategic asset that requires active investment. Community facilitation, social work practice, participatory design, trauma-informed care, trust-building and cultural competency are identified as "medium" urgency in the table because AI does not directly threaten them in the near term. But "medium urgency" does not mean "safe to neglect." As AI handles more of the information processing and administrative burden that currently consumes social economy practitioners' time, the comparative value of these distinctively human capabilities will increase. An organisation that invests in AI while disinvesting in relational practice will find that it has efficient systems and ineffective programmes. The sector's investment in human relational capability needs to keep pace with its investment in AI tools - not as a defensive measure, but as a deliberate strategy for preserving the sector's deepest source of value.

The research synthesis and knowledge use gap is both the most underappreciated skills deficit and the one where AI can most quickly add disproportionate value. South Africa's social sector generates substantial evidence - through evaluations, academic research, community consultations, programme reports and sector analyses. Almost none of this evidence makes it into practice in any systematic way. The research-practice gap is wide, well-documented, and enormously costly in terms of programmes that reinvent the wheel, interventions that repeat known failures, and communities that bear the cost of uninformed decision-making. AI can synthesise, translate and surface this evidence at a scale and speed that manual processes cannot match. But it requires practitioners who can commission AI-assisted synthesis intelligently, critically evaluate what it produces, and translate it into programme decisions. These are judgment capabilities, not technical ones - and the sector is not investing in developing them.

The Competitiveness Challenge

There is a structural dimension to the skills deficit that AI sharpens into urgency: **the social economy's ability to compete for talent in a market where AI competency commands a significant premium.**

The WEF's Future of Jobs Report 2025 projects that AI and big data competencies are the single fastest-rising skill category globally, with employers across all sectors competing aggressively for professionals who can work at the intersection of data, systems and human systems. In South Africa specifically, the demand for data-literate professionals is growing in financial services, healthcare, government and the technology sector simultaneously. The social economy is competing for precisely this talent from a structural disadvantage: lower salaries, weaker technology infrastructure, limited professional development investment, and an employer brand that rarely reaches the channels where talent-rich young professionals make career decisions.

The sector's response to this challenge cannot be salary competition alone - it cannot win that contest. Its response must be to offer something the corporate and technology sectors cannot: genuine mission, community embeddedness, the complexity and human significance of social systems work, and increasingly - if the right investments are made - the opportunity to work at the frontier of applied AI in the most consequential human contexts. This is a compelling proposition. But it requires the sector to make it credible, consistent and visible - on the channels where the talent it needs is actually listening.

> *The WEF Future of Jobs Report 2025 found that 63% of employers globally identify skill gaps as the primary barrier to digital transformation - up from 60% in 2023. This challenge is even more acute in the social economy, where AI-specific roles and professional development pathways barely exist.* ***The sector is starting this race from further back, with fewer resources, and against stronger competition. The urgency of investment in people cannot be overstated.***

A Multigenerational Workforce in Transition: The AI Generation Gap

One of the most significant and underappreciated dimensions of the AI transition in the social economy is that it is not happening to a single, uniform workforce. It is happening to a

multigenerational workforce in which different cohorts carry radically different relationships to technology, different motivations, different fears and different forms of capability. Managing this transition well requires understanding these differences with nuance - not to stereotype, but to design appropriate pathways for each group.

The data is now extensive. A 2025 survey by the London School of Economics found that AI usage at work varies from 83% among Gen Z workers to 52% among Baby Boomers - a 31-percentage-point gap that reflects not just age but fundamentally different relationships with technology, risk, trust and professional identity. Critically, this is not simply a gap in adoption rates. It is a gap in how AI is used, why it is used, what is trusted and what is feared.

Generation	Age Range (2026)	AI Adoption Rate	How They Use AI	Key Motivations	Key Barriers	Social Economy Implications
Generation Z	Born 1997-2012 (14–29)	76-83% use generative AI tools (Deloitte 2025; LSE 2025)	Experimentation, creativity, content creation, career development, learning - treat AI as a collaborative co-pilot	Innovation, speed, career competitive edge, personal development, digital nativeness	Overconfidence without critical evaluation; risk of accuracy blind spots; may undervalue relational depth	Crucial talent pipeline for the sector; drawn by mission and purpose; risk of non-attraction if sector perceived as tech-backward or low-paying
Millennials	Born 1981-1996 (30–45)	58-73% use AI; 62% report high AI expertise; highest workplace adoption rate (McKinsey 2025)	Productivity, problem-solving, research synthesis, workplace tasks - most proficient users across generations	Professional advancement, efficiency, keeping competitive, managing complex workloads	Work-life pressure; 'sandwich generation' stress; rapid pace of change creates fatigue	Current backbone of social sector leadership; need structured AI upskilling; positioned as inter-generational AI translators
Generation X	Born 1965-1980 (46–61)	36-60% use AI; mostly passive or productivity-oriented (Deloitte 2025)	Search enhancement, efficiency tools, selective professional use - strategic caution based on prior tech experience	Practical value, proven results, efficiency gains, reliability	Scepticism from prior tech disappointments; lower formal AI training access; culture of self-reliance	Occupy senior leadership and board roles; their buy-in is critical for AI governance; respond to ROI evidence not hype
Baby Boomers	Born 1946-1964 (62-80)	20-52% use AI; mostly embedded, familiar tools;	Familiar embedded tools (email AI, navigation,	Practical usefulness, simplicity, familiarity,	Trust deficits, privacy concerns, complexity	Hold governance roles, board seats,

Generation	Age Range (2026)	AI Adoption Rate	How They Use AI	Key Motivations	Key Barriers	Social Economy Implications
		71% have never used ChatGPT (Deloitte 2025, McKinsey 2025, LSE 2025)	voice assistants); rarely use standalone generative AI	clear value proposition	aversion, limited support	long-term funder relationships; need non-threatening AI exposure; bring institutional wisdom AI cannot replicate
Generation Alpha	Born 2013+ (under 13)	**Not yet in workforce; growing up as first truly AI-native generation – included merely for horizon scanning purposes**	Learning tools, creative exploration, AI-assisted education from early age	Natural curiosity; learning-first orientation	Digital safety and consent risks; critical thinking development essential	Will enter the sector from 2030+; their expectations will redefine what a technology-enabled social sector looks like

Table 14: Generational AI Adoption and Social Economy Implications: Motivations, Barriers and Workforce Strategy by Generation

Generation Z: The Mission-and-Technology Generation

Generation Z currently aged 14 to 29 represents the social economy's most critical talent pipeline. They are the most digitally native, the most likely to use AI tools experimentally and creatively, and crucially the most explicitly motivated by purpose and social impact in their career choices. The WEF's 2025 Possibilists Study found a 'strong sense of collective purpose' among youth, most of whom are driven by a desire to support their communities and create systemic change.

However, **Gen Z's relationship with the social economy is ambivalent and conditional.** They are drawn to mission, but they are also pragmatic about career development, technology quality and professional growth. South African Gen Z workers - shaped by loadshedding, remote learning disruption, high youth unemployment and digital disruption - are entrepreneurial, values-driven and acutely aware of their market value. They will not accept poor technology infrastructure, unclear career pathways or low salaries simply because the work is important.

The social economy risks losing this generation entirely if it fails to modernise its technology infrastructure, create clear AI-enabled career pathways, articulate a credible professional development proposition and offer competitive (or at least liveable) compensation. Gen Z makes career decisions through digital channels - 46% have secured jobs through TikTok, 76% use Instagram for career content, and 95% say a company's social media presence influences whether

they will apply. A sector that communicates its career opportunities primarily through sector journals and network word-of-mouth is invisible to this generation.

> ***The social economy has a genuine competitive advantage with Gen Z in one dimension: mission alignment.*** *Surveys consistently show Gen Z prioritises purpose, social impact and values alignment over salary when making career choices - to a greater degree than any previous generation. But this advantage is conditional. Mission alone is not enough if the work environment is technology-backward, the career path is unclear and the pay is impossible to live on. The sector must earn this generation's commitment, not assume it.*

Millennials: The AI-Proficient Leadership Generation

Millennials currently aged 30 to 45 are both the social economy's most AI-capable generation and its most critical leadership cohort. McKinsey's 2025 research found that 62% of employees aged 35–44 report high AI expertise, higher than Gen Z (50%), Gen X or Boomers. Deloitte found that 56% of Millennials use generative AI at work. They are simultaneously the generation managing AI's workplace introduction, leading multi-generational teams, and navigating the 'sandwich generation' pressures of caring responsibilities and career peak years.

In the social economy context, Millennials are the current and emerging generation of programme managers, MEL specialists, fundraisers, communications leads and increasingly, executive directors. Their AI capability is an enormous asset but it is largely self-taught, ungoverned and unrecognised. The sector has not built structured pathways for AI-capable Millennials to develop their skills, apply them strategically or be compensated for the additional value they bring.

Critically, Millennials are also the generation most likely to leave the social sector for better-paying, more technologically stimulating roles in the corporate or technology sectors. With 77% believing generative AI will impact their work within the next year, and many already using it extensively, they are aware of the gap between what they can do and what their organisations allow or support. Retaining this generation requires structured investment, recognition and career advancement not just mission alignment.

Generation X: The Sceptical Bridge

Gen X currently aged 46 to 61 occupies a uniquely important position in the social economy's AI transition. **They hold a disproportionate share of senior management, board and strategic leadership roles. Their buy-in is therefore essential for organisational AI adoption - they must approve AI policies, champion AI investments and model AI-positive culture. Yet they are among the most sceptical of AI's promises.**

This scepticism, however, is not simply resistance. Research suggests it is strategic caution earned through experience of technology transitions that promised transformation and delivered disappointment. Gen X lived through the introduction of the internet, the ERP revolution, social media's rise and the big data wave - all of which came with transformational promises that were partially realised at best. Their measured approach to AI is, in many ways, exactly what the social

economy needs - a counterweight to uncritical AI enthusiasm and a voice for governance, ethics and proven value.

The implication is not that Gen X needs to be converted into AI enthusiasts. It is that they need to be provided with credible, evidence-based exposure to AI applications that demonstrate concrete social value - and that the sector needs to create conditions where their scepticism functions as a governance asset rather than an adoption barrier.

Baby Boomers: Institutional Wisdom in the Age of AI

Baby Boomers hold significant power in the social economy - as board members, long-term funders, governance advisors and institutional memory holders. Their AI adoption is the lowest of all generations, but this framing can obscure their critical contribution. The deepest institutional knowledge about community relationships, sector history, stakeholder trust and governance practice lives in this generation. This knowledge is not replaceable by AI - and the sector would be impoverished if the AI transition caused Boomer leadership to disengage or feel irrelevant.

The strategic imperative is not to make Boomers into AI practitioners, but to design AI governance frameworks that draw on their wisdom: their understanding of what the sector has tried before, their relationships with funders and communities, their ability to ask hard questions about what AI is actually serving. The most effective boards in an AI-enabled social economy will be those that combine Millennial and Gen Z AI fluency with Boomer and Gen X institutional wisdom and ethical scepticism.

Future Careers in the AI-Enabled Social Economy

One of the most important and underdiscussed dimensions of AI's impact on the social economy is the **new career landscape it is creating**. The WEF Future of Jobs Report 2025 projects that AI will displace 92 million jobs globally while creating 170 million new ones - a net positive, but a deeply disruptive transition that will not be evenly distributed. **For the social economy specifically, the AI era is creating a category of roles that barely existed five years ago and are urgently needed now.**

These are not conventional 'tech jobs' transposed into a social context. They are genuinely hybrid roles: combining technical capability with deep sector knowledge, community competency, ethical grounding and systems thinking. **They represent a new professional identity for the sector: practitioners who are simultaneously socially rooted and technologically capable, who understand both the language of data and the language of community, who can translate between AI systems and human systems.**

The following table maps ten emerging career roles in the AI-enabled social economy, with specific reference to South African entry points and development portfolio relevance.

Role / Career	Description	Core Skills Required	Development Portfolio Relevance	SA Social Economy Entry Point
Social Impact Data Scientist	Analyses complex social and community datasets to identify patterns, measure outcomes, model scenarios and generate predictive insights for programme and portfolio decisions	Python/R, statistical modelling, data visualisation, impact measurement frameworks (IRIS+, SROI), community data ethics, POPIA compliance	All portfolios: especially health, education, food security, livelihoods	SA universities offering data science degrees; Explore Data Science Academy (EDSA); AI4D fellowships; intermediary-embedded roles
AI-Enabled MEL Specialist	Designs and manages AI-augmented monitoring, evaluation and learning systems; bridges quantitative dashboards and qualitative community voice; governs data ethics	MEL methodology, AI tools (b.world, DHIS2), qualitative analysis, systems thinking, learning facilitation, POPIA compliance	All portfolios: critical for health, education, community development and livelihoods programmes	NPOs, funders and government departments; currently a major gap in SA sector; immediate demand
Community Intelligence Analyst	Synthesises community-level data - surveys, feedback, participatory assessments, social media signals - into actionable ecosystem intelligence for funders, intermediaries and policymakers	Community research methods, participatory data tools (Ushahidi, Kobo), qualitative synthesis, spatial data, local language competency	Community development, GBV response, food security, youth development	Intermediaries, DSD, NDA, social research units; links to urban planning and social policy roles
Impact Investment Analyst (Social Economy)	Conducts due diligence on social enterprises and impact organisations; uses AI tools to assess impact evidence, financial sustainability and social return; manages impact reporting for DFIs and philanthropic funds	Financial modelling, IRIS+, SROI, impact due diligence, AI-assisted research, ESG frameworks, blended finance instruments	All portfolios across impact investing spectrum - education, health, livelihoods, climate	DBSA, IDC, NEF, Bertha Centre, Tshikululu and emerging impact investing firms in SA
AI Ethics and Governance Officer (Social Sector)	Develops and enforces organisational and	AI ethics frameworks (UNESCO, NIST),	Cross-cutting: essential for any organisation	Intermediaries, large NPOs, CSI departments,

Role / Career	Description	Core Skills Required	Development Portfolio Relevance	SA Social Economy Entry Point
	ecosystem AI governance frameworks; conducts algorithmic bias audits; manages community consent and data sovereignty; advises boards on AI risk	POPIA, King IV governance, bias detection, community engagement, policy literacy	deploying AI in beneficiary-facing programmes	foundations; currently absent from almost all SA social sector organisations
Social Innovation Technologist	Designs, adapts and deploys technology solutions for social sector challenges; bridges technical capability and community context; leads human-centred design processes	Human-centred design, low-code/no-code platforms, multilingual tech (Vulavula), community co-design, digital inclusion principles	Youth development, education, livelihoods, community development, entrepreneurship	Social innovation labs, tech NPOs, GSBI-affiliated social enterprises, university innovation hubs; links to SA tech startup ecosystem
Ecosystem Intelligence Manager	Leads the development and maintenance of cross-sector knowledge and data platforms; synthesises learning across multiple partners; produces ecosystem intelligence reports for funders and policymakers	Knowledge management systems, AI synthesis tools, stakeholder facilitation, ecosystem mapping, narrative communication	Primarily intermediary and funder roles; critical for education, health and food security ecosystem coordination	New Africa Capital Forum, Tshikululu, Southern Africa Trust, national coordination bodies; DSD and NDA at government level
Digital Inclusion and AI Literacy Facilitator	Builds AI literacy and digital capability across organisations and communities, particularly those at Readiness Levels 1–2; facilitates peer learning networks; bridges language and digital divides	Adult facilitation, digital literacy pedagogy, multilingual communication, AI tools training, community development	Education, youth development, community development, livelihoods; particularly CBOs and grassroots organisations	NGOs, SETAs, libraries, community centres, tech hubs; directly addresses the sector's equity gap in AI access
Participatory Research and AI Co-design Specialist	Leads community-centred research and co-design processes that integrate AI tools with participatory	Participatory Action Research (PAR), co-design facilitation, qualitative methods, AI tool	All portfolios where beneficiary voice is central: particularly health, GBV, child development, food security	Academic institutions, community-based research organisations, policy research

Role / Career	Description	Core Skills Required	Development Portfolio Relevance	SA Social Economy Entry Point
	methodologies; ensures community voice shapes AI system design and governance	integration, community ethics		units; currently rare in SA sector
Social Sector AI Programme Manager	Manages end-to-end AI adoption programmes within social organisations; coordinates technology, people, governance and change management; translates AI strategy into operational reality	Project management, change management, AI literacy, stakeholder management, M&E, organisational development	Cross-cutting: needed in any organisation moving from AI awareness to integration (Levels 2–4 on the readiness pathway)	Large NPOs, intermediaries, CSI departments, government social sector; growing demand as AI adoption accelerates

Table 15: Emerging Career Roles in the AI-Enabled Social Economy — Skills, Portfolio Relevance and SA Entry Points

The Portfolio-Specific Skills Landscape

Beyond these cross-cutting professional roles, the AI era is reshaping the skills requirements within each of the sector's major development portfolios. Practitioners working in education, health, food security, livelihoods, GBV response, child development, community development and entrepreneurship all face different AI-driven transformations in their daily work. The following table maps these portfolio-specific skills shifts.

