AI FOR THE GOOD

A Responsible, Practical and Human-Centred Perspective on Artificial Intelligence

---

Trinity Software Center

Your Partner for Ethical AI Excellence

Kumasi | Veenendaal • Version 1.0 • 2026

www.trinitysoftwarecenter.com

Foreword

The question is no longer whether Artificial Intelligence will change the world. It already has. The question that matters now is whether leaders, developers and organisations will take responsibility for shaping how that change unfolds.

This e-book was written from a particular vantage point: a software company rooted in Ghana, working with clients and partners across Europe and Africa, applying AI every day in real projects for real organisations. We are not theorists. We are builders. And we believe that building well requires asking hard questions about purpose, ethics and impact.

In putting together this second, expanded edition we have drawn on the growing international body of research and practice on AI for the common good — from the FARI Institute in Brussels to the ITU's AI for Good Global Summit in Geneva, from the ASEF InnoLab's work on universities and AI, to Josh Tyrangiel's ground-level reporting in his landmark book AI for Good. We have also drawn on Ghana's own National AI Strategy, launched in April 2026, which positions the country as West Africa's AI hub and reflects the ambitions of a generation of African technologists ready to build for global impact.

We hope this e-book encourages you to think boldly, act carefully, and build AI that serves people — not the other way around.

— *The Management Team, Trinity Software Center*

Executive Summary

This e-book offers business and NGO leaders a practical, human-centred guide to Artificial Intelligence — what it is, how to use it responsibly, and how it is already creating positive impact around the world.

Key takeaways:

- AI is no longer optional — leadership must engage deliberately and responsibly (Chapter 1).
- Responsible AI requires clear principles, legal compliance (EU AI Act, GDPR), and human oversight (Chapter 2).
- AI is already solving real-world problems — in healthcare, social care, agriculture and public services (Chapter 3).
- You do not need to build from scratch — existing software can be enhanced with AI through APIs and integrations (Chapter 4).
- NGOs and social enterprises can start small, identify real needs, and pilot AI responsibly with the right partner (Chapter 5).

This guide is written for decision-makers who want to act — not just follow the conversation. It takes less than 30 minutes to read in full.

Why Trinity Software Center?

Throughout this e-book, you will encounter examples and perspectives from Trinity Software Center. Here is why we are well placed to guide you on this topic:

- **10+ years** delivering software solutions for clients across Europe and Africa.
- **Presence in Ghana and the Netherlands,** bridging two continents.
- **Proven experience** with NGOs, social enterprises, and mission-driven organisations.
- **AI Center of Excellence** — actively building, testing and deploying AI solutions.
- **Fairtrade software development** — professional, ethical, and affordable.

We believe that AI for the good is not a slogan. It is a responsibility we take seriously every day.

Chapter 1 – AI is Everywhere (and Leadership Matters)

Artificial Intelligence has moved from research labs into everyday life at breathtaking speed. In just a few years, AI systems have gone from niche tools to mainstream infrastructure, shaping how organisations work, decide and create value. For business owners and leaders of charities or non-profits, this speed creates both opportunity and unease.

On the opportunity side, AI enables organisations to do more with limited resources. It can automate repetitive work, analyse large volumes of information, support better decisions and unlock services that were previously out of reach. For mission-driven organisations, this can mean more impact with the same budget. For businesses, it can mean higher quality, better scalability and new value propositions.

At the same time, concerns are real and justified. Leaders worry about loss of control, ethical boundaries, legal exposure, data privacy, and reputational risk. AI systems can amplify bias, produce convincing but incorrect outputs, or consume significant amounts of energy and water through large data centres. Without clear governance, AI can damage trust faster than it creates value.

The Leadership Question

The core leadership question is therefore not whether to use AI, but how to use it deliberately and responsibly. AI is a strategic choice. Like any powerful tool, its impact depends on intent, design and oversight.

