

Title	Responsible AI in philanthropy: Determinants of adoption across technology, organization and environment
Authors	Treacy, Stephen;Wang, Fan;Demesi, Hoperodney;Wang, Kemeng;Bhosle, Poornima;Twiss, Sean;Heavin, Ciara
Publication date	2026-04-09
Original Citation	Treacy, S, Wang, F, Demesi, H, Wang, K, Bhosle, P, Twiss, S & Heavin, C 2026, Responsible AI in philanthropy: Determinants of adoption across technology, organization and environment. in UK Academy for Information Systems Conference 2026 (UKAIS 2026), Sheffield, United Kingdom, 9-10 April 2026. Association for Information Systems Electronic Library, AISeL, pp. 1-28.
Type of publication	Conference item
Rights	© 2026, the Authors.
Download date	2026-09-22 04:31:29
Item downloaded from	https://hdl.handle.net/10468/18996

Responsible AI in Philanthropy: Determinants of Adoption Across Technology, Organisation and Environment

Dr. Stephen Treacy, stephen.treacy@ucc.ie
Fan Wang, 124101112@umail.ucc.ie
Hoperodney Demesi, 124116134@umail.ucc.ie
Kemeng Wang, 123107573@umail.ucc.ie
Poornima Bhosle, 124118803@umail.ucc.ie
Sean Twiss, 124101501@umail.ucc.ie
Prof. Ciara Heavin, C.Heavin@ucc.ie

Cork University Business School, University College Cork, Ireland

Completed Research

Abstract

Artificial Intelligence (AI) presents opportunities for philanthropy, but adoption remains constrained by ethical sensitivities, resource limitations, and mission-driven commitments. Existing frameworks such as TAM, UTAUT, and HOT-Fit lack the specificity required to capture the ethical, governance, and legitimacy challenges of this sector. This study adapts the Technology–Organisation–Environment (TOE) framework, enriched by the Context–Input–Process–Product (CIPP) model, to examine responsible AI adoption in philanthropic organisations. Drawing on fourteen semi-structured interviews with philanthropy leaders and AI practitioners, the findings reveal how technological barriers such as poor data quality and interoperability intersect with organisational challenges of cultural readiness, leadership, and ethical oversight, and environmental pressures of regulation, donor expectations, and reputational risk. The adapted TOE+CIPP framework highlights that trust, accountability, and mission alignment are as decisive as technical readiness. The study contributes theoretical extensions to IS adoption research and practical insights for mission-driven, resource-constrained sectors navigating responsible AI adoption.

Keywords: Responsible AI, Philanthropy, Technology–Organisation–Environment (TOE) framework, Adoption, Governance, Trust

1.0 Introduction

Artificial Intelligence (AI) systems are widely recognised as a transformative technology, reshaping economies and organisations through automation, analytics, and predictive capabilities. In the private sector, adoption is accelerating as firms seek efficiency gains, innovation, and competitive advantage (Venkatesh et al., 2016). By contrast, socially accountable sectors such as philanthropy face a different reality. Here, resource constraints, heightened accountability to donors and beneficiaries, and ethical obligations shape the conditions of adoption in ways that diverge sharply from profit-

driven enterprises (Ebrahim, 2005; Prentice, 2016). These contextual factors make philanthropy a compelling use case for exploring how AI can be integrated responsibly without eroding legitimacy or mission alignment. . Environmentally, the energy demands of training large-scale models have raised concerns about carbon emissions, prompting calls for “Green AI” approaches that prioritise efficiency alongside accuracy (Schwartz et al., 2020; Strubell et al., 2019). Societally, responsible AI emphasises fairness, accountability, transparency, and human-centred design (Floridi et al., 2018; Jobin et al., 2019). While both dimensions matter, it is the societal aspect that resonates most strongly with philanthropy. Because trust is central to sustaining donations and programme legitimacy, philanthropic organisations must ensure that AI adoption not only improves operations but also reinforces ethical accountability and stakeholder confidence.

IS research offers a variety of models to explain technology adoption, which provide useful but incomplete tools for this challenge. At the individual level, the Technology Acceptance Model (TAM) (Davis, 1989) and Unified Theory of Acceptance and Use of Technology (UTAUT) (Venkatesh et al., 2003) examine how perceptions shape behavioural intention. At the organisational level, the Technology–Organisation–Environment (TOE) framework highlights how internal capacity and external pressures interact with technological features (Baker, 2011; Tornatzky et al., 1990). Sector-specific models such as HOT-Fit in healthcare further underline the importance of alignment across domains (Yusof et al., 2008). Yet, while these approaches help explain readiness and fit, they do not sufficiently account for the governance, responsible, and ethical dimensions that are central to AI adoption in mission-driven contexts.

This study addresses that gap by adapting the TOE framework through the Context–Input–Process–Product (CIPP) evaluation model (Stufflebeam, 2000). Drawing on a comprehensive review of the relevant literature and insights from fourteen semi-structured interviews with experts in AI, philanthropy, and digital transformation, this study explores the technological factors influencing adoption decisions, the organisational capabilities and cultural dynamics that facilitate or hinder implementation, and the external pressures exerted by donors, regulators, and peer institutions. It also examines the additional governance and normative requirements, particularly transparency, accountability, and mission alignment that are necessary to

operationalise responsible AI in philanthropy contexts. These insights underpin an adapted TOE+CIPP framework that integrates the European Commission's *Ethics Guidelines for Trustworthy AI* (European Commission, 2019) and aligns with the requirements of the EU *Artificial Intelligence Act* (Regulation (EU), 2024).

The contributions are threefold. Firstly, the study extends TOE by embedding governance and ethical considerations. Secondly, it provides philanthropic organisations with a roadmap for adopting AI responsibly in resource-constrained environments. Thirdly, it demonstrates how high-level regulatory and ethical principles can be translated into actionable governance practices. While philanthropy is the empirical setting, the findings have wider implications for socially accountable sectors such as healthcare, education, and international development, where responsible AI adoption is pressing.

2.0 Literature Review

2.1 Operationalising Responsible AI

The ethics of AI has been shaped by a growing consensus around values such as fairness, accountability, transparency, and human agency (Floridi et al., 2018; Jobin et al., 2019). However, these principles have been repeatedly criticised for lacking actionable specificity, leading to what Mittelstadt (2019) terms a “principles-to-practices gap.” Further, Schiff et al. (2020) argue that organisations struggle to embed such values within day-to-day processes. For philanthropic organisations, this challenge is acute: legitimacy and trust are central to their survival, and ethical commitments that remain abstract risk being perceived as symbolic rather than substantive.