Development Portfolio	Current Skills Gap	AI-Augmented Role Emerging	New Competencies Required	Gen Z / Young People Entry Point
Education & Skills Development	Teacher training quality; ECD practitioner capacity; TVET curriculum relevance; data on learning outcomes	AI-enabled learning designer; adaptive curriculum specialist; ECD programme data analyst	Learning science + AI tool integration; multilingual content design; POPIA-compliant learner data management; inclusive pedagogy	ECD assistant with AI support tools; youth peer learning facilitators; digital skills coaches in TVET colleges
Health & Community Health	Community health worker (CHW) support and supervision; clinical data systems; rural diagnostic capacity	AI-assisted community health navigator; digital health programme coordinator; telehealth support specialist	mHealth tool operation; AI-assisted triage protocols; clinical data entry and validation; referral system management	Community Health Worker roles enhanced by AI decision-support; health data entry and quality roles; patient feedback system operators

Development Portfolio	Current Skills Gap	AI-Augmented Role Emerging	New Competencies Required	Gen Z / Young People Entry Point
Food Security & Agriculture	Smallholder advisory services; nutrition programme reach; early warning system use	Food systems data analyst; AI-enabled nutrition programme coordinator; early warning system operator	Food insecurity forecasting tools; GIS and spatial data literacy; community nutrition data collection; supply chain analytics	Data collectors for food security monitoring; community agri-advisors using AI crop advisory apps; food garden programme coordinators
Livelihoods & Youth Employment	Career guidance quality; SETA programme effectiveness; informal sector support; NEET tracking	AI-enabled career counsellor; livelihoods programme data specialist; SMME digital support advisor	Skills matching platform operation; labour market data analysis; SMME digital tools training; youth engagement through AI	Youth digital ambassadors; peer career coaches; SMME support roles in NPOs and SETAs; entrepreneurship hub data roles
Gender Equity & GBV Response	GBV case management systems; survivor data protection; legal aid access; cross-sector referral coordination	AI-enabled GBV case coordinator; survivor data protection officer; digital safety and security specialist	Trauma-informed AI tool design; POPIA-compliant survivor data management; digital safety protocols; AI-assisted legal navigation	Young women in community safety data roles; digital helpline operators; GBV awareness content creators using AI tools
Child & Youth Development	ECD quality assurance; child protection case management; youth mentorship reach; out-of-school tracking	AI-assisted child protection data analyst; youth programme impact specialist; ECD technology facilitator	Child protection data ethics; AI case flagging systems; youth engagement digital tools; POPIA consent for under-18s	Youth development coordinators using AI scheduling/ matching tools; young mentors trained in digital engagement; ECD programme trackers
Community & Social Development	Community needs assessment quality; social cohesion programme evidence; place-based planning capacity	Community intelligence analyst; social cohesion programme evaluator; place-based development data coordinator	Community data collection (Kobo, Ushahidi); participatory GIS; social network analysis; community feedback synthesis	Community data collectors; neighbourhood-level programme coordinators using digital tools; social cohesion facilitators with AI support
Entrepreneurship & Social Enterprise	Business support quality for social enterprises; impact measurement capacity; investor-readiness	Social enterprise impact analyst; AI-enabled business advisor; blended finance programme coordinator	Social enterprise MEL; impact investing frameworks (IRIS+); AI-assisted business planning; market analysis tools	Young social entrepreneurs with AI business tools; enterprise support coordinators; impact measurement

Development Portfolio	Current Skills Gap	AI-Augmented Role Emerging	New Competencies Required	Gen Z / Young People Entry Point
				associates in incubators

Table 16: AI-Era Skills Transformation Across Development Portfolios — Current Gaps, Emerging Roles and Young People Entry Points

A Competency Framework for AI-Era Social Economy Practitioners

The social economy needs a shared, sector-appropriate competency framework that describes what AI-era practitioners need to know, be able to do and embody at different levels of seniority and specialisation. The following framework draws on the WEF's Future of Jobs Report 2025, the Equip the Flip initiative's social sector leadership competencies, and the specific demands of South Africa's social economy context.

The framework is organised around seven competency domains, each with three tiers: core competencies expected of all practitioners regardless of role, advanced competencies for those in specialist or technical positions, and leadership competencies for those with governance, strategic or executive responsibilities.

Competency Domain	Core Competencies (All Practitioners)	Advanced Competencies (Specialists)	Leadership / Governance Competencies
AI Literacy & Digital Fluency	Understanding what AI is and is not; basic use of AI tools for communications, research and reporting; prompt engineering fundamentals; recognising AI-generated content limitations	AI tool selection and workflow design; data quality management; AI output critical evaluation and verification; low-code/no-code platform use	AI strategy development; organisational AI policy design; board-level AI risk oversight; technology partner management
Data Thinking & Evidence Culture	Basic data literacy: reading dashboards, understanding indicators; contributing to data collection; using evidence in programme conversations	Data analysis and visualisation; MEL system design; mixed-methods research; impact measurement framework application (IRIS+, SROI, GRI)	Evidence-based decision-making culture; learning organisation design; sector knowledge sharing leadership; research commissioning
Ethics, Equity & AI Governance	Understanding data privacy (POPIA); recognising algorithmic bias; community consent principles; digital safety basics	AI ethics framework application (UNESCO, NIST); bias auditing; data sovereignty design; community data governance	AI ethics policy development; governance framework design; regulatory engagement; sector ethics advocacy; board AI oversight
Systems & Ecosystem Thinking	Understanding the social ecosystem; recognising feedback loops; seeing	Systems mapping; ecosystem analysis; network thinking; theory	Ecosystem governance; cross-sector partnership leadership; systems

Competency Domain	Core Competencies (All Practitioners)	Advanced Competencies (Specialists)	Leadership / Governance Competencies
	how programme fits into broader system	of change with systems lens; adaptive programme management	change strategy; collective impact facilitation
Human-Centred & Relational Practice	Community facilitation; active listening; cultural competency; trust-building; Ubuntu-centred engagement; trauma-informed practice	Participatory design facilitation; co-design with communities; community-led research methods; cross-cultural leadership	Stakeholder convening; power-aware partnership design; community accountability mechanisms; social movement literacy
Adaptive Leadership & Organisational Resilience	Openness to change; learning orientation; resilience under uncertainty; collaborative problem-solving	Change management; organisational development; multi-generational team leadership; innovation facilitation; crisis management	Strategic AI adoption leadership; organisational transformation; multi-stakeholder governance; sector advocacy and thought leadership
Communication, Storytelling & Influence	Clear written and verbal communication; basic data storytelling; social media literacy; multilingual communication	Impact communication design; data visualisation; AI-assisted content creation; report writing; media and donor relations	Executive thought leadership; policy influence; sector voice and advocacy; publication and public discourse participation

Table 17: AI-Era Competency Framework for Social Economy Practitioners — Core, Advanced and Leadership Competencies Across Seven Domains

Several dimensions of this framework deserve particular emphasis given the South African context.

Protecting Ubuntu in an AI World

At the heart of African social practice is ubuntu - the understanding that personhood is constituted through relationship, community and mutual care. This is not a soft cultural nicety. It is the philosophical foundation of the sector's most effective practices: participatory design, community-led development, trauma-informed care, social cohesion work, and the relational trust that makes behaviour change possible. The competency framework must explicitly honour and protect these capabilities - not as alternatives to AI competency, but as the deeper orientation within which all AI tools must be deployed.

The social economy's competitive advantage against the technology sector is precisely this: it is not trying to optimise systems. It is trying to strengthen communities, restore dignity and create conditions for human flourishing. These are goals that require human wisdom, relational presence and cultural competency that no algorithm can provide. AI is a tool in service of these goals - never a replacement for them.

The Language Dimension

South Africa's linguistic diversity - 11 official languages plus the rich multilingual practice of most communities - creates a specific competency requirement that rarely appears in generic skills frameworks: the ability to work across languages, to commission and evaluate multilingual AI tools, and to ensure that the sector's AI adoption does not default to English in ways that systematically exclude Zulu, Xhosa, Sotho, Tswana and other language communities. This is both a technical requirement (choosing tools like Vulavula that support African languages) and a value requirement (insisting that language equity is non-negotiable in AI adoption decisions).

The Role That Holds Everything Together: The Human-Centred Technologist

Across all of these emerging roles, one profile emerges with particular importance for the social economy: the human-centred technologist. This is not primarily a technical role. It is a bridging role: someone who possesses enough technological fluency to work effectively with AI tools, data systems and digital platforms, but whose deepest orientation is toward the communities they serve, the social systems they work within, and the ethical responsibilities they carry.

The WEF projects that by 2026, nine out of ten jobs will require some combination of digital fluency, social influence and creative problem-solving. For the social economy, this means the sector's core capabilities: community facilitation, trust-building, empathy, participatory practice, systems navigation - are becoming more valuable, not less, in an AI-enabled world. The task is not to choose between technology and humanity. It is to cultivate both, deliberately and in service of the communities the sector exists to support.

> *The half-life of a learned skill is now estimated at approximately five years.* ***This means that today's social economy practitioners will need to actively refresh their capabilities at least twice before 2035.*** *The sector's investment in continuous professional development - and funders' willingness to pay for it - is not a peripheral nicety. It is a core component of organisational sustainability in the AI era.*

The CTO Seat: Why the Human-Centred Technologist Needs a Technical Partner

The human-centred technologist is a bridging role. That framing is precise and important: this practitioner holds enough technological fluency to work with AI tools, data systems and digital platforms, while their deepest orientation remains toward the communities they serve and the social systems they work within. This is the right profile for the integrative, mission-anchored AI work the social economy most urgently needs.

But the bridging role does not own the infrastructure. And infrastructure ownership matters.

A Chief Technology Officer function brings something categorically different: build-level technical assurance. Computer systems design, software development, AI infrastructure architecture, data security protocols, POPIA-compliant data residency configuration, vendor contract review, penetration testing oversight, and the technical governance that keeps sensitive

beneficiary data safe when an organisation is operating at scale. These are not capabilities the human-centred technologist is positioned to hold - and without them, an organisation's AI systems may function technically while failing the very governance standards that Section Nine of this guide establishes as non-negotiable.

- **The human-centred technologist is a bridging role. This integrative figure brings sector judgement and mission orientation.**
- **The CTO brings compliance and best-practice assurance.**

Both are necessary. Sector organisations need to know they need both: not as a counsel of perfection, but as a practical design principle for any organisation handling sensitive beneficiary data or deploying AI in programme-critical functions.

The Fractional CTO: Making the Seat Viable

CTO-level talent commands corporate market rates that most NPOs cannot approach. This is real. But it does not make the CTO seat optional - it makes the funding model for it the design challenge.

The fractional model is the solution. A fractional CTO does not require a full-time hire at market rates. The role flexes with the organisation's stage of AI adoption: heavier during architecture and system setup, when foundational decisions about data residency, security configuration and infrastructure design are being made; lighter once the build is stable and the human-centred technologist can manage day-to-day; scaling again when new initiatives require new build decisions.

More practically: a fractional CTO funded across three to six organisations within an intermediary's network stretches the role's cost to viability. Each organisation contributes a fraction of the total cost; the CTO provides the technical governance each needs, at the intensity each stage requires. This is how the sector can access CTO-level assurance without competing for full-time talent against the technology sector, and it is a model that intermediaries are institutionally well-placed to host and coordinate.

> ***Design principle for intermediaries:*** *A fractional CTO embedded within an intermediary's shared services function can serve as the technical governance layer for an entire network of grantees - reviewing vendor contracts before signing, configuring data residency settings, conducting quarterly security audits, and standing alongside human-centred technologists when AI systems need rebuilding or scaling. This is precisely the kind of shared infrastructure investment that transforms intermediaries from coordinators into genuine ecosystem intelligence and governance hubs.*

When Does an Organisation Need a CTO vs. a Human-Centred Technologist?

The answer depends on what the organisation is building and what data it handles. A CBO using Claude or Microsoft Copilot to draft reports and synthesise research findings needs an AI-literate staff member - a human-centred technologist at Readiness Level 2–3. It does not need a CTO.

An organisation deploying a beneficiary-facing chatbot, building a case management system with AI-assisted decision support, or integrating AI into a child protection flagging workflow needs both: a human-centred technologist to manage the programme logic and community accountability, and a CTO (fractional at minimum) to ensure the architecture is secure, the data is handled lawfully, and the system is built to a professional standard. The stakes are different, and the governance requirement scales accordingly. The table below provides a practical guide to when each seat becomes necessary.

AI Deployment Scenario	Human-Centred Technologist Needed?	CTO/Fractional CTO Needed?	Minimum Technical Governance
Staff using AI tools (ChatGPT, Claude) for drafting, research, communications - no personal data involved	Yes: AI literacy, prompt governance, use policy	**No:** not required at this stage	AI usage policy; staff training; tool selection criteria
AI-assisted MEL and reporting using aggregated/anonymised programme data	Yes: data thinking, tool configuration, quality review	**Optional:** helpful for data architecture review	Data classification policy; anonymisation protocols; DPA for any cloud tool
Community feedback system (surveys, SMS, WhatsApp) using names and contact data	Yes: consent design, tool management, community accountability	**Recommended:** data residency and DPA review required	POPIA lawful basis; DPA with tool provider; data residency confirmation; consent documentation
Beneficiary-facing chatbot or AI voice interface	Yes: community design, language equity, escalation protocols, trauma-informed review	**Required:** system architecture, security, API integration, crisis protocol integration	Full POPIA compliance; human escalation pathway; bias testing; security penetration test; incident response plan
AI-assisted eligibility screening or beneficiary matching	Yes: mission logic, community accountability, appeal processes	**Required:** algorithmic bias audit, explainability design, POPIAs.71 compliance	Automated decision compliance (POPIA s.71); human review pathway; bias audit; explainability documentation; appeal mechanism
Case management system with AI flagging (child protection, GBV, health)	Yes: sector practice, trauma-informed design, community trust	**Required:** full security architecture, Special Personal Information compliance, audit trails	POPIA Chapter 6 (Special Personal Information); Privacy Impact Assessment; security certification; incident response; regulatory notification protocols

Table 18: When Does an Organisation Need a CTO? A Decision Guide by AI Deployment Scenario

A Talent Strategy for the AI-Era Social Economy

Understanding the skills gap, the generational dynamics and the emerging career landscape is only meaningful if it translates into concrete talent strategy. The following table maps the sector's key

talent strategic priorities, the specific challenges each faces, the AI-era actions required and who must lead each.

Strategic Priority	Challenge in Current SA Social Economy	AI-Era Action Required	Who Must Lead
Attracting Young Talent	Low salaries, poor technology infrastructure, weak career pathways and limited AI exposure make the sector unattractive to digitally-native Gen Z talent relative to corporate or tech sector alternatives	Create visible AI-enabled career pathways; invest in technology infrastructure; develop competitive salary advocacy; build sector brand through digital channels (TikTok, Instagram - 76% of Gen Z use these for career discovery)	Intermediaries, SAIA/Inyathelo, sector umbrella bodies, CSI departments as talent investors
Retaining Millennial Leaders	Millennials are the sector's most AI-capable generation but face salary compression, limited advancement and high burnout - nearly half consider leaving roles within six months	Provide structured AI upskilling; create senior practice roles (not just management tracks); offer flexible working; invest in wellbeing; articulate mission-AI connection clearly	NPO boards, executive directors, HR-capable intermediaries
Engaging Gen X in AI Transition	Gen X leaders and managers are sceptical of AI based on prior tech experience; their buy-in is essential for organisational AI adoption but is rarely deliberately managed	Provide evidence-based AI case studies; focus on practical value not hype; create psychologically safe spaces for experimentation; frame as augmentation not replacement	Senior leadership teams, AI champions within organisations, professional development providers
Supporting Baby Boomer Transition	Boomers hold governance, board and senior advisory roles but have the lowest AI adoption; risk of governance without AI understanding	Develop board-level AI literacy programmes; simple, embedded AI exposure; draw on their institutional wisdom as complement to AI capability	Sector governance bodies, IoDSA, SAICA, NPO governance training providers
Building a Pipeline from Universities	Social work, development studies, public administration graduates enter the sector with minimal AI literacy, creating an immediate skills gap from day one	Advocate for AI literacy modules in social work, development studies and public administration degrees; create NPO internship programmes with AI components; fund bursaries for AI-adjacent roles	Universities, SETA system (ETDP SETA, HWSETA), intermediaries, CSI talent investment programmes
Professionalising the AI Function	No recognised AI roles, career pathways or salary benchmarks exist in the social sector; AI is treated	Create recognised professional roles (Social Impact Data Scientist, AI Ethics Officer, MEL Technology Specialist);	Sector bodies, NPO coalitions, HR-focused intermediaries, Trialogue benchmarking

Strategic Priority	Challenge in Current SA Social Economy	AI-Era Action Required	Who Must Lead
	as an add-on rather than a professional function	develop salary benchmarks; build professional development pathways	
Investing in CBO Capacity	Community-based organisations face the greatest talent access barriers; digital divides, low wages and infrastructure gaps make AI adoption structurally difficult	Fund shared AI facilitator roles within intermediary networks; support peer learning and community of practice models; prioritise multilingual tools (Vulavula, Masakhane); build offline-capable systems	Funders, intermediaries, Digital Access South Africa initiatives, GSMA AI for Development

Table 19: Talent Strategy for the AI-Era Social Economy — Strategic Priorities, Actions and Accountability

The Structural Investment That Makes Everything Possible

Underlying every element of this talent strategy is a financial and political commitment that the sector has historically failed to make: ***investing in people as a strategic asset rather than a programme cost.*** The social economy's chronic underinvestment in staff salaries, professional development, institutional capacity and technology infrastructure is not a secondary issue - it is the root cause of many of the sector's most persistent challenges, and the primary reason that AI adoption will struggle without structural change.

