



In collaboration with EY and Microsoft

Al for Impact: The PRISM Framework for Responsible AI in Social Innovation

WHITE PAPER JUNE 2024

Contents

Foreword	3
Executive summary	4
Introduction	5
1 Pathways for adoption: The PRISM Framework	8
1.1 The PRISM Framework	9
1.2 Case studies along the adoption pathways	12
2 Capabilities and risks	15
2.1 Ethics	16
2.2 Data	18
2.3 Business and organization	20
2.4 Technology	22
2.5 Costs and metrics	25
Conclusion	28
Contributors	29
Endnotes	31

Disclaimer

This document is published by the Schwab Foundation for Social Entrepreneurship in partnership with the World Economic Forum. The findings, interpretations and conclusions expressed herein are a result of a collaborative process facilitated and endorsed by the Schwab Foundation for Social Entrepreneurship or the World Economic Forum but whose results do not necessarily represent the views of the Schwab Foundation and World Economic Forum, nor the entirety of its Members, Partners or other stakeholders.

© 2024 Schwab Foundation for Social Entrepreneurship and the World Economic Forum. All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means, including photocopying and recording, or by any information storage and retrieval system.

Foreword



Jean-Philippe Courtois Executive Vice-President and President, National Transformation Partnerships, Microsoft



Gillian Hinde EY Global Corporate Responsibility Leader, EYGS



Cathy Li Head, Al, Data and Metaverse; Member of the Executive Committee, World Economic Forum

Now more than ever, businesses and organizations find themselves at a pivotal juncture – one that compels them to prioritize their communities and the environment at the core of their operations. The tumultuous landscape of ongoing global crises signals that business as usual is no longer an option. Against this backdrop of change and challenge, social innovators are demonstrating globally, across sectors and regions and in their communities, what the capitalism of the future can look like – a model where purpose and profit not only coexist but reinforce each other, thriving with profit.

Yet, as the 2030 deadline for the Sustainable Development Goals looms, significant gaps remain in the impact targets set by businesses and governments to ensure a just, equitable and sustainable world for all. Achieving these goals demands nothing less than collective action and a certain leapfrogging to accelerate the speed at which the targets are being met. The emergence of artificial intelligence (AI) presents a transformative opportunity for social innovators, non-profits, businesses and governments alike. The AI for Impact: The Role of Artificial Intelligence in Social Innovation report released in April 2024 examines the landscape of social innovators who are currently deploying AI worldwide. The key insights from this report underscore the immense potential of AI as a tool for driving social impact, particularly in critical domains such as healthcare, education and economic empowerment.

To realize impact on a grand scale, widespread AI awareness and the strategic, responsible deployment of the technology is necessary – whether integrated into operating models or as a product deployed as a means to effect change. Yet, internal and external barriers remain that prevent social innovators from adopting AI at scale, hindering their ability to harness its full potential. For social innovators, challenges such as technical complexity, skills gaps, access to AI technology along with balancing social and business objectives can be prohibitive to their understanding and deployment of the technology. On the other hand, technology organizations and developers have a responsibility to develop AI capabilities that are ethical and fair, address global disparities and do not exacerbate existing biases. This necessitates the creation of two-way dialogue between technology organizations and social innovators that integrates the voices of social innovators, enabling them to influence the technology roadmap and ensuring that AI capabilities are truly fit for purpose for impact enterprises.

This paper builds on earlier efforts to present a strategic roadmap for social innovators assessing how to include AI in their operations or models. It draws inspiration from thought leadership published by the World Economic Forum's AI Governance Alliance and is informed by the expertise of a network of social innovators and intermediaries on the frontline of change. Special thanks go to those involved in the working groups of the Global Alliance for Social Entrepreneurship AI for Social Innovation workstream, whose time and expertise have contributed to the development of this paper and whose lived experiences can be found as examples throughout the paper.

This paper will form the basis for work with intermediary partners and other ecosystem collaborators on the further development of training and skill-building toolkits for social enterprises. Our hope is that the frameworks and roadmaps presented herein, along with resources from corporate collaborators, can help social innovators realize the promise of Al-driven social change.

Executive summary

The PRISM Framework enables organizations to adopt artificial intelligence for impact and highlights technology shortcomings for ethical implementation.

The rapid development and dissemination of artificial intelligence (AI) capabilities presents a viable solution for social innovators to leapfrog structural gaps and scale their solutions. The first report in this series, *AI for Impact: The Role of Artificial Intelligence in Social Innovation*, was published by the Schwab Foundation's Global Alliance for Social Entrepreneurship in collaboration with EY and Microsoft. It offers a comprehensive analysis of the role of AI in social innovation, with insights from over 300 social innovators. It demonstrates how social innovators already embrace the technology for tangible outcomes, mostly in areas such as healthcare, environmental action, economic empowerment or education.

This second paper takes a deep look into best practices in the process of applying AI. It features social innovators who are spearheading the deployment of AI for impact, those who are implementing it as an enabler for scale and those who are just starting to consider the use of AI for their operations. Data shows that more than 50% of social innovators currently apply AI to enhance core products or services. Almost a third leverage AI to develop entirely new solutions and another 20% use it to enhance internal efficiencies or enable scale for their organization.

Building on this dataset, research and over 20 in-depth interviews, this paper introduces the PRISM Framework to guide social innovators and other organizations through the nuanced landscape of Al integration. The framework encourages organizations to start with lowrisk, low-cost AI applications and stresses the importance of organizational readiness over mere technological capability. It aligns with established methodologies like the AI Governance Alliance's Presidio Framework. It advocates an iterative but strategically aligned approach with specific evaluation gates for the implementation of AI.

The PRISM Framework outlines three layers of Al implementation: impact mission and strategy, adoption pathways, and capabilities and risks. Each layer addresses different elements of readiness and potential applications, from simple, internal use cases to complex, externally facing deployments. The modular approach allows organizations to tailor their Al strategies to their specific needs and capacities, ensuring that Al adoption is both impactful and sustainable. Rich case studies, such as SAS Brasil's cervical cancer screening and High Resolves Al-enhanced educational initiatives illustrate the framework's use.

The framework also highlights risks and technological shortcomings that need to be addressed for an equitable implementation. This includes, for example, data biases or technology structures that are not fit to allow for explainability of Al decisions – a key element to build trust among the communities that social innovators are trying to serve. Therefore, the paper calls for active engagement between Al for social innovation ecosystem, technology leaders and social innovators to jointly enable the ethical adoption of Al for positive impact.

↓ Image credit: LifeBank



Introduction

Artificial intelligence has the potential to scale impact in several domains but requires collaboration to help social innovators realize its maximum potential.