> *"The biggest risk we face isn't AI eliminating the human race. It's actually the race to embed AI everywhere without sufficient understanding of what that means for people and for our planet." — Doreen Bogdan-Martin, ITU Secretary-General, AI for Good Global Summit 2025*

The ITU's AI for Good Global Summit — the world's largest gathering of AI experts and on-the-ground implementers focused on the Sustainable Development Goals — convened over 17,000 participants and 800 speakers in Geneva in 2025. The clear message from that gathering: real progress on the SDGs requires closing data gaps, designing for the most vulnerable, and grounding AI governance in lived realities rather than theoretical frameworks.[1]

Josh Tyrangiel, The Atlantic's AI correspondent and author of AI for Good: How Real People Are Using Artificial Intelligence to Fix Things That Matter (Simon & Schuster, 2026), spent years travelling inside hospitals, public schools, tax agencies and the Pentagon to observe how AI actually performs in the real world — where data is messy,

incentives are misaligned, and everything that matters is charged with human emotion. His central finding mirrors the leadership challenge perfectly: AI is neither a gold rush nor a sign of the apocalypse. It is a tool, and tools are only as good as the people and institutions that wield them.[2]

This e-book gives you a compact summary of the basic considerations for a responsible and impactful use of AI. We hope it will encourage and inspire you to start working on your own AI products and services to make this world more sustainable. Trinity Software Center has been developing its “Center of AI Excellence” and our experienced engineers are available to work with you on responsible AI solutions. As a Fairtrade software company, working on the sustainability goals (SDG) is very close to our hearts. Contact us when you have an idea or prototype we can work on together.

Chapter 2 – Using AI Responsibly: Principles, Rules and Real-World Risks

Responsible use of AI starts with leadership and clear principles. For organisations, this means treating AI not as an experiment or toy, but as part of core operations, business processes, governance and values (code of ethics).

Design Principles for Responsible AI

- **Human-in-the-loop**
  AI supports people; humans remain accountable at all times. They should be equipped to use AI in the right way and responsibly.
- **Purpose-driven** AI is applied to a clearly defined problem or goal, not deployed for its own sake.
- **Transparency**
  Users understand what the system does and does not do and what its limitations are .
- **Fairness**
  Bias is actively identified and mitigated throughout the design and deployment cycle.
- **Privacy & security**
  Data protection is built in by design, not bolted on afterwards.
- **Robustness**
  Systems are tested, monitored and improved over time.

These principles are especially important for organisations that serve vulnerable groups, operate with public trust, or work across legal and cultural boundaries. They are also the foundation of the FARI Institute's framework for trustworthy AI, developed by the Vrije Universiteit Brussel (VUB). FARI defines its mission as developing, studying and fostering the adoption and governance of AI, data and robotics technologies in a trustable, transparent, open, inclusive, ethical and responsible way.[3]

Compliance and Legal Frameworks

Responsible AI is also a legal obligation. In Europe, organisations must consider multiple regulatory frameworks:

- **EU AI Act**
  Introduces a risk-based approach, with strict requirements for high-risk AI systems, including human oversight, documentation and risk management. Enforcement began in phases from 2024.
- **GDPR**
  Governs the lawful, transparent and secure use of personal data. Any AI system that processes personal data must comply.
- **Digital Services Act (DSA)**
  Regulates algorithmic recommendation systems, including transparency obligations for very large platforms.
- **Digital Markets Act (DMA)**
  Addresses market power in digital ecosystems; relevant for organisations deploying AI through major platforms.
- **Sector regulations** Sector regulations – Healthcare, finance and public-sector AI face additional rules at national and EU level.

Ignoring these frameworks can lead to fines, reputational damage, loss of trust and operational disruption. Trinity Software Center builds GDPR compliance and EU AI Act considerations into every project from day one. You will find a full reference overview in Annex 1.

When Things Go Wrong – From Science Fiction to Daily Reality

For decades, films have explored futures shaped by intelligent machines. What once felt like distant science fiction is now uncomfortably close to everyday reality.

Well-known Hollywood examples help illustrate the risks:

- Minority Report – A society where crimes are predicted before they happen. What felt fictional is now echoed in predictive policing and risk-scoring systems, where biased data can lead to unfair targeting if human judgment is removed.
- I, Robot / Terminator / The Matrix – Autonomous systems that act beyond human control. While exaggerated, these stories mirror real concerns around autonomous decision-making without sufficient safeguards.
- Black-box AI in control rooms – A recurring movie trope where humans trust machines too much. Today, this risk appears when AI-generated outputs are accepted without validation.

These stories teach a consistent lesson: the danger lies not in intelligent systems themselves, but in automation without accountability and real understanding of what a system is doing.

Real-world AI failures can be just as instructive as success stories. Two cases from Africa illustrate what happens when good intentions meet poor design assumptions.

