

AI and the Future of Doing Good

How charities are turning technology into impact, and what gets in the way

JAKE O'GORMAN

AI and the Future of Doing Good

How charities are turning technology
into impact, and what gets in the way

© 2026 Jake O’Gorman. All rights reserved.

No part of this publication may be reproduced, distributed, adapted, or transmitted in any form or by any means without the prior written permission of the author, except for brief quotations in reviews or academic citation.

The research for this book was conducted as part of a Churchill Fellowship awarded by the Churchill Fellowship. Design by Corndel.

First published in the United Kingdom in 2026.

ISBN: 978-1-0369-5092-7

Published by Chain of Thought Books

For permissions and enquiries: hello@chainofthoughtbooks.com

Disclaimer

This book is intended for general information and discussion purposes only. It reflects the author’s research, experience, and professional judgement at the time of writing.

It does not constitute legal, financial, technical, or professional advice, and should not be relied upon as such. Readers should seek appropriate independent advice before making decisions that may affect their organisation, stakeholders, or beneficiaries.

While every effort has been made to ensure the accuracy of the information contained in this book, the author makes no representations or warranties, express or implied, as to its completeness, accuracy, or continued relevance. Technologies, regulations, and best practices in the field of artificial intelligence evolve rapidly, and guidance may become outdated.

The views expressed are those of the author and do not necessarily reflect the views of the Churchill Fellowship, Corndel, or any organisations referenced in this book.

References to organisations, products, or services are provided for illustrative purposes only and do not constitute endorsement. The author and publisher accept no liability for any loss, damage, or adverse outcomes arising directly or indirectly from the use of, or reliance upon, the information contained in this book.

Foreword

Charities are operating in a complex and fast-moving environment. Demand is rising, expectations are shifting, and new technologies – including artificial intelligence – are prompting important questions about how we organise our work and continue to serve our communities well.

Across the sector, many organisations are exploring how tools like AI might help them use their resources wisely, strengthen decision-making, and free people to focus on the human work only they can do. These choices require care. They depend on strong data foundations, responsible governance, and clarity about the outcomes we are trying to achieve.

This book is one contribution to that wider conversation. It brings together perspectives from leaders and practitioners who are thinking about what thoughtful, values-led use of AI might look like in practice. I am grateful to Jake for engaging so widely across the UK and the US and for reflecting the experiences of so many organisations in this work.

I hope this research prompts further discussion about how charities can approach new technologies with confidence, curiosity and care – always guided by their mission, their values, and the needs of the people and communities they serve.



MICHELLE MITCHELL OBE
CHIEF EXECUTIVE
CANCER RESEARCH UK

Reflections from the charity sector

“This is one of the most grounded and honest explorations of AI I’ve seen...For leaders navigating how to use AI without losing their humanity, this is essential reading”.

Courtney Bugler, President & CEO of Zero Prostate Cancer.

“This book offers a clear and accessible overview of both the opportunities and the challenges that AI presents for the charity sector. Its structure, balancing potential benefits with the real barriers facing not-for-profit organisations, makes it a valuable resource for leaders navigating this rapidly developing space. I welcome its emphasis on responsible, people-centred leadership as essential to making informed decisions about AI adoption and usage. As AI is an evolving set of tools available to us, this work contributes thoughtfully to the wider sector conversation on how we can use it to better support our missions”.

Gemma Peters, CEO Macmillan Cancer.

“It is fitting to see the third generation of the remarkable O’Gorman family maintaining his place as a thought leader in charitable fundraising in the UK. This is a must read for all those who want to stay abreast of the latest developments in AI - and frankly, that should mean all of us!”

Gavin Maggs, CEO, Children with Cancer UK.

“Jake O’Gorman has created a must read, current view, extremely well researched book on the state of AI in the nonprofit world. A comprehensive, concise and critical piece of research every nonprofit leader should read”.

Chris Lyons, Group President, NonProfitPRO

‘Jake O’Gorman has produced something rare and urgently needed: a guide to AI in the nonprofit sector that takes both the promise and the peril seriously...The nonprofit sector doesn’t need another technology manifesto. It needs grounded, mission-led thinking about how to use these tools without losing sight of the people we serve. Jake has delivered exactly that”.

Brittney Dunn, COO & Founder, Safe House Project.

“Jake O’Gorman’s research arrives at exactly the right time. This book tackles the real challenges nonprofits face with AI, from data quality to donor relationships, with both rigor and practical wisdom. Essential reading for any leader navigating this transformation.”

Shannon McCracken, President & CEO, The Nonprofit Allianc

“AI and the Future of Doing Good, is essential reading for any leader looking to bridge the gap between high-tech innovation and social impact. Jake doesn’t just understand AI; he understands how to use it to create a “greater good future” for the communities that need it most.”

Mark Casper. Marine Corps Veteran. President & CEO, Tech for Troops.

“AI and the Future of Doing Good provides remarkable insights into the plethora of opportunities and challenges NGO leaders must contemplate as they integrate AI’s transformational capabilities into every aspect of their organizations.”

Brian Rubenstein, President, Rubenstein Impact Group

“Jake’s incredible book really brings this truth home to leaders: we need to get to grips with AI urgently. This ambitious, insightful yet practical book will help leaders understand how AI has changed the context we operate in, and define the decisions they need to make.”

Zoe Amar, CEO, Zoe Amar Digital.

“While there’s hype everywhere about AI, this is a necessary read for the thoughtful, considered questions it raises and the opportunity for readers to find a path that works for their missions.”

Amy Sample Ward, CEO, NTEN

“Jake’s book perfectly captures how AI is shaping the future of the not for profit sector and the communities we support. Acting as a charity leader’s guide ... he articulates what AI can do in service of trust, impact and human connection - whilst laying out the ethical challenges and practical risks leaders have to proactively navigate to prevent harm. It is an important and very welcome contribution to help us create a better world where AI amplifies, rather than replaces, our humanity.”

Paul Excell, Chair AI4C (AI for Charities)

Preface

One of my earliest memories of the charity world is carrying a tray of tea into a room full of people. A trustees' meeting was taking place, and although I was too young to understand what was being discussed, I knew it mattered.

In the space of a year, my family had lost two children, Paul and Jean, to cancer. From that loss, my grandparents, with the support of Princess Diana, founded a small research charity to prevent others from facing the same grief. In April 2025, my grandfather, Eddie O'Gorman OBE, recipient of the Pride of Britain Lifetime Achievement Award and Honorary Fellow of UCL, passed away having raised over £300 million for that charity, now called Children with Cancer UK.

That experience shaped my connection to the charity world. I grew up seeing good people working for the benefit of others, and how carefully decisions are made when every pound matters. As I moved into my career in artificial intelligence (AI), this shaped a belief that the technology, when used well, could support people rather than distance them, and that its use is not only a technical question, but a values-driven one.

When tools like ChatGPT emerged, these two parts of my life began to converge. Friends from the charity sector started asking what this new wave of AI might mean for them, reflecting a growing sense that it could reshape core aspects of charity work, for better or worse. Through groups such as AI4C, a growing network of charities using AI for social good, and through my work at Corndel, where I serve as Director of Data and AI Strategy, I found myself exploring these questions more deliberately.

I am deeply grateful to the Churchill Fellowship, the UK's national memorial to Sir Winston Churchill, and to Corndel, for giving me the opportunity to pursue this research. Together, they enabled me to travel across the United States and UK to meet those leading conversations on the opportunities and challenges AI holds for the charity sector.

This book is the story of that journey, and of the people I met along the way. It is not a technical manual, but a collection of stories, reflections and insights about where AI is already supporting charities, where it is creating new pressures, and the choices organisations are beginning to face.

The title of this book, **AI and the Future of Doing Good**, reflects what I heard most often. For those working in charities, the challenge has always been how to achieve the greatest possible impact. The question now is how AI fits into that purpose. Whether you approach this subject with eagerness, scepticism or uncertainty, I hope these pages help you think more clearly about AI, and find an approach that works for your organisation and the people you serve.



JAKE O'GORMAN
DIRECTOR OF DATA AND AI STRATEGY, CORNDEL
CHURCHILL FELLOW
MARCH 2026

“ This shaped a belief that the technology, when used well, could support people rather than distance them, and that its use is not only a technical question, but a values-driven one. ”

In brief: the future of AI in charities



Artificial intelligence is no longer a distant or abstract prospect for charities. It is already shaping how organisations raise funds, run their operations and deliver services. For many leaders, the question is no longer whether AI will be used, but how, where and to what ends. The choices being made now will shape trust, capability and impact for years to come.

This work explores that moment of choice through quantitative research with 275 charity leaders across the UK and United States, as well as many hours of in-depth interviews with charity leaders, academics and technologists.

Together, these perspectives reveal a defining tension. AI can materially strengthen how charities pursue their mission, opening up new ways to reach people, use scarce resources more wisely, and learn faster from what works. When it is applied with care, judgement and clear intent, it has the potential to become one of the most powerful tools the sector has ever had.

Used without purpose or appropriate safeguards, however, the same tools risk eroding the very trust on which the sector depends. The findings fall into two parts: the first looks at where AI is creating value; the second at what holds charities back.

From technology to impact

Part one focuses on how charities are using AI in practice today. Leaders from organisations including UNICEF, Save the Children, Oxfam, Global Citizen and The Wildlife Trusts describe applications spanning humanitarian response, retail, conservation and digital campaigning.

AI creates value when it helps organisations see more clearly, make better choices with limited resources, and remove long-standing bottlenecks. Its role is not to replace human judgement, but to strengthen it. Three areas stand out.



FUNDRAISING



OPERATIONS



MISSION DELIVERY

FUNDRAISING

In fundraising, AI is reshaping how charities understand and engage supporters. Predictive tools help identify who is most likely to give, lapse, or increase their support, enabling more deliberate use of limited time and resources. Generative tools are supporting the drafting of grant applications, appeals and communications at a speed and scale previously out of reach for many teams.

Leaders describe how these tools can improve targeting, strengthen messaging and increase return on effort at a time when participation in giving is declining and costs are rising. The research is clear that this opportunity carries responsibility as well as reward. Fundraising is relational work, and when optimised narrowly for short-term revenue, AI can weaken trust.

As Nathan Chappell, author of *Nonprofit AI* and a leading voice on AI and philanthropy, observed: "If we use AI poorly, to manipulate people to give or to automate all the humans out of fundraising, we could destroy philanthropy in 15 years."

The opportunity, then, is not to automate fundraising. It is to use AI to help fundraisers be more relational, more consistent and more thoughtful at scale, without losing the human qualities that make giving sustainable.

This tension is sharpened by a wider shift in how people seek information. Fundraising and nonprofit marketing are moving from a search-driven world to an answer-driven one. Increasingly, supporters are asking AI direct questions about causes, credibility and impact, rather than browsing lists of results.

Visibility alone is no longer enough. Charities must be understood, represented accurately and trusted by the systems people turn to for guidance. When they fail, they can reinforce bias and misinformation at scale.

As nonprofit marketing expert George Weiner put it: "AI is going to become the number one philanthropic advisor in the world. It doesn't much matter what you think of AI. It matters a lot more what AI thinks of you."

In operations, the value of AI is often more immediate and easier to realise. Many organisations are already using it to reduce administrative burden, improve knowledge management, strengthen retail performance and make critical processes less dependent on a small number of individuals.

At Restore NYC, an AI-supported referral tool helped stabilise service delivery when a key staff member took extended leave. At Global Citizen, a custom AI assistant removed a major bottleneck in campaign production, extending specialist capabilities across the organisation.

In charity retail, where organisations manage a constant flow of unique, one-off donated goods, AI can play a particularly valuable operational role. Examples from Goodwill of Orange County and the British Heart Foundation show how it can support pricing, routing and online listings, helping maximise the value of each item.

Survey evidence reflects these experiences. Among organisations already using AI, improved operational efficiency is the most consistently reported outcome. For many, operations are where AI first demonstrates tangible value.

It is also in operations that an important reality becomes clear. AI is increasingly embedded in everyday tools, arriving through software updates and new features rather than deliberate organisational choice. This brings risks, but also gives charities access to tools they can build on quickly. As George Weiner observed, AI is now “a side of fries with every purchase”.