Artificial intelligence (AI) offers transformative potential in solving complex societal challenges and is deployed across sectors from healthcare to agriculture. In April 2024, the Schwab Foundation's Global Alliance for Social Entrepreneurship's published the *AI for Impact: The Role of Artificial Intelligence in Social Innovation* report in collaboration with EY and Microsoft. It analyses AI implementations by over 300 social innovators and reveals that there is no gap in implementation between low-/middle-income and high-income countries but a strong diversity in the sectors of implementation: healthcare and environment had the highest share of AI implementations across the board. However, 80% of economic empowerment solutions were deployed by social innovators in low-/middle-income countries.

FIGURE 1 | Geographical summary of the social innovator dataset



Despite broad implementation, the report identifies gaps hindering wider adoption. Only 13% of Al initiatives focus on educational toolkits, primarily in the Global North. Gender disparities persist, with just 25% of women-led social enterprises using Al compared to half in the broader sector. Ethical and equity challenges also emerged. For example, most commercially available models are trained on data from high-income countries, delivering sub-par results for low- and middle-income countries.

Social impact domain



FIGURE 3

Social innovators deploying AI by gender

Woman founder or chief executive officer







This paper builds on previous insights and frameworks such as the Presidio Framework by the World Economic Forum's AI Governance Alliance to highlight avenues for broader AI adoption among social innovators. It combines data from previous studies with insights from 22 in-depth interviews and interactive workshops with social innovators and tech leaders, highlighting the nuanced decision-making processes in AI deployment and the ongoing challenges faced by social innovators.

Pathways for adoption: The PRISM Framework

Building on existing concepts, the PRISM Framework guides the responsible adoption of AI for impact.

The April 2024 *AI for Impact: The Role of Artificial Intelligence in Social Innovation* report emphasizes how social innovators are leading in AI adoption, notably in projects like SAS Brasil's cervical cancer screening and High Resolves's initiatives in AI-aided education despite connectivity challenges. This second paper uses their experiences to illustrate how AI can drive social and environmental benefits. It synthesizes insights from 300 social innovators and expert interviews into the PRISM Framework, a guide to implementing AI iteratively, responsibly and effectively.

The findings encourage social innovators to start small and execute low-cost/low-risk implementation when first approaching AI use cases. It stresses that social innovators and conventional companies alike need to consider their organizational readiness when designing AI implementation, as internal preparedness often outweighs technological and data considerations. The framework highlights adoption pathways based on organizational readiness and relevant capabilities and risks linked to these pathways.

The framework aligns with other concepts like the Presidio Framework by the World Economic Forum's AI Governance Alliance, which supports organizations in responsibly enhancing productivity and redefining industries through AI. The Presidio Framework starts by identifying and filtering the best generative AI use cases aligned with strategic objectives and funnels them through three common evaluation gates – business impact, operational readiness and investment strategy. These evaluation gates are designed to be applied in any sequence and iteratively.

FIGURE 5 | The Presidio Framework: Funnelling use cases through evaluation gates



Funneling use cases through evaluation gates

1.1 | The PRISM Framework

The PRISM Framework builds on the Presidio framework. It highlights how different adoption pathways are interdependent with different levels of organizational, technological and data readiness. It also emphasizes the need for alignment with an organizations impact ambition at every stage (see Figure 6). It highlights that social innovators are, by nature, impact-first organizations and their consideration of strategic alignment (as outlined in the framework by the AI Governance Alliance) goes beyond business alignment. Their considerations prioritize impact and also include core values and ethical principles when applying AI (such as prioritizing community voices in the design of solutions). That may not always be an extensive, highly designed and resource-intense strategy process but rather a constant reconciliation of planned use cases with impact objectives (e.g. early engagement of communities in product/service development, data ownership among beneficiaries or the prohibition of decision-making without human involvement).

Image credit:
 Education for
 Employment



FIGURE 6

6 The PRISM Framework



The second layer highlights the different adoption pathways. Innovators may choose different pathways depending on their readiness for AI, which will evolve over time. It showcases different approaches to AI implementation depending on the

organization's maturity, starting with low investment and low risk implementation. The examples below illustrate these pathways based on concrete case studies found in the dataset and interviews.

TABLE 1

Social innovator examples of adoption pathways

Adoption pathway	Focus	Example
Conscious tinkering Simple, low-risk implementation leveraging off-the-shelf solutions	Internal	Ashoka: The organization started its adoption of AI with experiments, leveraging open-source AI capabilities, to develop different internal tools while initially keeping the risk profile low. For example, the organization analyses the transcripts of its board meetings through a multi-query process with ChatGPT-4 Turbo to provide recommendations for process improvements and implementation of new AI tools.
Proof of value Implementation of tools with significant efficiency gains or new insights, demonstrating the value of AI to the organization	Internal	Ashoka: For promising applications, Ashoka sets out to prove AI value to the board and staff to encourage widespread internal use. For example, Ashoka has leveraged AI to analyse qualitative insights from its internal database of social innovators (changemakers). That way, it was able to highlight opportunities for new partnerships between them and collective action as a community at any point in time. It also created the foundation for a more diverse AI portfolio, which now counts 21 applications.
Efficiency enhancement Structured adoption of Al tools within the organization to streamline core processes	Internal	Education for Employment (EFE) promotes a culture of Al use by employees to save time and improve internal efficiency in departments such as marketing and grant writing. Al tools such as Microsoft Copilot has enabled team members to save an estimated 10-20% (about 1 to 3 hours) in time versus if they had not used Al at all. In other instances, the organization has deployed Al web scraping tools to predict job trends and inform internal strategy.
Delivery improvements Implement AI with externally facing processes and staff, improving product/service experience and impact	External	LifeBank's adoption of AI has enabled the company to leapfrog barriers presented by supply chain capital limitations. The enterprise developed its proprietary AI tools for supply chain management and to manage last-mile payments. It's AI-powered brainbox, Light, communicates with healthcare workers at the last-mile, improving supply chain efficiency by reducing human errors in decision-making.
Al as a core asset Leveraging Al to develop new solutions and services to achieve the impact objective	External	SAS Brasil is seeking to develop an AI-enabled diagnostic tool drawing on large- scale cervical cancer datasets to improve healthcare outcomes for women in vulnerable communities.
Al-first organization Organizational transformation to become data-driven and Al-enabled in almost every offering and service	External	Geekie made the decision to become a technology organization to fully leverage Al for education. It transformed its basic offering to be free-of-charge to be able to gather data at scale and build Al models. It transformed itself to embed technology into every facet of the organization – from training plans to suggestions for remedial work.

Al implementations insights



Among the 300 Al in social innovation use cases, most innovators (54%) leveraged Al to enhance product offerings (3. efficiency or 4. delivery), 30% created entirely new products or services and 12% leveraged Al to unlock scale for their solutions (5. core asset), 5% leveraged it for internal processes (1. tinkering and 3. efficiency).