Scholars have responded by designing artefacts that institutionalise responsibility. Gebru et al. (2021) propose “Datasheets for Datasets,” requiring structured documentation of data provenance, consent, and representativeness. This logic is extended through “Model Cards,” which standardise disclosures of intended use, subgroup performance, and limitations (Mitchell et al., 2019). These artefacts make transparency auditable by embedding documentation within workflows. Organisational approaches have evolved further, with Raji et al. (2020) proposing algorithmic auditing frameworks designed to systematically identify and mitigate risks throughout the

development and deployment lifecycle. Wieringa (2020) reframes algorithmic accountability itself, arguing that responsibility depends on clearly defined actors, forums, standards, and consequences. This highlights accountability as relational and institutional, not merely technical. The fairness literature reinforces this perspective. In a survey of bias mitigation techniques, Mehrabi et al., (2022) conclude that technical fixes alone are insufficient; effective fairness requires governance structures that resolve normative trade-offs across the lifecycle. These insights are especially relevant for nonprofits, where transparency and equity are integral to donor relations and public legitimacy. Embedding artefacts and accountability processes therefore serves not just compliance but also the cultivation of trust.

2.2 Governance, Regulation, and the Nonprofit Context

Governance research situates these artefacts within regulatory and institutional frameworks. The EU Artificial Intelligence Act, analysed by Veale & Zuiderveen Borgesius (2021), introduces a risk-based model requiring documentation, oversight, and monitoring for high-risk systems. These obligations elevate practices such as audits and documentation from voluntary guidelines to mandatory compliance. Hauer et al., (2023) consider transparency and review mechanisms, concluding that process-level evidence such as logs, datasets, and decision rationales—is more effective for accountability than explanation interfaces alone. Indeed, critical scholarship cautions against superficial implementations. Current practice may be characterised as “transparency theatre,” in which disclosure mechanisms signal openness without enabling genuine scrutiny (Ananny & Crawford, 2018). This warning aligns with nonprofit governance studies. Ebrahim (2005) warns of “accountability myopia,” where compliance-oriented metrics satisfy funder demands while undermining organisational learning. Prentice (2016) highlights the financial fragility of nonprofits, showing how short-term pressures distort strategic decision-making. These dynamics mean that an overemphasis on formal compliance could burden nonprofits without addressing deeper questions of equity or mission alignment.

Sector-specific studies further underscore these risks. Goldkind et al. (2025) analyse nonprofit human services, identifying how hype and harm accompany AI adoption. They propose heuristics for aligning adoption with organisational missions. Plaisance (2025) meanwhile examines philanthropy more directly, arguing that its reliance on

donor trust and vulnerable beneficiaries generates both suspicion and hope around AI. Similarly, adjacent research in the public sector provides further lessons. For example, Mikalef et al. (2022) demonstrate that organisational capabilities strongly mediate AI assimilation in government agencies, suggesting parallels with nonprofits, which face similar legitimacy pressures but with fewer resources.

The implication is clear, governance mechanisms must be thoughtfully adapted to resource-constrained environments (Cabayé, 2025), while simultaneously fostering trust and accountability. For nonprofits, adopting tools such as model cards or audits can strengthen legitimacy, but only if they are adapted to avoid compliance overload. The literature suggests that proportionate, evidence-based practices are more sustainable than wholesale importation of corporate or state models.

2.3 IS Adoption Theories and Their Limits for AI

IS research has produced a rich set of models to explain technology adoption and use. At the individual level, TAM (Davis, 1989) and its successor, UTAUT (Venkatesh et al., 2003), emphasise perceptions of usefulness, ease of use, social influence, and facilitating conditions as determinants of behavioural intention. These models remain influential, but their focus on user perceptions provides only a partial account of AI adoption in philanthropic organisations, where decisions are often driven less by end-user intention and more by strategic, governance, and accountability considerations.

At the organisational level, the TOE framework (Baker, 2011; Tornatzky et al., 1990) has become a cornerstone. TOE conceptualises adoption as shaped by three domains: the technological (e.g., relative advantage, complexity), the organisational (e.g., resources, managerial support), and the environmental (e.g., competition, regulation). Its value lies in flexibility: it has been applied across domains from supply chains to healthcare. However, TOE has two important limitations when applied to AI in nonprofits. First, it typically frames adoption as a static decision point, assessed through cross-sectional surveys, rather than as a process of ongoing assimilation. This makes it ill-suited for AI, where models evolve through retraining, redeployment, and monitoring, each introducing new risks. Second, TOE's environmental dimension emphasises market and regulatory pressures but under-specifies normative expectations such as fairness, accountability, and legitimacy. For philanthropic organisations, these

normative concerns are not peripheral, they are central to sustaining donor and beneficiary trust.

Extensions of TOE in sector-specific contexts help illustrate these limitations. The HOT-Fit model in healthcare emphasises alignment across domains and highlights that misfit undermines system success even where adoption occurs (Yusof et al., 2008). While valuable for its socio-technical perspective, HOT-Fit was not designed to capture external governance obligations such as algorithmic audits or lifecycle monitoring, which are increasingly central to AI adoption. Similarly, nonprofit accountability research cautions that adoption cannot be judged by efficiency alone: “accountability myopia” can distort priorities when organisations respond to narrow funder metrics at the expense of long-term learning (Ebrahim, 2005).

Beyond adoption, IS scholarship has also examined how technologies are assimilated into organisations. Orlikowski (1992) reconceptualised technology as enacted in practice, highlighting how ongoing use shapes organisational structures and routines. Lyytinen & Newman (2008) proposed a punctuated socio-technical change model, showing how adoption often unfolds in cycles of disruption and reconfiguration rather than linear progression. More recently, researchers have explored AI assimilation in government agencies to better understand how organisational capabilities and legitimacy pressures mediate outcomes (Mikalef et al., 2022). These perspectives underscore that AI adoption is best understood as a dynamic process requiring iterative alignment, not as a one-off decision.

Complementary insights also come from research on digital infrastructures. Henfridsson & Bygstad, (2013) describe the generative mechanisms through which digital infrastructures evolve, while Rai et al., (2012) show how interfirm IT capabilities enable value co-creation across organisational boundaries. Applied to AI in philanthropy, these studies highlight that adoption must be embedded within a broader ecosystem of donors, regulators, and beneficiaries. Therefore, decisions are not just organisational but infrastructural, shaping and shaped by the wider institutional environment.