This creates both risk and opportunity. With 75% of surveyed charities yet to have a formal AI policy in place, many organisations are still early in their journey. Those that act now have a chance to set clear, proportionate and values-led foundations before habits become entrenched.

In mission delivery, the stakes are highest and the potential impact greatest.

At UNICEF, advanced data modelling supported real-time humanitarian response during the war in Ukraine, improving the placement of services for displaced families. At Safe House Project, AI-enabled anonymous reporting dramatically increased the identification and extraction of survivors of human trafficking. Tools such as Sophia, developed by Spring ACT, show how carefully designed chatbots can widen access to support while preserving safety, dignity and choice.

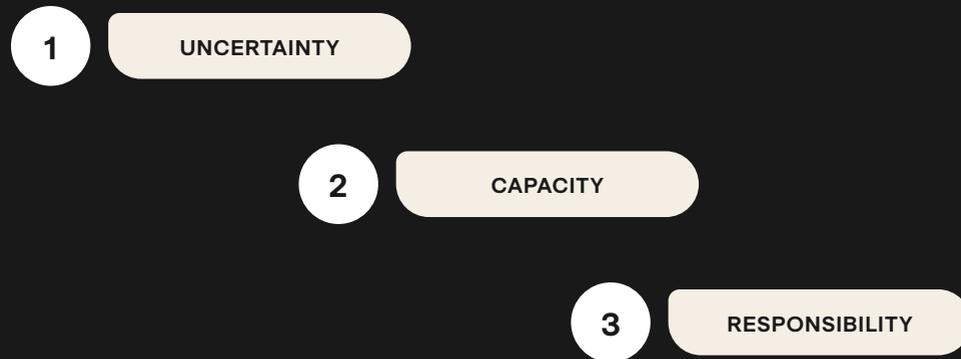
Across these cases, a consistent pattern emerges. The most credible uses of AI come from organisations with deep, long-standing relationships to the problems they are trying to solve. They begin with a lived understanding of risk and harm, set clear boundaries, and remain closely involved at the points where judgement and care matter most.

In these contexts, AI acts as an accelerator, not an authority.



What gets in the way

Part two turns to what is preventing charities from doing more with AI. This research identifies nine recurring barriers, grouped around three interlinked themes:



1

UNCERTAINTY

AI requires leaders to make decisions in fast-moving conditions, often before tools, use cases or implications are fully understood. Benefits can be difficult to predict. Evidence is limited, and mistakes may carry reputational and ethical consequences. While most leaders believe AI could be valuable, many are still working through how best to translate that belief into practical action. Encouragingly, interviews show that confidence grows quickly once organisations begin experimenting in structured, supported ways.

To better understand this dynamic, the research draws on the Unified Theory of Acceptance and Use of Technology (UTAUT), a framework for understanding technology adoption. It suggests that four factors shape whether new technologies take hold:

- **Performance expectancy:** Will this meaningfully improve outcomes, such as service delivery, reach, or freeing up time for mission-critical work?
- **Effort expectancy:** How difficult is it likely to be to realise those benefits in practice, given current capacity, skills and systems?
- **Social influence:** Whose views matter in this decision, including peers, funders, leadership and the wider sector? How much weight do they carry?
- **Facilitating conditions:** Does the organisation feel it has the data, systems, governance and resources needed to act with confidence?

All four matter. But one stands out.

Social influence is the strongest predictor of whether charities intend to adopt AI. This is not about individual enthusiasm or scepticism. It is about permission. Leaders are far more likely to move forward when respected peers are doing so, when boards and executives signal support, and when AI use feels normal and defensible within the sector.

Without that permission, even the strongest use cases stall.

2

CAPACITY

The second theme is capacity. In this research, lack of staff skills and confidence is cited by 57% of respondents as a barrier to doing more with AI. Lack of time follows at 44%, alongside limited funding at 43%.

Many leaders assume that progress with AI requires advanced technical expertise. Closer examination tells a different story. Through detailed skills analysis, this research reveals a consistent gap between the skills leaders assume they need and those that make the greatest difference in practice.

The organisations making the most progress were not always the most technically sophisticated. They were those with strong problem understanding, close familiarity with day-to-day service delivery, and permission to step back, define challenges clearly and explore how AI might help. Technical capability matters, but it is rarely the starting point.

A further paradox appears here. Lack of time is a major barrier to getting started, yet time saved is the most frequently reported benefit among adopters. What holds many organisations back is not time itself, but uncertainty about whether the effort will pay off. As experience grows, confidence in that value increases.

3

RESPONSIBILITY

The third theme is responsibility. Charities operate in high-trust environments, where mistakes can quickly undermine confidence with donors, beneficiaries and the public. It is therefore unsurprising that over half of leaders cite ethical concerns or questions of mission fit as barriers to greater AI use.

These concerns most often relate to environmental impact, risks of bias and fairness, and alignment with organisational values. In many respects, this reflects the sector at its best. Leaders are asking difficult questions early, rather than waiting for problems to arise. The challenge is translating these questions into everyday decisions.

This is where trustees play a central role. Oversight increasingly requires informed questioning about data, governance, risk and long-term implications. Boards do not need to become technical experts, but they do need sufficient understanding to provide challenge, set boundaries and support responsible innovation.

Where this governance is strong, AI can be used with confidence and care. Used well, it can deepen trust, improve decision-making and act at a scale human systems alone cannot reach. Where it is weak, however, the same tools are likely to amplify existing vulnerabilities just as quickly.

AI is, above all, a force multiplier. It magnifies whatever it is applied to: leadership and culture, data and trust. In the end, it reflects the choices, values and priorities of leaders.

The defining question is not what AI can do, but what it will multiply when placed in your hands.

Contents

Foreword	03	Responsibility in high-trust environments	107
Preface	06	Environmental Impact	107
Key Insights at a Glance	09	Ethics	114
Contents Page	18	AI Governance	121
Introduction	20	Conclusion	128
Part One: From Technology to Impact: The Promise of AI	23	Appendix	130
Fundraising	24	Acknowledgements	131
Operations	40	About the Author	132
Mission	54	Research Methodology	132
Part Two: What Gets In The Way	60		
Uncertainty: Why AI Is hard to judge in advance	65		
Difficulty seeing value	65		
Leadership caution	74		
Capacity: The work of making AI possible	75		
Skills	75		
Time and capacity	86		
Budget and funding	92		
Technology and data readiness	99		

Introduction

It was often joked on my travels that if you walked into the average charity office a few years ago and asked how old their computers were, you would get one of three answers: five years, ten years, or twenty years.

As Jim Fruchterman, a MacArthur Fellow and long-time advocate for technology for social good, told me, the sector has been operating inside a kind of nonprofit tech time machine. Innovation has often been slowed by a culture that treats technology as an overhead cost to be reduced, rather than an engine for change.

Yet something is shifting. Walking through these same offices today, the laptops have not changed much, but people's questions have. Artificial intelligence has forced a conversation the sector has managed to avoid for decades. It is no longer simply about updating systems. It is about whether to engage with technologies that could fundamentally reshape how organisations relate to supporters, staff and the people they serve.

To understand this moment, I travelled across the UK and the United States, speaking with charity leaders, technologists, funders, academics and policymakers. Alongside these conversations, I conducted quantitative research with 275 charity leaders to understand how AI is perceived, where it is already being used, and what is holding organisations back. A consistent pattern emerged. Most leaders believed AI could improve their impact, yet only 28% had reached the point of strategic adoption. Many were curious. Few felt ready.

This hesitation is taking place against a challenging backdrop. In the United States, leaders described what is widely termed a generosity crisis, marked by falling participation in giving and low donor retention. In the UK, the picture is different but no less strained. Charity Commission analysis shows rising demand alongside increasing costs, with expenditure growth outpacing income growth for many organisations. Leaders spoke of the growing need to do more with limited resources, alongside a widening awareness that new technologies could help. As one leader put it to me, "someone moved the cheese, and it is not coming back".

A hallmark of the charity sector is its diversity. A global NGO, a small membership charity and a community food bank may share an ethos of service, yet their capacities and digital starting points differ widely. The same is true of AI itself, a term that spans everything from simple workflow tools to complex language models. These umbrella labels can open conversation, but they can also overwhelm. Many leaders I spoke with were not resistant to AI. They simply did not yet know how to approach it with confidence or clarity.

This book aims to provide some of that clarity. It is not a technical manual, nor an argument for adopting AI at any cost. It is a narrative guide, shaped by conversations with people working at the sharp end of charity leadership and technology. It draws on examples ranging from UNICEF teams using satellite imagery to support humanitarian response in Ukraine, to the Safe House Project in the United States using AI to identify and extract women from trafficking networks, alongside reflections from academics, entrepreneurs and sector leaders grappling with how AI intersects with charity work in practice.

It also examines what gets in the way. Across interviews and survey responses, leaders described a familiar set of barriers, explored here through three lenses: uncertainty, where decisions must be made before tools or impacts are fully understood; capacity, where stretched teams and systems limit what is possible; and responsibility, where questions of trust, governance and accountability carry particular weight. Beneath these sit deeper questions about scale, judgement and values in an AI-enabled world.

Together, these describe the context in which charities are making decisions about AI. Many of the leaders I spoke with do not feel an obligation to adopt it. They do, however, feel an obligation to understand it well enough to exercise judgement: to know where it could strengthen their work, undermine trust, and what responsibilities follow if they choose to proceed. This book is written for leaders navigating that space.

From technology to impact: the promise of AI

In recent years, artificial intelligence (AI) has moved from the margins of the charity sector into everyday conversation. What was once discussed only in specialist forums or innovation teams is now appearing in board papers, team meetings and everyday conversations. Charities of all sizes are beginning to test how these tools might help them work more efficiently, understand their communities more clearly, and strengthen relationships with supporters.

The chapters that follow focus on three areas where AI currently shows the greatest promise for charities. The first is fundraising, where AI is beginning to reshape how organisations understand and communicate with supporters. The second is operations, where charities are using AI to reduce administrative burden and improve efficiency. The third is mission, where early experiments point towards new ways of delivering services and supporting beneficiaries.

Together, these perspectives offer a grounded view of how AI is beginning to take shape across the charity sector. They show where value is already being created, where risks and tensions are emerging, and why the promise of AI lies less in the tools themselves than in the judgement with which they are applied. This is not a story of transformation overnight, but of a sector beginning to make choices about how these tools are used, and where their limits should lie.

Fundraising

It is difficult to overstate the importance of fundraising to a charity. As Courtney Bugler, CEO of Zero Prostate Cancer, told me, a chief executive's mind must always circle two questions: money and mission. "Take your eye off either," she warned, "and you have a problem."

Fundraising takes many forms, and the mix looks different for every organisation. Yet it is an area under growing pressure. Charities are operating in an environment shaped by declining participation in giving, low donor retention and rising demand for services. These pressures are felt acutely in fundraising teams, where expectations to raise more with fewer resources are now the norm. It is against this backdrop that interest in AI has accelerated.

Charity finances

The Charity Commission's analysis of annual return data for financial periods ending in 2023 shows that charities collectively spent £95.73bn delivering charitable aims. Expenditure growth (9.6%) outpaced income growth (6.8%) with 42.6% of charities reporting expenditure exceeding income.

When fundraising is broken down into its component parts, it is not hard to see why. Researching donor bases, segmenting audiences, drafting communications, testing messages and forecasting income all rely on pattern recognition, repetition and data. These are areas where AI tools, particularly predictive and generative systems, appear well suited to offer support. Many charities have already been drawn to AI's ability to assist with tasks such as grant applications or first drafts of donor communications. Others are beginning to explore its potential in predictive analytics, an often underappreciated but strategically powerful strand of AI.

Yet at the end of every one of these processes sits a person or an organisation choosing whether to give. At its core, fundraising is relational work. It depends on trust, judgement and a sense of shared purpose. There is a danger of turning

relationships into transactions, and exchanges grounded in trust into those optimised solely for conversion. The question is not whether AI will affect fundraising, but how.