Bridgespan's 2025 research on South African CSI explicitly identifies investing in grantee institutional capacity - including staff capacity - as one of the six key practices that distinguish transformational from transactional social investment. The implication for funders is direct: ***if AI is to deliver value in the social sector, funders must be willing to fund the human development that makes AI adoption possible, not just the programme costs where AI outputs will eventually appear.***

Academic and Training Institution Partnerships

The gap between what universities produce and what the social economy needs in an AI era is significant and growing. Social work graduates, development studies majors, public administration professionals and community development practitioners are entering the sector with minimal AI literacy, limited data competency and no exposure to the emerging roles that AI is creating. Closing this gap requires **active engagement between sector bodies and academic institutions** - advocating for AI literacy modules in social science curricula, creating NPO internship programmes with AI components, and funding bursaries for AI-adjacent roles within the sector.

South Africa's SETA system - particularly the Education, Training and Development Practices SETA (ETDP SETA) and the Health and Welfare SETA (HWSETA) - has a direct mandate to address skills gaps in the social sector. AI literacy, data thinking and digital inclusion competencies should be urgently integrated into the learnerships, skills programmes and workplace learning frameworks administered through these bodies.

The Wellbeing Dimension

A talent strategy for the AI-era social economy cannot ignore practitioner wellbeing. Social Current's 2025 workforce analysis found that social sector employees are navigating change, stress, political polarisation, isolation and vicarious trauma simultaneously - making the sector's already high burnout risk even more acute. The AI transition adds a further layer of uncertainty: will my role be automated? Do I have the skills to remain relevant? Is the sector I have devoted my career to capable of navigating this transformation?

These anxieties are real and must be named and addressed - not dismissed with reassuring platitudes about AI as augmentation. The sector needs to invest in psychological safety, transition support, peer learning networks and honest communication about what AI adoption means for different roles. Practitioners who feel threatened and unsupported will resist AI adoption regardless of its potential value. Those who feel prepared, supported and respected will become the sector's most powerful AI change agents.

People First: The Non-Negotiable Foundation

The social economy's AI journey will ultimately succeed or fail not at the level of technology but at the level of people. The platforms, tools, frameworks and standards discussed throughout this guide are only as valuable as the human capability that can deploy them thoughtfully, govern them ethically and ground them in genuine community accountability.

This means the sector needs to make a set of commitments that it has historically been reluctant to make: ***to pay practitioners competitively, to invest in their professional development continuously, to create career pathways that are visible and achievable, to build technology infrastructure that makes their work more powerful rather than more frustrating, and to protect the relational and community-embedded capabilities that are the sector's deepest source of value.***

It means designing an AI adoption pathway that is explicitly multigenerational - drawing on Gen Z's digital nativeness and Boomer's institutional wisdom, on Millennial AI proficiency and Gen X strategic caution, in ways that build collective capability rather than creating new forms of exclusion within the workforce.

And it means telling a compelling story about why the social economy is where the most interesting, most consequential and most human work in an AI-enabled world will take place. Because that story is true - the challenge is making the sector capable of living up to it.

> *The future social economy practitioner is not a social worker who has learned to use AI, nor an AI professional who has learned about social issues.* ***They are something genuinely new: a human-centred technologist, a community-rooted data thinker, a relational systems navigator who can hold the complexity of human need and the possibility of intelligent tools in productive, ethical and community-accountable tension.*** *Developing this practitioner - in their thousands, across the continent - is perhaps the most important investment the sector can make in its own future.*

SECTION THIRTEEN

Toward an Intelligent Social Economy: Recommendations

Who this section is for: All readers. These recommendations are addressed to the sector as a whole, with named accountabilities for funders, intermediaries, government, NPOs and sector bodies. They are organised in four clusters and intended to be read as a connected set. Pick up the ones where you have agency - and hold the others as a framework for the conversations you need to have.

The future of Africa's social and impact economy increasingly depends on its ability to become more intelligent, more adaptive, more collaborative and more interconnected. AI may significantly accelerate this transformation. But technology alone will not solve systemic inequality. The future social economy will still depend on trust, ethics, stewardship, local legitimacy, participatory governance and human relationships.

This is particularly important in Africa. Because Africa's social economy is not merely a smaller version of Global North systems. It is more relational, more adaptive, more networked, more hybrid and more community-embedded. These characteristics are not deficits to be overcome - they ultimately become strategic advantages in an AI-enabled future.

The following recommendations are addressed to the sector as a whole, with specific implications for funders, intermediaries, NGOs, government and communities. The recommendations are organised in four clusters: People and Capability, Governance and Ethics, Ecosystem Intelligence and Infrastructure, and Policy and Systemic Change - and are provided to specific actors with named time horizons. They are intended to be read as a connected set, not a menu.

Cluster A: People, Capability and Talent — Investing in the Human Foundation

No AI strategy survives contact with an under-capacitated workforce. The most sophisticated platform deployed in an organisation without AI literacy, data governance capacity or technical leadership will produce poor outcomes - or worse, ungoverned ones. These recommendations address the people and capability preconditions that underpin everything else.

RECOMMENDATION A1

Invest in AI Literacy as a Sector-Wide Infrastructure Priority

Funders, intermediaries and CSI departments should treat AI literacy investment not as a project-specific line item but as sector infrastructure, analogous to the investment in MEL capacity that the sector made over the past two decades. This means funding AI literacy programmes at every level: board governance briefings for Boomer and Gen X leadership, structured upskilling for Millennial programme staff, and digital inclusion pathways for CBO practitioners in languages other than English.

The ETDP SETA and HWSETA have a direct statutory mandate to address skills gaps in the social sector. AI literacy, data thinking and human-centred technology competencies should be urgently integrated into their learnership and skills programme frameworks. This requires active advocacy from sector bodies including Inyathelo (South African Institute for Advancement - Inyathelo) and sector intermediaries such as IPASA (Independent Philanthropy Association of South Africa).

Time horizon:
- *6 months: First AI literacy cohorts and board briefing programmes.*
- *12 months: SETA integration advocacy and intermediary-delivered AI training networks.*
- *3 years: Sector-wide AI competency baseline visible in workforce data and funder reporting requirements.*
- *Who leads: Intermediaries, DSD, Inyathelo CSI departments, ETDP SETA, HWSETA.*

RECOMMENDATION A2

Build the Pipeline — Develop Human-Centred Technologists from Within the Sector

The social economy needs a new professional profile urgently: the human-centred technologist, someone who holds deep sector knowledge, community competency and sufficient AI fluency to build and govern AI systems. This profile does not exist in significant numbers yet, and salary advocacy alone will not create it. What is needed is deliberate pipeline investment: structured development programmes that take experienced sector practitioners and build their AI and data capabilities, combined with pathways that bring technically skilled young people into the sector with proper mission formation.

This is not primarily a recruitment problem. It is a development problem and it requires funders to invest in the time, structure and mentorship needed to develop practitioners who hold both sides of the equation. Intermediaries are well-positioned to host such programmes across their networks, creating shared development infrastructure that individual NPOs cannot afford alone.

Time horizon:
- *12 months: First cohort of human-centred technologist development programmes in at least two intermediary networks.*
- *3 years: Recognised professional pathway with salary benchmarks and a community of practice of 200+ practitioners.*
- *Who leads: Intermediaries, universities, sector bodies, forward-looking CSI departments as talent investors.*

RECOMMENDATION A3

Establish the Fractional CTO Model as a Sector-Funded Infrastructure

The sector needs technical leadership, not just AI literacy. A Chief Technology Officer function brings the build-level assurance that the human-centred technologist role cannot provide alone: POPIA-compliant infrastructure architecture, data security standards, AI system auditing, vendor

assessment and the technical governance that keeps sensitive beneficiary data safe. This is not a role most NPOs can hire full-time at market rates. The solution is a fractional model: one CTO shared across three to six organisations in an intermediary's network, funded collectively, scaling in time with each organisation's stage of AI adoption.

Intermediaries are the natural anchor for this model. A fractional CTO embedded within an intermediary's technical infrastructure can serve multiple grantees simultaneously, providing architecture oversight during setup, lighter-touch maintenance once systems are stable, and scaling back in as new initiatives are built. This is how the sector can access technical leadership at a viable cost without bidding against the corporate tech market for full-time talent.

Time horizon:

- *6–12 months: First fractional CTO pilots in two or three intermediary networks.*
- *2–3 years: Established shared CTO infrastructure across major intermediary networks; documented cost-sharing models available for replication.*
- *Who leads: Well-resourced intermediaries with technology mandates; DFIs and foundations willing to fund shared technical infrastructure.*

RECOMMENDATION A4

Address Scarcity, Not Just Salary — Pool Technical Talent Across the Sector

The talent challenge facing the social economy in AI is not primarily a salary problem, though salaries matter. It is a scarcity problem. Human-centred technologists are rare because few people hold the combination of sector depth, community competency and AI capability that the role requires. Fractional CTOs are expensive even on a fractional basis because demand for tech leadership in the South African market is high. Bidding up sector salaries will not solve either constraint.

What solves the scarcity constraint is building supply: more human-centred technologists developed from within the sector (Recommendation A2), more shared CTO capacity across networks (Recommendation A3), and a deliberate secondment and loaning culture where larger, AI-capable organisations temporarily second technical staff to smaller organisations. The corporate sector has precedent for this in pro bono consulting models; the social economy needs an equivalent for technical capacity.

RECOMMENDATION A5

Make the Sector Visible to Gen Z — Build a Compelling Digital Talent Brand

The social economy's greatest competitive advantage with the most digitally capable generation in history is mission and purpose, and it is squandering this advantage through invisibility on the channels where Gen Z makes career decisions. TikTok, Instagram and LinkedIn are not optional communication channels for the sector's talent strategy; they are the primary arena where this generation decides whether a career in the social economy is worth pursuing. A sector that communicates its opportunities primarily through sector journals, word-of-mouth and sector events is invisible to the talent it most needs.

This requires a coordinated sector brand investment: intermediaries, umbrella bodies and well-resourced NPOs should co-invest in a compelling digital narrative about careers at the intersection of social impact and AI. This is not marketing; it is a survival strategy for a sector that will be increasingly competing with technology companies that offer both mission framing and superior technical environments.

RECOMMENDATION A6

Redefine "Overhead" — Fund Institutional Intelligence as a Grant Condition

The social economy's most persistent funding pathology is the treatment of institutional capacity: data systems, staff development, technology infrastructure, governance investment - as overhead to be minimised rather than as the foundation on which programmatic impact depends. This pathology predates AI. AI makes it existentially urgent.

A CSI manager who funds a community nutrition programme but refuses to fund the data system that would tell the implementing NPO whether the programme is working, is not making a cost-saving decision. They are making a decision to remain ignorant of their own impact. A foundation that funds an NPO's services but will not fund the MEL specialist who could synthesise learning across their portfolio is not being fiscally responsible. They are preventing the accumulation of evidence that would make future funding decisions better.

The practical recommendation is direct: funders - CSI departments, philanthropic foundations, bilateral donors, government DFIs - should establish a minimum institutional investment threshold as a condition of grant-making. A credible threshold, based on sector evidence, is between 20-32% of grantmakers and philanthropists annual budget. (Source: Candid). That includes grants for institutional capacity, including staff AI development, data governance, technology infrastructure and learning systems. This is not generosity. It is a precondition for impact.

Specifically, funders should:

- Permit AI tool subscriptions and staff AI training as eligible grant expenditure
- Fund shared AI infrastructure through intermediaries as a portfolio-level investment, not a line item in individual grants
- Require AI readiness assessment as part of due diligence, and fund the development of readiness where gaps are identified
- Commit to multi-year funding cycles that give organisations the stability to invest in AI adoption rather than surviving month to month

The Bridgespan 2025 research on South African CSI identifies investing in grantee institutional capacity as one of six practices that distinguish transformational from transactional social

investment. The sector has known this for years. AI simply raises the stakes of continuing to ignore it.

According to Candid (2024 Foundation Giving Survey) operational funding standards are evolving, grantmakers typically deploy unrestricted support through two distinct mechanisms:

AS A % OF GRANTMAKING

Instead of funding isolated projects, many grantmakers now allocate a defined median of **32%** of their total annual grant dollars toward unrestricted support. This varies heavily by foundation, ranging anywhere from 0% to 100% depending on the institution's charter.

AS AN OVERHEAD RATE

When providing restricted project funding, global grantmakers traditionally imposed a tight 10% ceiling on administrative or "overhead" costs. Today, "Pay-What-It-Takes" philanthropy advocates push for realistic coverage of an organisation's true infrastructure. The industry standard now typically lands between **15% and 30%** — with major global funders such as the Bill and Melinda Gates Foundation using a **15% flat rate** for indirect cost.

Time horizon:

- *6 months: First funders adopt explicit institutional capacity thresholds and permit AI investment in grant budgets.*
- *12 months: Sector bodies publish guidance on intelligent overhead as a grant norm.*
- *3 years: Institutional investment in AI readiness is a standard component of CSI and philanthropic due diligence. Who leads: Trialogue, IPASA, NACF, Bertha Centre, DBSA, IDC - with advocacy from Inyathelo and sector umbrella bodies.*

Cluster B: Governance, Ethics and Responsible Deployment

Governance and ethics are not constraints on AI adoption, they are the preconditions for responsible adoption that the sector can stand behind publicly, defend to its beneficiaries, and sustain over time. These recommendations address the governance infrastructure the sector must build.

RECOMMENDATION B1

Develop Sector-Wide AI Governance Standards — Urgently

The vast majority of South African NPOs, funders and intermediaries operate without any AI governance policy. This is not a criticism, the tools are new and the sector has been responding in real time. But the window for getting ahead of governance rather than catching up to it is closing. As AI tools become embedded in programme delivery, beneficiary communication and funding decisions, the absence of governance frameworks creates growing legal, ethical and reputational exposure.

Sector bodies, SAMEA, Inyathelo, Independent Philanthropy Association, sector-specific umbrella organisations, should convene urgently to develop a Social Economy AI Governance Code: a

lightweight, practical framework drawing on King IV, POPIA, the AU Continental AI Strategy, UNESCO's AI Ethics Recommendation and ISO/IEC 42001. This should be accompanied by a simple, accessible AI policy template that any NPO can adapt within a day's work.

The governance framework must address eight non-negotiables:

- *POPIA compliance protocols for AI tools; data residency decision rules; community consent standards for AI-enabled programmes; algorithmic bias audit requirements; vendor assessment criteria; community appeals and grievance mechanisms; board-level AI oversight accountabilities; and mandatory human review pathways for AI-assisted decisions about individuals.*

Time horizon:

- *6 months: First draft code.*
- *12 months: Adopted by leading sector bodies.*
- *3 years: Included in NPO compliance frameworks and funder due diligence.*

Who leads: SAMEA, Inyathelo, IPASA, intermediaries, in partnership with legal and governance experts.