When implementing AI across the adoption pathways, different capabilities and risks are considered with varying intensity. This is described in the third layer of the PRISM framework:

- 1. Ethics: Beyond the consideration of general values, principles and the impact mission in layer 1, this capability and risk consideration includes elements such as considering bias in AI models, ensuring transparency of decision-making, accountability for the decisions made and integrating ethical AI as a business principle into the organization at a later stage of the organizational deployment.
- 2. Data: Data strategies may be applied in different facets during the implementation of AI, including an analysis of existing data assets, the cleaning and preparation of data, managing data intake as well as data privacy considerations when using data especially in highly sensitive impact areas such as healthcare.
- 3. Business and organization: Organizational readiness is a key component of implementing AI. This may include setting a clear vision and leadership for the use of AI (especially when leveraging AI as a core asset or transforming the organization to an AI-first organization), a continuous learning strategy, the development of skills and talent as well as change management and the development of an AI or data culture. For social innovators, stakeholder management is particularly important to create the necessary buy-in and trust among key constituents.

- 4. Technology: Many social innovators struggle with the identification of the right technology for their use cases. In early adoption cases, social innovators may simply apply off-the-shelf solutions with low implementation complexity and costs. Later on, model complexity, the need for reinforcement training (especially for use cases with the need for contextual data), interoperability with ecosystem actors or internet access issues may surface as well.
- 5. Costs and metrics: Given the broad availability of AI technology, costs are not much of an issue in early stages of AI adoption. The number of requests with an application programming interface (API) remains low, data volumes tend to be restricted and organizational readiness may not need to be considered. But in laterstage implementations, variable costs quickly ramp up and indirect expenses mount to ensure that AI implementations remain impactful and ethically responsible. Social innovators may need to consider and monitor costs that are not limited to software or licensing fees but also include data acquisition, data preparation or hardware costs.

The following section explores the abovementioned capabilities and risks in more detail. The PRISM Framework offers a modular approach to considering the implementation of AI for impact. On one end of the spectrum, social innovators with strong organizational readiness - solid internal AI skills and capabilities, a healthy enterprise-wide attitude to the technology and a clear articulation and understanding of their data assets and ethical concerns - may apply capabilities and manage risks across the entirety of the framework. Organizations that have not built such internal capabilities may leverage the technology to tinker and therefore use a low-risk entrance point to implement the technology at low cost and with minimal impact on the business strategic goals and ethical concerns.

1.2 Case studies along the adoption pathways

The following use cases illustrate the different intensity levels of AI adoption across the different pathways.

CASE STUDY 1 Ashoka's tinkering approach to internal AI implementation

At Ashoka, the development of AI prioritizes internal deployments to ensure thorough learning and refinement before opening applications to a wider audience. This hesitation stems from a conscious decision driven by the organization's values, ethical considerations and governance protocols.

Among the array of techniques employed, one stands out as particularly effective: pattern identification in large volumes of qualitative data. Leveraging its database of fellows, Ashoka uses this technique to extract valuable insights and generate tailored summaries. For instance, by querying the database for fellows with a connection to intellectual property, the organization classifies and tags these fellows, enabling swift analysis and summary generation.

Self-described as "tinkering", this approach leverages existing organizational data and deploys off-the-shelf solutions (initially ChatGPT Turbo, now ChatGPT 4o and the suite of Claude 3 models) to provide proof of value to the organization. Ashoka's Fellow database has been analysed from a diversity, equity and inclusion (DEI) perspective, so the organization is well aware of the biases. This makes it easier to address risks like structured discrimination by the Al tools. Internal legal and ethics guidelines shape the design of applications and prevent AI developments for certain use cases like automated decision-making or social scoring. And because Ashoka builds solutions internally, it is easy to provide transparency for the recommendations provided by highlighting the steps of the multi-query process. The initial implementation of AI had low requirements for organizational preparedness and featured a rather simple technology stack. Yet it provided the basis and buy-in for what is now a rich portfolio that spans 21 internal applications, stewarded by the Ashoka AI Lab. These applications are now more complex and require General Data Protection Regulation (GDPR) considerations, ethical terms of use, they affect important business processes like fellow selection, and involve higher costs at scale. considerations, ethical terms of use, they affect important business processes like fellow selection, and involve higher costs at scale.

FIGURE 8 | PRISM profile of a "tinkering" application at Ashoka



CASE STUDY 2 LifeBank's "Light" staff recommendation system

Recognizing the limitations of solely human-driven decisions, LifeBank transitioned to leverage AI for enhanced decision support. A brain box, internally referred to as "Light", is a semi-automatic system that offers delivery recommendations for its field staff while allowing human override for more nuanced decisions.

LifeBank's implementation approach is an example of the iterative nature of technological integration aligned with evolving business needs. Ahead of deploying Light, the organization analysed its business requirements and closely engaged stakeholders to identify pain points and explore Al implementation possibilities. It then catalogued its available data to evaluate insights and identified gaps for ongoing collection as

part of its data strategy. These approaches laid the foundation for informed decision-making and resource allocation.

Through early buy-in from its field staff and stakeholders, LifeBank has ensured a smooth transition, mitigating resistance from users accustomed to existing systems. Furthermore, the successful deployment of the product presents an important use case in how integrating human expertise into technology can facilitate seamless adoption and ensure ethical use. This process design makes sure to enable handoff points between AI recommendations and human decision-making. Figure 9 outlines the PRISM profile of LifeBank's implementation of the AI recommendation system.

FIGURE 9 | PRISM profile of LifeBank's "Light" solution



CASE STUDY 3 SAS Brasil's cervical cancer screening

The social enterprise SAS Brasil is in the process of developing an external-facing AI diagnostic tool to predict the risk of cervical cancer and speed up its diagnosis in local communities in Brazil as a response to a need that the organization has identified. Its approach to developing the technology filters through all aspects of capabilities and risks of the PRISM framework (level 3).

The social enterprise prioritizes addressing the ethical implications of AI applications in healthcare and has constituted an internal ethical framework with its partners to guide the tool's development. It also addresses a major challenge of data bias, as datasets are mostly culled from sick patients based in the Global North, with limited data from healthy patients to train the AI. SAS Brasil has prioritized ethical data collection and research integrity as central to its impact mission to bring quality healthcare to communities, safeguarding the awareness that beyond data points, each database entry represents a human life.

SAS Brasil has partnered with local academic and health institutions and embarked on creating expansive datasets representative of diverse populations. These efforts to constitute equitable data are expensive and time-intense, with a strong resource focus on partnerships and ecosystem development.