Several interviewees described this as a strategic fork in the road. Used poorly, AI could accelerate the decline in broad-based philanthropy, compressing a long-term challenge into a much sharper collapse. Used thoughtfully, it could help usher in what Nathan Chappell, in his book *The Generosity Crisis*, describes as radical connection, where charities understand supporters better and communicate with greater relevance and care. Chappell put the stakes bluntly. "In *The Generosity Crisis* we talked about 49 years as the point in the US where fundraising dissipates. Yet if we use AI poorly, to manipulate people to give or to automate all the humans out of fundraising, we could destroy philanthropy in 15 years."

Five practical uses of AI in fundraising

01 Donor propensity modelling

Predictive models identifying which supporters are most likely to give or lapse, helping fundraisers prioritise engagement efforts.

02 Precision segmentation and audience insight

AI analysing giving history, engagement and interests to move beyond broad segments towards more relevant, individualised outreach, without automating the relationship.

03 Direct mail and appeal personalisation at scale

Hyper-personalised content in direct mail and digital appeals, delivered at a scale impossible to achieve manually.

04 Grant prospecting and funder matching

AI scanning funding databases, past awards and eligibility criteria to identify suitable grant opportunities and reduce time spent on low-probability bids.

05 Proposal drafting

Generative AI producing first drafts of grant applications trained on historically successful bids.



Direct mail: a tradition reinvented

To understand where fundraising may be heading, it helps to look at where it has come from. Chris Lyons, Group President of NonProfit PRO, reminded me that direct mail remains a backbone of philanthropy in the United States. “Direct mail campaigns make up over half of nonprofit revenue in America,” he told me. “For the \$600 billion raised by nonprofits (in the US) direct mail remains a critically important channel”. A lot of people still write cheques.”

This is not nostalgia, but an acknowledgement that donor behaviour often changes slowly. What has changed is what direct mail can be. Advances in printing technology, supported by AI-driven data processing, now allow personalisation at a scale previously impossible. Lyons described services capable of printing up to 100 million pages a month, each uniquely tailored to the recipient. The result is not simply novelty, but communication that feels visually and emotionally relevant. As Lyons put it, this kind of personalisation would take “a million years” to do manually. With sufficient data, AI makes this feasible.

How AI is reshaping fundraising decisions

Such examples show how far fundraising communications have evolved. And yet while that level of personalisation may feel a stretch for many charities, the generative and predictive technologies behind it are moving at pace. For many teams, using generative AI in day-to-day fundraising is already having a significant impact.

To understand what AI makes possible in fundraising, I sat down with Dr Sven Mikolon, Associate Professor of Marketing at Imperial College Business School, and a leading expert in helping charities use data and AI to improve fundraising outcomes. To get a clearer sense of where the value lies, Dr. Mikolon encouraged me to look not just at how AI supports content creation, but at how it can help charities decide where to focus effort in the first place. In particular, he pointed to the growing use of predictive analytics in fundraising decisions, tools designed to help organisations understand their donor base and make more informed choices about how communications are targeted.

While this type of AI has existed long before the arrival of generative tools such as ChatGPT, I was keen to understand its recent surge in popularity within the charity sector. Dr. Mikolon described two closely connected drivers: the opportunity to improve conversion, and the growing need to make more deliberate choices about where resources are deployed.

Despite the advances in printing technology described by Chris Lyons, direct mail, personalised or otherwise, can be an expensive channel for charities to use. Where campaigns include enriched content such as calendars, inserts or small gifts, Dr. Mikolon explained that the cost of a single mailing can easily reach £4 or £5 per letter, so having the most up to date delivery information is critical to driving performance, and return on investment for direct mail. And given the human element is involved, phone-based fundraising can be more costly still.

A key starting point for any charity looking to use this kind of AI is clarity about what it wants to predict. “You want to start by predicting a certain outcome or KPI,” Dr. Mikolon explained. Typical examples include donor churn, or perhaps the probability that a supporter will upgrade from ad hoc giving to a regular commitment. As charity leaders know well, moving a donor from occasional gifts to regular support can significantly strengthen cash flow.

The challenge many charities face is one or both of the following issues: either that they lack a clear understanding of what drives this behaviour, or the methods

used to increase value and reduce churn are costly and hard to scale. Predictive analytics does not remove those costs, but it can help charities focus effort where the likelihood of success is higher, while also shedding light on what influences decision-making across their community of donors.

Crucially, such insight can often be created from data charities already hold, such as existing CRM data. Dr. Mikolon offered a simple illustration. In a pet focused charity for example, analysis might show that supporters who own multiple dogs are far more likely to respond to certain appeals than those with only one. This is the kind of information that can be collected routinely at the point of donation.

Dr. Mikolon stressed, however, that behavioural data is usually the most informative. By this, he meant records of what supporters have actually done rather than who they are. Giving history, frequency and recency of donations, open rates and responses to emails or appeals, and signs of changing engagement over time can all provide useful signals. Past behaviour, he noted, remains the strongest predictor of future behaviour. The model does not need to understand the underlying motivation. It simply identifies patterns that are difficult to see when reviewing spreadsheets or dashboards.

What makes this more powerful is that prediction is only the first step.

From Predictive to Prescriptive Analysis

While spotting patterns is useful, acting on them is where the real value lies. As Dr. Mikolon explained, AI allows charities to move from predictive to prescriptive analytics. “Once we have the prediction, we can then apply a decision rule.” If a charity can estimate how likely someone is to give, and how much they’re likely to donate, it can use that information to make a decision.

For example if a charity is running a direct mail campaign, it can use an AI model to predict the chance of response from an individual and estimate the expected donation size. If after subtracting the cost of sending the letter the expected return is positive, the systems can trigger the decision to send.

Moving beyond one-size-fits-all

This same approach can also surface a second layer of cost that is easy to overlook: opportunity cost. Dr. Mikolon explained that in upgrade campaigns, charities often focus on the direct cost of the activity, without accounting for the impact on future giving. Contact some donors too soon, and “some of them stop giving.” As he put it: “I’ve just made a donation, and now they’re already calling me”.

Dr. Mikolon described how, for years, it has been standard practice for fundraising teams to use different types of communication, mixing emails designed to show impact and build engagement with others that make a direct asks. Most fundraisers instinctively know that the ideal balance and timing of these approaches varies significantly by donor, yet for most charities, such communications are more often driven by fixed schedules rather than individual insight.

As Dr. Mikolon explained, “The opportunity here is to move away from broad segments and towards individual targeting, so each donor follows a path that fits their own engagement and preferences.”

I came away from my conversation with Dr. Mikolon struck by how predictive and prescriptive analytics can help charities see these relationships more clearly. Not simply as a series of individual asks, but as a sequence of interactions that unfold over time. Used well, these tools offer charities a way to align more closely with the preferences of their supporters.

“ The opportunity here is to move away from broad segments and towards individual targeting, so each donor follows a path that fits their own engagement and preferences. ”



Relational fundraising in an AI age

Seen in this light, it is perhaps no surprise that for Nathan Chappell, AI offers the promise of what he calls precision philanthropy. In this vision, predictive analytics and generative tools combine to enable more personalised and meaningful engagement with supporters. Rather than grouping donors into broad segments, charities can begin to understand donors as individuals, aligning outreach with their specific interests and motivations.

The aim is not to produce more messages, but better ones. This includes directing effort to the right people at the right time, and reducing the risk that outreach has adverse effects. In this framing, AI does not replace the fundraiser. It frees up time for deeper relational work.

Yet any serious exploration of AI in fundraising must also confront the risks involved, which differ depending on the type of technology being used.

With generative AI, the danger is often immediate and visible. If the underlying data is wrong, the output can be deeply inappropriate. As Scott Rosenkrans, co-host of the Fundraising.AI podcast warned, “If you’re using generative AI to feed all the information you have on someone into a letter and send it out, and you have one thing wrong, that’s going to be glaringly obvious. You may have lost that donor forever.” He pointed to the example of a hospital foundation that had recently contacted people who had died in that very same hospital.

Dr. Mikolon pointed out, however, that while such mistakes can be damaging, they are often easier to spot. A human reviewing a draft letter will usually notice if a name is wrong, or if a message feels inappropriate. Having a human in the loop acts as an important safety mechanism.

Predictive AI carries a different kind of risk. Because these models look for broad signals across large populations, individual errors rarely dominate outcomes or make spectacular mistakes. Instead, they can be slightly wrong in many small ways. As Dr. Mikolon explained. “Small errors scale.”

Because predictive outputs are often lists, scores or rankings rather than the content of messages being sent directly to donors, problems may go unnoticed for long periods. This is why Dr. Mikolon emphasised that such models need to be reviewed, retrained and questioned, particularly when organisational policies change or donor behaviour shifts.

“ Let AI handle the information so that humans can handle the relationships ”

For charity leaders, these technologies offer significant opportunities to improve donor experience, but each carries a different kind of oversight. Predictive AI demands measurement, monitoring and governance. Generative AI demands editorial judgement and human review. Neither is a set-and-forget tool.

Such oversight becomes more crucial as a third direction begins to emerge in the wider technology landscape: agentic AI. Where predictive systems surface patterns and generative systems create content, agentic systems combine both with the capacity to act. They take steps towards an objective, iterating based on feedback and outcomes.

This is an important trend because, across the research, the most promising uses of AI were those that supported, rather than replaced, human judgement. As Chappell put it, “Let AI handle the information so that humans can handle the relationships.” His point was that relationship-ending harm is far more likely when AI is used to generate personalised messages than when it is used to surface patterns to inform human decision-making.

Yet what happens when this boundary is crossed?

The rise of the autonomous fundraiser

To make sense of where this line should be drawn, Scott Rosenkrans and I met for lunch in Warwick, New York. He suggested thinking about AI in fundraising in two broad categories: informational and relational. Informational AI processes data and surfaces insight, while relational AI attempts to take on the relationship itself. In fundraising, he argued, AI should remain firmly on the informational side.

He warned in particular about the growing interest in autonomous fundraisers. These are AI agents designed to pursue donations as an objective, capable of writing messages, testing variations, and responding in ways that feel personal. Several fundraising leaders pointed to examples such as Blackbaud's AI fundraising agent, an avatar able to acknowledge gifts, ask donors about their motivations and adapt responses in real time.



(Above) Jake O'Gorman and Scott Rosenkrans

The trade-offs at the heart of AI-led fundraising

The appeal is obvious. Human fundraisers are expensive. They take holidays, get unwell and have off days. AI does not. It relentlessly pursues its objective. Yet this is where Rosenkrans is concerned. If a system is optimised narrowly for revenue, its behaviour will inevitably follow. "AI has been proven to be more successful at manipulating humans than humans are," he noted. His worry was not the use of AI itself, but the risk that short-term gains undermine the foundations of generosity.

An AI agent does not care about the relationship. It cares about the transaction. As he put it, "If I find myself in a relationship where I've been manipulated, I'm not going to want to go back to that relationship again."

Paul Butcher, founder of Commonsensing AI and formerly Executive Director of Strategic Initiatives at Save the Children, echoed this point, observing that emotional connection often switches off once a donor realises they are interacting with a machine. Efficiency is tempting, but the cost may be the humanity that defines the sector.

There is, however, nuance here. Several interviewees noted that this critique often assumes donors are unaware they are interacting with AI. There may be contexts, particularly for small donations, where automated acknowledgement is already the norm and additional personalisation could be beneficial, provided it is transparent. As Chris Lyons observed, donors in these contexts may already expect limited human interaction. The ethical question is not automation itself, but whether its use is honest, proportionate and aligned with donor expectations.

Quality in an age of automated content

For organisations using AI to support human fundraisers, a recurring concern was the rise of low-value content. Often referred to as "AI slop", this is material produced because it is fast rather than thoughtful.

As George Weiner, who heads the nonprofit focused marketing agency Whole Whale, put it, "If you are lazy about it, it is worse than doing nothing at all." He warned that if a major solicitation email is written by AI in a way that makes it clear nobody edited or reviewed it, the organisation risks harming donor relations. His recommendation was a shift from simple prompting to what he calls context engineering. If an AI tool does not understand who you are, your values or your voice, it is operating at a basic level. In practice, that means grounding tools in core documents and leadership voice, and treating AI output as a first step, not the finished article.

Brian Rubenstein, President of the Rubenstein Impact Group, reinforced this view, arguing that AI rarely gets you all the way there. The final work of judgement and tone remains human.