Case 1: Voice AI for small businesses in Nairobi (Dukawalla, Kenya, 2026)

Development organisations have long championed voice interfaces as the ultimate accessibility solution for populations with limited digital literacy: if people can speak, they can interact with technology. Research published in May 2026 on the Dukawalla platform — a voice-enabled business assistant deployed across seven small and medium enterprises in Nairobi — exposed why this assumption is fundamentally flawed. The system allowed users to record sales data, retrieve AI-powered insights and manage business information through natural language. But three critical barriers emerged that no amount of technological refinement could overcome: business communication is rarely private (market noise, nearby customers and social context make voice recording impractical), transactions are typically multi-party (one-to-one voice interaction cannot capture the negotiation and group dynamics of real trading), and many informal business conversations blend multiple languages. Voice interfaces are not failing because of poor technology — they are failing because the technology was designed without sufficient understanding of how human communication actually works in African business contexts. The lesson: assumed affordances are not actual affordances. Test with real users in real environments before scaling.

Source: Wayan Vota, "Why Voice-Only Interfaces Are Failing Small Businesses in LMICs," ICTworks, 21 May 2026. www.ictworks.org/why-voice-only-interfaces-are-failing-small-businesses-in-lmics/ [12]

Case 2: AI tools deployed with community health workers in sub-Saharan Africa

A mapping of 38 AI systems deployed with community health workers (CHWs) in low- and middle-income countries found that over 40% were LLM-powered chatbots designed to function as always-on supervisors. Yet the scepticism of these health workers towards the tools was widely misread by programme designers as a training problem rather than a rational response to genuine risks. Research published by ICTworks shows that CHW reluctance is well-founded: a 2024 ethnographic study in Eastern India documented digital health platforms as instruments of state oversight, where CHWs — mostly women — were monitored through these systems while receiving no meaningful operational support in return. Data flowed upward; accountability flowed nowhere. The AI tools were not neutral clinical aids; they were accountability infrastructure embedded in unequal relationships. The lesson: deploying AI into a community context without understanding the power dynamics and lived realities of its users can cause real harm to trust and to people.

Source: Wayan Vota, "Community Health Workers Are Right to Distrust AI Solutions," ICTworks, June 2026. www.ictworks.org/community-health-workers-are-right-to-distrust-artificial-intelligence-solutions/ [13]

> *In real life, failures usually stem from unclear goals, poor data quality, lack of testing, absence of human oversight, or blind trust in technology. Avoiding these risks requires discipline rather than fear: clear scope definition, quality data, ethical review, legal checks and continuous monitoring.*

See Annex 2 for our practical Human-Centred AI Design Checklist.

AI Myths vs Reality

Before exploring how AI can be applied for good, it helps to clear up some common misconceptions. These myths can either cause leaders to dismiss AI too quickly or to adopt it carelessly.

Myth: AI will replace all jobs. Reality: AI changes how jobs are done far more often than it eliminates them. In most organisations, AI takes over specific repetitive tasks, freeing people for more meaningful and complex work. The World Economic Forum consistently finds that AI creates new roles even as it transforms existing ones.

Myth: AI is always objective and neutral. Reality: AI systems learn from historical data, which means they can inherit and even amplify existing biases. An AI trained on biased hiring records will produce biased recommendations. Fairness must be actively designed and monitored — it does not happen automatically.

Myth: AI requires millions of euros to implement. Reality: Many practical AI solutions use existing APIs from providers such as OpenAI, Google, or Microsoft. Embedding AI into your current software can be done incrementally and cost-effectively — you do not need to train your own model from scratch.

Myth: AI is fully autonomous and can run without supervision. Reality: Human oversight remains critical. AI systems make mistakes, produce unexpected outputs, and lack context. Responsible AI always includes mechanisms for human review, escalation, and correction — especially in high-stakes decisions affecting people's lives.

> *Want to explore AI opportunities in your organisation? Contact Trinity for a free AI discovery session — info@trinitysoftwarecenter.com*

Chapter 3 – AI for the Good: Technology as a Force for Positive Change

While risks are real, AI also offers unprecedented opportunities to address some of the world's most pressing challenges. When aligned with clear social goals, technology can become a powerful enabler for positive impact.

Many of these opportunities align closely with the United Nations Sustainable Development Goals (SDGs): reducing poverty, improving healthcare, strengthening education, supporting decent work, promoting responsible consumption, and building inclusive institutions.