Taken together, this literature suggests that while IS adoption theories provide valuable foundations, they require extension to capture the realities of AI in nonprofit contexts. TAM and UTAUT highlight user acceptance but neglect governance; TOE provides a flexible structure but treats adoption as static and under-specifies normative pressures; HOT-Fit brings in socio-technical alignment but omits governance; assimilation and digital infrastructure studies highlight dynamics and interdependencies but have yet to be integrated with adoption frameworks. For philanthropy, where accountability and legitimacy are paramount, these gaps are especially salient. What is required is an approach that extends TOE with explicit governance and evaluation dimensions into every stage of adoption.

3.0 Methodology

3.1 Research Design

This research adopts an interpretive, qualitative design to investigate how philanthropic organisations engage with AI. Qualitative approaches are well suited to examining phenomena that are socially embedded, emergent, and insufficiently theorised (Klein & Myers, 1999). The adoption of AI within nonprofit and philanthropic contexts represents such a phenomenon: organisational practices are still in formation, governance frameworks remain unsettled, and the implications of responsibility and sustainability are only beginning to be articulated.

This approach provided flexibility to elicit rich, detailed accounts while ensuring that core themes relevant to the study's theoretical framing were addressed consistently across participants. Interviews allowed respondents to reflect on organisational strategies, governance practices, and perceptions of legitimacy in their own words, while also permitting the researcher to probe issues that emerged during discussion. The resulting data offer insight into both common sectoral dynamics and the organisation-specific contexts in which AI adoption decisions were made.

3.2 Theoretical Framing

The study is informed by the TOE framework (Baker, 2011; Tornatzky et al., 1990), which conceptualises adoption decisions as shaped by technological factors, organisational readiness, and environmental conditions. While TOE is a widely used

and adaptable lens, it has notable limitations when applied to AI. Specifically, TOE does not adequately capture the governance accountability dimensions that have become central to AI adoption. To address this, TOE was integrated with the Context–Input–Process–Product (CIPP) evaluation model (Stufflebeam, 2000). The CIPP model emphasises the importance of assessing context, inputs, processes, and products to generate comprehensive insights. When combined with TOE, this approach enabled the study to address not only the structural factors influencing adoption, but also the evaluative dimensions of governance, and accountability. Context in this study encompassed regulatory requirements, sectoral expectations, and donor demands. Inputs included resources, capabilities, and technological readiness, but were extended to capture sustainability-related concerns such as environmental footprint and efficiency. Processes referred to the practices through which AI systems were implemented, including auditing, documentation, and stakeholder engagement. Products covered outcomes, both in terms of performance and efficiency, and in terms of legitimacy and accountability. This integration of TOE and CIPP ensured that adoption was studied as a lifecycle process shaped simultaneously by organisational, environmental, and normative factors.




Figure 1. Integrated TOE x CIPP Model of Responsible AI Adoption in Philanthropy

This model illustrates adoption as a dynamic lifecycle in which Technology, Organisation, and Environment jointly shape adoption context, inputs, governance processes, and outcomes. Solid arrows depict the primary influence pathways from

TOE conditions through the CIPP stages to adoption outcomes (legitimacy, capability, and implementation), while the solid feedback loop indicates that outcomes reshape future environmental expectations and organisational capacity over time. The dashed feedback arrow represents an indirect pathway through which adoption experiences may prompt technological reconfiguration (e.g., infrastructure upgrades or vendor switching), typically mediated through organisational investment decisions. We present this as explanatory pathways grounded in qualitative evidence, intended to support analytic generalisation and inform future hypothesis testing.

3.3 Data Collection and Analysis

Data were generated through 14 semi-structured interviews carried out between May and July 2025. Participants were purposively sampled (Patton, 1990) to represent a range of philanthropic organisations varying in size, funding model, mission, and degree of technological sophistication. This approach was chosen to capture diversity in adoption experiences and to illuminate both commonalities and differences across contexts. The target group included mid-to senior-level professionals and academics with experience in philanthropy, AI, or digital transformation. To ensure relevance and diversity across the three focal domains, fourteen experts were recruited: six from philanthropy, and four each from AI and digital transformation (See Table 1).

Interviewee Role
1. Senior Digital Transformation Consultant
2. Educational Technology Expert
3. Information Systems Researcher
4. Data Analytics & Business Intelligence Manager
5. Third Sector Consultant
6. Technology Consultant
7. Philanthropic Legislation Lobbyist
8. Director of Philanthropy
9. Senior Digital Transformation Consultant
10. Digital Strategist
11. Head of Transformational Philanthropy
12. AI Ethics Ambassador
13. Head of Philanthropic Giving

Table 1. Overview of Interviewees

Recruitment began with professional networks and was supplemented through snowball sampling as participants recommended colleagues with relevant expertise. All participants provided informed consent before the interviews commenced. Interviews were conducted via secure videoconferencing platforms to accommodate geographical spread, and each lasted between 45 and 75 minutes. Ethical approval for the study was obtained, in line with the study's ethical approval and participant consent, the data are securely held and used exclusively for this research to ensure confidentiality and protect sensitive organisational information.

The interview protocol was informed by the TOE+CIPP framework (See Appendix A). These explored perceptions of AI's benefits and risks, organisational readiness and resources, the influence of external pressures such as regulation and donor expectations, the governance mechanisms in place to ensure accountability, the consideration of sustainability factors, and the outcomes experienced in terms of both efficiency and legitimacy.

Interview transcripts were coded manually and collaboratively using thematic analysis to generate practice-informed insights grounded in real-world contexts (Braun & Clarke, 2006; Ando et al., 2014). A hybrid coding strategy was employed, combining inductive open coding with deductive alignment leveraging the core dimensions of the TOE framework (Fereday & Muir-Cochrane, 2006; Gale et al., 2013) and CIPP frameworks as sensitising concepts rather than rigid analytical templates (Blumer, 2006). This process resulted in a formal codebook comprising 22 themes across technological, organisational, and environmental domains. Thematic categories were refined through iterative group discussion and agreed by majority consensus, enhancing analytical consistency and reliability (MacQueen et al., 1998; Nowell et al., 2017).