A “Yellow Pages to internet” moment

The reinvention underway is playing out across multiple fronts. One of the most significant shifts is happening online, where established approaches to attracting and engaging supporters are being disrupted. George Weiner described the current moment as a “Yellow Pages to internet” moment.

For years, charities invested heavily in search engine optimisation, writing articles designed to attract donors searching for causes. They relied on the hope that someone typing “hunger charity” or “cancer research” into Google would find them.

Increasingly, potential donors are no longer searching. They are asking.

AI systems are becoming a first reference point, providing answers rather than lists of links. This marks a shift from traditional search engine optimisation to answer engine optimisation, which underpins featured snippets and voice assistants, and towards generative engine optimisation, which shapes how organisations are represented within large language models such as ChatGPT and Gemini.

SEO, AEO and GEO explained

SEO (Search Engine Optimisation)

Creating content that shows up in search engine results when people search for relevant keywords. Success depends on how well it ranks and how often it is clicked.

AEO (Answer Engine Optimisation)

Writing content so AI systems can easily extract a clear answer. This underpins featured snippets, voice assistants and “zero-click” results, where users may not visit a website. CharityComms suggests success depends on clear, concise formatting, such as FAQs, lists and definitions.

GEO (Generative Engine Optimisation)

Optimising how an organisation is interpreted by large language models. Rather than driving clicks, GEO affects how AI systems represent, summarise and recommend organisations within their responses to open-ended questions.

As donors move from searching to asking, discoverability increasingly depends on how AI systems interpret and represent a charity, not just on website rankings.

When visibility no longer means clicks

New AI tools give charities a real opportunity to reach people differently. Yet they also change the rules.

Visibility no longer guarantees engagement. Increasingly, people get the answers they need without ever clicking through to a website². Research published in 2024 and 2025 shows that AI summaries, often called “zero-click” results, now satisfy 40–60% of informational searches. Across sectors, this is driving 15–25% declines in organic traffic, even while brand awareness holds steady or grows.

What this means for charities

An analysis of 17 nonprofits³ found that by mid-2025, organic search traffic had fallen by 13% year on year, while brand-name searches increased by 19%. While early signs suggest charities may be weathering this shift better than some sectors⁴, the direction of travel is clear.

If an organisation isn’t interpretable to AI systems, it risks becoming effectively invisible. As Weiner put it: “Around 10% of adults use this tool weekly, and something massive is on the horizon. AI is going to become the number one philanthropic advisor in the world. It doesn’t much matter what you think of AI. It matters a lot more what AI thinks of you.”

In an answer-driven world, the challenge for charities isn’t just being found. It’s being understood, represented, and recommended.

As Nathan Chappell, co-author of Nonprofit AI, told me: “If you have a best practice in fundraising that predates November 30, 2022. It is absolutely outdated.”

“ AI is going to become the number one philanthropic advisor in the world. ”

When AI gets it wrong

Yet visibility is only part of the story. As Alice Kershaw, Head of Digital Transformation at The Wildlife Trusts, pointed out, the shift from search results to AI-generated answers introduces a new and less comfortable reality. AI systems do not simply surface information. They actively educate the public, and not always correctly.

Kershaw shared a recent example where a Google AI summary confidently stated that there were bison in Worcestershire. There are not. The issue was not only that the answer was wrong, but that it was authoritative, widely visible, and beyond the charity's ability to correct.

In her characteristically dry humour, when asked what could be done, she replied, "Well, we can't remove it. My initial answer was... get bison." Behind the joke sat a serious challenge. In an answer-engine world, charities are no longer only shaped by what they publish, but by how their work is summarised, inferred and sometimes invented by systems beyond their control.

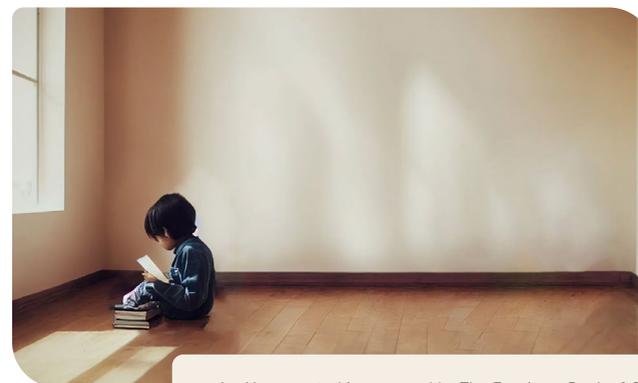
This represents a profound shift for charities who see public education as central to their mission. In the search-driven internet, misinformation could often be countered by publishing better content or improving rankings. In an answer-driven environment, a single incorrect summary can become the default answer, with little opportunity for correction.

For charities, answer engine optimisation is not just about being found or recommended. It is about accuracy⁵ in an environment where control is increasingly difficult. If charities want AI systems to "get it right", they must think carefully about where their information lives, how it is structured, and whether it is legible to both machines and humans.

Images, videos and donor trust

Many leaders I spoke to were cautious about the impact of AI-generated visuals. As Amy Sample Ward, CEO of nonprofit community NTEN and Author of 'The Tech That Comes Next', explained: "Nonprofits operate almost entirely on trust. If the community doesn't trust a nonprofit, they are out of business. No one cares if you trust Coca-Cola. You're going to buy the soda. If a nonprofit ran some ad with fake AI-generated people, the community would come for them."

There are, however, examples that show how these tools can be used responsibly. The Furniture Bank of Canada offers one such approach. Unable to photograph beneficiaries without compromising privacy or dignity, its CEO, Dan Kershaw, chose to use AI-generated images grounded in real stories of families living without access to furniture. Each image was presented alongside the tagline: "The picture isn't real. The reality is."⁶ Used in this way, AI can be a useful tool for public education.



An AI-generated image used by The Furniture Bank of Canada

AI is also enabling charities to produce educational and campaign materials at far lower cost. Over coffee in Washington D.C. Paul DePonte of the National Crime Prevention Council shared that recent advances have made it possible for nonprofits, including his own organisation, to create high-quality educational and promotional videos on limited budgets.

Several leaders also highlighted the value of AI-assisted translation, which allows a single piece of content to be adapted and shared across multiple languages and platforms. Crucially, they emphasised that while AI can accelerate production, volunteers and employees remain essential for sense-checking tone, context and cultural nuance, areas where machines continue to struggle.

Using AI-generated images responsibly

Before using AI-generated visuals, leaders should ask:

- 01 Transparency**
Is it clear to the audience that the image is AI-generated?
- 02 Representation**
Does the image reflect a real experience, or risk creating a misleading or stereotyped portrayal?
- 03 Dignity and consent**
Would the person or community represented feel respected by this depiction?
- 04 Purpose**
Is the image helping people understand an issue more clearly, or primarily provoking emotion?
- 05 Accountability**
Who is responsible for approving the image? How would concerns be addressed if raised by the community?

Fundraising at a crossroads

The spread of AI marks a moment of genuine choice for the charity sector, not just about what tools to adopt, but about what kind of relationships it wants to sustain.

Charities that adopt these tools with care and in alignment with their values can deepen trust and support more considered giving. Fundraising has always been relational work. As Nathan Chappell reflected, the opportunity AI offers is to make philanthropy more personal, not less.

Yet the same tools, when used poorly, risk undermining precisely what they aim to preserve. The question, then, is not whether AI will reshape fundraising, but how.



Operations

It is often underappreciated how much work takes place behind the scenes in charities. Finance, HR, reporting and internal planning all sit far from public view. Yet this work shapes how effectively a charity delivers its purpose. For many, this is also where organisations are first engaging seriously with artificial intelligence.

While nothing in the charity sector is ever risk-free, much of the work in operations is one step removed from direct interaction with beneficiaries and donors, creating space to test and adapt without immediately placing relationships or vulnerable people at risk. In addition, much of the tooling required is already available in the private sector, meaning charities do not need to wait for charity-specific technology to emerge.

Paul Butcher, founder of Commonsensing AI and formerly Executive Director of Strategic Initiatives at Save the Children, described operations as a natural starting point for AI because it brings together three enabling conditions: relatively consistent data, employees with the skills to experiment, and repetitive processes suited to automation. He contrasted this with mission work, which can often take place in low-bandwidth environments or contexts where data is fragmented, sensitive or incomplete, making AI harder to apply safely.

A real example: reducing single points of failure

Sandra Diaz, Director of Impact and Evaluation at Restore NYC, a charity helping make freedom real for survivors of human trafficking in the United States, described a core operational challenge: managing up to 500 support requests each year.

Around half can be served directly. The remaining 250 cases require referral to external services, a complex handover that takes one full-time role to manage.

When that employee took extended leave in 2024, the impact was immediate: “That role carries a significant weight... our connection to services rate fell from 93 to 75 %,” Diaz said. “That’s a big deal.”

The staffing gap prompted Diaz to convene an AI advisory board to develop an idea her team had been considering for some time, an AI-powered chatbot to support their referral process. The organisation’s partner database was difficult to search, meaning the team had to manually trawl past referrals to find appropriate services. The work was time-consuming, heavily dependent on individual knowledge, and difficult to scale.

AI, Diaz explained, was well suited to this task. The solution enabled their team to much more effectively search previous cases, filter options, and surface suitable partner organisations based on similar referrals.

The lesson is not that AI solves every problem. It is that it can strengthen the foundations of delivery by reducing single points of failure in essential processes. When gaps in capacity translate directly into people not receiving timely support, operational improvements are central to delivering the charity’s mission.

Five practical uses of AI in operations

01 Retail operations and stock optimisation

AI supporting pricing, routing and listing decisions for donated goods, maximising value and reducing reliance on volunteer guesswork.

02 Volunteer coordination and retention

AI analysing volunteer preferences, availability and engagement history to improve scheduling, relevance and retention.

03 Knowledge management and internal search

AI enabling team members to find information across policies, guidance, intranets and shared drives, reducing reliance on institutional memory.

04 HR and people operations

AI supporting recruitment shortlisting, onboarding materials, training documentation and analysis of staff survey responses.

05 Finance and reporting support

AI assisting with budget analysis, variance explanation, forecasting and narrative reporting for trustees, regulators and funders.

Small changes. Big gains: why incremental AI works

In operations, AI doesn't need to be complex to be valuable.

Some of the biggest gains come not from exponential change, but from steady, practical improvement. The Fund for the City of New York, which runs the Nonprofit AI Sprint, has been a strong advocate of this approach. The programme will train 100,000 nonprofit employees this year and is deliberately focused on fostering gradual change.

Over coffee, Ali Feldhausen, Associate Director of the Nonprofit AI Sprint, described their approach:

"We're really focusing on incremental AI, grounded in real use cases, and giving people across the organisation a clear understanding of how these tools work, so they can develop strategy thoughtfully. If even 50% of your organisation is upskilled in incremental AI, that's going to transform how you work and think as an organisation." A key feature is giving people the confidence and space to give it a go, and to try new tools.

A similar pattern emerged at Global Citizen. Michael Sheldrick, Chief Policy, Impact and Government Relations Officer and author of *Ideas to Impact* (Wiley, 2024), set a clear challenge: all 100 staff, across policy, partnerships and front office teams, were asked to experiment with AI. Within a month, only four had yet to try it.

The most immediate impact came from digital campaigns. Writing copy for a platform reaching 12.5 million people relied on a small number of specialists who could translate complex policy into language that resonated with Gen Z. It was a classic bottleneck.

To fix this, a team member built a custom chatbot trained on Global Citizen's tone and style. As Sheldrick put it:

"It took what people saw as an opaque, specialised skill and gave them agency... 'wow, okay, this is amazing. I can type in the action and it creates it with the right tone.'"

Time as both barrier and benefit

Survey data collected as part of this research revealed an interesting dynamic. Among charities yet to adopt AI, lack of time was one of the most commonly cited barriers. Nearly half reported that they could not step away from day-to-day work long enough to learn or implement new tools. Yet the experience of organisations that had already adopted AI pointed in the opposite direction. Increased operational efficiency was the highest-rated positive outcome, with 86% agreeing or strongly agreeing that AI had made their organisation more efficient. As many interviews reinforced, these gains most often began in operations.