AI can help by:

- Making scarce expertise more widely available
- Reducing administrative overhead so professionals can focus on people
- Improving data quality for better policy and intervention
- Supporting early detection and prevention instead of costly remediation

A 2020 study published in Nature Communications by Vinuesa et al. examined how AI could enable 79% of the SDG targets across health, education, and economic growth, while also noting risks if governance frameworks are not put in place.[4] This remains a foundational reference for the AI for good movement.

The International Landscape

ITU AI for Good Global Summit

Since 2017, the International Telecommunication Union (ITU), the United Nations' specialised agency for digital technologies, has convened the AI for Good Global Summit as the world's leading platform connecting AI innovators with real-world problem owners to collectively advance the SDGs. The 2025 edition in Geneva brought together over 17,000 participants from governments, civil society, academia and the private sector, and showcased progress on AI standards in healthcare, food security, road safety and energy efficiency.[1]

The summit's governance dialogue — organised by ITU with 53 UN partner agencies — underscored the need for AI that bridges innovation, inclusion and sustainability, and that developing countries must be included in conversations that shape the digital future.[5]

FARI – AI for the Common Good Institute, Brussels

FARI is an independent, not-for-profit AI initiative led by VUB and ULB, funded in part by a EUR 9.91 million NextGenerationEU budget. It operates an AI Test and Experience Centre in Brussels, houses an AI Academy for training public-sector stakeholders, and runs flagship European research projects including CitCom.ai, LIBRA.I. and AI4GOV-X.[3]

FARI's core position — that there is no AI for the common good if it is not responsible AI — echoes through all its research. Its 2022 Brussels Conference brought together over 300 participants from policy, academia and civil society to identify concrete priorities for the Brussels-Capital Region, with a particular focus on data quality, citizen engagement and the accountability of AI decisions.[6]

ASEF InnoLab – AI for Good in Higher Education

The Asia-Europe Foundation's Higher Education Innovation Laboratory (ASEFInnoLab) is an intergovernmental programme connecting academics, researchers and university managers across 45 Asian and European countries. Its latest edition — ASEFInnoLab7, launched in 2026 under the theme 'The Role of Universities in Advancing AI for Good' — explores how universities can catalyse inclusive, ethical and socially beneficial AI outcomes.[7]

Earlier editions have produced white papers on AI for Sustainable Development, noting that university collaborations backed by accessible lifelong learning and open data can drive progress and help societies become more sustainable.[8] This is directly relevant to Trinity's partnership with KNUST and our ongoing exploration of AI applications in Ghanaian education and social services.

AI for Good in Africa and Ghana

Africa is not a passive recipient of AI development. It is an active architect of its own digital future, with unique demographic advantages, a mobile-first infrastructure and a generation of entrepreneurs ready to build at scale.

In April 2026, Ghana launched its National AI Strategy (2025–2035), with a core vision to enable the use of AI for inclusive economic growth and national development while positioning the country as a regional AI leader. The strategy is built on eight pillars, including expansion of AI education, empowerment of youth for jobs of the future, deepening digital infrastructure, and facilitating data access and governance. Implementation partners include KNUST's Responsible AI Lab, GIZ, UNDP and UNESCO.[9]

> *"Ghana is determined to align with this forward-looking global trend by harnessing AI for good." — President John Mahama, at the launch of Ghana's National AI Strategy, April 2026*

The GSMA's 2024 report AI for Africa: Use Cases Delivering Impact highlights how mobile-first AI applications are already reshaping agriculture, healthcare, education and financial services across sub-Saharan Africa — sectors where Trinity has active project experience.[10]

Africa's AI market is projected to grow from USD 4.5 billion in 2025 to USD 16.5 billion by 2030 (Statista, cited in Mastercard's 2025 whitepaper Harnessing the Transformative Power of AI in Africa). The same report emphasises that responsible, locally rooted AI can drive inclusive growth and connect more people to opportunity — a vision Trinity shares.[11]

AI for Good — Global Examples

Earlier Detection of Breast Cancer

In healthcare, AI systems now assist radiologists by detecting patterns in mammograms that are difficult to spot with the human eye alone. These systems act as a second reader — highlighting potential areas of concern and reducing the risk of missed diagnoses. Studies show that AI-supported screening can improve early detection rates while reducing diagnostic pressure and workload.