RECOMMENDATION B2

Require POPIA Compliance as a Precondition, Not an Afterthought

The practical AI tool selection guide provides the operational detail. The governance principle is simple: no AI tool that processes identifiable beneficiary data should be adopted without a documented data processing agreement, a data residency determination, and a POPIA lawful-basis assessment. This is not optional, it is a legal obligation. And it is one that the sector routinely overlooks when the excitement of a useful new tool overtakes the discipline of compliance.

Funders should make POPIA-compliant AI governance a condition of grant funding for any programme that includes AI-enabled data collection or analysis. This is not punitive; it is protective - of beneficiaries, of organisations, and of the sector's public legitimacy.

RECOMMENDATION B3

Close the Gender Digital Divide as an AI Equity Non-Negotiable

An ImpactHER 2024 survey found that 86% of women across 52 African countries lack basic AI proficiency, with internet access 37 percentage points lower than men and device ownership at 84% absent in rural areas. Any AI strategy for the African social economy that does not explicitly address this divide will deepen gender inequality rather than reduce it, contradicting the sector's core values.

Every AI literacy investment, every tool selection decision, every chatbot or community interface deployment must be reviewed through a gender equity lens. Are tools available in languages women in this community actually use? Do women have the devices and connectivity to access them? Are women practitioners receiving AI upskilling at the same rate as men? These are not rhetorical questions, they require data, monitoring and accountability.

RECOMMENDATION B4

Centre Community Voice in AI Governance — Not Just Programme Design

The sector has made genuine progress in participatory programme design over the past decade. AI governance requires the same participatory discipline: communities whose data AI systems process should have a voice in how those systems are designed, what data is collected, how decisions are made and how errors are corrected. Community data trusts, participatory AI audits, and regular community review of AI-generated insights are not aspirational, they are the governance mechanisms that make AI legitimate in the eyes of those it most affects.

Organisations working in high-sensitivity areas - GBV, child protection, health, legal aid - have an especially acute obligation here. A survivor whose case data is processed by an AI system she was not informed about has had her dignity and agency violated, regardless of the good intentions of the organisation involved. Consent-first, community-accountable AI design is the only design that is ethically defensible in this sector.

Cluster C: Ecosystem Intelligence, Infrastructure and Measurement

The transformation from fragmented to intelligent social economy requires shared infrastructure, not just improved individual organisations. These recommendations address the collective investments that no single actor can make alone.

RECOMMENDATION C1

Invest in Shared Ecosystem Intelligence Infrastructure

Funders and intermediaries should jointly invest in the shared data systems, learning repositories and coordination infrastructure that give the entire social ecosystem better visibility of itself. This means mapping who is operating where, identifying geographic service gaps (only Gauteng receives concentrated CSI; Limpopo, Northern Cape and North West remain persistently under-served), flagging where duplication is creating waste, and making learning accessible across the system.

South African intermediaries - NACF, Tshikululu, the Southern Africa Trust and others - are well-positioned to anchor this infrastructure. But positioning is not the same as resourcing: funders must be willing to invest in the intelligence layer of intermediaries as a strategic priority, not a grudge cost. The Bridgespan 2025 research identifies this as one of the six key practices that distinguish transformational from transactional CSI.

RECOMMENDATION C2

Integrate AI Across the Full Programme Lifecycle — Not Just Visible Functions

The AI Value Chain (Table 1) maps ten stages from design and diagnosis through to exit and sustainability. Organisations consistently over-invest in AI adoption in high-visibility, low-complexity functions - communications drafting, annual report generation, social media content - while the highest-value applications remain underused: real-time adaptive MEL, exit

planning and sustainability modelling, ecosystem diagnostic intelligence, community voice systems, and theory of change synthesis.

Funders should actively encourage and fund AI adoption at MEL and learning stages, not just operational and communications stages. The goal is not efficiency; it is adaptive intelligence - organisations that learn faster, respond more accurately to community need, and accumulate evidence that makes future programmes better.

RECOMMENDATION C3

Connect AI to the Sector's Measurement Standards Infrastructure

AI does not replace IRIS+, SROI, GRI, SDG alignment or the Impact Management Project's Five Dimensions. It makes them more powerful, more accessible and more applicable to organisations that currently cannot afford the consultancy costs of full framework implementation. Funders and implementing organisations should invest in connecting AI tools to the sector's existing measurement infrastructure: using AI-native platforms like b.world and community feedback tools like 60 Decibels as practical entry points.

The goal is not more reporting. It is genuine adaptive learning. Organisations that use evidence to improve programmes in real time, not organisations that produce better-formatted annual reports for funders. This distinction matters, and funders' measurement requirements should reflect it.

RECOMMENDATION C4

Create Protected Space for Experimentation — and Fund Failure

AI adoption in the social sector requires protected space for experimentation, failure and learning, not the performance of AI capability for funders. This means funders explicitly making room in grant terms for AI pilots that do not work. It means intermediaries creating peer learning forums where organisations share failures as openly as successes. It means sector bodies convening regular AI learning exchanges where what has not worked is treated as evidence as valuable as what has.

The social economy has a deeply embedded culture of reporting success to funders and concealing struggle. AI adoption will be slower, more expensive and less effective if this culture persists. Protected space for honest experimentation is not a soft aspiration, it is a practical precondition for collective learning.

Cluster D: Policy, Systems Change and African AI Sovereignty

The social economy does not exist in a policy vacuum. The frameworks governing AI - nationally, continentally and globally - will shape what is possible for the sector for decades. These recommendations address the sector's responsibility to shape those frameworks, not merely comply with them.

RECOMMENDATION D1

Engage Actively with the South African AI Policy Process

South Africa's national AI policy framework is in development. The social economy sector, including NPOs, intermediaries, funders and social enterprises, has both the standing and the obligation to participate in shaping it. The sector serves the communities most likely to be affected by AI-related risks: those with least power to contest algorithmic decisions, least access to legal recourse, and least representation in the technology design processes. This makes the sector's voice in policy formation a community accountability issue, not just an advocacy opportunity.

Sector bodies should engage with the DTIC's AI policy process, the DCDT's digital economy strategy, and the Parliamentary committees overseeing digital rights and social development. They should do so with clear positions on data sovereignty, algorithmic accountability, digital inclusion and the specific regulatory needs of community-facing AI applications.

RECOMMENDATION D2

Prioritise African-Built AI Solutions — Strategically, Not Just Symbolically

Choosing Lelapa AI's Vulavula over a US-based voice platform is not just an ethics statement, it is a data quality decision, a data sovereignty decision, and a sector development decision. African-built AI tools understand African languages, social dynamics and community contexts in ways that tools built on Global North data cannot. They keep data in African hands. And they develop the African AI ecosystem that the sector will eventually need to sustain itself.

This means sector procurement policies should actively prefer African-built solutions - and where none exist for a specific use case, funders should consider co-investing in their development. Masakhane's open-source model library, the Deep Learning Indaba's research community, and Lelapa AI's language infrastructure are all sector investments, not just tech startup stories.

RECOMMENDATION D3

Invest in the African AI Ecosystem as Social Infrastructure

The sector should regard investment in the African AI ecosystem: in Masakhane, in the Deep Learning Indaba, in AI4D Africa, in African language datasets, as social infrastructure investment, equivalent to investing in community health infrastructure or school buildings. Because an African AI ecosystem built on African data, in African languages, governed by African ethical frameworks, is the precondition for an AI-enabled social economy that serves African communities rather than extracting value from them.

This is a long-horizon investment, and it requires foundations, bilateral donors and government DFIs to fund it as such - not as a once-off project, but as a sustained commitment to building the infrastructure of African digital sovereignty.

A Consolidated View of the Recommendations

The narrative above develops each recommendation with context and rationale. Table 20 below provides a reference layer for funders, intermediaries, NPOs, sector bodies and government to locate the recommendations that apply to them and the time horizon involved. Owners are indicative; most recommendations require coordination across multiple actors.

Recommendation	Primary Owner	Time Horizon	Cluster
A1. Adopt a sector-wide AI accountability framework	IPASA, Trialogue, CSI Forum	6–12 months	A
A2. Publish portfolio-level AI usage and outcomes data	Corporate funders	12 months	A
A3. Fund intermediary AI capacity as portfolio infrastructure	Corporate and foundation funders	12–24 months	A
A4. Require POPIA and community-data governance evidence in due diligence	All grant-making funders	6 months	A
A5. Replace one-year extractive grants with multi-year learning grants	Corporate funders	12–24 months	A
A6. Move the institutional-overhead floor to 20–25%	Corporate and foundation funders	12 months	A
B1. Reposition from coordinators to intelligence platforms	Intermediary leadership	12–24 months	B
B2. Establish fractional CTO / AI-lead capacity hosted at intermediary level	Intermediaries with funder support	12 months	B
B3. Build sector-wide AI talent pipeline and exchange programmes	Intermediaries, sector bodies	24 months	B
B4. Develop shared infrastructure for POPIA-compliant tooling	Intermediary consortia	24 months	B
C1. Conduct AI readiness self-assessment as a board-led process	All NPOs and CBOs	6 months	C
C2. Establish AI governance committee with community representation	NPOs with community-facing programmes	12 months	C
C3. Invest in AI literacy across all staff before tool adoption	All NPOs and CBOs	6–12 months	C
C4. Adopt ethical procurement principles for AI tools and vendors	NPOs and intermediaries	12 months	C
D1. Submit a coordinated sector input to the SA national AI policy process	Sector bodies and intermediaries	6 months	D
D2. Build cross-sector AI ethics communities of practice	Sector bodies, academia, civil society	12 months	D

D3. Engage government as both regulator and adopter	Sector bodies, intermediaries	12–24 months	D
D4. Position African voices in global AI governance dialogues	Academia, government, sector bodies	Ongoing	D

Table 20: Consolidated Recommendations Summary

CONCLUSION

The Transition from Fragmentation to Intelligence

The future of Africa's social economy depends not only on funding, programmes or technology - but on our collective ability to build systems capable of learning, adapting and collaborating intelligently. **The transition from fragmentation to intelligence is not primarily a technological challenge. It is a governance challenge, an equity challenge, and - most fundamentally - a question of what kind of social economy we choose to build.**

AI offers tools for seeing patterns we could not previously see, for connecting knowledge that has been siloed, for giving voice to communities that have been under-heard, and for enabling smaller organisations to access the institutional capability that has historically been the privilege of the well-resourced. These are not trivial possibilities. In a sector characterised by chronic resource scarcity and compounding social need, the ability to work more intelligently - to learn faster, coordinate better, adapt more responsively - could translate into millions of lives better supported.

But the opportunity is not guaranteed. AI deployed without African agency will replicate existing inequalities. AI deployed without community consent will deepen the extractive dynamics that the social sector exists to challenge. AI deployed without governance will create new vulnerabilities in systems that are already fragile. And AI deployed as a substitute for adequate and sustained funding will ultimately disappoint.

What is required is intentionality: deliberate choices about whose AI we build, whose data we govern, whose voice we centre, and whose benefit we prioritise. Africa's social economy has always been fundamentally relational - built on trust, on community legitimacy, on the wisdom of context and lived experience. These are not liabilities in an AI-enabled world. They are the deepest and most durable assets the sector possesses.

> *The most effective social economy institutions of the coming decade may not be those with the most advanced technology.* ***They will be those that have most successfully combined AI-enabled intelligence with African relational wisdom, ethical stewardship and genuine community accountability.*** *In that combination lies the possibility of a social economy that is not only more intelligent, but more just.*

The choices made now - about infrastructure, governance, ethics, inclusion and design - will shape the trajectory of this transition for a generation. The moment to make them intentionally is not tomorrow. It is now.

PART FOUR

Appendices: Supplementary Tools Pack

IN THIS PART

> *Organisations cannot govern what they cannot afford to staff. They cannot learn from what they cannot afford to measure.*

APPENDICES

Supplementary Tools Pack

Eight audience-specific resources developed alongside the guide

This section contains eight supplementary resources developed to accompany the main practitioner guide. Each is designed to stand alone as a practical reference for a specific audience. All are grounded in the guide's analysis and frameworks - readers of the full guide will find these tools a practical distillation; readers encountering these tools first will find pointers to the full guide for deeper context.

Tool	Audience
Funder's AI Diligence Checklist	CSI managers and corporate funders
Board's Ten Questions on AI	Board members and governance leads
CBO Companion — Plain Language Guide	Community-based organisations and grassroots practitioners
How to Budget for AI Adoption	All organisations; especially NPO directors and funders
Disintermediation Risk and Opportunity	Intermediaries and ecosystem coordinators
Government as AI Adopter	Policymakers, DFIs and government social sector leads
A Worked Example for MEL Practitioners	MEL practitioners and impact analysts
AI for Social Enterprises	Social enterprise founders and operators

TOOL 1: FUNDER'S AI DILIGENCE CHECKLIST

For CSI managers, philanthropic foundations and grant-making institutions

This checklist supports funders who want to move from passive to active engagement with AI in their grantee portfolios. It covers three moments in the funding relationship: before making a grant (due diligence), during implementation (active grant management), and at the portfolio level (strategic oversight). It is modelled on the guide's Seven-Decision Framework and Governance Checklist in Section Eleven, translated into a funder-facing format.

Why this matters for funders

The single most powerful lever available to funders who want AI to create value in the social economy is not selecting the right AI tools - it is funding the institutional capacity that makes responsible AI adoption possible. Organisations cannot govern what they cannot afford to staff. They cannot learn from what they cannot afford to measure. This checklist helps funders ask the right questions to ensure their grants are enabling intelligent organisations, not just funded ones.

Part A: Pre-Grant Due Diligence — Assessing AI Readiness

Use these questions when evaluating a new grant application or renewing an existing relationship.

A1. Organisational AI Readiness

☐ **Readiness baseline:** Has the applicant completed an AI Readiness Self-Assessment (or equivalent)? Ask them to share it.

☐ **Data governance:** Does the organisation have a documented data governance policy, or is one in development?

☐ **Accountability:** Is there a named staff member responsible for technology and data decisions - even if not a dedicated role?

☐ **Use case clarity:** Has the organisation identified specific, concrete use cases for AI - or is the interest still general and exploratory?

☐ **Infrastructure:** Does the organisation's current technology infrastructure (internet, devices, power) support the AI applications they are considering?

A2. POPIA and Legal Compliance

☐ **Personal data involvement:** Does the applicant's work involve the collection, storage or analysis of personal information about beneficiaries?

☐ **Lawful basis:** If yes - does the organisation have a lawful basis for processing personal information (POPIA s.11), and can they articulate it?

☐ **Cross-border data transfer:** Has the organisation assessed which, if any, AI tools they currently use process data outside South Africa (POPIA s.72)?

☐ **Automated decision-making:** If AI will make or inform decisions about individual beneficiaries - does the organisation have a human review pathway and appeals process (POPIA s.71)?

☐ **Community consent:** Has the organisation obtained or planned for community consent to AI-assisted data analysis of their information?

A3. Ethical and Equity Considerations

☐ **Language equity:** Has the applicant considered whether the AI tools they are using or planning to use work effectively in the languages their communities speak?

☐ **Bias review:** Has any AI tool used in beneficiary-facing work been reviewed for potential bias against the communities being served?

☐ **Error and escalation:** Does the applicant have a plan for what happens when AI makes an error that affects a beneficiary? Is there a human fallback?

☐ **High-sensitivity contexts:** For organisations working in GBV, mental health, child protection or other high-sensitivity areas - has a trauma-informed review of AI tools been conducted?

Part B: During the Grant — Active Portfolio Management

Use these prompts in quarterly check-ins, narrative reporting reviews and site visits.

B1. Implementation Tracking

☐ **Progress against plan:** Is the organisation's AI adoption progressing as planned? If not - what is the barrier (capacity, infrastructure, governance, tools)?

☐ **Policy implementation:** Has the organisation's AI usage policy been implemented and communicated to all staff?

☐ **Human oversight:** Are AI outputs being reviewed by a human before they influence programme decisions or beneficiary interactions?

☐ **Incident tracking:** Has the organisation encountered any AI-related errors, unexpected outputs or governance concerns? How were they handled?

B2. Learning and Adaptation

☐ **Learning documentation:** Is the organisation documenting what works and what does not in its AI applications - and sharing this learning?

☐ **Staff impact:** Has AI adoption created any new staff capacity pressures, resistance or wellbeing concerns that need to be addressed?

☐ **Data-to-decisions:** Has the AI application produced any data or insights that have influenced programme design or funder reporting?

Part C: Portfolio-Level AI Strategy — The Funder's Own Obligations

These questions are for the funder's internal use - not for grantees. They reflect the guide's direct challenge to funders about their own role in enabling or constraining AI adoption.

☐ **Eligible cost categories:** Does our grant budget template allow grantees to include technology investment, AI tool subscriptions and staff AI training as eligible costs?