SAS Brasil has managed to build internal AI skillsets to deploy this technology. The organization's collaboration with local universities includes working with teachers, PhD and graduate students in computer science tracks to develop the diagnostic tool. Figure 10 outlines SAS Brasil's PRISM profile and its intense approach to implementing AI as a core asset.

FIGURE 10 | PRISM Profile of SAS Brasil's cervical cancer screening



Capabilities and risks

Different adoption pathways require different levels of risks and capability assessments.

The five capabilities and risks outlined in the PRISM framework outline potential considerations for social innovators. Table 2 and the interactive assessment complementing this paper allow organizations

to evaluate which of these capabilities and risks are relevant for them, based on their adoption pathways and use case.

TABLE 2 Capabilities and risks by adoption pathway

				Adoption	pathway		
Capability	Details	Conscious tinkering	Proof of value	Efficiency enhancement	Delivery improvements	Al as a core asset	Al-first organizations
Ethics	Fairness and avoiding bias Ensuring transparency Maintaining accountability Ethical Al framework			1			
Data	Cataloging data assets Creating a data strategy Organizing data through architecture Managing data Securing and protecting data Consent and control Anonymization and privacy by design Data governance						
Business and organization	Leadership and vision Stakeholder engagement Continuous learning and adaptation Skill development and acquisition Change management Culture						
Technology	Type of technology Model complexity Training data requirements Interpretability and explainability Latency and real-time processing Scalability and deployment Model updating and maintenance Integration and ecosystem compatability Internet access						
Costs and metrics	Development costs Data preparation costs Data acquisition costs Hardware costs Infrastructure costs Licence fees Personnel and change management costs Training costs Integration costs Legal and compliance costs Metrics and measurement						

Not applicable for this adoption pathway 🔵 Relevant risk or capability for this adoption pathway



2.1 | **Ethics**

Considering ethics capabilities and risks ensures that AI systems are intelligent and equitable at the same time. As social innovators often work with vulnerable and marginalized groups, this is one of the key considerations. The PRISM Framework

Fairness and avoiding bias

Al is a mirror reflecting the data fed into it. When the reflection is distorted by biased data, the outcomes can perpetuate societal disparities. Unfortunately, many Al models have been developed using Western data, which means that their outputs are either not specific to the context of social innovators in low- and middle-income countries. or – even worse – highly biased and can produce results with negative consequences. Organizations implementing Al for delivery, as a core capability or as an Al-first organization therefore need to ensure that their Al systems act as lenses that correct, rather than amplify, biases. This correction

highlights the following ethical aspects that are relevant to different Al implementations. Table 2 allows the read to cross-reference each risk and capability with the adoption pathways.

begins with the careful design of algorithms and the curation of datasets. It is essential to draw data from diverse sources and rigorously test self-developed algorithms across varied scenarios, demographics and use cases. In some cases, this means retraining Al solutions based on more diverse data.

In applications like facial recognition, where historical data may have underrepresented certain demographics, diverse data samples may need to be included. In sentiment analysis, nuances and context may need to be integrated to understand culturally specific sentiments.



SAS Brasil is investing time and resources in the deployment of its AI-enabled diagnostic tool because datasets around health are largely based on Northern Hemisphere data and did not reflect the local communities that SAS Brasil hopes to serve.

Ensuring transparency

The enigmatic nature of Al's decision-making is often compared to a "black box" where inputs and outputs are visible but the processing remains opaque. For organizations dedicated to social impact, peeling back the layers of this black box can be imperative. The decision-making process of Al, especially in complex systems like predictive analytics and recommendation engines, should be made clear wherever possible.

For social innovators developing their own algorithms, transparency involves documenting and explaining the algorithm decision paths in a language that is accessible to all stakeholders. For example, if a predictive model identifies a high-risk group for a certain disease, the factors contributing to this prediction should be clear and comprehensible to medical professionals and patients alike. Open-source tools like LIME (Local Interpretable Model-Agnostic Explanations) can be integrated into models and help explain predictions in high-stakes areas.

MapBiomas trained its machine learning and prediction algorithms to protect the Amazon rainforest against fires on hundreds of million sampled pixels gathered by organizations across different biomes in Brazil from 1985-2022.¹ The data generated in such way is made available free of charge and fully transparent. This is possible through MapBiomas' network of more than 30 organizations, who collaborate on building every dataset. The data is available via Google Earth Engine in order to allow stakeholders across the world to replicate the process for other geographies.²

However, transparency tends to remain an ambition for many social innovators using off-the-shelf, proprietary models (such as ChatGPT or Gemini). These proprietary models do not grant their customers insights into the algorithm and hence render explainability almost impossible. There is a clear need for technology leaders to engage in conversations with social innovators to allow them to highlight such issues and, ideally, influence the technology roadmap so that AI solutions are fit for purpose. For now, in the absence of full transparency, innovators tend to provide "chain of thought" explanations that highlight the multi-step queries fed into the proprietary models.



Dimagi collaborates with developer teams at OpenAI to preview upcoming AI capabilities and showcase how to deploy them for impact. In 2024, Dimagi built a tool based on Voice Engine, a new offering from the technology organization for the creation of custom voices, and ChatGPT-4 to upskill local community health workers at the last mile by providing instructions and training in local languages such as Swahili and Sheng.³

Haqdarshak is deploying AI at the forefront of economic empowerment in India to enable citizens in marginalized and low-income communities to access government resources and initiatives that would otherwise be out of reach. The social enterprise is using AI capabilities such as chatbots and agent management systems to improve products and services in marginalized communities. The solutions complement the efforts of frontline agents, rather than seeking to replace them, making them more efficient and productive. Echoing other social innovators with similar sentiments, Haqdarshak has expressed a desire to share its AI use cases with technology organizations so that, in turn, developer teams can study, understand and ultimately design AI capabilities with an impact-focus built in.

Maintaining accountability

Establishing accountability within AI systems sets up guardrails to ensure that there is a clear path to determine errors and bias. For an impact-first organization, this accountability translates into mechanisms that allow – wherever possible – tracing back through the AI decision-making process.

These mechanisms involve technical tools that track decisions to their algorithmic origins and governance policies that outline who is responsible for tracking, managing and – if needed – reversing the AI's decisions. These human-AI interaction processes with clear handoffs between AI and human operators tend to take a significant amount of time in the solution design process. A high level of accountability fosters trust and ensures that stakeholders feel secure in the knowledge that the AI systems are operating within the bounds of ethical responsibility.