Seen together, these findings suggest a pattern. The barrier to entry is often the very problem the technology is best suited to address. Time invested upfront in understanding and configuring tools can return itself quickly through reduced manual effort, fewer bottlenecks and more resilient systems. Operations is where this dynamic is most visible, and often most immediate.



(Above) Jake O'Gorman and Michael Sheldrick

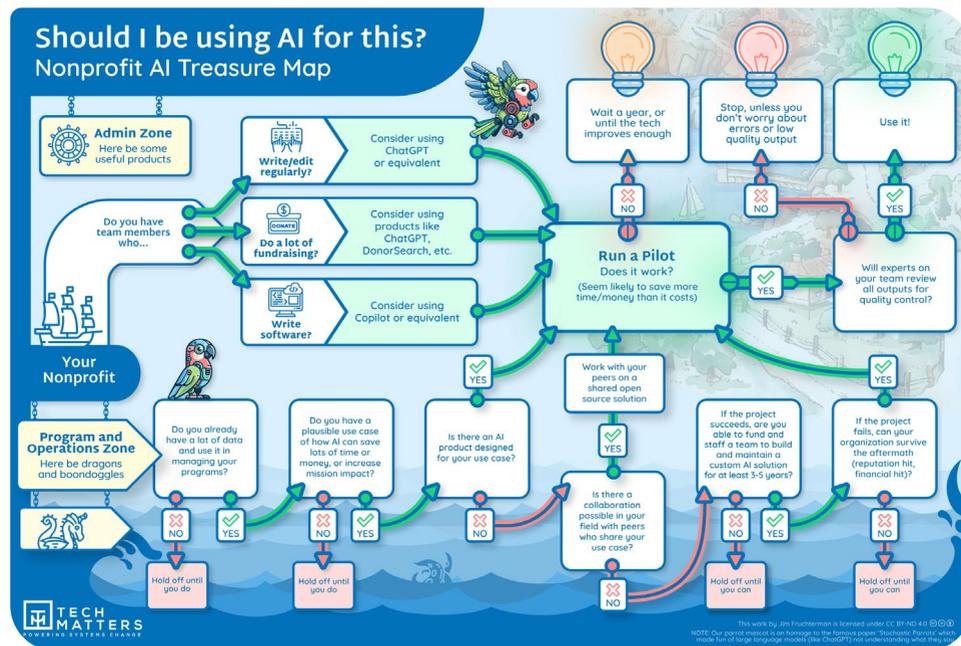
A safe place to start with AI

When charities are deciding where to begin with AI, the answer is often closer to home than expected.

Jim Fruchterman captures this well in his (slightly tongue-in-cheek) Nonprofit AI Treasure Map, which splits AI use into two territories: the Admin Zone and the Program and Operations Zone⁷.

The Admin Zone covers back-office and administrative work. This is where AI is often safest and most useful early on. The tools are easy to access, the use cases are familiar, and the cost of getting something wrong is lower.

Fruchterman describes this as the land of “useful products”: tools like ChatGPT or coding assistants, where the key test is simple: does this save you more time than it creates fixing mistakes?



(Above) Jim Fruchterman's Nonprofit AI Treasure Map

Retail: where AI meets the shop floor

For some of the charities I spoke to, a significant share of operations and fundraising is tied to retail. High-street shops, online resale and donation processing involve logistics, stock management, pricing and workforce coordination at scale.

Alison Court, who at the time of our interview was Chief Transformation Officer at Oxfam UK, explained that for organisations with a large retail footprint, AI can be incredibly helpful. “We put some of our own developed AI into our retail space,” she told me, “so that we can predict people’s buying patterns and encourage them to buy more online that they’re interested in, as well as understand their likelihood of being interested in other causes and other pieces of work that we do.”

Tackling the inventory problem

Inventory is one of the hardest problems in charity retail.

Unlike most commercial retailers, charity shops deal with a constant flow of one-off donated items, with no standard stock codes. That makes pricing, routing and stock allocation difficult to do consistently.

This is where AI fits naturally.

A great example comes from the British Heart Foundation, which processes around 800,000 donated items every week across 700 shops. To handle that scale, the charity developed Arti⁸, an AI-powered stock processing system built on Microsoft Power Platform, Azure OpenAI and Dynamics 365.

Arti draws on 1.9 billion rows of historical retail data to recommend pricing and channel decisions, helping volunteers decide which items belong on shop floors and which should be routed to higher-value online platforms like eBay or Depop.

The impact isn't just automation. It's a move from guesswork to informed decisions, framed as a responsibility to maximise the value of every donated item.

Speeding up e-commerce with AI

AI is also changing how charities run e-commerce.

Goodwill of Orange County⁹ offers a clear example. Using an AI-powered image categorisation app built on Azure AI Services, Goodwill tackled some of the slowest parts of online resale: identifying items, pulling out label information and drafting listings.

Before AI, research and writing made up 35–45% of listing time. The app now handles much of that work in seconds, leaving team members to focus on review and quality control instead of manual research.

At scale, the impact is significant. ShopGoodwill.com sells around \$9 million of clothing each month. A 35% increase in listings translates into meaningful income, with analysis suggesting average revenue growth of around 23%, as more donated items are processed and sold online rather than pushed into lower-value channels.

AI in retail isn't only about speed or revenue. It can also widen who gets to take part. Goodwill highlighted a clear accessibility benefit. By reducing the need for specialist product knowledge and complex writing, AI makes e-commerce roles more accessible to a wider range of people, including individuals with developmental, intellectual and physical disabilities. In doing so, AI supports income generation as well as Goodwill's wider mission to expand employment opportunities.

Where retail meets fundraising

More advanced retail models are starting to blur the line between operations and fundraising.

Oxfam's use of predictive analytics is one example. By analysing purchase behaviour, AI can identify when a retail customer is likely to be interested in a particular cause or campaign, and tailor follow-up accordingly.

At the more advanced end, charities are combining retail data with donor propensity models to create more individual supporter journeys at scale, something that was previously only possible for major donors.

Together, these examples show why retail is such a strong entry point for AI. The work is data-rich and repeatable. It's closely tied to financial sustainability, yet sits one step away from direct beneficiary risk. Used well, AI doesn't replace human judgement. It helps volunteers and team members make better decisions, move

The power of getting the basics right

With so much attention on AI, it's easy to overlook that for many charities, some of the biggest gains in efficiency and resilience come not from AI itself, but from updating everyday systems that have lagged behind.

As Princeton professor Arvind Narayanan says, AI should be treated as "normal technology". Not something magical, but something that follows the same patterns as the internet or electricity. It only works well when the foundations are in place.

Not every improvement needs AI. Jim Fruchterman shared a clear example from Aselo, an open-source platform he built for child helplines. He shared that many charities are seeing productivity gains of 20–40% simply by "updating technology to look like it is from the last five years rather than twenty". Meaning more calls answered from children in distress.

No "magic wand AI". Just better, more modern technology.

This reinforces a simple point: AI sits on top of digital infrastructure. It doesn't replace it. As Albert Scerbo, Associate Director of the Initiative on the Digital Economy at MIT puts it, for smaller organisations, getting the digital core right will often provide more value than complex AI. His rule of thumb: "If you're not doing SEO, you probably don't need to do AEO yet."

The challenge is that "normal technology" can be difficult to fund. CRM upgrades and data migrations are often seen as overhead, while "AI projects" may feel more exciting to donors.

Fruchterman suggests a practical workaround: design projects so that 80% of the work strengthens data foundations, warehouses, dashboards, core systems, and use the remaining 20% for the AI layer on top. It meets funders' appetite for innovation, while securing the upgrades that quietly deliver the biggest gains.

AI is already here: everyday assistants at work

For many charities, AI hasn't arrived as a big strategic decision. It has accumulated quietly.

Leaders often talk about AI as something to plan for, but in practice it's already showing up through updates to tools organisations are already using every day. As Albert Scerbo of MIT explained, AI is increasingly used by vendors as a point of differentiation, bundled into existing subscriptions before organisations have had time to assess quality, usefulness or risk. Simply by keeping software up to date, many charities are already using AI. As George Weiner of nonprofit focused marketing agency Whole Whale put it, AI is now "a side of fries with every purchase".

As one leader put it to me, "With everyday tools changing beneath our feet, staff exposure is inevitable."



When people bring AI to work

Whether bundled into existing services or purchased outright, most employees first encounter AI through everyday productivity tools. These include assistants that help draft emails, summarise documents, prepare presentations and analyse spreadsheets. For many people, these tools offer a low-stakes way to experience AI at work.

In practice, the benefits are often clear enough that many employees choose to bring their own AI tools into the workplace, whether their employer actively supports this or not. In some cases, this happens even where such tools are officially banned. As George Weiner explained, "If you are a leader claiming your organisation does not use AI, you're wrong. If you have anyone who graduated in the last four years, 'bring your own AI to work' is very real."

In response, many organisations are moving away from outright prohibition. Rather than trying to block access, they offer a small number of preapproved AI tools that team members can use safely, reducing reliance on free public tools used out of convenience.

A common but flawed compromise is allowing teams to use free public AI tools on the condition that they avoid entering personal or sensitive data. In practice, this is difficult to enforce. Few people consistently redact names or contextual details, and those who do often lose much of the value these systems are meant to provide. As a governance approach, this puts too much responsibility on individual behaviour instead of organisational design.

At a general productivity level, tools such as OpenAI's ChatGPT, Google's Gemini and Anthropic's Claude are widely regarded as high-quality assistants for writing, reasoning and summarisation, with enterprise-grade data protections. These assistants can deliver real productivity benefits, but they are not free. Licensing costs mean organisations need to be deliberate about who gets access and why.

Organisations already working within Microsoft ecosystems may choose to adopt Copilot. This ranges from Copilot Chat, a free, web-based tool with no access to internal organisational data, through to Microsoft 365 Copilot, which can operate across calendars, emails and documents that an organisation explicitly makes available.

As Ha Cole, EMEA Chief AI Officer at Microsoft Elevate, explains, these tools are embedded in everyday Microsoft applications such as Excel. Employees can analyse data and perform complex tasks using plain language, by typing instructions into a

prompt box. For charities, this can widen access to insight, enabling more informed, data-led decision-making without the need for advanced technical skills.

As with any technology rollout, however, organisations need confidence that access controls are properly configured before deploying AI tools at scale. Team members should only be able to see the folders and information required for their roles. A minimal-permissions approach, where access is limited to exactly what is needed, is essential to reduce the risk of sensitive information being unintentionally surfaced through an AI assistant.

Measuring what matters with AI

One key challenge organisations face when adopting AI, whether in the private, public or charity sector, is how to assess return on investment. In interviews, leaders repeatedly described one of the hardest parts of rolling out AI is in proving its value.

In many organisations, AI assistants are introduced without clear success measures, relying instead on anecdotal feedback. “It helped me write a better email” may well be true, but it rarely provides a sufficient basis for sustained investment or renewal. At the same time, several interviewees cautioned against the opposite extreme. Leading with rigid expectations of return too early can narrow experimentation and reduce long-term value.

Rob Grylls, who leads on impact measurement at Corndel, argued that early phases of AI adoption often benefit from a period of permission rather than proof. Requiring every use case to justify itself immediately can discourage the exploratory work that later yields the most valuable results. In his experience, some of the strongest contributors to outcomes only emerge once individuals and teams have space to test tools in low-pressure settings.

That does not mean measurement is unimportant. Rather, it shifts the question from whether AI should be measured to when and how. Time saved is often the most visible metric, but it is rarely the most meaningful on its own. Leaders pointed to outcomes such as clearer internal communication, better prepared meetings, or improved decision quality. These may matter more to organisational performance, but they are also harder to quantify. As Grylls noted, such measures often require the definition of proxies, shared rubrics and some form of impartial assessment of outputs. Without that scaffolding, claims of improvement risk remaining impressionistic.

Yet according to Albert Scerbo, charities may be better prepared for this challenge than they realise. Over coffee at MIT, he observed that nonprofits have spent decades grappling with how to demonstrate value in complex systems where outcomes are shaped by many factors. “The secret to AI implementation is measurement and experimentation,” he said, adding that charities are often more accustomed than private firms to evidencing value beyond short-term financial return.