☐ **Institutional capacity threshold:** Are we applying the principle that 20–25% of grant funding should be available for institutional capacity - which includes AI readiness?

☐ **Readiness as due diligence:** Do we require our grantees to demonstrate AI readiness as part of due diligence - and if not, should we?

☐ **Grant duration:** Are we funding multi-year grants of sufficient duration to allow organisations to invest in AI adoption (minimum 3 years)?

☐ **Shared infrastructure:** Are we investing in shared AI infrastructure through intermediaries in our portfolio — or expecting each small organisation to figure it out alone?

☐ **Portfolio-level AI mapping:** Have we assessed our own portfolio for AI opportunity - which of our grantees could benefit most from AI-enabled MEL, communications, or programme design?

☐ **Funder transparency:** Are we prepared to be transparent with our grantees about how we ourselves are using AI in grant management, reporting analysis and portfolio oversight?

The cost of not asking these questions

Funders who do not actively support AI readiness in their portfolios are not neutral: they are choosing the status quo. The status quo is: underfunded institutional capacity, manual reporting systems, fragmented data, and a sector that cannot learn at the speed the challenges require. AI does not solve this. But funders who fund it well can.

See Section Thirteen, Recommendation A6 of the guide for the full funder accountability framework.

Quick Reference: Key guide sections for funders	
Your stakeholder row in the ecosystem matrix	Section Six, Table 2
Thematic AI applications by portfolio focus	Section Seven
POPIA compliance obligations for funders handling grantee data	Section Nine
AI Readiness Self-Assessment (to share with grantees)	Section Ten
Funder recommendations and accountability framework	Section Thirteen, Cluster A

TOOL 2: THE BOARD'S TEN QUESTIONS ON AI

For board members, trustees and governance leads in NPOs, foundations and social enterprises

A board that has never discussed AI is not a board that has decided against it: it is a board that is governing blind. AI is already influencing how the organisations you govern collect data, produce reports, engage beneficiaries and manage staff. The question is not whether to have a governance conversation about AI. It is whether to have it before something goes wrong.

These ten questions are designed for a board agenda - not as an audit, but as a structured governance conversation. Each question comes with the context you need to ask it well, and a note on what a good answer looks like.

Who should lead this conversation?

The executive director or CEO should present to the board on questions 1–4. Questions 5–8 should involve the finance and audit committee. Questions 9–10 are governance questions for the full board. Boards that have a technology subcommittee should use that forum, but not as a substitute for full board discussion - AI has strategic, ethical and legal dimensions that require the full board's engagement.

Question 1: What AI tools are we currently using — and does the board know?

It is common for AI tools to enter an organisation through individual staff members using free consumer tools before any formal decision has been made. Board members are often unaware this is happening. This question asks management to provide a complete, honest account of current AI use - not just planned or approved use.

What a good answer looks like: A list of tools currently in use, whether formally adopted or informally used by staff, with the use case for each. 'We don't know' is an acceptable starting point - but it should be followed by a commitment to find out.

Question 2: Are any of these tools processing personal information about our beneficiaries?

Under the Protection of Personal Information Act (POPIA), as well as King V guidance on corporate governance and directors responsibilities - the board has ultimate accountability for how the organisation handles personal information. If AI tools are processing beneficiary data, even for internal analysis, the board needs to know. This is not a technical question. It is a governance and legal accountability question.

What a good answer looks like: Clear identification of which tools handle personal data, what the lawful basis for processing is, and whether a Data Processing Agreement is in place with each tool provider. Uncertainty here should trigger an immediate governance review.

Question 3: Is AI being used in any decisions that affect our beneficiaries' access to services?

When AI recommends, screens or ranks beneficiaries, even as a support tool for human decisions, it creates legal and ethical obligations. POPIA section 71 gives individuals the right to human review of automated decisions that significantly affect them. If your organisation's AI tools are influencing who receives services, who is prioritised, or who is referred where, the board needs to understand and govern this explicitly.

What a good answer looks like: A clear account of where AI influences beneficiary-affecting decisions, with confirmation that human review and an appeals pathway exist for every such decision.

Question 4: Do we have an AI usage policy — and has the board approved it?

An AI usage policy is a governance document, not just an operational one. It should specify which tools are approved, what data may be used with each, what is prohibited, and who is accountable for compliance. Boards that leave policy development entirely to management are abdicating a governance responsibility.

What a good answer looks like: A board-approved policy, reviewed within the last 12 months, that covers approved tools, data use permissions, prohibited uses, staff responsibilities and incident reporting. A policy that exists but has not been communicated to all staff is not a functioning policy.

Question 5: Are we budgeting for AI adoption — or expecting it to happen for free?

Responsible AI adoption is not free. It requires staff training, tool subscriptions, governance review time, data governance infrastructure, and - as readiness increases - specialist technical support. Boards that do not create budget space for these costs are implicitly deciding that AI adoption will happen unsupported, which typically means it happens ungoverned.

What a good answer looks like: A budget line or identified capacity for AI development, even if modest. The guide recommends that organisations approaching funders request that at least 20–25% of grant budgets be available for institutional capacity - including AI readiness. Boards should actively support this request.

Question 6: How would we know if something went wrong with an AI tool we are using?

AI tools produce errors. They can generate incorrect information, reproduce biases embedded in their training data, or make recommendations that are inappropriate for a specific community context. Boards should ask: do we have an incident reporting pathway? Would staff feel safe

raising an AI-related concern? Would a beneficiary who was harmed by an AI decision know how to raise it?

What a good answer looks like: A clear incident reporting pathway for AI-related concerns, integrated into existing risk management systems. A culture where staff are encouraged to flag AI errors as learning opportunities, not performance failures.

Question 7: Are our staff equipped to use AI responsibly — and do they know what is expected of them?

AI literacy is not the same as AI enthusiasm. The board should ask whether staff understand what AI can and cannot do reliably, how to critically evaluate AI outputs, what the governance requirements are, and what to do when AI produces something problematic. This is a people question, not a technology question.

What a good answer looks like: Evidence of structured AI literacy investment, even light-touch training on responsible use, with clear expectations communicated to all staff. Not 'our younger staff are good with technology,' which conflates digital fluency with AI literacy.

Question 8: Are we thinking about AI in our strategic planning — or treating it as an IT issue?

AI is not an IT issue. It is a strategic question about how the organisation learns, decides, engages with communities and accounts for its impact. Boards that delegate AI entirely to management's technology decisions are missing the strategic dimension. The next three-year strategy should include a clear position on AI - not a technology roadmap, but a statement of intent about what the organisation will and will not use AI for, and why.

What a good answer looks like: AI is named as a strategic consideration in the current plan, with a clear link to mission, values and the communities the organisation serves. The board has discussed and endorsed this position.

Question 9: Are we confident that our AI use reflects our values?

Every organisation has stated values: dignity, community agency, equity, accountability. AI applications can either reinforce or undermine these values, depending on how they are designed and governed. A board that has not explicitly tested its AI use against its organisational values has not completed its governance responsibility.

What a good answer looks like: A deliberate values-alignment review of AI applications, not a once-off tick-box, but a recurring governance question. The board should be able to explain, if asked by a funder or community member, why each AI application reflects rather than contradicts the organisation's commitments.

Question 10: Do the communities we serve know we are using AI with their data — and have they consented?

This is the question that most boards have not yet asked. Community consent to AI-assisted analysis of their data is both a legal obligation under POPIA and an ethical commitment that reflects the social economy's foundational principle of community agency. It is not sufficient to include AI in a privacy policy that beneficiaries do not read. Genuine informed consent requires that communities understand, in accessible language, what AI is being used for and what they are agreeing to.

What a good answer looks like: Evidence that community consent processes have been updated to include AI-specific language, tested for comprehension in the languages communities speak, and genuinely offered as a choice - not a condition of service access. See the Community Questions Appendix in Section Nine of the guide for the questions communities have the right to ask.

For the board agenda
These ten questions can be structured as a 60-minute board agenda item: Questions 1–4 as a management presentation (20 minutes), Questions 5–8 as a finance and audit discussion (20 minutes), Questions 9–10 as a full board deliberation on values and community accountability (20 minutes). The output should be a documented board resolution on AI governance - including whether a dedicated subcommittee, policy review or external governance review is required.

Quick reference: Key guide sections for boards	
Ethics, risks and governance framework	Section Nine
Minimum Viable AI Governance Checklist	Section Eleven, Table 12
POPIA compliance obligations	Section Ten
Community Questions Appendix	Section Nine, Appendix
Board member competency requirements	Section Twelve

TOOL 3: CBO COMPANION — A PLAIN LANGUAGE GUIDE TO AI

For community-based organisations, grassroots practitioners and home-based care workers

How to use this companion

This guide is written for you - the person running a community crèche, a home-based care programme, a youth club, a soup kitchen or any community organisation. You do not need a technology background to use it. It is in plain language. If anything is unclear, that is our failure, not yours.

The full practitioner guide (From Fragmentation to Intelligence) is long and detailed. This companion gives you the parts that matter most for your organisation right now.

What is AI, in plain terms?

Artificial Intelligence - AI - is a type of computer technology that can do tasks that used to require a human being. It can read and write text, answer questions, analyse information, translate languages, recognise faces and images, and make recommendations.

You have probably already used AI without knowing it. The autocorrect on your phone is a form of AI. The suggestions that come up when you search something online. The WhatsApp messages from MomConnect that pregnant women receive. These are all AI at work.

AI is not magic. It is not a robot. It is not going to replace the work you do with communities. What it can do is take some of the administrative, writing and information-processing tasks off your plate — so you can spend more time on the work that actually requires a human being.

What can AI actually do for a small community organisation?

You do not need to be a big, well-funded organisation to benefit from AI. These are things you can do today, with a basic smartphone and a data connection:

- **Write documents and reports:** AI tools like ChatGPT or Claude can help you draft a funding proposal, an annual report, a letter to a funder or a newsletter to community members. You describe what you want to say, and the tool helps you write it. You review and correct it. This can save hours of writing time.
- **Summarise and make sense of your data:** If you collect data on paper forms - attendance registers, home visit records, beneficiary information - AI tools can help you design better forms and summarise what the data shows.
- **Translate materials into local languages:** You can use AI tools to translate a document from English into isiZulu, isiXhosa, Sesotho or other South African languages. The translation may not be perfect - always have a community member check it - but it is a starting point.
- **Plan and prepare:** AI can suggest agendas, help you write minutes, draft job descriptions, create training materials or help you prepare for a difficult conversation with a funder.

- **Access guidance and support:** Tools like ECD Connect (available on WhatsApp) and Zuzi are designed specifically for community practitioners - they give you access to guidance on child development, nutrition and care in the languages your community speaks.

What AI cannot do

AI cannot replace the trust you have built with your community. It cannot sit with a grieving family. It cannot understand the specific history of your neighbourhood. It cannot make decisions about people's lives on its own - or at least, it should never be allowed to.

When AI produces something that does not feel right for your community, trust that instinct. You are the expert on your community. AI is a tool.

Before you start: three questions to ask yourself

- **Question 1:** Am I going to use any personal information about community members with this tool? If yes: stop and read the section below on protecting people's data before you continue.
- **Question 2:** Does the tool work in the language my community speaks? If not, consider using Vulavula or Masakhane-based tools instead.
- **Question 3:** Do I have reliable internet access? If connectivity is a barrier, look for offline-capable tools - or connect with your intermediary or support organisation about shared access options.

Protecting your community's data — what you must know

South Africa has a law called the Protection of Personal Information Act - POPIA. It applies to your organisation, even if you are small. Here is what it means in practice:

- **Personal information is protected:** Names, ID numbers, health information, addresses, photographs of community members - these are all personal information and are protected by law.
- **Do not share personal data with free AI tools:** You must not upload personal information about community members into free online AI tools. This includes ChatGPT, Google Gemini, or any similar free tool. These tools may store and use your data in ways you cannot control.
- **People have rights over their data:** Community members have the right to know if you are using AI to analyse information about them. They have the right to say no.
- **When in doubt, ask:** Before you use any AI tool that involves community members' information, ask your intermediary or support organisation for guidance.

The safe rule of thumb
If you are using AI only with general information - writing a report that does not name individuals, drafting a proposal, summarising publicly available research - you are in safe territory with most AI tools.

If the information involves real people from your community - their names, their health, their circumstances - do not use a free online tool. Ask for guidance first.

Your first steps: a simple starting plan

You do not need to do everything at once. Here is a safe, simple way to begin:

Step	What to do
Week 1–2	Try one thing. Use a free AI tool (ChatGPT or Claude.ai) to draft one document: a proposal, a report introduction, a newsletter. Use only general information, no personal data. See what comes out.
Week 3–4	Evaluate honestly. Did it save time? Was the output usable? What did you have to correct? Share what you learned with one colleague or peer organisation.
Month 2	Choose one regular task where AI could help - and commit to using it consistently for that task for 30 days. Note what changes.
Month 3	Talk to your intermediary or network about what tools other organisations like yours are using. You do not need to figure this out alone.
Month 4–6	If you work with beneficiary data, talk to your support organisation about POPIA-compliant tools before using AI for data-related work.

Questions your community has the right to ask you

If your organisation uses AI to analyse data about community members, the community has the right to ask you these questions - and you must be able to answer them honestly:

- ☐ What information about us are you collecting?
- ☐ Why are you collecting it, and what will you use it for?
- ☐ Will you share it with anyone else?
- ☐ Is AI being used to analyse information about us?
- ☐ Will AI be used to make decisions about our access to services?
- ☐ If we say no to our data being used, what happens?
- ☐ Who do we contact if something goes wrong?

You are not alone in this
The organisations that have the most to gain from AI - and the most to lose if it goes wrong - are community organisations like yours. The tools, the training, the governance support and the infrastructure investment that would make AI work for you are things that funders and intermediaries must provide. If they are not providing them, that is a gap worth naming.
Use the questions above to start a conversation with your funders and support organisations about what they are doing to make AI accessible and safe for community organisations.

TOOL 4: HOW TO BUDGET FOR AI ADOPTION

A cost framing guide for NPOs, funders and social enterprises - not a price list

Why this guide avoids specific tool prices

AI tool pricing changes rapidly. A subscription that costs X today may cost 2X or 0.5X within six months. Enterprise pricing varies by organisation size, usage volume and negotiation. Listing specific prices would make this section obsolete within months of publication.

What does not change is the structure of what AI adoption costs. This section gives you the framing to build a real budget for your organisation - whatever the current price of specific tools turns out to be.

Part A: The Five Cost Categories of AI Adoption

When organisations budget for AI, they typically think only about tool subscriptions. This is the smallest part of the real cost. A complete budget for AI adoption covers five categories:

Cost category	What it includes	Who often forgets to budget for it
1. Tool costs	Subscriptions, API usage fees, enterprise licensing, per-user costs. Consumer tiers may be free; enterprise tiers that protect your data are not.	Everyone: but this is the category most often included
2. People costs	Staff time for learning, experimenting, reviewing AI outputs and building new workflows. Training and upskilling. Time for governance review. Fractional CTO or human-centred technologist support if needed.	Most organisations: people time is the largest hidden cost
3. Governance costs	Time to develop and implement an AI usage policy. Data mapping (understanding what personal data you hold). POPIA compliance review. Legal advice on data processing agreements.	Small organisations: governance is seen as overhead, not investment
4. Infrastructure costs	Reliable internet access. Capable devices. Power stability (relevant in load-shedding contexts). Data storage that is South Africa-resident if handling personal information.	Rural and peri-urban organisations: most significantly
5. Learning costs	Time to document what works and what does not. Peer learning and knowledge sharing. Iteration - the cost of trying something, finding it does not work, and trying again.	Everyone: iteration is funded in almost no AI budgets

Part B: Indicative Budget Scales by Readiness Level

The following framing is indicative, it describes the relative scale and composition of AI budgets, not specific amounts. Actual costs depend on organisation size, sector, existing infrastructure and the specific tools chosen.

Readiness level	Budget composition	What the budget mostly funds
Level 1 — Awareness (0–6 months)	Mostly people time; minimal tool costs. Consumer-tier tools are often free. The dominant cost is staff time to experiment, learn and evaluate.	Staff time for learning and experimentation. Possibly one externally facilitated AI literacy session.
Level 2 — Experimentation (6–18 months)	People and governance costs dominate. Tool costs begin to appear as paid tiers become necessary for POPIA-compliant enterprise use. Governance development requires dedicated time or external support.	Governance policy development. Paid tool subscriptions for approved enterprise tools. Possible fractional human-centred technologist support (part-time, project-based).
Level 3 — Integration (18 months–3 years)	Infrastructure and people costs grow significantly. Multiple tool subscriptions. Possible data infrastructure investment. Ongoing staff development. A dedicated AI champion or part-time data role.	Staff capability development. Multiple enterprise tool subscriptions. Possible data governance infrastructure. Fractional CTO-level support for architecture decisions.
Level 4 — Institutionalisation (3+ years)	Full cost structure across all five categories. Permanent staff roles. Custom integrations. Organisation-wide governance framework. External evaluation of AI impact.	Permanent AI-enabled roles. Custom tool configurations. External impact evaluation. Contribution to sector-wide learning.