While many participant responses mapped directly onto TOE and CIPP constructs, the data also revealed sector-specific dynamics characteristic of philanthropy. These included tensions between donor priorities and organisational missions, staff scepticism regarding algorithmic transparency, and challenges in balancing operational efficiency with institutional legitimacy. Such themes were retained and integrated into the analysis

to avoid constraining interpretation solely within existing theoretical categories. Attention was given not only to convergent experiences of adoption but also to cases of resistance, delay, or abandonment, as these provided critical insights into the limits and boundary conditions of responsible and sustainable AI adoption in the philanthropic sector.

4.0 Findings

The findings are organised using the TOE framework, which provides the structural lens for understanding AI adoption in philanthropic organisations. Each dimension of TOE is examined through the evaluative categories of the CIPP model, which brings additional clarity to how readiness, implementation, and outcomes unfold in practice. This combined framing allows the analysis to move beyond adoption as a single decision event, instead highlighting the lifecycle dynamics of responsibility and sustainability.

4.1 Technology

The technological environment of philanthropic organisations emerged as a foundational determinant of AI adoption. Interviews revealed fragmented infrastructures, data quality concerns, interoperability challenges, compliance bottlenecks, and acute shortages of in-house expertise. Together, these factors illustrate how *inputs, processes, and products* in the CIPP model condition the technological dimension of the TOE framework.

Technological inputs: fragmented infrastructures and weak data foundations.

Participants repeatedly described their core systems as outdated and poorly integrated. As one remarked, *“many systems still offer a basic list view or Excel tracker,”* while another admitted, *“we’re at the early stages of data maturity... people are using SharePoint and spreadsheets.”* Such infrastructure fragmentation makes experimentation difficult and reflects a wider pattern of limited digital readiness across the nonprofit sector (Rathi & Given, 2010; Salamon & Anheier, 1998). Even where organisations sought real-time analytics, they found existing systems unable to cope. As one manager put it: *“It’s no longer just, ‘what did we do last week?’ It’s more like, ‘what’s happening now, in real-time?’... but our systems aren’t built for that.”* The

absence of standardised APIs further compounds the problem, forcing costly, bespoke integrations and discouraging innovation (Gangwar et al., 2014).

Processes: governance gaps, compliance risks, and interoperability. Across organisations, data was scattered across CRMs, grant management systems, and local drives, often with significant gaps or duplicates. One participant was explicit: *“Good data hygiene is what makes any AI viable.”* The lack of shared taxonomies and centralised repositories undermined confidence in outputs, as illustrated by one hospital where mismatched patient arrival times rendered dashboards unreliable. Interviewees noted that up to 60 per cent of AI budgets could be consumed by data cleansing alone. Despite these challenges, few organisations had formalised governance mechanisms. Several pointed to the absence of stewardship roles or accountability systems, leading to persistent errors and risks of bias. Adoption was further constrained by integration barriers with external systems. One participant noted: *“We’re still unsure about how to secure it when integrating third-party AI APIs. That’s the next bottleneck.”* In regulated domains such as healthcare, compliance frameworks like HL7 limited interoperability: *“Any system we implement has to be compliant with HL7... otherwise, it’s a non-starter.”* Comparable issues were reported in finance and education, where data privacy and auditability requirements raised the threshold for experimentation (Stahl, 2021; Wirtz et al., 2019).

Processes also reflected strong risk aversion. Data Protection Officers frequently intervened to halt pilots over compliance and reputational concerns. One participant explained: *“We’ve flat-out banned ChatGPT on internal data unless we can prove isolation.”* This reflects broader anxieties about breaches and misuse, particularly after high-profile incidents involving generative AI platforms. As another noted, *“When it comes to tech, it’s all about compliance... making sure systems work together... and the institution is protected.”* This defensive posture channels innovation into low-risk domains while discouraging high-impact applications, consistent with prior research that finds nonprofits often treat innovation as liability rather than opportunity (Stahl, 2021).

Technological products: outcomes shaped by scale, skills, and safeguards. Where pilots did proceed, scalability was a frequent barrier. Systems that functioned under controlled conditions often faltered in real-world use. One participant recalled how a pipeline handling 10 GB in two hours expanded to eight hours when scaled, delaying

outputs and frustrating decision-makers. Such performance gaps meant proof-of-concept systems rarely matured into dependable tools.

At the human level, skills shortages severely limited capacity to exploit AI. *“There’s always a capacity issue... we’re not a big team,”* reflected one manager. Others highlighted that even where staff had basic programming knowledge, AI-specific skills were missing. As one interviewee put it: *“Most of our staff haven’t even heard of prompt engineering... a massive digital skills gap.”*

Some organisations experimented with lightweight governance mechanisms to mitigate these challenges. For instance, traffic-light categorisation was used to assign risk levels to proposed use-cases, enabling faster decision-making for non-experts. One participant explained that such systems already existed in other parts of philanthropy, such as fundraising approvals, and could be adapted for AI. These tools offered a pragmatic way to reconcile the tension between enabling innovation and ensuring accountability. Taken together, the findings show that the technological dimension of AI adoption in philanthropy is shaped less by enthusiasm for innovation than by the constraints of infrastructure, data quality, compliance, and expertise. Through the lens of TOE, these factors demonstrate limited technological readiness, while the CIPP model reveals how inputs (fragmented systems and data), processes (governance and compliance mechanisms), and products (scalability and legitimacy outcomes) interact to determine what is technologically possible. Without foundational upgrades and robust safeguards, AI risks remaining theoretical: a possibility acknowledged by multiple participants who described their organisations as “stuck” at the pilot stage.

4.2 Organisation

The organisational dimension of the TOE framework emerged as a decisive factor in shaping how philanthropic institutions adopted, resisted, or adapted to AI. Across the interviews, participants highlighted uneven readiness, gaps in expertise, the importance of leadership and change champions, cultural resistance, and the need for alignment with mission and values. Viewed through the lens of the CIPP model, organisational inputs, processes, and products reveal how structural and cultural conditions either enabled or inhibited responsible and sustainable AI adoption.

Inputs: readiness, expertise, and leadership support. A consistent theme was the uneven organisational readiness for AI. Participants stressed that technical expertise alone was not decisive; rather, process maturity, leadership engagement, and

institutional coherence were critical enablers. As one put it, *“You don’t need technical expertise to implement technology. What you really need is process knowledge and leadership.”* This underscores that readiness is as much about organisational processes and leadership posture as it is about technical infrastructure.