Establishing an ethical AI framework

An ethical AI framework is the blueprint upon which organizations can construct a responsible AI strategy. This framework is not a one-size-fitsall solution; rather, it must be as unique as the mission and context of the organization itself. It encompasses a set of guidelines that navigate the ethical challenges specific to different sectors and AI applications. For instance, an organization using Al in healthcare will need to consider guidelines for patient consent and medical data privacy. In contrast, an organization using Al to support judicial decisionmaking will focus on fairness and avoiding reinforcing historical injustices. Developing this framework involves consulting with experts, studying best practices and engaging with the communities affected by the Al systems.



RECODE's mission to democratize generative AI in the slums and favelas of Brazil began with establishing a coalition to ensure the promotion of digital inclusion and the development of a manifesto to guide the ethical use of the technology and safeguard data generated from its communities.



2.2 | **Data**

Data powers the engine of artificial intelligence. This chapter explores the multi-faceted realm of data in the context of impact-first organizations, outlining

Cataloguing data assets

Understanding data assets entails an inventory of what information is held, where it originates and how it can be expanded in the future. This could range from quantitative data like transaction logs to qualitative insights gathered from direct customer

Creating a data strategy

A data strategy supports the organization's mission. It includes protocols for data collection that prioritize ethical considerations and align with the intended social impact. It supports collecting data that is easy to acquire and seeking information that can

Organizing data through architecture

The data architecture outlines how data will be organized and integrated within the organization. For social innovators, this can mean building systems that are capable of handling diverse data types from various sources. It entails creating data models that accommodate complex social metrics

Managing and preparing data

Effective data management ensures that the right data is available, consistent and usable when needed. It builds on the data architecture and establishes processes for data entry, updates and maintenance. For social innovators, this often means dealing with data that is unstructured or feedback or community engagements. The aim is to identify and catalogue these assets, evaluating how

they can serve immediate needs and forecasting

their potential to address future objectives.

the roadmap to harnessing data for the greater

good while maintaining the utmost respect for

individual rights and societal norms.

fill knowledge gaps and drive the social mission forward. The strategy can include a plan for analysing data to gain insights into the effectiveness of the organization's actions and the needs it aims to address.

and support efficient data retrieval for analysis and reporting. Data modelling tools can help visualize the dependencies of data (e.g. the connections between transaction data, demographic user information and insights about their experience with the solution).

comes from unconventional sources, necessitating a more nuanced approach to categorization and use. Oftentimes, data needs to be digitalized and translated from local languages to be usable with Al tools.



Often, social innovators operate in environments where information is predominately kept on paper. Dimagi, for example, has developed its own Optical Character Recognition (OCR) tool to digitize paper-based information in local languages.

Securing and protecting data

Social innovators often work with vulnerable populations and, in some cases, must rely on offline solutions that require data storage on local systems. Data security ensures that data is protected from unauthorized access and breaches. It involves implementing technical controls like encryption and network security and administrative controls such as access policies and staff training.

Consent and control

Consent and control are about ensuring that individuals have a say in how their data is used. This involves creating clear consent forms, providing users with options to opt-in or opt-out of data collection, and being transparent about data use policies. For applications that involve personal data, these considerations are ethical and legal necessities. Social innovators often face the challenge that their users are not accustomed to technology and/or have low literacy rates, which prohibits them from making informed decisions about the use of their data and either revert to the standards proposed by the social enterprise or deny any access to their data.

Anonymization and privacy-by-design

Anonymization techniques are used to prevent personal data from being tied back to individuals. This is especially relevant for datasets used in training AI, where the removal of identifiers can help to protect privacy while still allowing for valuable insights to be gained. It is important to note that a surprisingly low number of data points is needed to identify a person. In a historic survey, MIT researcher Latanya Sweeney showed that 87% of US citizens can be uniquely identified by only using their gender, postcode and date of birth⁴ – information readily available in healthcare data. Privacy-by-design tries to address these issues by integrating privacy into the development process of new products, systems, or processes.

Data governance

Data governance refers to the overarching management of data availability, usability, integrity and security in organizations. It involves policies, procedures and controls that ensure the effective and ethical management of data. For social innovators, data governance is critical in ensuring that data practices align with the social mission and comply with regulatory requirements.



↓ Image credit: High Resolves



2.3 | Business and organization

One of the key elements of Al implementation is the preparation of the business and organization. Structuring a business and organization for Al integration demands a concerted effort across various facets.

Leadership and vision: Towards an AI-enhanced future

For social innovators implementing AI as a core asset or even transforming to an AI-first organization, leadership commitment is essential to fostering an environment where AI can thrive. Leaders must communicate a clear vision of how AI can enhance the organization's impact, where it is off limits and how AI initiatives align with the organization's overarching social objectives. Consider the case of a non-profit aiming to leverage AI for environmental conservation. Leaders there should ensure that AI solutions, such as predictive analytics for poaching activities, are both technologically sound and aligned with the mission to protect wildlife. It can also be helpful to implement a blocklist – use cases and areas for which the organization prohibits the use of AI. For example, educational organizations may decide to rule out the use of AI to grade students.

Stakeholder engagement: The foundation of community-centric AI

Stakeholder engagement is the practice of involving those who will be affected by AI – the users, community members, ethical watchdogs and

policy influencers. It is about actively listening and incorporating their perspectives into the AI design and deployment process.



For example, RECODE's AI coalition consists of representatives from schools, slums, community centres and government, as well as corporate partners such as Microsoft and Accenture and ecosystem enablers.

Continuous learning and adaptation: Evolving with communities

The AI landscape, especially its ethical and privacy facets, is in constant flux. Mature organizations may establish mechanisms for ongoing education and policy adaptation. This can be achieved through regular training sessions, ensuring constant knowledge inflow through, for example, AI ethics journals, or participating in industry consortia focused on responsible AI. For instance, a social enterprise working with AI in education might stay abreast of new research on educational equity to ensure its AI tools do not inadvertently widen achievement gaps.

Skill development and acquisition: Building AI capability

An honest assessment of existing skills against the needs of AI initiatives becomes critical as the group of users expands. Impact-first organizations might partner with academic institutions to access cutting-edge AI research or provide staff with training to understand AI tools. For instance, a charity looking to use AI to analyse donor data for better engagement strategies may implement training programmes to upskill their marketing team in data analytics.



Image credit: Dimagi

Change management: Managing the AI transition

Al can transform job roles and operations that can often come with fears among staff – some very concrete such as job loss, others more abstract such as lack of technological understanding. These tensions necessitate a change management strategy as technology applications trickle through the organization. Success factors include

Culture: Fostering an AI ethos

Al-first organizations will need to develop an Alinformed culture that embraces change, is ready to learn from both successes and setbacks and critically reviews any further use of Al based on ethical and performance concerns. One example of encouraging a culture of innovation where trial communicating clearly about changes, involving employees in the AI transition journey and providing support such as mentorship programmes. An organization using chatbots for customer service will need to transition some employees from traditional support roles to chatbot management and oversight roles, which requires planning and support.

and error are part of the journey is organizations that hold "fail fairs" where teams share lessons from unsuccessful projects. This cultural shift is vital for staff to feel comfortable working alongside AI systems and viewing them as tools to enhance their work rather than threats to their roles.