Part C: The Cost of Inaction

The cost of not adopting AI is real: it is simply less visible than a tool subscription invoice. Organisations and funders who are calculating whether AI adoption is 'worth it' should factor in what inaction costs.

The inaction cost frame

For every hour a programme officer spends manually compiling a donor report that AI could draft in minutes, that is an hour not spent on community relationships, programme design or learning.

For every evaluation that sits unread because synthesising it manually is too time-consuming, that is learning the sector paid for and never used.

For every small NPO that cannot demonstrate its impact because it cannot afford MEL infrastructure, that is an organisation at risk of losing funding - not because its work is poor, but because it cannot make its work visible.

For every community member who cannot access health guidance in their first language because the tools were only built in English, that is a person the sector has failed to reach.

Inaction cost	What it looks like in practice
Administrative burden	Staff spending 20–30% of their time on tasks AI could assist with: proposal drafting, report compilation, data entry, translation
Learning loss	Evaluation reports produced at significant cost and never synthesised or shared: learning trapped in unread PDFs
Impact invisibility	Smaller organisations unable to demonstrate their impact to funders because they lack MEL infrastructure: creating competitive disadvantage unrelated to programme quality
Capability inequality	Large, well-funded organisations gaining further advantage through AI while smaller organisations fall further behind: widening the institutional capability gap
Language exclusion	Beneficiaries excluded from digital services, health guidance and information because tools are English-only: continuing the pattern that has always characterised under-investment in African language infrastructure
Missed funding	Organisations that cannot produce AI-assisted evidence of impact losing out to competitors who can: regardless of comparative programme effectiveness

Part D: How to Ask Funders for AI Budget

Many organisations know they need to budget for AI but are unsure how to raise it with funders. Here is a framing that is grounded in the guide's evidence base:

- **Frame AI investment as institutional capacity, not technology cost**. 'We are requesting support to build our data and learning infrastructure' is a more fundable proposition than 'we want to subscribe to AI tools.'
- **Be specific about the use case and the expected return.** 'AI-assisted MEL synthesis will allow us to analyse 12 months of community feedback data in hours rather than weeks, improving our programme responsiveness' is a concrete, fundable claim.
- **Reference the guide's recommendation that 20–25% of grant budgets should be available for institutional capacity - which includes AI readiness.** This gives you an evidence-based reference point for the conversation.
- **Ask funders explicitly whether AI tool subscriptions and staff AI development time are eligible costs in their grant budgets.** Many funders have not yet decided, your question may prompt them to open this up.
- **If a funder says no to AI-specific budget, ask whether they fund organisational learning and adaptation.** AI-assisted learning infrastructure is a reasonable interpretation of that category.

TOOL 5: DISINTERMEDIATION — RISK AND OPPORTUNITY FOR INTERMEDIARIES

For intermediaries, ecosystem coordinators, umbrella bodies and capacity-building organisations

Disintermediation

The guide identifies disintermediation as a primary risk for South African intermediaries in an AI-enabled social sector. This tool develops that analysis in full - describing what the risk is, where it is real versus overstated, and what the strategic response looks like for an intermediary that wants to be essential to the ecosystem rather than bypassed by it.

What disintermediation means — and does not mean

Disintermediation is the removal of an intermediary from a transaction or relationship that previously required it. In financial services, digital banking disintermediated the bank teller. In the media sector, the internet disintermediated the newspaper. In logistics, platforms like Uber disintermediated the dispatcher.

In the social economy, the disintermediation risk for intermediaries is specific: if AI tools allow funders to assess grantee impact directly, and implementing organisations to access strategic capability independently, the traditional intermediary functions of data aggregation, capacity building and knowledge brokering may become less necessary.

This risk is real. But it is not uniform, and it is not inevitable.

WHERE DISINTERMEDIATION RISK IS HIGHEST

Functions that are primarily information-passing - grant reporting aggregation, compliance checking, standardised capacity building content - are most exposed to AI-enabled disintermediation. If an intermediary's primary value is consolidating reports that AI can now process directly, or delivering training content that AI can now personalise on-demand, that value proposition is at risk.

WHERE DISINTERMEDIATION RISK IS LOWEST

Functions that require trusted relationships, contextual knowledge, community embeddedness and systems navigation are least exposed. No AI tool can replicate an intermediary's decade of relationships with a funder, its understanding of the political economy of a specific sector, or its capacity to convene actors who would not otherwise sit in the same room.

A Strategic Risk Map for South African Intermediaries

Intermediary function	Disintermediation risk	Strategic response
Grant reporting aggregation	**HIGH**: AI can aggregate and analyse grantee data at scale, directly	Move from data aggregation to data intelligence: synthesise, interpret and act on data, not just collect it
Standardised capacity building and training delivery	**HIGH:** AI can deliver personalised, on-demand training content at zero marginal cost	Shift from content delivery to learning facilitation: AI delivers content; intermediaries facilitate application, reflection and peer learning
Compliance checking and M&E verification	**MEDIUM:** AI can automate rule-based compliance checks; judgment-dependent cases still need human review	Become the governance authority: set the standards, train the reviewers, adjudicate edge cases
Knowledge brokering - connecting practitioners to research and to each other	**MEDIUM:** AI can surface relevant research; it cannot build the trust that makes peer learning work	Invest in community intelligence: curated, contextualised knowledge for specific ecosystems, not generic content aggregation
Strategic advice and organisational development	**LOW**: contextual, relational, adaptive; hard to automate	Double down on this as the core value proposition; use AI to free up staff time for higher-value advisory work
Ecosystem convening and systems navigation	**LOW:** trust, relationships and system knowledge are non-automatable	Position as the ecosystem intelligence hub: the place where AI-generated insights become collective action
Grant-making and fund management	**LOW:** judgment, relationships and accountability remain human	Integrate AI into the diligence and portfolio analysis functions, while maintaining human decision authority

The Opportunity Frame: From Coordinator to Intelligence Hub

The guide argues that disintermediation risk is real for intermediaries that stay in the coordinator role, **but that AI creates a genuinely new and more valuable role: the ecosystem intelligence hub.** This is not a minor rebranding of what intermediaries already do. It is a substantive repositioning.

An ecosystem intelligence hub:

- **Operates shared AI infrastructure that individual NPOs cannot afford or justify alone**: providing access to enterprise-grade tools, data governance frameworks and technical support across a portfolio of organisations
- **Synthesises learning across the ecosystem in real time**: not producing an annual report, but maintaining a live, AI-assisted picture of what is working, what is not and where the gaps are
- **Functions as the POPIA governance anchor for smaller organisations in its network**: providing data processing agreement templates, consent framework guidance and compliance support that individual CBOs cannot develop independently

- **Trains and develops the AI literacy of practitioners across its ecosystem**: not as a once-off workshop, but as an ongoing learning infrastructure
- **Translates ecosystem intelligence into advocacy**: using AI-generated evidence to make the case for systemic change with funders, government and the public

What this repositioning requires

The transition from coordinator to intelligence hub is not automatic and is not cheap. It requires:

Requirement	What it means in practice
Mandate clarity	An explicit board decision that the intelligence hub function is core to the intermediary's mission - not a side project or a funder-pleasing initiative
Data infrastructure investment	A shared data platform capable of aggregating, anonymising and analysing data across the NPO network - requiring capital investment and ongoing maintenance
Technical capacity	At minimum a human-centred technologist; at scale, a data analyst or data scientist with social sector knowledge
Governance frameworks	Data sharing agreements with NPO partners. POPIA compliance architecture. Community consent frameworks that apply across the portfolio.
Funder commitment	Multi-year core funding that supports the intelligence hub function itself - not just the programme delivery of member organisations
Staff capability	An AI literacy programme for intermediary staff that moves from awareness to proficient use across all functional areas

Questions intermediaries should be asking right now

☐ Which of our current functions are most exposed to AI-enabled disintermediation in the next three years?

☐ Which of our functions depend on relationships, contextual knowledge and trust that AI cannot replicate?

☐ Are we currently positioned to become a shared AI infrastructure provider for smaller organisations in our network?

☐ Do we have a mandate from our board and funders to invest in the intelligence hub function - or are we implicitly expected to stay in the coordinator role?

☐ What would it take to develop a shared data platform for our portfolio? Who would fund it?

☐ Are we having honest conversations with our funders about the cost of staying relevant in an AI-enabled ecosystem?

The existential version of this question

Intermediaries that do not actively reposition toward intelligence hub functions risk becoming the sector's most expensive middle layer: collecting reports that AI can aggregate, delivering training that AI can personalise, and brokering connections that platforms can now facilitate directly.

The social economy needs intermediaries. But it needs intermediaries that are essential - that provide genuine intelligence infrastructure, not just administrative coordination. The AI era is a forcing function for that repositioning. Organisations that treat it as a threat will be disintermediated. Organisations that treat it as an invitation will become more valuable than ever.

TOOL 6: GOVERNMENT AS AI ADOPTER — A POLICYMAKER'S REFERENCE

For policymakers, government social sector leads, DFIs and government departments with social mandates

Government

The government is both the social economy's largest funder and its most influential regulator. How the government approaches AI in the social sector - as funder, as service deliverer, as regulator and as employer - will shape the trajectory of the entire ecosystem. This reference addresses four distinct government roles, the specific risks and obligations each carries, and the actions that would most accelerate responsible AI adoption across South Africa's social sector.

Role 1: Government as Social Service Deliverer

Government departments: Health, Social Development, Basic Education, Human Settlements, deliver services to millions of South Africans through systems where AI is already present or imminent. The SASSA case in the guide (Section Six) is the clearest example: biometric verification, automated means testing and digital life certification at national scale, affecting more than 26 million grant recipients.

The governance stakes in this role are the highest in the sector. When government AI goes wrong at scale, the consequences fall on the country's most vulnerable households: people with limited ability to seek recourse and limited awareness that an algorithm has shaped their access to services.

AI application area	Primary governance obligation
Beneficiary eligibility and means testing	Human review pathway for all adverse automated decisions (POPIA s.71). Accessible appeals process in all 11 official languages. Bias audit before deployment and annually thereafter.
Biometric identity verification	Fallback pathway for individuals whose biometrics fail - elderly, disabled, workers with physical disabilities. Privacy impact assessment. Data retention limits.
Service allocation and prioritisation	Algorithmic transparency: affected individuals must be able to understand the basis of decisions affecting them. Community consultation before deployment in specific geographies.
Community health worker deployment and routing	Human override authority for CHWs. Explainability of routing recommendations. Clinical supervision of AI-assisted triage.
School performance analytics and learner tracking	Parental consent for individual-level tracking. Prohibition on using AI outputs as a substitute for teacher judgment in assessment. Protection of minor data under POPIA.

Role 2: Government as Grant-Maker and CSI Regulator

The government shapes how corporate social investment flows through B-BBEE legislation, the Companies Act and the NPO Act. This regulatory role gives the government the ability to create or remove incentives for responsible AI adoption in the social sector - without a single rand of direct expenditure.

- **B-BBEE alignment:** Amend B-BBEE Skills Development and Socio-Economic Development scorecards to explicitly recognise AI literacy training and AI infrastructure investment as qualifying expenditure. Currently, most AI investment falls into definitional gaps that make it difficult for corporates to claim against their B-BBEE commitments.
- **NPO Act guidance:** Issue guidance to NPO sector funders clarifying that technology investment, AI tool subscriptions and staff AI development costs are legitimate grant expenditure - removing the ambiguity that causes many funders to refuse these budget lines.
- **National framework:** Develop a National Social Sector AI Framework - a policy document that sets minimum standards for AI governance in government-funded social programmes, analogous to the National Evaluation Policy Framework. This would create a governance floor for the entire sector.
- **Tax incentive alignment:** Establish a Section 18A precedent for AI infrastructure donations - enabling corporates to receive tax benefits for funding shared AI infrastructure for the NPO sector (shared data platforms, AI literacy programmes, governance frameworks), as they can for direct charitable donations.

Role 3: Government as Employer of Social Sector Practitioners

The government is the largest employer in South Africa's social sector: employing social workers, community health workers, teachers, nurses, public servants and ECD practitioners at national scale. The AI transition for government employees is a workforce development challenge of significant scale and consequence.

The guide's workforce analysis (Section Twelve) has direct application here:

- **Frontline practitioners:** Social workers and community health workers need AI literacy: specifically the ability to critically evaluate AI outputs, maintain human judgment in high-stakes casework, and use AI for administrative tasks without delegating clinical or social assessment to algorithms.
- **Data and MEL practitioners:** Government MEL and data practitioners need training in AI-assisted evidence synthesis, automated data quality checking and AI-native reporting tools - without which the government's own evidence-based policymaking capacity will fall behind that of well-resourced private sector actors.
- **Leadership:** Senior officials and political principals need AI governance literacy: not technical expertise, but sufficient understanding to ask the right questions of system designers, procurement officials and technology vendors. The Board's Ten Questions tool (Tool 2 in this pack) is relevant here.

- **Education and training pipeline:** TVET colleges and public universities producing social sector graduates need AI-integrated curricula - producing practitioners who enter the workforce with baseline AI literacy as a standard professional competency.

The relevant SETAs - ETDP SETA for education and training practitioners, Health and Welfare SETA for health and social development workers - have a direct mandate to fund AI literacy development for these workforces. Sector Skills Plans for the 2026–2031 period should explicitly prioritise AI literacy as a scarce skill.

Role 4: Government as AI Policy Actor

South Africa does not yet have a comprehensive national AI policy framework. The Presidential Commission on the Fourth Industrial Revolution (PC4IR) produced recommendations in 2020, but implementation has been uneven and the social sector has been largely absent from national AI governance conversations. This is a gap with significant consequences.

Policy gap	Consequence for the social sector
No national AI ethics framework	Each department and NPO invents its own governance approach - or adopts none. No floor of minimum standards.
No procurement standards for AI in social services	Government departments procuring AI tools for social service delivery have no sector-specific guidance on ethical requirements, community consent standards or bias testing obligations.
No requirement for algorithmic transparency in government social programmes	Beneficiaries of government AI systems cannot understand how decisions affecting them were made - contrary to the spirit of POPIA s.71 and constitutional rights to just administrative action.
No national AI literacy investment framework	Skills development for the public sector's AI transition is fragmented, unfunded and uncoordinated - with no national standard for what AI literacy means for different roles.
No community participation mechanism in government AI governance	Communities most affected by government AI systems have no formal channel for input into their design, governance or evaluation - reproducing the extractive dynamic that has characterised government technology procurement.

Five actions that would accelerate responsible AI adoption across the social sector

- **Within 12 months:** Develop a National Social Sector AI Governance Framework: setting minimum standards for AI use in government-funded social programmes, developed through consultation with the NPO sector, intermediaries and community representatives.
- **Within 6 months:** Direct ETDP SETA and HWSETA to include AI literacy as a priority scarce skill in their 2026–2031 Sector Skills Plans, with associated learnerships and bursary programmes.
- **Within 12 months:** Issue POPIA guidance specific to AI in the social sector: clarifying how consent, data minimisation, cross-border transfer and automated decision-making obligations apply to the most common AI use cases in government and NPO social service delivery.

- **Within 6 months:** Establish an AI in Social Services Interdepartmental Working Group: convening Social Development, Health, Basic Education, the Presidency's Data Science Unit and the Information Regulator to coordinate AI governance across the social sector.
- **Annual commitment:** Commission and publish an annual Social Sector AI Readiness Report: tracking AI adoption rates, governance compliance, community participation levels and equity gaps across the NPO and government social sectors. This should be produced with and for the sector, not about it.

A note on surveillance risk

Government's AI governance obligations are not only about efficiency and service delivery. They are also about power. AI tools deployed by the government in social service delivery create the technical infrastructure for surveillance of vulnerable populations - particularly in the context of grant eligibility verification, community health monitoring and school performance tracking.

The social economy's experience of government data systems includes a long history of data collected for service purposes being used for control, exclusion or policing. AI amplifies this risk significantly. A national social sector AI governance framework must include explicit prohibitions on secondary use of beneficiary data, meaningful community consent mechanisms, and independent oversight of government AI systems affecting vulnerable populations.