Staff expertise and training were nonetheless important, especially for building confidence and reducing resistance. One participant emphasised: *“If you don’t get those people on board early, forget it... You need to find a ‘what’s in it for me.’”* Here, adoption hinged on whether employees perceived tangible benefits in their day-to-day work. Participants also pointed to a disconnect between high-level strategy and frontline realities. When AI was not aligned with workflows or values, projects stalled despite formal resourcing.

Leadership engagement was consistently identified as a critical input. Effective leadership was not about formal authority but about demonstrating commitment, securing resources, and aligning projects with organisational goals. *“A project like that would typically start at a lower level and then work its way up to the board,”* one participant explained, underlining how board-level sponsorship could shape momentum. In this regard, AI adoption required visible executive backing to move beyond experimentation.

Processes: champions, planning, governance, and inclusivity. Change champions played an essential role in translating strategic interest into practical uptake. As one interviewee described, *“You need a change champion... someone inside the organisation who understands the system deeply and can support others.”* Yet without institutional authority or alignment with formal policies, even passionate champions found it difficult to sustain impact. Mid-level coordination roles were equally important: *“There should be someone in the middle who takes overall responsibility... who understands what each team member does.”* This echoes the CIPP model’s emphasis on process discipline, without structured roles, initiatives lost coherence.

Several participants noted that philanthropic organisations embraced AI enthusiastically but often without clear strategies. Projects initiated in response to sectoral trends or peer pressure lacked robust planning and risk management. *“You need a roadmap, clear stages... timelines, training, data migration plans,”* one participant urged, warning that vague goals and shifting expectations undermined sustainability. Planning was seen not only as a technical necessity but also as an ethical one. As one participant stressed, *“You need to be doing research and making sure the*

organisation's ethics are being followed." Steering committees, phased vendor engagement, and internal consultations were cited as practices that anchored projects in both governance and ethics.

Organisational processes also reflected the influence of regulatory scrutiny. One interviewee highlighted how accountability structures were expected by regulators: *"That's why whistleblower protections exist... regulators expect firms to know that and not exploit customers' cognitive biases."* This demonstrates how process design was informed by both internal strategy and external oversight, reinforcing the role of CIPP in evaluating alignment with ethical and regulatory expectations.

Products: cultural readiness, trust, and alignment with mission. The outcomes of organisational practices were most evident in terms of cultural readiness and trust. Nearly half the participants identified scepticism, fear, and disengagement among staff as barriers to adoption. AI was often perceived as a threat to job security. One participant described how *"AI discussions were seen as a precursor to job cuts."* Such mistrust was exacerbated by weak communication and exclusion from decision-making. Even well-funded initiatives faltered when implemented in cultures resistant to change or unclear about AI's purpose.

Participants stressed that successful outcomes required deliberate trust-building and inclusion. *"Aligning AI tools with mission and values requires thoughtful, inclusive design,"* one explained. Generational divides further complicated adoption: *"In a traditional culture with stakeholders from different generations... the culture gap is something senior leadership must consider seriously."* These insights emphasise that products of adoption are not limited to technical efficiencies but include cultural shifts, employee trust, and organisational cohesion.

The interviews revealed frustrations with superficial AI adoption. *"Too many organisations merely bolt AI onto their systems... They hope for the best and chase profit rather than aligning every AI initiative with clear ethical standards and corporate purpose,"* observed one participant. This reflects the risk that AI projects, if not rooted in mission and values, can erode legitimacy rather than enhance it.

Inclusivity extended beyond internal culture to the design of tools for external stakeholders. Several interviewees criticised the lack of user-centred approaches. *"The end user gets forgotten... cost, technical capability, and profitability tend to take priority over user needs,"* one noted. Another reflected on exclusion: *"I've been in meetings where they're discussing the app... I know it'd be brilliant for people like my*

parents, but they'd never go near it." In some cases, organisations responded with bespoke solutions: *"These are custom, homegrown tools... they work well because they were built for our environment."* Yet participants also acknowledged structural design biases: *"Developers often share the same education, background, and beliefs... what they create ends up reflecting their own perspectives."* For mission-driven organisations rooted in equity, such homogeneity presented ethical challenges that directly impacted legitimacy.

The overarching lesson from the organisational dimension is the need for alignment with mission, values, and strategy. Rather than chasing trends, participants advocated for measured, mission-driven adoption supported by structured planning and ethical safeguards. As one participant summarised, *"It's the job of the change manager... to challenge the technology implementation, understand it, and identify where the pain points are going to be. There's no such thing as a challenge-free implementation."* Others underscored philanthropy's unique dependence on trust: *"Charities don't sell products... we sell hope, visions, good deeds... So, trust is hugely important."* For these organisations, accountability was non-negotiable. *"Take accountability... You must be able to answer the question: who is responsible if something goes wrong?"*

Through the TOE+CIPP lens, organisational readiness (inputs), governance and planning (processes), and cultural trust and legitimacy (products) emerge as decisive factors. Adoption was not determined by technical expertise alone but by whether organisations could embed AI into their missions, values, and cultures in ways that sustained trust. Without this alignment, even promising projects risked stalling or generating reputational harm.

4.3 Environment

The environmental dimension of the TOE framework captures the external pressures that shape how philanthropic organisations pursue AI adoption. Across interviews, participants highlighted the influence of regulation, donor expectations, sectoral norms, competitive pressures, geography, funding constraints, and sustainability concerns. When examined through the CIPP lens, these pressures appear as contextual drivers, resource inputs, procedural safeguards, and legitimacy outcomes.

Context: regulatory change and donor expectations.

Regulation emerged as one of the most salient external forces shaping adoption trajectories. The forthcoming EU AI Act was repeatedly cited as a decisive influence

on pace and scope. As one participant explained, *“We’re watching what’s coming out of the EU... it’ll probably decide how fast we move on this stuff.”* While some welcomed clear principles-based rules around accountability, others viewed stringent requirements, such as those imposed by GDPR, as costly barriers that slowed progress. One interviewee captured this sentiment bluntly: *“The red tape... shapes what’s possible.”* Donor expectations added a parallel layer of pressure. Interviewees stressed that funding often hinged on value alignment and demonstrable impact. As one noted, *“If your project doesn’t align with those values... it doesn’t matter what your peers are doing.”* This reflects how accountability to funders and beneficiaries’ conditions not only whether AI projects are resourced but also how they are framed.

Inputs: sectoral trends, funding constraints, and geographic variation.