Ashoka encourages this culture and enables teams to tinker with a wide range of AI capabilities to explore solutions that can improve internal efficiencies.



2.4 | Technology

Selecting the right AI technology is a pivotal decision and one of the most complicated for organizations at an early stage of their adoption. The choices can significantly influence an organization's ability to achieve its social mission effectively. Organizations need to consider the technology selection process, ensuring the chosen AI aligns with the specific needs and constraints of socially driven initiatives.



↑ Image credit: High Resolves	Type of technology
	 Textual data: Natural language processing (NLP), sentiment analysis and chatbots are valuable in processing textual information. An example is a mental health support platform that uses NLP to interpret user messages and provide personalized support. NLP and chatbot solutions like OpenAI's ChatGPT have been leveraged to create responsive chatbots for mental health, providing real-time, conversational support.⁵
	BarefootLaw uses a chatbot to disseminate legal information and provide first line legal guidance. Sentiment analysis tools like IBM Watson Tone Analyzer can assess customer feedback on social programmes, determining the public's sentiment towards initiatives. ⁶
	Visual data: Computer vision and image recognition technologies are critical to analysing visual data. For instance, an AI-driven system that helps farmers detect plant diseases from images could use image recognition.
	Microsoft's Computer Vision application programming interface (API) has aided conservation efforts by analysing satellite images for deforestation. ⁷ Google's Cloud Vision API has been used by healthcare apps to analyse medical imagery to detect diseases. ⁸
	Auditory data: Speech recognition and translation are indispensable in interpreting and interacting with auditory data. A mobile app translating spoken language in real-time is a typical application aiding communication in multilingual communities.
	Speech-to-text services like Google Speech-to-Text enable real-time transcription services for community consultations. Al-powered services like DeepL Translator can assist NGOs in breaking language barriers during international relief efforts.

User behaviour data: Recommendation systems and predictive analytics can enhance services by analysing user behaviour data.

For example, Geekie, an e-learning platform, uses these technologies to suggest courses to learners based on their past interactions.

Haqdarshak is extending the scope of the metrics of its AI solution beyond mere functionality to evaluate its influence on user behaviour. While monitoring its government benefit support programmes, the team scrutinizes whether its solution is inadvertently triggering disadvantageous reliance on credit and debt versus more prudent financial habits. Continuous monitoring provides behavioural data that enables the team to close down certain programmes or otherwise make adjustments for responsible deployment.



Micro-lending platform Kiva partnered with DataRobot to create and deploy models that predict which loans should be promoted based on their unlikeliness to be funded (prioritizing promotion on hard-to-fund projects)

Model complexity: Balancing accuracy with resources

The complexity of AI models ranges from simple decision trees to complex deep learning networks. The choice depends on the required accuracy versus available computational resources.

For example, a start-up providing diagnostic services may use a complex model for high accuracy, while a non-profit with limited resources may opt for a simpler model that still provides valuable insights.

MapBiomas, for instance, opted for a highly complex, Deep Neural Network classification model for classifying some rainforest areas where the simple analysis and classification of the pixeled imagery itself was not possible and required contextual analysis. This classification model is based on mathematical calculations that can perform machine learning and visual pattern recognition of the target image itself in relationship to surrounding areas.

Training data requirements: Quantity vs quality

The amount and type of training data required by AI technologies vary. Where deep learning needs large datasets, some machine learning algorithms can operate on smaller, less resource-intensive data. While a system like OpenAI's GPT requires extensive data to fine-tune language models, smaller, custom-built models can function with limited data. A social enterprise in the educational sector might use a small language model (SLM) to create localized educational content, requiring far less data than its larger counterparts.

Interpretability and explainability: Trust through transparency

As outlined, the interpretability of AI models is vital in applications where AI decisions have a significant impact, such as in credit scoring for underbanked populations. Transparent models can help build trust within the community by allowing stakeholders to understand and challenge AI decisions. Tools like local interpretable model-agnostic explanations (LIME) can help explain predictions from custombuilt models in high-stakes areas such as healthcare diagnostics. This transparency is key in building trust among end-users, such as patients and doctors, who need to understand the basis of Al-driven recommendations. Unfortunately, as outlined earlier, broadly available, proprietary models tend to lack this explainability.

Latency and real-time processing: Efficiency in action

Real-time processing is essential for applications like emergency response systems, where AI must deliver prompt results. Organizations should seek technologies that balance real-time processing with the constraints of their operating environment, especially in low-connectivity areas. TensorFlow Lite allows AI solutions to be deployed on mobile devices with limited internet access, process data locally and upload information when connectivity is reestablished. Solutions like these could be crucial for applications like real-time language translation devices used by field workers in remote locations with poor connectivity.

Scalability and deployment environment: Growth and accessibility

Al solutions should scale with the organization's growth and be compatible with their deployment environment. A health surveillance system, for

example, needs to function across different regions, adapting to the varying scales of data and environment constraints.



TensorFlow Lite allows for the deployment machine learning models on mobile and internet of things (IoT) devices, facilitating scalability and ease of updates, which is crucial for health-tracking apps in low-resource settings.

Model updating and maintenance: ensuring longevity

The ability to update and maintain AI models with new data ensures their relevance over time. Organizations might favour models that support incremental learning to stay current with evolving social phenomena.

Integration and ecosystem compatibility: Seamless operation

Al technologies must integrate with existing systems and fit within the larger technology ecosystem. This is critical for organizations that rely on various data sources and require analytics platforms to operate harmoniously. An example is integrating a new Al module into an existing platform for community feedback, which requires compatibility with the current database structures. Organizations with a need to make data and data events available to multiple stakeholders might use solutions like Apache Kafka to ensure that the data flows efficiently between their AI applications and systems in the ecosystem, ensuring consistency and reliability of the services provided.

Internet access: Connectivity considerations

In areas with limited internet access, Al technologies need to operate effectively offline or with minimal data transmission. Models capable of

running on local devices or efficiently syncing data when connectivity is available are preferable.

High Resolves, a social enterprise working with young people to help them build future-ready skills, has factored limited connectivity into its AI strategy, allowing its AI-powered training system to be deployed in communities with low access to technology in Asia, Africa and Australia.



Ushahidi, an open-source platform, has successfully enabled local devices to collect data, which is synced whenever the connection permits, providing critical information during crises in low-connectivity areas.