This is not a theoretical risk. It is a governance requirement that must be built in from the outset — not retrofitted after deployment.

Quick reference: Key guide sections for policymakers	
Government row in the stakeholder matrix	Section Six, Table 2
SASSA case study	Section Six
National AI policy landscape and gaps	Section Two
Thematic applications across government social portfolios	Section Seven
Workforce development for government social sector	Section Twelve
Recommendations addressed to government	Section Thirteen, Cluster D

Tool 7: A Worked Example for MEL Practitioners: From Manual Synthesis to AI-Assisted Outcome Mapping

Why this tool exists

The case for AI in monitoring, evaluation and learning is easy to make in the abstract: synthesis at scale, pattern detection across years of qualitative data, faster reporting cycles, more time for learning. The case is much harder to translate into a Monday-morning workflow. MEL practitioners in this sector live with mixed data of variable quality, donor reporting timelines that do not pause for innovation, and the entirely reasonable suspicion that any tool described as transformational probably means more work, not less.

This tool meets that suspicion head-on. It walks one composite South African NPO — Imbasa Community Trust through the Seven Decision Framework introduced in Section Eleven, then through a side-by-side manual-versus-AI-assisted SROI on one of its flagship programmes. The numbers, dates and direct quotes are illustrative. The workflow is not. It mirrors the actual sequencing, time costs and judgement calls a mid-sized NPO M&E team would face in 2026.

Imbasa is a fiction. The dilemmas are not.

Setting the scene: who Imbasa is

Imbasa Community Trust is a composite mid-sized social-economy organisation: Eastern Cape-headquartered, with delivery partners in Gauteng and the Western Cape; fifteen full-time staff; an annual operating budget of approximately R8 million. It runs two flagship programmes: **The Imbasa ECD Practitioner Programme**, which delivers site-based training and mentoring for early childhood development practitioners across 200 community ECD sites, and the **Imbasa Youth Civic Action Programme**, which supports youth-led community organising in three peri-urban districts. The M&E function consists of two people: a Senior M&E Lead and a part-time research assistant. Funding comes from three large foundation grants, one corporate CSI partnership and small recurring grants from public foundations.

Imbasa's M&E challenge is the one most of the sector recognises. Three years of mixed-method data - quantitative practitioner tracking, qualitative practitioner interviews, parent and community feedback sessions, programme observation forms, donor narrative reporting - sit across SharePoint, KoBo Toolbox, three donor reporting templates and approximately 400 audio recordings from community feedback sessions, two-thirds of which have never been transcribed. The M&E Lead estimates that around 60% of her time and 80% of the research assistant's time goes to extracting, formatting and synthesising data for donor reports - not to analysing it for programme improvement or sharing it back with the communities and practitioners who generated it. The donors she serves include three different reporting frameworks: one IRIS+-aligned, one SROI-style, one bespoke. She is, by her own admission, drowning in evidence and starved of insight.

This is the situation in which AI has been suggested. The board chair sent her an article. The corporate funder is asking whether the team has an AI strategy. The Executive Director wants to know what is realistic before the next reporting cycle begins in eight weeks.

STEP 1

Running Imbasa Through the Seven Decision Framework

The Seven Decision Framework in Section Eleven is meant to be used precisely in this situation — when the pressure to do something with AI arrives before any consideration of whether AI is the right answer to the actual problem. The gates, applied to Imbasa, look like this.

1

GATE 1

Is the problem clearly defined?

Imbasa's problem is that it does not have enough data. It has too much data the team cannot synthesise fast enough. The clearly defined problem is synthesis of mixed-method evaluation data into formats that serve both donor reporting and programme learning, with a current bottleneck of approximately 1,800 person-hours per year tied up in extraction and formatting. Naming the problem this precisely changes what AI is being asked to do — not generate insights, but accelerate synthesis.

PASS

2

GATE 2

Is AI actually the right answer?

For synthesis of large volumes of qualitative text and audio, large language models are demonstrably effective. For pattern detection across structured tracking data, machine learning approaches have a stronger evidence base than generative AI. Imbasa's problem is primarily a synthesis problem, with a secondary pattern-detection layer. Both are addressable, but with different tools, and an honest answer is that the synthesis half is mature and ready; the pattern-detection half is exploratory.

PASS, WITH SCOPING

Address synthesis first; treat pattern detection as a learning project, not a deliverable.

3

GATE 3

Do we have the data foundation?

Imbasa has more raw data than it can use, but the data foundation has known weaknesses. The 400 audio recordings are unevenly tagged. The practitioner tracking data has changed schema twice in three years. Community feedback session transcripts use inconsistent identifiers. AI synthesis tools amplify whatever data structure they receive — garbage in, structured garbage out. Before any AI

synthesis happens, roughly six weeks of data preparation work is required: standardising identifiers, completing transcription, tagging consents from the recording sessions.

PASS, WITH PREREQUISITE

Six-week data preparation phase precedes any AI synthesis.

4

GATE 4

Have we done the POPIA and ethics work?

This is the hardest gate. The audio recordings contain identifiable voices of practitioners, parents and — in some cases — children. The original consents covered programme evaluation, not AI processing. Imbasa cannot send these recordings to a third-party cloud-based AI tool without either obtaining additional consent or using on-premise / private deployment models that keep the data inside South Africa. The Esethu Framework principles referenced in Section Four become directly relevant: who benefits from the synthesis, and how does the community see the result of its own data?

CONDITIONAL PASS

Synthesis proceeds with structured text data immediately; audio synthesis is paused pending consent renewal and a POPIA-compliant deployment decision.

5

GATE 5

Do we have the human capacity?

The M&E Lead has good prompt-engineering instincts but no formal AI training. The research assistant is willing but junior. Neither has time to learn a new tool from scratch in eight weeks. The realistic capacity strategy is to bring in a fractional MEL-AI specialist for an initial 12-week engagement covering tool selection, prompt design, quality assurance routines, and team upskilling — not to hire a new full-time role.

PASS, WITH FRACTIONAL CAPACITY

This recommendation aligns with Recommendation B2 in Section Thirteen.

6

GATE 6

Have we set realistic success criteria?

A reasonable first-year target: reduce extraction and formatting hours by 40% (from ~1,800 to ~1,080) while maintaining donor reporting quality scores. Generate one quarterly programme learning brief

that the team would not previously have had time to produce. Share two community feedback summaries back with the communities that generated them. Note what success does not include: discovering new programme insights AI cannot do that reliably, and any goal that depends on it is set up to fail.

PASS, WITH EXPLICIT NON-GOALS

7

GATE 7

Have we planned for failure?

What happens if the AI summarises something inaccurately and Imbasa reports it to a donor? The team needs a quality-assurance protocol — every AI-generated synthesis is reviewed against source documents by the M&E Lead before it leaves the building. What happens if the fractional MEL-AI specialist becomes unavailable? Documentation of prompts and protocols must live in Imbasa, not in the consultant's head. What happens if a donor objects to AI involvement in their reporting? Disclosure language goes into reporting templates from day one. What happens if a community member objects to their words being processed by AI? An opt-out pathway is documented and respected.

PASS, WITH DOCUMENTED PROTOCOLS

Project proceeds.

STEP 2

Manual versus AI-Assisted SROI on the ECD Practitioner Programme

With the framework cleared, the M&E Lead chooses one specific deliverable as the pilot: the upcoming six-monthly Social Return on Investment review of the Imbasa ECD Practitioner Programme. SROI is a useful test case because it requires both quantitative analysis (input costs, financial proxies, outcome quantification) and qualitative synthesis (practitioner narratives, parent feedback, observed practice change). It also tends to consume disproportionate evaluator time — the M&E Lead estimates this SROI takes around 14 person-days of effort when done manually.

What follows is a side-by-side comparison of the same SROI delivered through Imbasa's current manual workflow and through an AI-assisted workflow following the protocols above. Both workflows produce the same final deliverable. They differ in where human time is spent.

STAGE 1

Defining scope and stakeholders

MANUAL	AI-ASSISTED
1 day	1 day

Both workflows require the same upfront thinking: which stakeholders are included, what outcomes count, what time horizon. AI adds nothing here — this is judgement work.

STAGE 2

Gathering and structuring data

MANUAL	AI-ASSISTED
3 days	1.5 days

Manual workflow: extract practitioner data from KoBo, pull observation records from SharePoint, transcribe priority community feedback sessions, reconcile inconsistent identifiers — all by hand. AI-assisted workflow: structured text data is fed to a POPIA-compliant deployment, with prompts that extract, reconcile and structure the data against a predefined schema. The M&E Lead spot-checks 10% of records. The remaining audio synthesis is parked pending consent renewal.

Saving: ~1.5 days. Caveat: requires structured prompts built once, then reusable. The first cycle takes longer because the prompts have to be designed.

STAGE 3

Outcome mapping and financial proxy selection

MANUAL	AI-ASSISTED
3 days	2.5 days

AI can suggest financial proxies from established SROI databases and surfaces where similar programmes have used which proxies. The M&E Lead retains the decision. AI reduces the search time; it does not replace the judgement.

Saving: modest but real — around half a day.

STAGE 4

Qualitative narrative synthesis

MANUAL	AI-ASSISTED
4 days	1 day

This is where the biggest time saving sits. AI synthesis of 60+ practitioner interview transcripts into themed narratives is fast, and — with careful prompt design — high quality. The M&E Lead spends one day reviewing the synthesis against source documents, identifying where the AI has flattened nuance or missed a minority voice, and adding the editorial judgement that makes the narrative resonate with a South African reader rather than a generic one.

Saving: ~3 days. Critical risk: AI flattens minority voices and consensus-confirms majority views. The review step is non-negotiable; the cycle saves 3 days only if quality review is built in.

STAGE 5

Calculation, sensitivity analysis and write-up

MANUAL	AI-ASSISTED
3 days	2 days

Calculation itself is a spreadsheet job. AI helps with sensitivity analysis prompts and write-up drafting. The M&E Lead retains final write-up, but the first-draft load is reduced.

MANUAL TOTAL	AI-ASSISTED TOTAL	SAVED PER CYCLE
14 person-days	8 person-days	6 person-days

Across four cycles per year, this is 24 person-days — around 12% of the M&E Lead's annual capacity, redirected from extraction to learning and community return..

What this worked example actually proves

It does not prove that AI transforms MEL. It proves that, when applied with discipline to a clearly defined synthesis bottleneck, AI redirects approximately 12% of a senior evaluator's annual time from extraction to judgement work. That is not a transformation. It is a meaningful operational saving in an environment where MEL teams are systematically under-resourced. It is also entirely contingent on six things: structured data, completed POPIA work, human review of every output,

transparent disclosure to donors and communities, fractional specialist support in the first cycle, and a set of success criteria that does not include things AI cannot reliably do.

The single biggest lesson is the one the framework is designed to surface: most of the value sits in the gate work that happens before any tool is opened. Imbasa's eight-week pre-AI preparation - data structuring, consent renewal, governance documentation, fractional specialist scoping, success criteria definition - is where this project will succeed or fail. The tools are interchangeable. The discipline is not.

Tool 8: AI for Social Enterprises: Revenue, Sustainability and the Mission Tension

Why this tool exists

Social enterprises occupy an uncomfortable position in this guide. They are mentioned in the value chain. They appear in the stakeholder matrix. Aerobotics and GROW ECD are referenced as case material. But until this point, the guide has spoken primarily to organisations whose revenue model is grant-dependent and whose accountability flows upwards to funders. Social enterprises sit differently. Their accountability flows upwards to investors, sideways to customers, and inwards to a social mission that has to survive contact with commercial reality.

The AI opportunity and risk look different from this position. The opportunity is sharper: AI can move commercial metrics — unit cost, conversion rates, demand forecasting accuracy, customer service throughput — in ways that materially extend runway and reach. The risk is more concentrated: the same tools that improve commercial performance can, applied without care, accelerate mission drift. A social enterprise that uses AI to optimise route efficiency at the expense of serving the customers nobody else will is no longer a social enterprise; it is a business with a logo.

This tool is for organisations whose model requires both sides of that balance to hold. It addresses where AI genuinely moves the commercial needle, where it threatens the mission, and what discipline is needed to capture the first without paying the second.

Where AI moves the commercial needle for South African social enterprises

The honest commercial applications fall into five categories. Each is supported by demonstrated practice in comparable South African or African operating environments, with caveats for context.

01

Demand Forecasting and Inventory

Social enterprises operating in agriculture, food systems, last-mile distribution or skills training all face demand-pattern problems that machine learning approaches address well — the data is structured, the patterns are real, and the cost of getting it wrong is concrete. Aerobotics has built its position partly on this terrain. For smaller social enterprises, the entry point is not to build proprietary models but to use the demand-forecasting features now standard in operational software (from inventory tools to logistics platforms), and to redirect freed-up working capital into mission delivery.

02

Customer Service and Engagement at Scale

Social enterprises serving informal or peri-urban markets often face a customer service problem that traditional call centres cannot economically solve: small ticket sizes, multiple

South African languages, communication via WhatsApp rather than email. Large language models with strong multilingual capability — including models like InkubaLM where the languages are supported — now make AI-augmented WhatsApp customer service economically viable at price points that were not available even two years ago. The mission application is critical: this is not about reducing human contact, but about extending service to customers who were previously priced out of reach.

03

Marketing Copy, Content and Design

For enterprises with thin marketing budgets, generative tools cut content production costs substantially. The discipline required: brand voice and cultural authenticity are not features the tools handle on their own, and a social enterprise that ships AI-generated marketing copy without rigorous review will produce work that sounds identical to every other enterprise using the same tools. The mission value is to make brand work affordable enough to be done at all; the mission risk is sounding generic.

04

Pricing Optimisation and Segmentation

This is the most ethically charged application on the list. AI pricing tools can identify willingness-to-pay patterns and segment accordingly. For a commercial enterprise, this is straightforward. For a social enterprise whose mission includes serving customers who are price-sensitive by definition, it is fraught. The legitimate use case is identifying customer segments who could pay more to subsidise those who cannot. The illegitimate use case is charging vulnerable customers more because the model identifies they have fewer alternatives. The governance line between the two is the social enterprise's mission statement, and any pricing optimisation work without an explicit, board-approved framing of where that line sits is a mission risk waiting to materialise.

05

Back-Office Cost Reduction

The least exciting and most reliable application. Bookkeeping automation, document processing, contract review, recruitment screening — the back-office tasks that consume disproportionate management time in small enterprises are precisely the tasks AI tools handle well in 2026. The mission relevance is indirect: every hour of management time saved on bookkeeping is an hour available for the customers the enterprise exists to serve.

The scale-versus-mission tension AI introduces

The hardest question for a social enterprise considering AI adoption is not technical. It is whether the efficiency gains AI delivers are compatible with the friction the mission requires.

Consider three examples.

Childcare Logistics

A social enterprise providing childcare logistics in low-income communities can use route optimisation to lower per-trip cost. But the route that maximises efficiency may bypass the most isolated households — the ones the enterprise exists to reach.

Skills-Training Enterprise

A skills-training enterprise can use AI screening to reduce drop-out rates by selecting candidates with higher predicted completion probabilities. But the candidates most likely to complete are also, statistically, the candidates with the most existing support — the opposite of who the enterprise was founded to serve.

Agricultural Enterprise

An agricultural enterprise can use AI-powered yield optimisation to advise farmers, but the smallholders most in need of advice are the ones with the least data, and the AI will inevitably serve those it can model better.

In each case, the AI is doing exactly what it was asked to do. The problem is what was asked. The discipline a social enterprise needs is to define the mission constraint before the optimisation target - to specify which customers must be reached, which communities cannot be deprioritised, which thresholds of inclusion must be met - and to build those as hard constraints in the brief given to any AI tool, not as soft preferences the model is free to relax in pursuit of efficiency.

A four-question test before any AI adoption decision

Before any commercial AI application is deployed, the social enterprise leadership team should be able to answer four questions on the record.

Question 1: Mission Preservation

The mission test. If this AI application succeeds on every commercial metric, will it still be true that the enterprise is reaching the customers it was founded to reach? If the honest answer is uncertain, the application is not ready for deployment.

Question 2: Equity Safeguards

The equity test. Are there customers, communities or geographies whose access will materially degrade if this application is optimised for commercial efficiency? If so, what is the operational guarantee that protects them — not a hope, an operational guarantee.

Question 3: Customer Transparency

The disclosure test. Customers and beneficiaries have a right to know when AI is in the loop on decisions that affect them. What does the enterprise tell them, in plain language, and through what channel?

Question 4: Data Dignity

The data test. South African social enterprises serving low-income customers are processing personal information from people whose alternatives are limited. POPIA compliance is a floor, not a ceiling. What additional safeguards are in place because the customers cannot easily go elsewhere?