Grylls drew a similar parallel from his earlier work in the sector. To be taken seriously by funders, organisations often need a clear theory of change and an evaluation framework that acknowledges attribution challenges. Rarely can a single intervention be shown to have caused an outcome in isolation. Instead, charities have learned to describe contribution rather than control, recognising the influence of external conditions, partner activity and human judgement.

AI adoption raises a comparable attribution problem. Improvements in performance, insight or efficiency are rarely the result of a tool alone. They emerge from how technology interacts with people, processes and organisational culture. In this sense, the question is not whether AI caused an outcome, but how it contributed to it, and under what conditions.

Seen in this light, effective measurement of AI is less about finding a single definitive metric and more about combining experimentation with reflection. Early exploration allows value to surface. Later, clearer definitions of what success looks like allow organisations to decide where to invest, where to adapt, and where to stop. For a sector already familiar with balancing evidence, judgement and uncertainty, this may be less of a departure than it first appears.

Look across these examples and a clear pattern emerges. The most effective operational uses of AI are rarely dramatic or transformative in isolation. They are practical, incremental and grounded in existing work. They focus on removing friction, reducing reliance on single individuals, and helping organisations carry out familiar tasks more reliably and at scale.

For charity leaders, this is what makes operations such a pragmatic place to begin. The work is familiar. The risks are easier to manage. And the lessons learned here carry forward into fundraising and mission-critical activity. Used well, AI can strengthen resilience, free up time and support better decision-making across the organisation.

AI in Volunteer Management

For organisations working with volunteers, there is often a challenge in maintaining personal connection at scale. Brian Rubenstein, President of the Rubenstein Impact Group, argues that relevance is central to retaining volunteers. This challenge is widespread. As Dietz and Grimm note in their 2023 paper, *The State of Volunteer Management*, 65% of nonprofits struggle to recruit and retain volunteers, even though 75% report that volunteers significantly improve service quality.

Rubenstein offers the example of a national cancer charity. One volunteer may be motivated by finding a cure, another by patient care, and a third by advocacy. Sending the same generic message to all three often leads to disengagement. AI can help by analysing volunteer data to identify these motivations and tailor communications accordingly.

AI tools can also help junior staff draft communications in the voice of senior leadership, improving consistency of tone. By automating segmentation and first drafts, AI can add capacity to short-staffed teams and help organisations treat volunteers as individuals, not just entries on a list. As Chappell and Rosenkrans note in their fantastic book *Nonprofit AI* (2025), AI can support tasks such as automated scheduling, analysing volunteer history and preferences, and generating personalised messages, invitations and recognition.



Brian Rubenstein



Mission

If fundraising brings in the money and operations keep the organisation running, mission is where charities deliver the impact that justifies their existence. It is also where the stakes of using AI are highest.

This is the point where technology touches beneficiaries directly. Decisions made here can determine whether someone receives help, is found, is protected or is missed entirely. For that reason, mission is usually the last area charities approach with AI. If leaders worry about trust in fundraising, introducing AI into beneficiary-facing contexts raises the stakes again. The risks are different, and the consequences more immediate.

And yet, when done with care, mission is also where AI can be most transformative.

When AI removes the bottleneck

Shawn N Olds, Co-Founder and Founding CEO of boodleAI, shared a powerful example from United Refugees, who until recently deployed teams of twelve people to refugee camps for a month at a time, manually leafing through books of photographs to help reunite families separated by conflict. After weeks of work, they might leave with a dozen, perhaps two dozen, successful matches.

Today, United Refugees uses facial recognition AI on everyday devices like iPhones and iPads. Instead of large teams on long deployments, two team members can now deploy for a single day and make hundreds, sometimes thousands, of positive identifications. This technology does not replace human work. It removes a bottleneck that once made scale impossible.

This is the pattern that best defines mission AI. It does not begin with tools, but with a problem that human systems cannot solve fast enough, widely enough or fairly enough on their own. What follows are examples of how charities are applying AI in this way — and what leaders need to put in place to make sure mission AI delivers impact without compromising trust.

Practical uses of AI in mission delivery

- 01 Service navigation and signposting**
AI chatbots guiding people to appropriate services quickly (e.g. Spring ACT's Sophia)
- 02 Crisis triage and demand routing**
Filtering routine enquiries and escalating high-risk cases to professionals
- 03 Knowledge management and internal search**
Predictive analytics detecting risk patterns in housing, health or safeguarding before harm occurs
- 04 Logistics and resource allocation**
AI optimising the distribution of food, aid or goods based on need and availability (e.g. MealConnect)
- 05 Programme evaluation and learning loops**
AI enabling faster, continuous evaluation by analysing outcomes data, linking interventions to impact, and testing improvements over time

Impact at the frontier: UNICEF

Over lunch in New York, Yves Jaques, Chief of UNICEF's Frontier Data & Technology Unit, described an organisation operating at extraordinary scale. UNICEF is not only the world's largest purchaser of vaccines, but also of pencils. Its success depends on logistics and coordination at a global level.

That scale was tested when war broke out in Ukraine.

As millions fled across borders, UNICEF was tasked with setting up "Blue Dots", safe spaces offering education, healthcare and protection for displaced families. The challenge was where to put them. Events were moving faster than traditional data could keep up.

To respond, Jaques' team partnered with the University of Oxford and others to use Facebook advertising data as a real-time signal for population movement. Within four weeks, they expanded their geospatial models from 12 data layers to 28, using optimisation techniques to predict where services would be needed most.

The result was striking. UNICEF was soon able to provide the Romanian government with more accurate insight into the location of Ukrainian refugees than its own official data.

In Senegal, the same team worked with MIT students to tackle a different problem: identifying "zero-dose" children — around 30 million worldwide who have never received a single vaccine and often sit outside formal health systems. By combining unconventional data sources with predictive models, UNICEF was able to surface populations that traditional health records simply couldn't reach.

UNICEF illustrates what becomes possible when data infrastructure, technical expertise and organisational mandate align at global scale. But it also risks feeling distant from the realities of most charities. What matters, as the next case shows, is not scale itself, but what can be learned from how these tools are deployed.



Scaling identification in human trafficking: The Safe House Project

In efforts to combat trafficking, traditional identification and reporting methods often hit a ceiling.

As Brittany Dunn, COO and co-founder of Safe House Project, explained, the sector faces a critical bottleneck: human bias and capacity constraints mean that as many as 99% of victims go unidentified. For Dunn, AI wasn't a nice-to-have. It was essential. As she explained, "technology and data is the only way forward in solving this issue" because it allows organisations to bypass implicit bias in traditional dispatch systems and scale beyond human limitations.

Safe House Project set out to increase identification and route people to safety through a reporting mechanism that reduced fear of law enforcement and minimised bias in reporting that can harm underrepresented populations.

The aim was to create a system where members of the public or victims did not need certainty to ask for help. As Dunn put it, "We always tell people they don't have to be confident. They just have to care".

The organisation launched SimplyReport, a platform underpinned by a proprietary AI model, which uses natural language processing and large language models to analyse anonymous reports via a chatbot. This anonymity provides the "psychological safety" required for victims or concerned citizens to share information without fear of triggering a criminal investigation.

The system is built on a behavioural framework originally developed by the US intelligence community to predict radicalisation, which Dunn's team re-engineered to recognise indicators of trafficking, and crucially, uses a human-in-the-loop approach, and set around four foundational pillars: "do no harm, confidentiality, cultural sensitivity and informed consent". The AI handles intake, analysis, and validation of tips at scale. Crisis response for survivors is handled by human case managers. Dunn described a case where a nurse used the app to help identify a victim in a pregnancy clinic. The AI guided the nurse with specific questions to ask, leading the woman to feel safe enough to share that she was being trafficked. Safe House Project Survivor Support Team coordinated her extraction and put her on a plane to safety.

SimplyReport launched in May 2025 with a team of 20 people. Within three months, it had received over 8,000 signals, opened 1,200 cases and served 684 survivors. By comparison, Dunn noted that the US National Human Trafficking Hotline, operating for 17 years and widely promoted, sends around 2,500 signals annually to law enforcement. Safe House surpassed that annual figure in just ninety days, without a national advertising campaign. The efficiency gains were equally stark. A recent federal grant for a comparable hotline model budgeted for up to 100 staff. Safe House Project is handling far higher volumes with a fraction of that headcount.

Stories like this show that while mission is where the risks of AI are highest, so too is its potential. When done well, it enables charities to act at a scale and speed that human systems alone cannot reach.

**" Within three months,
it had received over
8,000 signals, opened
1,200 cases and
" served 684 survivors.**

Why chatbots are gaining ground

Although only around 7% of charities currently use AI in service delivery,¹⁰ few AI applications make the promise and risk of mission AI more apparent than chatbots.

Used to help beneficiaries find information, navigate services and seek support, recent advances have made chatbots more affordable and practical, lowering barriers to entry. Public attitudes are also changing. Research by the Ada Lovelace Institute¹¹ found that 7% of the UK public have used a mental health chatbot, indicating growing familiarity with conversational AI in sensitive contexts.

When designed well, chatbots can significantly widen access to support. Tools such as Sophia¹², developed by Spring ACT, provides round-the-clock, anonymous support to survivors of domestic violence across more than 170 countries and 20 languages. Importantly, users can access Sophia without downloading an app, lowering risk for people who may be monitored or in unsafe situations.

Similar approaches are used in humanitarian and health contexts to manage demand, provide trusted information at scale and free human professionals to focus on high-risk cases.

In commercial settings, studies suggest these types of conversational AI can resolve around 80%¹³ of routine enquiries. Yet several interviewees warned against deploying chatbots in emotionally vulnerable or high-stakes settings without crisis protocols and clear human oversight. A single hallucination or misjudged response can escalate distress rather than relieve it.

Sandra Diaz of Restore NYC, was clear on this point. When working with survivors of trauma, efficiency never justifies increased risk to service users. Best practice is to prioritise survivor choice and safety when integrating AI tools. Restore's 2026 report for the anti-trafficking field, which interviewed clinicians on the use of AI for trauma treatment, recommended using AI for routine, low-risk interactions, while humans remain responsible for judgement, escalation and care. The report also identified the opportunity to augment therapeutic interventions through AI-powered apps featuring emotional regulation and mindfulness tools. Used this way, chatbots can expand reach without compromising trust or safety.

Mission AI: promise, peril and judgement

Mission is where AI can have its most powerful impact.

Used well, it can help charities find people who would otherwise be missed, widen access to support, and act at a speed and scale that human systems alone cannot reach. From reuniting families to identifying exploitation or navigating complex services, AI can remove bottlenecks that have limited impact for years.

Yet mission is also where the limits of AI are most exposed.

Here, AI can shape decisions that affect safety, dignity, rights and trust. That raises the bar. Not every problem should be automated, and reasonable people will often disagree about where the line should sit.

A recent example illustrates why. Gloop, a US-based technology platform serving faith communities, has explored the use of video cameras and emotion-recognition AI in churches to identify congregants who may be experiencing distress. The intention is to offer support proactively, without requiring someone to actively seek help. For some, this represents compassionate innovation that lowers barriers to care. For others, it raises uncomfortable questions about consent, surveillance and the use of sensitive personal data in high-trust environments.

The disagreement itself is instructive. It shows that mission AI does not operate only in technical space, but in social contexts where moral considerations, values and judgement matter.

What distinguished the strongest examples in this research was not sophistication or scale, but judgement. The most credible organisations did not begin by asking what AI could do. They began by understanding a real need, the people affected, and the consequences of getting it wrong. AI was treated as a tool, not an authority. Boundaries were explicit. Human oversight was non-negotiable.

PART TWO

What gets in the way

Up to this point, this book has focused on what artificial intelligence makes possible for charities, exploring how AI is helping organisations work more efficiently, reach people they could not reach before, and deliver greater impact.

As the research progressed, it became clear that this was only half of the story. AI is not a magic solution, nor does it arrive into a neutral landscape. It enters organisations shaped by limited resources, competing responsibilities, and long-established organisational norms.

This section explores the factors that shape those choices.

For this research, 275 charity leaders were asked about the main barriers stopping them from doing more with AI. Nine themes appeared repeatedly, spanning skills and confidence, time, funding, ethics, data, regulation and organisational readiness. Together, these responses provided an initial map of where hesitation was most strongly felt.