Several participants described how perceptions of sectoral trends, peer influence, and competition provided both momentum and anxiety. One remarked, *“There’s pressure to adopt the newest tech... organisations jump on board fast.”* Another added, *“Competitors are making noise about AI, but it’s hard to tell what’s real and what’s hype.”* These comments illustrate how mimetic pressures often substituted for clear strategy, creating a reactive, hype-driven posture. In contrast, some organisations practiced what participants termed *“strategic patience,”* emphasising caution and deliberate pace. *“Slow adopters were being diligent... it gives them a real edge,”* observed one executive, linking patience to more sustainable outcomes. Such organisations often implemented phased pilots, commissioned external audits, and introduced feedback loops before scaling adoption.

Geography and sectoral variation further shaped environmental inputs. European organisations, operating under strict data laws and regulatory scrutiny, tended to proceed cautiously. In contrast, counterparts in parts of the Global South, where formal safeguards were weaker, sometimes prioritised efficiency gains over governance. Funding structures also influenced scope and appetite for risk. As one participant explained, *“Matching the size of the project with the size of the donor... would take out a massive amount of work.”* These dynamics demonstrate how organisational environment, in combination with financial realities, dictates both ambition and restraint.

Processes: culture, due diligence, and risk management.

Culture within organisations, and especially among boards, significantly influenced how external pressures were interpreted. As one participant observed, *“We have a few*

champions at the top... but the board is divided. Some see AI as a gimmick.” These cultural divides often determined whether external pressures translated into adoption or rejection. Due diligence emerged as a central environmental process. One interviewee described it succinctly: *“Due diligence is everything... you want to be absolutely sure that [partners] are morally and ethically sound.”* Such practices reflected strong traditions of accountability in philanthropy, yet participants admitted that sustainability criteria were rarely integrated into these assessments. This created a gap where partners could be vetted for ethical probity but not for environmental impact.

Participants also expressed anxieties about risks inherent in third-party AI systems. These included potential data misuse, unclear accountability, and systemic bias. One interviewee highlighted ethical concerns: *“If I wrote: Here is our Technology Strategy; please advise on how to improve adoption, that would be highly unethical.”* Others flagged privacy risks and the commodification of user data, warning that *“biases in AI can replicate and spread, affecting users or even whole populations.”* These insights underline how environmental processes, such as partner selection and risk management, were often underdeveloped in relation to the specific challenges of AI.

Products: reputation, trust, and sustainability outcomes.

The outcomes of environmental pressures were most evident in how organisations perceived legitimacy and reputational risk. Many participants stressed that philanthropy is uniquely dependent on public trust. As one put it: *“Charities don’t sell products, we sell hope, we sell visions, we sell good deeds.”* Alumni donors and volunteers were described as particularly attentive, amplifying reputational exposure. Organisations often sought to protect trust through communication strategies, charters, and ethical standards. Participants cited instruments such as Donor Charters and CASE’s Bill of Rights as examples of sector-wide tools that helped embed ethical commitments.

Sustainability concerns also emerged as a significant though unevenly integrated product of environmental awareness. Several participants highlighted the energy, and water demands of large-scale AI systems, with one remarking, *“One ChatGPT prompt uses surprising amounts of water... sustainability is becoming key.”* Yet sustainability criteria were not consistently embedded in decision-making or partner evaluations, creating a disconnect between awareness and implementation.

5. Discussion

The findings reveal a dynamic interplay between technological readiness, organisational capacity, and external pressures, such as donor expectations, regulatory demands, and peer influence, that collectively shape adoption decisions. By reflecting on these insights in relation to existing literature, the study extends theoretical understanding and reveals sector-specific considerations with broader relevance for responsible AI.

5.1 Extending IS Adoption Theories

The findings confirm and extend insights from established IS adoption frameworks. The TOE framework provided a useful structure for analysing adoption and the empirical evidence highlights three ways in which it requires adaptation for AI in philanthropy. First, TOE's technological dimension, while attentive to readiness factors such as integration and data quality (Gangwar et al., 2014), does not adequately capture the governance and compliance constraints that participants identified as central to adoption decisions. In this sector, technological adoption is not merely a matter of capability but of legitimacy, as organisations balance donor trust and regulatory obligations. Second, the organisational dimension of TOE is often interpreted in terms of resources and managerial support (Baker, 2011), but here it was evident that cultural readiness, trust-building, and mission alignment were decisive. Participants repeatedly stressed that AI projects failed not for lack of technical knowledge but for lack of inclusive engagement, ethical safeguards, and leadership commitment. This extends adoption research by showing that in philanthropy, organisational "inputs" cannot be reduced to resources: they must encompass values and cultural factors that condition legitimacy. Third, the environmental dimension of TOE typically emphasises competitive and regulatory pressures (Tornatzky et al., 1990). While these were present, the findings show that donor expectations and reputational risk are at least as important. Unlike private firms, philanthropic organisations are structurally dependent on public trust, alumni donors, and volunteers. This means that "environment" must be theorised as including not only market and regulatory forces but also the moral economy of legitimacy. The integration of CIPP was therefore critical. By structuring the analysis around context, inputs, processes, and products, the study moved beyond static adoption models to capture lifecycle dynamics. This addresses critiques of TOE as overly cross-sectional (Lyytinen & Newman, 2008) and responds to calls for adoption research to consider assimilation and long-term accountability (Mikalef et al., 2022).

5.2 Responsible AI in Philanthropy

Much of the literature on responsible AI has focused on corporate or governmental contexts (Jobin et al., 2019; Raji et al., 2020). In philanthropy, responsibility is shaped less by market competition and more by donor influence, reputational exposure, and accountability to vulnerable populations.

The technological findings underscore that without robust data governance, interoperability, and security safeguards, AI initiatives are unlikely to move beyond pilots. Yet resource constraints mean that philanthropic organisations often lack the expertise and infrastructure to develop such safeguards. This creates a paradox: the very organisations most committed to equity and responsibility are often the least equipped to ensure them. This highlights the importance of external support mechanisms, whether from regulators, sectoral bodies, or collaborative networks, to build capacity for responsible AI. Organisationally, the findings show that leadership and cultural readiness are critical but fragile. While executive sponsorship can accelerate adoption, mistrust among staff, especially fears of job displacement, can derail even well-funded projects. This aligns with IS research emphasising the role of user perceptions in adoption (Davis, 1989; Venkatesh et al., 2003) and extends it by showing that in philanthropy, perceptions are bound up with mission and values, not only with perceived usefulness or ease of use. Responsible AI in this context therefore requires inclusive, participatory approaches that foreground staff and stakeholder voices.