Three South African pathways for social enterprise AI adoption

Different starting points require different sequences. The patterns below are not prescriptive - they are starting templates for three common South African social enterprise positions.

PATHWAY A

Early-stage enterprise with limited revenue

Start with back-office cost reduction and marketing content. These have low risk to mission, deliver quick operational savings, and build internal AI literacy without committing the enterprise to vendor relationships or operational dependencies it cannot exit. Defer demand forecasting, customer-facing AI and any pricing work until at least 12 months of operational AI experience exists in the team.

PATHWAY B

Established enterprise scaling delivery

Demand forecasting and logistics optimisation are the natural entry points. The mission constraint must be specified upfront and built into the brief - not retrofitted after optimisation has happened. Customer service automation follows in cycle two, with a hard requirement that human escalation pathways remain available in all customer-facing languages.

PATHWAY C

Enterprise with strong commercial position seeking deeper impact

The opportunity here is the inverse of the usual question. Rather than asking how AI improves commercial performance, the enterprise asks how the operational data and customer insight already generated by commercial AI can be redirected to deepen the mission - identifying underserved segments the enterprise could reach, surfacing the patterns of exclusion in the current customer base, or making operational data available to partner organisations working in adjacent mission areas. This is where AI starts to genuinely advance social enterprise practice rather than simply improve its margins.

What this section does not say

It does not say that AI is a sustainability solution for social enterprises that lack a viable commercial model. A social enterprise that is not financially sustainable will not be made so by AI - efficiency tools amplify the gap between revenue and cost; they do not close it. It does not say that every social enterprise should adopt AI. Many of the most effective South African social enterprises operate in deeply relational segments where the operational gains AI offers are immaterial relative to the trust capital that defines the enterprise. It does not promise that any of this is easy. The discipline required - specifying mission constraints before optimisation targets, holding the equity test on every deployment, treating disclosure as non-negotiable - is meaningfully harder for a social enterprise than for either a pure commercial business or a pure NPO. That is the price of holding both sides of the model. AI does not remove that price; it raises the stakes on getting it right.

APPENDICES

Key References and Data Sources

South African Context

- Trialogue Business in Society Handbook 2025: CSI expenditure data and NPO sector analysis
- Bridgespan Group (2025): 'Delivering on the Promise of Corporate Social Investment in South Africa': Six transformational CSI practices based on 36 expert interviews
- Inyathelo / SAIA (South African Institute for Advancement): NPO fundraising, governance and sustainability resources
- National Planning Commission: National Development Plan 2030: Social development goals and civil society role
- Protection of Personal Information Act (POPIA, 2020): Data governance obligations for AI in the sector
- NPO Act 71 of 1997 and NPC/Section 21 company framework: Legal context for civil society organisations
- The King Codes: The King Codes (officially the King Report on Corporate Governance) do not reside under any specific legislation. They are a set of voluntary, non-statutory principles and leading practices developed by the Institute of Directors in South Africa. While the Code itself is not a law, it heavily intersects with and is referenced by several key pieces of South African legislation and market regulations:
    - **The Companies Act (No. 71 of 2008):** The courts and Companies Act use the King Code as a benchmark to measure the fiduciary duties and standard of conduct expected of company directors. (Under this Act – the Social and Ethics committee requirements drive much of the CSI landscape in South Africa.)
    - **JSE Listings Requirements:** For companies listed on the Johannesburg Stock Exchange, applying the King Code is a mandatory requirement rather than a voluntary one.
    - **Sectoral Laws:** The King Code (including the latest King V Code) aligns its governance requirements with statutory mandates like the National Environmental Management Act (NEMA) and various financial regulations.
- King V, effective from 2026, reflects a broader shift in governance thinking: technology is no longer a support function but a core element of leadership accountability. A key development is the explicit inclusion of artificial intelligence within information and technology governance. Under King V, boards are expected to take responsibility for how data, systems and emerging technologies are used across the organisation. This includes providing strategic direction and ensuring that technology enables, rather than undermines, organisational objectives. Importantly, the Code shifts accountability directly to the board. Directors can no longer defer AI oversight to IT or compliance functions. They are expected to understand the implications of AI, ensure appropriate human oversight, and align its use with ethical principles such as fairness, transparency, and accountability. AI-related risks, including bias, privacy breaches, cybersecurity threats and reputational damage, must now be integrated into enterprise risk

management. At the same time, boards are required to balance these risks against the opportunities AI presents for innovation, efficiency and competitive advantage. Ultimately, King V reframes from the role of directors as stewards of both value and trust in a digital economy. As AI becomes embedded in decision-making, boardrooms are now at the forefront of ensuring that innovation is guided by responsibility, ethics and long-term sustainability. For South African organisations, this marks a crucial shift from asking if artificial intelligence (AI) should be governed, to understanding how to govern it well.

African AI Ecosystem

- African Union Continental AI Strategy (endorsed July 2024, Accra, Ghana): Africa's AI governance framework aligned with Agenda 2063
- McKinsey Global Institute (2025): 'Leading, Not Lagging: Africa's Gen AI Opportunity': Economic potential and adoption analysis
- GSMA (2024): 'AI for Africa: Use Cases Delivering Impact': Nigeria deep dive and pan-African use case analysis
- Mastercard (2025): 'AI in Africa': Infrastructure, adoption challenges and sectoral opportunities
- Brookings Institution (2024): 'How AI is Impacting Policy Processes and Outcomes in Africa'
- Nature Africa (2025): 'Africa Leading the Global Effort for AI That Works for All': Masakhane, Lelapa AI, Deep Learning Indaba
- MIT Technology Review (2024): 'What Africa Needs to Do to Become a Major AI Player': InkubaLM, Lelapa AI, language barriers
- TechCabal / Jade Abbott profile (2025): 'Lelapa AI's Jade Abbott is Building AI That Speaks African Languages'
- AI4D Africa Programme: Call for research on socio-economic impacts of AI in Africa: At its core AI for Development (AI4D) seeks to contribute to two critical outcomes: Shaping policy: Establishing responsible AI policies and regulations that minimise harm and promote inclusivity, ethics, sustainability, and human rights. Scaling innovation: Supporting and expanding impactful AI applications that address key development challenges, such as gender inequality and environmental sustainability.
- UNESCO (2025): 'Growing Africa's AI Garden: A Future Rooted in People': AU AI for Africa Initiative
- TechCabal Insights (2025): 'The Building Blocks Africa Needs for AI Adoption': Five infrastructure gaps
- ImpactHER (2024): Bridging the Digital Divide: Empowering Women & Girls in Africa through Digital Literacy and AI Education: Survey data on women and AI proficiency across 52 African countries
- WEF Future of Work Report 2025: Labour market projections including South Africa digitisation impact
- Stanford AI Index 2024: African AI adoption data including 27% daily ChatGPT use in Kenya
- Microsoft AI Diffusion Report 2025: No African country has reached 20% AI adoption

Impact Measurement and Standards

- GIIN / IRIS+ System (iris.thegiin.org): Global impact investing measurement standards and 5 Dimensions framework
- Social Value UK: SROI methodology, principles and practitioner guidance
- Impact Management Project: 5 Dimensions of Impact framework
- GRI Standards (globalreporting.org): Sustainability reporting framework for environmental, social and governance performance
- b.world: AI-native impact measurement, reporting and storytelling platform
- 60 Decibels / Lean Data: Rapid beneficiary feedback methodology and community voice tools
- SPAR4D: Systems Practice, Adaptive Response for Development

AI Ethics and Governance

- UNESCO Recommendation on the Ethics of Artificial Intelligence (2021): Global AI ethics framework
- NIST AI Risk Management Framework (AI RMF 1.0): Risk governance standard for AI systems
- ISO/IEC 42001:2023: AI Management System Standard
- AU Data Policy Framework: African data sovereignty and governance principles
- King V Report on Corporate Governance for South Africa (2026): Governance principles applicable to AI oversight - Institute of Directors South Africa (IoDSA): AI governance guidance for boards

Future workforce: Skills and Competencies

- World Economic Forum Future of Jobs Report 2025: Skills projections, AI displacement and creation figures, top skills to 2030
- Deloitte Digital Consumer Survey 2025 : Generational AI adoption rates by standalone and embedded AI tools
- The London School of Economics (LSE) 2025 GENERATIONS Survey (conducted in collaboration with Protiviti): A global study exploring how AI is reshaping the multigenerational workforce: AI at Work Survey 2025 — 83% Gen Z, 73% Millennial, 60% Gen X, 52% Boomer workplace AI use
- McKinsey Global Institute: The state of AI in 2025: Agents, innovation, and transformation, November 5, 2025: AI expertise by age group — 62% of Millennials vs 50% Gen Z vs 22% Boomers report high expertise
- Pew Research Center (2025): 30% of Millennials use ChatGPT at work vs 18% of Gen X and Boomers
- Mather Institute Gen Xperience Study 2026: Gen X as multigenerational workplace bridge; AI scepticism as strategic caution
- Menlo Ventures State of Consumer AI 2025: 45% of Baby Boomers have used AI in past 6 months
- PNP Staffing Group - 2024 Nonprofit Salaries and Staffing Trends report 2024: 59% of NPOs found it harder to fill roles in 2024
- Social Current (2025): Navigating Workforce Challenges: Equip the Flip leadership competency framework

- ManpowerGroup World of Work for Generation Z (2025): Gen Z will be one third of global workforce by 2030

Broader Research Themes

- NPC – State of the Sector Report 2024: Ready for Reset - 40% of registered NPOs unable to meet financial obligations
- The African Next Voices (ANV) project — a pan-African initiative to address the deep underrepresentation of African languages in AI. With over 9,000 hours of recorded speech in 18 languages across South Africa, Kenya, and Nigeria, this work is the most extensive open dataset creation effort of its kind on the continent. The Data Science for Social Impact group (DSFSI) - University of Pretoria, coordinated the South African arm of ANV, which focuses on seven languages: isiZulu, isiXhosa, Sesotho, Setswana, Sepedi, isiNdebele, and Tshivenda.
- The Possibilists Study 2025: A research study focused on young people aged 14 to 35 who are currently leading initiatives that address pressing social and environmental challenges. Initiated by ChangemakerXchange, The Possibilists is a global alliance of the world's leading youth-focused social-innovation networks.
- Candid Insights 2023: 4 things we learned about foundations and general operating support. Funders gave anywhere from 0% to 100% of total grant dollars as general operating support. The median foundation in 2022 awarded 32% of total grant dollars as GOS. In 2021, the median was 30%. According to the Center for Effective Philanthropy (CEP), in the decade prior to the COVID-19 pandemic, 21% of grants were for GOS. Candid's U.S. Social Sector Dashboard indicates that in 2019, 23% of private and community foundation grant dollars were designated as general operating support.

Glossary of Practitioner Terms

This glossary provides plain-language definitions of the AI, data and social-economy terms used throughout this guide. It is designed for practitioners, not specialists, and prioritises usefulness in the South African sector context over technical precision.

- **Algorithmic decision-making.** Decisions made or substantially influenced by software that follows predefined rules or learned patterns. POPIA gives data subjects rights regarding decisions based solely on automated processing.
- **CBO.** Community-Based Organisation. A grassroots organisation operating in and accountable to a specific local community.
- **Code-switching.** Alternating between languages within a single conversation — common in South African contexts. A meaningful test of whether multilingual AI tools actually work in practice.
- **CSI.** Corporate Social Investment. South African term for what is called corporate philanthropy or corporate social responsibility elsewhere.

- **Data colonialism.** The extraction of value from communities through data collection that benefits external actors more than the community itself. A central concern in the African AI ethics conversation.
- **Data dignity.** The principle that communities and individuals retain meaningful interest in data generated about them — not just consent to collection, but a share in the value created.
- **Disintermediation.** The removal of middle-layer actors when technology enables direct connections. The risk and opportunity AI creates for intermediary organisations.
- **ECD.** Early Childhood Development. A core focus area in South African social-sector practice, from birth to school-ready age.
- **Ecosystem intelligence.** The connective layer that allows actors in a social economy to coordinate, learn and adapt across organisational boundaries. The framing this guide proposes for AI's most valuable role in the sector.
- **Esethu Framework.** A community-data governance framework developed in South African context that specifies how communities benefit when their data is used by AI systems. Referenced throughout this guide as a sector reference point.
- **Fine-tuning.** The process of taking a pre-trained AI model and training it further on a specific dataset to specialise its outputs. More efficient than training from scratch.
- **Foundation model.** A large AI model trained on broad data that can be adapted to many downstream tasks. The technical category includes GPT-4, Claude, Gemini and similar systems.
- **Fractional CTO.** A senior technical leader engaged on a part-time or shared basis across multiple organisations — a cost-effective way for smaller NPOs to access strategic AI and data capability.
- **Generative AI.** AI systems that produce new content — text, images, audio, code — rather than only classifying or predicting. The category most relevant to this guide's discussion of synthesis applications.
- **GBV.** Gender-Based Violence. A priority sector in South African development work; appears throughout the guide where AI applications intersect with sensitive personal data and survivor support.
- **Hallucination.** When an AI system generates fluent, confident output that is factually incorrect. The reason every AI output in evaluation, reporting or governance contexts must be human-verified.
- **Hugging Face.** A widely used platform for hosting and sharing open-source AI models. The infrastructure layer behind much of the African AI development ecosystem.
- **Human-in-the-loop.** A system design in which AI generates suggestions or drafts but human judgement makes final decisions. The default architecture this guide recommends for most social-economy applications.
- **IMP.** Impact Management Project. A widely adopted framework for measuring and managing social impact.

- **Intelligence deficit.** The framing of this guide proposes what holds back the sector — not lack of effort or commitment, but lack of the connective intelligence that lets actors learn from each other and act in coordination.
- **IRIS+.** A widely adopted impact measurement framework developed by the Global Impact Investing Network (GIIN). Referenced in Section Eight.
- **LLM.** Large Language Model. The AI architecture behind most generative text applications relevant to this guide.
- **Machine learning.** A category of AI in which systems learn patterns from data rather than following pre-programmed rules. Distinct from — and often more appropriate than — generative AI for pattern detection and prediction tasks.
- **MEL.** Monitoring, Evaluation and Learning. The sector's preferred term for the function elsewhere is called M&E or impact measurement, with the L emphasising that data should improve practice, not just satisfy reporting.
- **NEET.** Not in Employment, Education or Training. A widely used indicator of youth disengagement; in South Africa the expanded NEET rate sits around 45% for the 15–34 cohort.
- **NPO.** Non-Profit Organisation. The formal South African registration category for what is elsewhere called a nonprofit, charity or NGO.
- **On-premise / private deployment.** Running AI systems on infrastructure controlled by the user organisation rather than on a vendor's cloud. Often required for POPIA compliance when processing sensitive personal data.
- **POPIA.** The Protection of Personal Information Act, 2013 — South Africa's data protection legislation. The regulatory foundation against which all AI applications processing personal data must be assessed.
- **Portfolio intelligence.** A funder's capacity to learn across the grantees in its portfolio — surfacing patterns, sharing what works, identifying systemic gaps. A function AI can materially strengthen.
- **Prompt engineering.** The practice of designing instructions to large language models to produce reliable, useful output. A learnable skill, not magic; this guide treats it as a core practitioner competency.
- **RAG.** Retrieval-Augmented Generation. An architecture in which an AI model draws on a specific document store at the time of generating output, rather than relying only on its training. More reliable for organisation-specific synthesis.
- **SROI.** Social Return on Investment. A framework for quantifying social value created relative to inputs.
- **Theory of change.** A structured articulation of how a programme's activities lead to its intended outcomes. The conceptual foundation for most impact measurement work.
- **Vector database.** A specialised database optimised for storing and searching the embeddings AI systems use to find semantically similar content. The infrastructure layer enabling RAG.

Key Acronyms used throughout the document

- **POPIA** Protection of Personal Information Act
- **CSI** Corporate Social Investment
- **NPO** Non-Profit Organisation
- **CBO** Community-Based Organisation
- **MEL** Monitoring, Evaluation and Learning
- **M&E** Monitoring and Evaluation (used interchangeably with MEL in some sections)
- **IMM** Impact Management and Measurement
- **IRIS+** Impact Reporting and Investment Standards
- **SROI** Social Return on Investment
- **IMP** Impact Management Project
- **DFI** Development Finance Institution
- **SMME** Small, Medium and Micro Enterprise
- **GBV** Gender-Based Violence
- **ECD** Early Childhood Development
- **NLP** Natural Language Processing
- **LLM** Large Language Model
- **API** Application Programming Interface