2.5 | Costs and metrics

The deployment of Al in impact-first organizations is an investment in technology and the organization's very ethos and mission. Understanding the costs associated with this deployment and the metrics by which its success is measured is

Development costs

Expenses in AI development can vary widely. Tools and platforms like TensorFlow, PyTorch or Keras offer no-cost entry points but the associated expenses come from the need for specialized personnel. For instance, the development of AI for a healthcare app that predicts patient

Data preparation costs

Before data can train or be used by AI, it must be cleansed and structured – a process that often represents a significant portion of the total AI investment. For example, organizations like J-PAL have gathered extensive field data for socioeconomic research that can be leveraged for AI but needs to be prepared accordingly.

This means hiring data scientists and investing in the time-intensive tasks of data labelling and

Data acquisition costs

Where organizations either lack data or need to acquire complimentary, high-quality data, this can become a major cost driver. For example, an urban planning organization might purchase satellite imagery data to analyse city expansion patterns or invest in IoT sensors to gather real-time environmental data. fundamental. Organizations need to understand the pragmatic aspects of AI implementation from a cost perspective and through metrics that ensure alignment with social impact goals.

treatment outcomes might involve AI experts and collaboration with clinicians to provide medical insights. Cloud-based development environments such as Google Colab can also help launch with low cost but may introduce costs at scale as development grows.

cleaning, which can be facilitated by tools like Amazon Mechanical Turk or Figure Eight for crowdsourced data annotation. This step is particularly critical for applications in which AI is expected to interpret complex and nuanced social data, such as understanding patterns in urban mobility or educational attainment. Non-profits like DataKind provide volunteer expertise for tasks like cleaning and labelling data, which can otherwise be a significant expenditure.

Collaborations with universities or public data initiatives can be a cost-saving measure, providing access to rich datasets at reduced or no cost. For example, MapBiomas datasets are freely available and have been developed in collaboration with universities across all countries of operation.

Image credit: LifeBank



Hardware costs

The computational demands of AI can necessitate robust hardware. While cloud computing options like Microsoft Azure, AWS EC2 instances or Google AI Platform offer scalability, they may lead to high ongoing expenses – especially at scale. Social innovators can leverage programmes such as Entrepreneurs for Positive Impact to access free credits and technical support.

Infrastructure costs

Infrastructure goes beyond mere hardware. It encompasses the entire ecosystem needed to support AI workloads. This includes networking equipment to ensure fast data transfer and storage solutions for the vast amounts of data processed

Licence fees

For certain proprietary Al solutions, licence fees can be substantial. However, vendors often negotiate these fees and some may offer tiered pricing or discounts for non-profits. An example is In contrast, one-time investments in in-house servers and GPUs for organizations like those analysing large-scale genomic data for public health research can be more cost-effective in the long run but require significant upfront capital.

by Al systems. Cloud services, such as Microsoft Azure, Amazon S3, Google Cloud or IBM Cloud, provide scalable solutions but can introduce variable costs based on usage levels.

the use of IBM Watson in the education sector for personalized learning experiences, where licensing costs are adjusted for institutional budgets.

engineers and data specialists need to work in

tandem to integrate AI into existing systems.

Moreover, reorganizing staff and operations to

integrate AI – like the World Bank's restructuring to

incorporate data-driven decision-making - can incur

Personnel and change management costs

The transition to Al-centric operations can be complex. It is not just about hiring data scientists but also about the change management experts who guide the transformation. These costs are seen in initiatives like a municipality adopting Al for traffic flow optimization, where urban planners,

Training costs

Staff training is an ongoing expense, vital for maximizing the benefits of Al tools. Online courses from platforms like Coursera or edX on Al and machine learning can be part of a cost-effective strategy to upskill staff. Additionally, dedicated workshops on ethical Al use and data privacy add consultancy and training costs.

depth to the team's expertise. Beyond technical training, costs include educating the broader team on AI benefits and workflow integration, similar to the training Amnesty International provides its staff when adopting new data tools.

Integration costs

Integrating AI into existing IT systems is a bespoke process, often requiring custom solutions. For instance, a non-profit integrating AI into its donor management system to predict fundraising trends must ensure compatibility with the current customer relationship management (CRM) API, which might entail custom API development or middleware.

Legal and compliance costs

The AI legal landscape is complex and everchanging and varies by jurisdiction. Ensuring compliance, particularly for global organizations like Doctors Without Borders using AI to manage patient data in different countries, can result in significant costs for legal reviews and technology audits. For instance, ensuring AI complies with the GDPR in the European Union (EU) may involve expenses for both technological adjustments and legal expertise.

Metrics and monitoring

When it comes to evaluating the impact and performance of AI solutions in the social sector, the metrics extend beyond traditional business key performance indicators (KPIs) to reflect the organization's social mission. Some metrics and indicators to consider:

- Efficiency gains: Quantifying the reduction in time or resources, like the reduced hours spent by caseworkers on data entry thanks to Alpowered CRM systems.
- Cost-efficiency: Hard currency savings achieved by automating tasks or reducing errors, exemplified by AI systems that help optimize energy use in community projects.
- Accuracy improvement: Measuring the increased accuracy of predictions, analysis or diagnostics (compared to human performance).
- Reach: The number of incremental individuals directly benefiting from the AI solution, which can be seen in the expanded reach of disaster response efforts aided by AI-driven logistics planning.

Legal advisory fees can accumulate when navigating the contractual nuances of Al partnerships. Engaging with Al vendors, navigating data sharing agreements and protecting the intellectual property of Al-generated content necessitate expert legal input, as seen with social platforms moderating content through Al.

- Outcome improvement: Measurable improvements in the targeted social issue, akin to the reduction in recidivism rates when AI is used for personalized re-entry programmes.
- Learning and adaptation: The AI system's ability to improve over time, reflected in more accurate predictions or recommendations, as observed in adaptive learning platforms for education.
- Stakeholder satisfaction: Feedback from users and beneficiaries on the AI solution's effectiveness and usability, often collected through surveys and direct feedback.
- Data-driven insights: The new and actionable insights provided by AI, which can inform policy decisions or the creation of new services.



The non-profit Crisis Text Line uses AI to analyse message data and improve service. Its metrics include the number of crises averted and response time improvements.

Conclusion

Social innovators are leading the AI revolution, applying it in sectors such as healthcare, education and environmental conservation to significantly boost their impact on complex social challenges. Their approach, as outlined in the PRISM Framework, sets a standard for ethical AI integration across sectors, emphasizing the balance between organizational readiness, ethical considerations and potential benefits.