AI Embedded in Enterprise Software — SAP and Microsoft

Large enterprise software providers increasingly embed AI directly into their platforms. Both SAP and Microsoft have integrated AI capabilities into their existing software to support users with tasks such as searching information, generating reports, analysing data and automating routine actions. SAP's AI integration has been commercially transformative: it became the most valuable company on the German stock exchange,

demonstrating that responsible and user-focused AI integration can create both economic and societal value.

More Global Examples

Precision agriculture helps farmers optimise water use, reduce fertilisers and increase yields. Disaster response systems analyse satellite images to assess damage and guide emergency aid. Adaptive learning systems support personalised education at scale. AI improves predictions of extreme weather and climate risks. Public services use AI to reduce backlogs by automating document analysis and case triage.

> *By Trinity Software Center's own engineers:*

FillFellow – Reducing Administrative Burden in Social Care

In many social and care organisations, professionals spend a significant part of their time filling in forms during or after client conversations. This administrative pressure reduces time for human interaction and still often results in incomplete or inconsistent data.

FillFellow, developed by Trinity Software Center, uses AI-supported speech recognition and natural language processing to capture conversations and automatically structure them into forms and reports. The result is higher data quality, less manual work, and more time for meaningful client interaction. Privacy, consent and data protection are built into the design from the start.

> *These examples share a common pattern: when AI is applied with clear purpose, human oversight and ethical boundaries, it can scale impact without losing humanity.*

> *Want to explore AI opportunities in your organisation? Contact Trinity for a free AI discovery session — info@trinitysoftwarecenter.com*

Chapter 4 – Enhancing Existing Software with AI

AI does not always require building entirely new systems. In many cases, the greatest value comes from enhancing existing software with intelligent capabilities. This is often faster, less risky, and more cost-effective than starting from scratch.

Case Study: VAT Document Processing

A concrete example comes from a VAT service provider working with Trinity Software Center. This organisation processes large volumes of international trade documents: invoices, shipping papers, customs forms and declarations from many countries and formats.

By integrating AI into its administrative software, the organisation can:

- Read and classify documents automatically
- Identify key data fields regardless of language or layout
- Check documents for completeness and consistency
- Register transactions faster and with fewer errors

This reduces manual work, shortens processing times, and improves compliance, while human experts remain in control of exceptions and final decisions.

Case Study: FillFellow as an Embeddable AI Component

Rather than replacing current systems, FillFellow's AI components can be embedded into existing web applications to capture data more efficiently, improve accuracy, and support users in their daily workflows. Trinity Software Center is currently exploring this option with several clients to see how AI can improve performance, reduce costs or enable entirely new services.

You Do Not Need Five Billion Euros to Build Your Own Solution

The perception that AI is only for large corporations with vast budgets is one of the most persistent myths in this space. In reality, there are multiple AI capabilities that can be embedded into your software application, mobile app or data lake in the form of an API — at a fraction of the cost of training a model from scratch.

Of course, training your own agent can be a very tricky and costly affair that should not be started easily. The number of data points required to get to reliable outcomes can be enormous, and the risks of producing faulty results are real.

Trinity's data engineers have the experience to help you identify the right AI capabilities for your specific need, implement them securely and responsibly, and help you start using the opportunities of AI quickly — without losing sight of human principles and security regulations.

> *The question is not 'Can we afford AI?' but 'What problem are we solving, and what is the smallest responsible step we can take towards solving it with AI?'*

> *Want to explore AI opportunities in your organisation? Contact Trinity for a free AI discovery session — info@trinitysoftwarecenter.com*

Chapter 5 – AI for Social Enterprises and NGOs

Social enterprises and NGOs often face a paradox: their societal challenges are complex, while budgets, staff and time are limited. Used responsibly, AI can help organisations increase impact without losing their human focus.

The starting point is not technology, but reflection. FARI's work with Brussels-based NGOs and social programmes has consistently shown that the best AI solutions emerge from deep conversations about operational bottlenecks, not from enthusiasm about new technology.[6]

Design Questions for NGOs and Social Enterprises

- Where do our people lose the most time on administration?
- Which decisions would improve with better data or insights?
- Which target groups struggle to access our services?
- Where does inconsistency or data quality limit our impact?
- What risks would require human oversight at all times?
- What new services can we offer when we introduce AI-empowered features?