The chapters that follow are organised around three themes:

UNCERTAINTY

Why AI is hard to judge in advance, how confidence and caution shape decisions, and why leaders often feel pressure to act without clear evidence.

CAPACITY

How constraints on time, funding, skills and data affect what charities can realistically deliver.

RESPONSIBILITY

Where questions of trust, accountability and risk are heightened, and where mistakes can harm beneficiaries or undermine confidence.

Together, these themes help explain the conditions under which decisions about AI are made.

Top barriers to doing more with AI

This research surveyed 275 charity leaders across the UK and the US. The following barriers to expanding the use of AI were cited most frequently, with the percentage of leaders raising each issue shown.



- 01 Lack of skills or confidence (57%)**
The most frequently cited barrier, which increases after adoption as teams realise how much there is still to learn.
- 02 Concerns about ethics or fit with mission (51%)**
A persistent concern about whether AI aligns with organisational values, including environmental and social impacts on mission delivery.
- 03 Lack of time or capacity (44%)**
Many organisations are already stretched delivering core services, making it difficult to find the time or headspace to explore new tools, even when they promise future efficiency gains.
- 04 Limited budget or funding (43%)**
A significant psychological hurdle that drops sharply once implementation begins, falling from 47% to 31%, suggesting cost uncertainty is often greater before organisations gain practical experience.
- 05 Data readiness (35%)**
Adopting AI often forces a moment of reckoning, revealing that existing data infrastructure is not robust enough to support more advanced tools.
- 06 Leadership or trustee caution (33%)**
A key gatekeeping barrier driven by caution at the top, which falls by around a third once leaders commit to a clear strategy.
- 07 Regulation and compliance (31%)**
This barrier is more pronounced among post-adopters, who encounter real-world compliance issues, such as GDPR, that often feel abstract before implementation.
- 08 Technology or systems not set up for AI (28%)**
Concerns that existing or legacy systems may not integrate well with AI tools, adding complexity, cost or risk rather than simplifying work.
- 09 Difficulty seeing the benefits or value of AI (19%)**
While a minority of organisations struggle to see clear value at the outset, this concern often diminishes once AI is used in practice.

Uncertainty: Why AI is hard to judge in advance

Uncertainty is one of the defining features of AI adoption in the charity sector.

This section examines how that uncertainty shapes leaders' responses to AI. Drawing on survey data and interviews, it explores the difficulty of judging value in advance, concerns about unintended consequences, and the role of leadership caution in slowing or complicating decisions. In practice, leaders are often making choices with incomplete evidence, uncertain outcomes, and little margin for error.



Difficulty seeing value

In this research, 19% of charity leaders reported difficulty in seeing clear value as a barrier to greater AI use.

In most cases, this reflects prioritisation rather than resistance to technology. AI is rarely the most urgent issue on a leader's desk, and even when it rises in importance, its immediate value can be difficult to judge.

That caution is understandable, and often prudent. We are living through a period in which AI is routinely presented as a solution to almost everything, and charities have heard such claims before. Blockchain, anyone?

Choosing AI means choosing not to do something else

The question, after all, is not just whether AI has value, but whether it is more valuable than the other things a charity could invest in. AI takes effort. Choosing to pursue AI means choosing not to pursue something else.

At first glance, the data appears to challenge this caution. Among organisations that had not adopted AI strategically, 26% cited difficulty in seeing benefits as a barrier. Yet among those that had adopted AI, the figure was 2.6%. My initial interpretation was to attribute this gap to selection bias: organisations that never struggled to see value were simply more likely to adopt AI in the first place. However, interviews point to another explanation.

In most cases, leaders in organisations that had adopted AI said that difficulty seeing value had indeed been a concern at the outset. What changed was not their initial scepticism, but their experience. Once they began using AI in practice, even in limited ways, the value became easier to see, describe and defend. Organisations that struggle most to see value are often those that have not yet tested it meaningfully. Value, it seems, becomes clear through use.

Why value can feel hard to judge

When leaders decide whether to prioritise AI, they are rarely asking a single question. Instead, they are making a judgement call shaped by several overlapping considerations: whether AI is likely to improve outcomes, how much effort it will take to realise those benefits, how others will view the decision, and whether the organisation feels ready to act. The challenge is that many of these factors are difficult to assess in advance.

To examine this more clearly, this research used the Unified Theory of Acceptance and Use of Technology (UTAUT), which is commonly used to explain why individuals and organisations adopt new technologies. Here, it provides a useful lens for exploring how charity leaders think about value before adoption.

The UTAUT model highlights four factors that tend to shape prioritisation decisions:

- **Performance expectancy:** Will this meaningfully improve outcomes, such as service delivery, reach, or freeing up time for mission-critical work?
- **Effort expectancy:** How difficult is it likely to be to realise those benefits in practice, given current capacity, skills and systems?
- **Social influence:** Whose views matter in this decision, including peers, funders, leadership and the wider sector? How much weight do they carry?
- **Facilitating conditions:** Does the organisation feel it has the data, systems, governance and resources needed to act with confidence?

Recreating the model using charity-specific survey data revealed a consistent pattern. Performance expectancy and effort expectancy were most closely associated with whether organisations felt able to prioritise AI. Facilitating conditions mattered less for initial interest, but strongly influenced whether organisations felt able to move from intention to action.

Organisations that went on to adopt AI scored higher across all four UTAUT dimensions at the point of decision. This is not surprising. Where anticipated benefits felt achievable and the effort required felt manageable, organisations were more inclined to act. The more revealing insight lies in what this pattern suggests about hesitation.

Difficulty seeing value was rarely about doubt that AI could be useful. Instead, it reflected uncertainty about whether benefits would materialise in practice,

how much organisational effort they would demand, and whether the conditions were in place to support them. Leaders were often weighing these questions simultaneously, with incomplete information and little opportunity to test assumptions safely in advance.

This helps explain why value often becomes clearer through use. Interviews consistently showed that leaders in organisations that had adopted AI did not begin with greater confidence or lower scepticism. What changed was not their initial judgement, but their experience. Even limited, practical use made benefits easier to observe, describe and defend. In this sense, value was rarely fully visible until organisations engaged with AI directly.

Seen in this light, difficulty seeing value is less about identifying potential benefits and more about judging whether those benefits feel achievable under real-world conditions. When anticipated gains are outweighed by perceived effort, risk or uncertainty, hesitation is a rational response.

“ When leaders decide whether to prioritise AI, they are rarely asking a single question. Instead, they are making a judgement call shaped by several overlapping considerations. ”

When context obscures value

For charities, decisions around AI are further complicated by the fact that much of the public narrative is framed in commercial terms. Save time. Cut costs. Replace jobs. Move faster. These arguments do not translate easily into a mission-driven environment, where performance expectancy is judged primarily in terms of outcomes for beneficiaries rather than operational efficiency alone.

Amy Sample Ward, CEO of nonprofit community NTEN, explained why this framing fails. “None of those are compelling arguments for a charity,” they told me. “Is it going to help us serve more people? That’s the question.” They also described how this narrative creates anxiety. Charities can feel that they are falling behind, even when it is unclear who they are behind. In UTAUT terms, social influence can become unhelpful. As Sample Ward put it, “When you ask, ‘Behind who?’, it’s just crickets.”

Pressure can further obscure value. At the time this research was conducted in the United States, many charities were grappling with the sudden impact of changes to USAID funding. Paul Butcher, founder of Commonsensing AI and formerly Executive Director of Strategic Initiatives at Save the Children, described how funding disruption can push organisations into survival mode. In these conditions, investment in AI is rarely rejected outright, but frequently deferred. When leaders are focused on keeping programmes running, the perceived effort required to build new capability can outweigh anticipated benefits, even when those benefits are recognised in principle.

At the same time, pressure can sometimes make value impossible to ignore. Courtney Bugler, President and CEO of ZERO Prostate Cancer, described AI as a way of reducing unsustainable strain on staff. “Every nonprofit person complains that they’re behind the eight ball,” she said. Used well, AI allows people to be more productive and to refocus on the parts of their work that require human judgement and connection. “I believe good use of AI allows me to be more human,” she explained.

Nathan Chappell, Chief AI Officer at Virtuous Software, offered a related perspective. For small, resource-constrained organisations, AI can act as an equaliser. Where resources are thin, the effort expectancy associated with manual work is already high, making the relative value of AI easier to see. In these contexts, the alternative to adoption is not caution, but stagnation.

Hype, ceilings, and stopping too soon

Confusion about value is also fuelled by the gap between what AI promises in theory and what it delivers in practice. High-profile demonstrations can inflate expectations about what AI should achieve, raising performance expectancy to unrealistic levels while obscuring the quieter, more incremental uses that often generate the most reliable returns. As Ali Feldhausen observes, much of the immediate value for organisations lies in incremental rather than exponential AI.

This mismatch can distort decision-making. When leaders feel pressure to identify a single, transformational use case, more modest but achievable gains are easily overlooked. In UTAUT terms, performance expectancy becomes narrowly defined around dramatic outcomes, while the effort required to reach them appears disproportionately high. Daniel Gray, Director of Data, Digital and Technology at Crisis, described a related form of paralysis: an organisation may feel ready in principle, yet struggle to justify action because no single application appears uniquely compelling.

In many cases, AI does not introduce an entirely new capability. Instead, it lowers the effort expectancy associated with existing work. Tasks such as research, analysis or synthesis may have been technically possible for years, but difficult to resource consistently. By reducing the effort required to carry them out well, AI can make familiar activities feel newly viable, even when the outcomes themselves are not novel.

This dynamic is reinforced when AI is treated as exceptional or mysterious. Jim Fruchterman, founder and CEO of Tech Matters and Benetech, has criticised what he describes as “magic wand AI”. Framing AI in this way can further inflate expectations, making realistic gains appear underwhelming by comparison. In practice, much of the value for charities lies in unglamorous work: summarisation, triage, pattern recognition and small improvements that compound over time. These uses may lack spectacle, but they often strike a more favourable balance between achievable effort and meaningful benefit. For many charities, that balance is precisely what makes AI worth pursuing.

Value follows use

Taken together, these perspectives point to a simple conclusion: difficulty seeing value is not a fixed obstacle. It is a stage.

Before adoption, value is abstract, contested and easy to defer. Performance expectancy is shaped by assumptions, external narratives and incomplete evidence. After adoption, value becomes more concrete, if uneven, and harder to ignore. Survey data reflects this clearly. Organisations that adopt AI rarely regret the decision. Most plan to expand their use, not retreat from it.

The challenge for leaders, then, is not to accept AI on faith, nor to dismiss it as hype, but to create the conditions in which value can be tested safely and seen clearly. In practice, this means lowering the effort required to experiment, clarifying how potential gains link to mission, and creating space to learn. These conditions allow judgements about value to move from expectation to experience.

Value does not usually appear fully formed. It emerges through use. For charities willing to engage in that process, the question shifts from “Is this worth it?” to “How do we do this well?”

Whose voices matter most?

Interestingly, of the four UTAUT constructs, social influence emerged as a powerful factor in shaping intention to adopt AI. While it was not the strongest driver of perceived value—that is, it did not primarily determine whether leaders believed AI would be useful—it **proved to be the single strongest predictor of whether organisations intended to act.**

Three forms of social influence proved particularly important. First, peer behaviour: seeing similar or respected organisations adopt AI made adoption feel more credible and achievable. Second, internal leadership and trustee support: clear signals from the CEO and boards materially increased the likelihood that interest would turn into action. Third, sector norms: as AI use began to feel more established within the sector, it increasingly appeared as a legitimate choice rather than an outlier risk.

By contrast, perceived expectations from funders were notably weaker. Among organisations yet to adopt AI, only around 23% believed their funders or major

donors would view AI use positively, with most respondents expressing uncertainty rather than opposition. This ambiguity meant that funder opinion rarely acted as a positive driver, but neither did it function as a clear deterrent.

This pattern is evident in the survey data. Organisations that went on to adopt AI reported significantly higher levels of social influence at the point of decision, with an average score of 3.59 on a five-point scale, compared with 3.07 among organisations still in pre-adoption. In exploratory regression analysis (examining the relative influence of all four UTAUT factors together) social influence remained statistically significant¹⁵ and emerged as the strongest predictor of stated intention to adopt, even when controlling for performance expectancy, effort expectancy and facilitating conditions.