These trailblazers demonstrate that Al's reach extends beyond commercial uses, significantly enhancing how organizations tackle societal issues when aligned with a clear social mission. The PRISM Framework captures these practices, offering an iterative approach to Al implementation and a robust methodology to align Al initiatives with impact goals along the varying degrees of adoption. It shares best practices that can inform Al adoption regardless of industry and encourages organizations at all stages of readiness to deploy Al for impact in a way that fits their current understanding of the technology. Social innovators' progressive use of AI is a model for responsible and impactful technology use. Looking ahead, incorporating their approaches into AI strategies will be vital to harnessing AI's full potential to benefit humanity. Technology leaders are encouraged to engage with these innovators to prepare the future roadmap of AI for impactful application, jointly addressing issues such as transparency, explainability or interoperability. This paper highlights accomplishments in AI, addresses gaps and encourages ongoing innovation and collaboration in the AI for social good sector.

↓ Image credit: High Resolves



Contributors

Lead authors

Daniel Nowack

Head, Global Alliance for Social Entrepreneurship, World Economic Forum

Adedoyin Omidiran Initiatives Lead, Global Alliance for Social Entrepreneurship, World Economic Forum

Yucca Reinecke Manager Technology Consulting, Financial Services, EY

Sreenivasa Sunku Director, Global Client Technology, EY

EΥ

Steven Deng Consultant, EY Australia

Hunter Hanneman Consultant, EY US

Jun Hao Senior Associate, EY Singapore

Gillian Hinde EY Global Corporate Responsibility Leader, EYGS

Laura Turkington EY Global Commercial and Innovation Leader, EYGS

David Weiss Consultant, EY Canada

Emily Will EY Global Impact Entrepreneurship Leader, EYGS

Microsoft

Emma Urrutibehety Business Manager to the Executive Vice-President; President, National Transformation Partnerships

Anthony Virapin Leader, Entrepreneurship for Positive Impact

World Economic Forum

François Bonnici Director, Schwab Foundation for Social Entrepreneurship; Head of Foundations

Connie Kuang Lead, Metaverse Value Creation

Benjamin Larsen Lead, Artificial Intelligence and Machine Learning

Cathy Li Head, Artificial Intelligence, Data and Metaverse; Member of the Executive Committee

Acknowledgements

The authors thank the following members of the Schwab Foundation, representatives of the Forum's network of Young Global Leaders, and members of the Al for Social Innovation workstream of the Global Alliance for Social Entrepreneurship for their insights and invaluable contributions to this paper:

Gerald Abila

Founder and Chief Executive Officer, BarefootLaw, Uganda

Rodrigo Baggio President and Founder, Recode, Brazil

Mehrdad Baghai Co-Founder and Chief Executive Officer, High Resolves, Australia

Andrew Baird Chief Executive Officer, Education for Employment (EFE), USA

Sabine Zink Bolonhini Chief Executive Officer and Co-Founder, SAS Brazil, Brazil

Sangu Delle Chairman and Chief Executive Officer, CarePoint, Ghana

Aniket Doegar Co-Founder and Chief Executive Officer, Haqdarshak Empowerment Solutions, India

Jim Fruchterman Founder and Chief Executive Officer, Tech Matters, USA

Temi Giwa-Tubosun Founder, LifeBank, Nigeria

Njideka U. Harry Founder and Member of the Executive Board, Youth For Technology Foundation (YTF), USA

Production

Danielle Carpenter Editor, Eagle Eye Communications

Bianca Gay-Fulconis Designer, 1-Pact Edition

Martha Howlett Editor, Studio Miko Mark Horoszowski Chief Executive Officer and Co-Founder, MovingWorlds, USA

Jonathan Jackson Founder and Chief Executive Officer, Dimagi, USA

Vukosi Marivate Chair of Data Science; Associate Professor of Computer Science, University of Pretoria, South Africa

Odin Mühlenbein Co-Founder, Ashoka Al Lab, Germany

Savita Muley Chief Product Officer, Haqdarshak Empowerment Solutions, India

Ayo Olufemi Chief Technology Officer and Co-Founder, LifeBank, Nigeria

Marta Perez Cuso Economic Affairs Officer, UN Economic and Social Commission for Asia and the Pacific, Bangkok

Claudio Sassaki Entrepreneur-in-Residence, Stanford University, USA

Adriana Mallet Toueg Chief Medical Officer and Co-Founder, SAS Brasil, Brazil

Endnotes

1.	Tasso Azevedo, MapBiomas. (2023, 22 February). Interview. (S. Reid, Solutions Insights Lab, Interviewer). <u>https://sjn-</u> static.s3.amazonaws.com/solutions-insights-lab/Complete+Interview+with+Tasso+Azevedo+(MapBiomas).pdf.
2.	MapBiomas. (2023). Algorithm Theoretical Basis Document (ATBD). <u>https://mapbiomas-br-site.s3.amazonaws.com/</u> ATBD_MapBiomas_Fogo_Cole%C3%A7%C3%A3o_1.pdf.
3.	OpenAI. (2024). Navigating the Challenges and Opportunities of Synthetic Voices. <u>https://openai.com/blog/navigating-</u> the-challenges-and-opportunities-of-synthetic-voices.
4.	Sweeney, L. (2002). k-Anonymity: A model for protecting privacy. International Journal on Uncertainty, Fuzziness and Knowledge-based Sytems, vol. 10, no. 5, pp. 557-570.
5.	Hale, E. (2023). ChatGPT is giving therapy. A mental health revolution may be next. <i>Al Jazeera</i> . <u>https://www.aljazeera.</u> com/economy/2023/4/27/could-your-next-therapist-be-ai-tech-raises-hopes-concerns.
6.	Alvino, N. (2020). Sentiment Analysis and Visualization for Citizen Opinion in Social Media. [Paper presentation]. Proceedings of The National Conference On Undergraduate Research (NCUR) 2020 Montana State University, Bozeman, MT. <u>https://libjournals.unca.edu/ncur/wp-content/uploads/2021/01/3269-Alvino-Nicholas-FINAL-FINAL.pdf</u> .
7.	Microsoft. (n.d.). Applications - Planetary computer. https://planetarycomputer.microsoft.com/applications.
8.	Google Health. (n.d.). Al-enabled imaging and diagnostics previously thought impossible. <u>https://health.google/health-</u> research/imaging-and-diagnostics/.
9.	lbid.



COMMITTED TO IMPROVING THE STATE OF THE WORLD

The World Economic Forum, committed to improving the state of the world, is the International Organization for Public-Private Cooperation.

The Forum engages the foremost political, business and other leaders of society to shape global, regional and industry agendas.

World Economic Forum

91–93 route de la Capite CH-1223 Cologny/Geneva Switzerland

Tel.: +41 (0) 22 869 1212 Fax: +41 (0) 22 786 2744 contact@weforum.org www.weforum.org