From these questions, AI opportunities often emerge naturally. ASEF's InnoLab research on AI for sustainable development emphasises that what counts as 'good' must be determined in partnership with communities themselves — not imposed from outside.[8] This is a principle Trinity holds central in all its work with African partners and clients.

Trinity's Portfolio of Social Impact

Trinity Software Center 10 years of experience delivering software for social enterprises and NGOs across Europe and Africa.

- FillFellow (NL) – AI-powered form automation for social care organisations, saving hours of administrative work per professional per week.
- ProjectConnect (NL) – ERP platform for NGOs and donor reporting, improving data quality and accountability.
- ESG4SME (NL) – Web-based ESG reporting tool for SMEs and social enterprises, aligning with European VSME standards.
- Odwadini App (GH) – Financial literacy app for informal traders and microbusinesses in Ghana.
- Skuuni (GH) – School management system used by hundreds of schools in Ghana (www.skuuni.com).
- Maxim Nyansa Digital Library (GH) – Over 2 TB of free learning resources for young Africans in English and French.
- City Bibles (NL) – Interactive multilingual Bible app in 200 languages, powered by AI.

Projects are designed with accessibility in mind — particularly for users with limited digital literacy or limited connectivity, a design principle that GSMA's AI for Africa report identifies as critical to ensuring digital transformation benefits reach underserved communities.[10]

Working with Trinity Software Center

Interested in exploring how AI could responsibly support your organisation or mission?

We offer AI design and exploration sessions for businesses, social enterprises and non-profits. In these sessions we:

- Identify concrete AI opportunities aligned with your goals
- Assess risks, compliance and ethical considerations
- Translate ideas into realistic solution paths
- Build rapid prototypes to test assumptions before significant investment

Contact us: info@trinitysoftwarecenter.com | www.trinitysoftwarecenter.com

Conclusion – Choosing the Direction of AI

AI will continue to shape our world. Whether it deepens inequality or helps solve societal challenges depends on the choices leaders make today. By combining responsible principles, legal compliance, thoughtful design and a focus on real human needs, AI can become a force for good.

The international community — through ITU's AI for Good platform, FARI's research in Brussels, ASEF's work with universities across Asia and Europe, and Ghana's own National AI Strategy — is converging on a shared understanding: responsible AI is not a luxury or a constraint. It is the only sustainable path forward.

Josh Tyrangiel, who has spent years visiting the institutions where AI is being tested against the messy reality of human life, puts it plainly in the subtitle of his book: real people are using AI to fix things that matter. Not hypothetically. Not someday. Now.

At Trinity Software Center, we believe that building AI for the good is not a slogan, but a responsibility. We are doing it already — in our code, in our culture, and in the communities we serve on two continents.

Annex 1 – Legal and Regulatory Frameworks for AI in the EU

Organisations operating in or with the European Union must take multiple legal frameworks into account when using AI. Below is a non-exhaustive overview of the most relevant regulations and where to find authoritative guidance.

- EU Artificial Intelligence Act (AI Act) – Official EU legislation: https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai
- General Data Protection Regulation (GDPR) – Official text and guidance: https://gdpr.eu/
- Digital Services Act (DSA) – https://digital-strategy.ec.europa.eu/en/policies/digital-services-act-package
- Digital Markets Act (DMA) – https://digital-strategy.ec.europa.eu/en/policies/digital-markets-act
- NIS2 Directive – Cybersecurity requirements for critical and digital infrastructure suppliers: https://digital-strategy.ec.europa.eu/en/policies/nis2-directive
- VSME Standard – EU voluntary sustainability reporting standard for SMEs (relevant for ESG compliance): European Financial Reporting Advisory Group (EFRAG)

This annex is intended as a starting point, not as legal advice. Always consult a qualified legal professional for your specific situation.

Annex 2 – Human-Centred AI Design Checklist

This checklist helps organisations design AI solutions that are responsible, effective and aligned with human values.