Environmentally, the findings highlight how donor expectations and reputational risk condition adoption. Unlike commercial organisations, where shareholder value is paramount, philanthropic organisations “sell hope, visions, good deeds.” This reliance on legitimacy amplifies the reputational risks of irresponsible AI. It also explains the cautious posture many organisations adopt: as one participant put it, slow adopters gained an edge by protecting mission integrity. These insights extend responsible AI research by emphasising that sectoral context shapes not only how responsibility is defined but also how it is operationalised.

5.3 Environment as a Dual Imperative

Sustainability emerged as both an enabler and a constraint. On one hand, participants acknowledged growing concerns about the environmental costs of AI, including energy

and water use (Verdecchia et al., 2023). On the other hand, the findings highlight that this was rarely integrated into due diligence or partner selection processes, creating a disconnect between awareness and action. This resonates with broader critiques that sustainability is often treated as aspirational rather than operational (Bolón-Canedo et al., 2024).

At the organisational level, sustainability was also understood in social terms. Inclusive design, accessibility, and alignment with mission were repeatedly emphasised as conditions for legitimacy. This extends debates on “Green AI” by demonstrating that sustainability must be understood holistically, encompassing environmental stewardship, ethical accountability, and cultural inclusivity. For philanthropy, sustainability is not a secondary concern but part of the sector’s identity: failure to embody it risks undermining trust with donors and beneficiaries alike.

5.4 Implications for Theory and Practice

The study offers several implications. Theoretically, it demonstrates the value of integrating TOE with CIPP to capture the lifecycle dynamics of responsible AI adoption. This combination begins to address limitations in existing adoption models and provides a framework that can be applied beyond philanthropy to other mission-driven sectors such as healthcare, education, or social services.

Practically, the findings highlight the need for philanthropic organisations to invest in AI governance infrastructure, leadership development, and cultural readiness alongside technical tools. Early planning, inclusive engagement, and ethical safeguards are not optional add-ons but central to responsible adoption. Funders also have a role to play by aligning incentives with responsible practices, they can mitigate the pressure to adopt AI prematurely. Regulators, meanwhile, should consider how to provide clarity without imposing compliance burdens that smaller organisations cannot meet.

5.5 Limitations and Future Research

While the study provides new insights, it also has limitations. The focus on philanthropic organisations means findings may not be generalisable to other sectors. The sample size, though sufficient for qualitative depth, cannot capture the full diversity of the sector globally. Future research could build on this study by examining

responsible AI adoption across different mission-driven sectors, by exploring the role of cross-organisational collaborations in building capacity, and by conducting longitudinal studies to trace how adoption trajectories evolve over time. Indeed, there is an opportunity to investigate how philanthropic organisations can act as ethical stewards of AI, influencing broader norms and practices globally.

References

- Ananny, M., & Crawford, K. (2018). Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society*, 20(3), 973–989. <https://doi.org/10.1177/1461444816676645>
- Baker, J. (2011). The Technology–Organization–Environment Framework. In *Information Systems Theory* (pp. 231–245). https://doi.org/10.1007/978-1-4419-6108-2_12
- Blumer, H. (2006). What Is Wrong with Social Theory? In *Sociological Methods*. Routledge.
- Bolón-Canedo, V., Morán-Fernández, L., Cancela, B., & Alonso-Betanzos, A. (2024). A review of green artificial intelligence: Towards a more sustainable future. *Neurocomputing*, 599, 128096. <https://doi.org/10.1016/j.neucom.2024.128096>
- Cabayé, J. (2025). Enhancing Impact Measurement of Philanthropic Organisations: A Human-AI. Last Accessed 25/01/26 <https://www.utupub.fi/bitstream/handle/10024/182660/Thesis1.pdf>
- Davis, F. D. (1989). Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology. *MIS Quarterly*, 13(3), 319–340. <https://doi.org/10.2307/249008>
- Ebrahim, A. (2005). Accountability Myopia: Losing Sight of Organizational Learning. *Nonprofit and Voluntary Sector Quarterly*, 34(1), 56–87. <https://doi.org/10.1177/0899764004269430>

- European Commission. (2019). *Ethics guidelines for trustworthy AI | Shaping Europe's digital future*. <https://digital-strategy.ec.europa.eu/en/library/ethics-guidelines-trustworthy-ai>
- Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations. *Minds and Machines*, 28(4), 689–707. <https://doi.org/10.1007/s11023-018-9482-5>
- Gangwar, H., Date, H., & Rao, A. D. (2014). Review on IT adoption: Insights from recent technologies. *Journal of Enterprise Information Management*, 27(4), 488–502. <https://doi.org/10.1108/JEIM-08-2012-0047>
- Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Iii, H. D., & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, 64(12), 86–92. <https://doi.org/10.1145/3458723>
- Goldkind, L., Ming, J., & Fink, A. (2025). AI in the Nonprofit Human Services: Distinguishing Between Hype, Harm, and Hope. *Human Service Organizations: Management, Leadership & Governance*, 49(3), 225–236. <https://doi.org/10.1080/23303131.2024.2427459>
- Hauer, M. P., Krafft, T. D., & Zweig, K. (2023). Overview of transparency and inspectability mechanisms to achieve accountability of artificial intelligence systems. *Data & Policy*, 5, e36. <https://doi.org/10.1017/dap.2023.30>
- Henfridsson, O., & Bygstad, B. (2013). The Generative Mechanisms of Digital Infrastructure Evolution. *MIS Quarterly*, 37(3), 907–931. <https://www.jstor.org/stable/43826006>

- Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
- Klein, H. K., & Myers, M. D. (1999). A Set of Principles for Conducting and Evaluating Interpretive Field Studies in Information Systems. *MIS Quarterly*, 23(1), 67. <https://doi.org/10.2307/249410>
- Lyytinen, K., & Newman, M. (2008). Explaining information systems change: A punctuated socio-technical change model. *European Journal of Information Systems*, 17(6), 589–613. <https://doi.org/10.1057/ejis.2008.50>
- Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2022). A Survey on Bias and Fairness in Machine Learning. *ACM Computing Surveys*, 54(6), 1–35. <https://doi.org/10.1145/3457607>
- Mikalef, P., Lemmer, K., Schaefer, C., Ylinen, M., Fjørtoft, S. O., Torvatn, H. Y., Gupta, M., & Niehaves, B. (2022). Enabling AI capabilities in government agencies: A study of determinants for European municipalities. *Government Information Quarterly*, 39(4), 101596. <https://doi.org/10.1016/j.giq.2021.101596>
- Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. (2019). Model Cards for Model Reporting. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 220–229. <https://doi.org/10.1145/3287560.3287596>
- Mittelstadt, B. (2019). Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence*, 1(11), 501–507. <https://doi.org/10.1038/s42256-019-0114-4>