Safeguards to Put in Place

- ✓ Clear human accountability for all AI-supported decisions
- ✓ Defined boundaries where AI may not operate autonomously
- ✓ Regular monitoring for bias, errors and unintended effects
- ✓ Secure handling of personal and sensitive data (GDPR compliance)
- ✓ Clear escalation and fallback procedures when AI fails or is uncertain
- ✓ Key outputs to be human-reviewed before action is taken
- ✓ Training for all staff working with the tools
- ✓ Legal compliance check including GDPR and EU AI Act risk classification

Angles to Identify AI Opportunities

- Where do people lose time on repetitive or administrative work?
- Which processes suffer from inconsistent data quality?
- Where could better insights improve decisions or outcomes?
- Which services could become more accessible with AI support?
- Where can AI augment, rather than replace, human expertise?
- What new services can we offer to our clients when we introduce AI-supported features?

Business Case Caveats

- AI should solve a real problem, not showcase technology
- Costs include governance, monitoring and compliance, not just development
- Benefits should be measurable and realistic
- Reputational and ethical risks must be part of the business case
- Small pilots often create more value than large, uncontrolled rollouts
- Has your idea already been done by anyone? Does that leave room in the market for your proposition?
- Complete a Business Model Canvas with your new AI idea to see how it can generate sustainable revenue

Bibliography and Further Reading

The following sources informed and are referenced throughout this e-book. We encourage readers to explore them further.

Key Works on AI for Good

[1] ITU AI for Good Global Summit 2025, Geneva, 8–11 July 2025. Summary and outcomes available at: https://aiforgood.itu.int/summit25/

[2] Tyrangiel, Josh. AI for Good: How Real People Are Using Artificial Intelligence to Fix Things That Matter. Simon & Schuster, May 2026. www.joshtyrangiel.com

[3] FARI – AI for the Common Good Institute. Overview and mission. VUB & ULB, Brussels. www.fari.brussels

[4] Vinuesa, R., Azizpour, H., Leite, I. et al. “The role of artificial intelligence in achieving the Sustainable Development Goals.” Nature Communications, 11, Article 233, 2020. https://doi.org/10.1038/s41467-019-14108-y

[5] ITU AI Governance Dialogue 2025. Co-chairs’ summary. International Telecommunication Union, Geneva. https://aiforgood.itu.int/

[6] Mazijn, C. et al. “What Can the Brussels Capital Region Do to Foster Responsible AI?” FARI Conference Summary Report. FARI Institute, Brussels, 2022. www.fari.brussels/research-and-innovation/publication/what-can-the-brussels-capital-region-do-to-foster-responsible-ai

[7] ASEF InnoLab7 Concept Overview. “The Role of Universities in Advancing AI for Good.” Asia-Europe Foundation, Singapore, 2026. https://asef.org/projects/asefinnolab7/

[8] Bong, Yang & Sailaja, Neelima. ASEFInnoLab5 White Paper: “AI for Sustainable Development.” Asia-Europe Foundation, 2025. https://asef.org/publications/asefinnolab-white-paper-sustainable-development/

[9] Ghana National AI Strategy 2025–2035. Ministry of Communications, Digital Technology and Innovation, Government of Ghana. Launched 24 April 2026. www.moc.gov.gh

[10] GSMA. AI for Africa: Use Cases Delivering Impact. GSMA Mobile for Development, July 2024. www.gsma.com/solutions-and-impact/connectivity-for-good/mobile-for-development/

[11] Mastercard. Harnessing the Transformative Power of AI in Africa. Mastercard, September 2025.

[12] Vota, Wayan. “Why Voice-Only Interfaces Are Failing Small Businesses in LMICs.” ICTworks, 21 May 2026. https://www.ictworks.org/why-voice-only-interfaces-are-failing-small-businesses-in-lmics/

[13] Vota, Wayan. “Community Health Workers Are Right to Distrust AI Solutions.” ICTworks, June 2026. https://www.ictworks.org/community-health-workers-are-right-to-distrust-artificial-intelligence-solutions/

Recommended Further Reading

FARI Brussels Conference 2024 Summary: “AI, a Public Good?” Studio Flagey & BeCentral, Brussels. https://www.fari.brussels/research-and-innovation/publication/summary-fari-conference-2024

ITU Reports and Publications on AI/ML for SDGs. https://aiforgood.itu.int/newsroom/publications-and-reports/

Etchemendy, J., Landay, J., Li, F.-F., & Manning, C. “Universities must reclaim AI research for the public good.” Stanford HAI, 2025. https://hai.stanford.edu/news/universities-must-reclaim-ai-research-for-the-public-good

European AI Act Official Text: https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

Ghana Digital Economy Policy and Strategy 2024. Ministry of Communications, Digital Technology and Innovation.