- Orlikowski, W. J. (1992). The Duality of Technology: Rethinking the Concept of Technology in Organizations. *Organization Science*, 3(3), 398–427. <https://doi.org/10.1287/orsc.3.3.398>
- Patton, M. Q. (1990). *Qualitative Evaluation and Research Methods* (2nd ed.). London: Sage Publications.
- Plaisance, G. (2025). Artificial Intelligence (AI) in the Context of Nonprofits and Philanthropy: Suspicion and Hope for Researchers and Organizations. *Journal of Philanthropy*, 30(2), e70022. <https://doi.org/10.1002/nvsm.70022>
- Prentice, C. R. (2016). Understanding Nonprofit Financial Health: Exploring the Effects of Organizational and Environmental Variables. *Nonprofit and Voluntary Sector Quarterly*, 45(5), 888–909. <https://doi.org/10.1177/0899764015601243>
- Rai, Pavlou, Im, & Du. (2012). Interfirm IT Capability Profiles and Communications for Cocreating Relational Value: Evidence from the Logistics Industry. *MIS Quarterly*, 36(1), 233. <https://doi.org/10.2307/41410416>
- Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 33–44. <https://doi.org/10.1145/3351095.3372873>
- Rathi, D., & Given, L. M. (2010). Research 2.0: A Framework for Qualitative and Quantitative Research in Web 2.0 Environments. *2010 43rd Hawaii International Conference on System Sciences*, 1–10. <https://doi.org/10.1109/HICSS.2010.317>
- Regulation (EU). (2024). *Regulation—EU - 2024/1689—EN - EUR-Lex*. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>

- Salamon, L. M., & Anheier, H. K. (1998). Social Origins of Civil Society: Explaining the Nonprofit Sector Cross-Nationally. *Voluntas: International Journal of Voluntary and Nonprofit Organizations*, 9(3), 213–248. <https://doi.org/10.1023/A:1022058200985>
- Schiff, D., Rakova, B., Ayes, A., Fanti, A., & Lennon, M. (2020). *Principles to Practices for Responsible AI: Closing the Gap* (Version 1). arXiv. <https://doi.org/10.48550/ARXIV.2006.04707>
- Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green AI. *Communications of the ACM*, 63(12), 54–63. <https://doi.org/10.1145/3381831>
- Stahl, B. C. (2021). *Artificial Intelligence for a Better Future: An Ecosystem Perspective on the Ethics of AI and Emerging Digital Technologies*. Springer International Publishing. <https://doi.org/10.1007/978-3-030-69978-9>
- Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and Policy Considerations for Deep Learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645–3650. <https://doi.org/10.18653/v1/P19-1355>
- Stufflebeam, D. L. (2000). The CIPP Model for Evaluation. In D. L. Stufflebeam, G. F. Madaus, & T. Kellaghan (Eds.), *Evaluation Models: Viewpoints on Educational and Human Services Evaluation* (pp. 279–317). Springer Netherlands. https://doi.org/10.1007/0-306-47559-6_16
- Tornatzky, L. G., Fleischer, M., & Chakrabarti, A. K. (1990). *The processes of technological innovation*. Lexington books.
- Veale, M., & Zuiderveen Borgesius, F. (2021). Demystifying the Draft EU Artificial Intelligence Act—Analysing the good, the bad, and the unclear elements of the

- proposed approach. *Computer Law Review International*, 22(4), 97–112.
<https://doi.org/10.9785/cri-2021-220402>
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User Acceptance of Information Technology: Toward a Unified View. *MIS Quarterly*, 27(3), 425–478. <https://doi.org/10.2307/30036540>
- Venkatesh, V., Thong, J. Y. L., Chan, F. K. Y., & Hu, P. J. H. (2016). Managing Citizens' Uncertainty in E-Government Services: The Mediating and Moderating Roles of Transparency and Trust. *Information Systems Research*, 27(1), 87–111. <https://doi.org/10.1287/isre.2015.0612>
- Verdecchia, R., Sallou, J., & Cruz, L. (2023). *A Systematic Review of Green AI* (Version 3). arXiv. <https://doi.org/10.48550/ARXIV.2301.11047>
- Wieringa, M. (2020). What to account for when accounting for algorithms: A systematic literature review on algorithmic accountability. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 1–18.
<https://doi.org/10.1145/3351095.3372833>
- Wirtz, B. W., Weyerer, J. C., & Geyer, C. (2019). Artificial Intelligence and the Public Sector—Applications and Challenges. *International Journal of Public Administration*, 42(7), 596–615.
<https://doi.org/10.1080/01900692.2018.1498103>
- Yusof, M. Mohd., Kuljis, J., Papazafeiropoulou, A., & Stergioulas, L. K. (2008). An evaluation framework for Health Information Systems: Human, organization, and technology-fit factors (HOT-fit). *International Journal of Medical Informatics*, 77(6), 386–398. <https://doi.org/10.1016/j.ijmedinf.2007.08.011>

Appendix A – Interview Guide

1. Please share your experience with AI and technology adoption in the philanthropic or non-profit sector?
2. What AI tools or technologies do you see as most practical or promising for philanthropic organisations today?
3. From your perspective, what are the biggest technical roadblocks non-profits face when trying to adopt AI?
4. How do issues like poor data quality or weak digital infrastructure affect an organisation's ability to use AI effectively?
5. What kind of in-house technical expertise do you think is necessary to sustain AI efforts in the non-profit sector?
6. What are the AI technologies and associated risks that philanthropic organisations should be actively preparing for?
7. How does leadership and organisational culture influence whether AI is adopted—and how responsibly it's used?
8. What role do staff skills, training, or capacity play in supporting or slowing down AI adoption?
9. What internal resources and ethical considerations are required to make AI adoption viable and responsible in a non-profit setting?
10. How do donor expectations, partnerships, regulations, and public trust shape decisions around AI adoption in philanthropic work?