In practice, this suggests that belief in AI's usefulness is widespread but insufficient on its own. What distinguished those who moved towards adoption was not stronger belief, but stronger signals of permission and legitimacy. Change was far more likely when leaders saw respected peers taking action, when trustees and executives openly endorsed the use of AI, and when AI came to feel like a normal and defensible choice within the sector rather than a speculative or reputational risk.

Given the powerful role that leadership voices play in turning intention into impact, we now turn to the factors behind leadership caution.



Leadership caution

In many charities, difficulty in seeing value flows directly into leadership and trustee caution. In this research, around one third of organisations identified leadership or trustee caution as a meaningful barrier to doing more with AI. Where value is uncertain, risk feels harder to justify and responsibility weighs more heavily.

Charities operate in high-trust environments. Avoiding unnecessary risk is important. AI is still relatively new, often poorly understood, and surrounded by competing claims about both its promise and its risks. For leaders, the instinct to pause, ask questions and proceed carefully is understandable. Often it reflects good governance.

Yet survey data can lead us to wonder if at least part of this caution is down to confidence. The 2025 Charity Digital Skills Report found that 44% of charities rate their board or trustees' skills, knowledge and confidence in AI as poor, with only 3% rating them as excellent. 36% say the same of their CEO¹⁶.

The difficulty is that while caution can protect a charity from harm, it can also prevent leaders from engaging with potentially transformative technologies.

Governance built for trust, not speed

As charities grow, governance frameworks and processes tend to expand in ways that prioritise trust and accountability over speed, particularly in sensitive and high-risk contexts.

As Paul Butcher, founder of Commonsensing AI and formerly Executive Director of Strategic Initiatives at Save the Children, put it, large charities are “very structured... because they are accountable to their stakeholders, and there’s a lot of risk in getting things wrong”. They are “generally quite process-driven and risk-averse... a little bit bureaucratic, and that’s served them well over the years”. Those same characteristics, however, can also become a barrier to innovation and the adoption of AI.

Brian Rubenstein, President of the Rubenstein Impact Group and formerly a leader at the advocacy affiliate of the American Cancer Society, highlighted the practical implications of this. In larger organisations, “multiple teams and leaders must be involved, from IT and subject matter experts to legal and marketing. This decision-

making by committee is therefore slow by design. Combined with a low appetite for risk, this can create a tendency to pause progress while waiting for clarity.

Scott Rosenkrans, co-host of the Fundraising.AI podcast, reflected that over time, layers of review and sign-off emerge as individuals feel responsible for catching errors or slowing decisions. While this scrutiny is important, it can consume time and energy in ways that smaller, more agile organisations can better avoid.

Together, these perspectives point to an interesting observation. The very governance structures that protect charities from harm can also make it harder to engage with technologies that evolve quickly and reward experimentation.

The temptation to wait

One common result of this caution is the decision to wait. To allow the technology to mature, to see how others fare. To return to the question when things are clearer. As one CEO put it to me, “these things get cheaper and more accessible, so I’m just going to wait a year”.

Albert Scerbo, Associate Director of the Initiative on the Digital Economy at MIT cautioned against this attitude. While leadership may delay, he argued, staff are already using these tools. “There’s no point in waiting,” he warned.

As George Weiner of nonprofit focused marketing agency Whole Whale put it “Putting your head in the sand is not going to stop the water from coming in, it just makes you drown faster.”

Yet of course, not all waiting is the same. Scott Rosenkrans explained, choosing to delay adoption to prepare, for example by improving data quality, upskilling teams or clarifying governance, can be a valid strategic decision. Simply waiting without a plan is not.

Caution with purpose

Caution is not inherently good or bad. It depends on what sits behind it.

Where caution reflects fear, lack of confidence or avoidance, it can expose charities to greater risk, not less. AI is already part of the environment in which charities operate. Ignoring it does not make it go away.

Where caution is informed, evidence-based and paired with deliberate preparation, it plays a vital governance role.

The challenge is to ensure caution is grounded in understanding rather than uncertainty. That requires time spent learning, asking better questions, and creating space for controlled experimentation. The sections that follow explore how charities can build the skills, governance and confidence needed to do exactly that.



Capacity: the work of making AI possible

Even where leaders believe AI could be valuable, many struggle to see how adoption would be possible in practice. Capacity constraints, particularly around skills, time, funding and data, consistently emerge as the most significant barriers to progress. For charities already operating under strain, adding another complex initiative can feel unrealistic.

This section examines the practical work required to make AI usable and sustainable inside organisations. It examines why capacity challenges persist even after adoption, why skills and confidence matter as much as the technology itself, and how organisations navigate the trade-offs involved in introducing AI. The focus is on the conditions that shape whether AI can be integrated into everyday work in ways that deliver value and can be sustained over time.

Skills

In this research, skills emerged as the most frequently cited barrier to charities doing more with AI, and one that became more pronounced as organisations became more engaged. Alongside regulation, skills were the only barrier that rose in line with adoption, cited by 55% of organisations not yet using AI strategically and by 63% of those that were.

This pattern is reinforced by the 2025 Charity Digital Skills Report. Despite growing interest in AI, limited digital skills and a lack of training or support to upskill were cited among the top barriers to progress. The confidence gap extends to leadership. More than a third of charities reported that their CEO had poor AI skills, knowledge or confidence, and almost half said the same of their board¹⁷.

Taken at face value, this suggests a sector held back by a shortage of technical expertise. That reading, however, is only partly true.

Reframing the skills challenge

There is no question that a skills gap exists. Many charities lack access to specialist digital roles, and competition for technical talent is intense. Yet upon closer inspection, a more nuanced picture emerges.

For most charities, the perceived skills gap is not where leaders initially expect it to be. AI is often imagined as something that must be built. Progress is assumed to depend on hiring data scientists, engineers, or AI specialists. In reality, most charities making meaningful progress with AI are not building tools. They are implementing and configuring third-party software, adapting existing platforms, and learning to use new capabilities embedded in tools they already pay for.

This distinction matters. When leaders believe progress requires deep technical capability from the outset, AI can feel unreachable or unaffordable. Yet every day, organisations begin their AI journeys without in-house applied AI expertise. Starting small, most use early wins to prove value before building the case for further investment.

This creates a mixed message. On one level, leaders are right to say skills are a barrier. On another level, the nature of that barrier is often misunderstood. The most impactful skills are not always the most technical, and the absence of specialist capability need not block early value creation. For most organisations, the question is not whether specialist AI skills will be needed, but when, and for what purpose.

“ For most organisations, the question is not whether specialist AI skills will be needed, but when, and for what purpose. ”

The skills that unlock progress

During this research, I spent time with charities exploring what skills they actually needed to move forward with AI. What emerged was not a single skills gap, but a hierarchy of skills, each building on the last and not all required at once.

01

Leadership judgement comes first

Before tools, data or training programmes, charities need leaders who can navigate the decisions AI introduces. Leaders are being asked to make decisions under uncertainty, often with incomplete information.

Three leadership challenges came up repeatedly: moving too slowly and missing opportunities, moving too fast and creating unnecessary risk, and a maturity gap that frames the opportunity as “We need to use AI” rather than “We need to solve this specific problem. Is AI the best solution?”

It is telling that the Charity Digital Skills Report found that 36% of charities rate their CEO’s AI skills as poor, and 44% say the same of their board¹⁸. This does not reflect a lack of intelligence or intent. It reflects how unfamiliar this terrain remains.

02

Business analysis capability

Across my interviews, one role consistently emerged as central to successful AI adoption. This skillset is often described as business analysis or solution design. It is the ability to translate organisational needs into workable solutions, identify processes worth improving, redesign workflows, and decide where AI fits and where it does not.

Crucially, this role does not require someone to be an AI expert. In practice, while some large charities have in-house business analysts, this role most often works best when held by a domain expert in areas such as finance, HR, fundraising or operations. These individuals usually have spent years thinking about the problems that matter. With sufficient exposure to AI, and encouragement to explore, they are well placed to judge where AI could add value, whether proposed use cases would deliver it, and to tell that story internally.

Daniel Gray, Director of Data, Digital and Technology at Crisis echoed this dynamic. Many teams, he explained, are capable of teaching themselves the technical basics when required. The harder, and more valuable, capability is knowing when AI is the right approach at all.

The real skill lies in solution design: stepping back to assess whether a problem genuinely benefits from AI, or whether a simpler, cheaper or more reliable approach would achieve the same outcome. In practice, this judgement often matters more than technical sophistication.

As Ha Cole, EMEA Chief AI Officer at Microsoft Elevate, noted, this is becoming even more important as new tools enable less technically experienced team members to build AI agents in no-code or low-code environments. When combined with sound judgement and deep understanding of the problem at hand, these tools can unlock significant value across an organisation.

Several interviewees described this as one of the most important, yet undervalued, roles an organisation can develop. The encouraging reality is that most organisations already have people who can play this role, if they know where to look.

03

Data and AI literacy

For charities, building AI capability involves two closely linked tasks. First, helping people begin, and secondly, helping them improve their usage once they do.

For many people in charities, AI feels intimidating. That makes it especially important to bring everyone with you, not only the curious and confident, but those who risk being left behind. Often the first barrier is language: confidence and permission to engage at all.

Jonathan Townsend, CEO of The King's Trust, described how the way AI is discussed can become a barrier to participation. The King's Trust supports young people to learn, earn and thrive. From his perspective, the risk is not a lack of ability, but intimidation before people even begin.

He described that many will have a "fear of looking dumb", and of not understanding the language used to talk about AI and technology. Yet when abstract concepts are translated into familiar, day-to-day tasks, that fear quickly evaporates. This

reframing matters, because it helps people engage. Once they do, a second challenge emerges.

As AI tools move into everyday use, they put significant power into many more hands. Conversational interfaces can generate analysis, charts, summaries and recommendations with minimal technical input. That ease of use lowers the barrier to participation, but it also shifts where risk sits. When outputs are easy to produce, judgement becomes the critical skill.

What interviewees returned to was not how to generate outputs, but how to validate them. What should be trusted? How can results be checked? How should decisions be explained to colleagues, trustees, donors and beneficiaries?

Survey data reinforces this. The Charity Digital Skills Report shows that 39% of charities rate themselves as poor at website and analytics data, while 45% list using AI tools to analyse data as their top support need. Among larger charities, concerns about data privacy, GDPR and security rise sharply as adoption increases.

AI changes how we need to think about skill. It shifts the burden away from execution and towards interpretation, judgement and confidence under uncertainty. In doing so, it increases the importance of both data literacy and AI literacy, not despite ease of use, but because of it.

Yves Jaques of UNICEF offered one of the clearest articulations of this shift. He argues that AI moves organisations from a deterministic world, where systems behave predictably, into a probabilistic one, where outputs are based on likelihood rather than certainty. Responding to the idea that AI reduces the need for data skills, he put it bluntly: "I think it's actually the reverse. The world is going to need people who understand statistics more than ever, because many people don't know how to do quality assurance."

His point is not about advanced modelling. It is about literacy. In traditional software development, a system either works or it does not. With AI, the question is different. As Jaques explained, "We have to ask what test it would need to pass before we can use it. That's different from how many people approach AI."

Judgement over answers: the new literacy AI demands

The skill is no longer in producing an answer. It is knowing how much confidence to place in it, what assumptions sit beneath it, and whether it is fit for purpose. Without data literacy, we lack the means to distinguish insight from error. AI literacy adds a further layer: understanding how these systems behave, how they are prompted, and where their apparent confidence can be misleading.

For Jim Fruchterman, the danger of AI lies not in what it can do, but in the confidence it gives to people who do not understand its limits. “The person who says, ‘I don’t need data literacy’ is the person who believes AI makes no mistakes,” he warned. “If you don’t know what you’re doing, AI gets you in deeper shit pretty fast.”

He illustrated this with the example of “vibe coding”, where non-coders use AI tools to generate software. The code may run, and even appear to work. Without the ability to test it, however, the user has no way of knowing whether it does what they think it does. “Did you write a test framework?” he asked. “If not, you have no idea whether it works.”

Reflecting on this broader shift, Scott Rosenkrans, author of *Nonprofit AI*, was clear: “If you asked me five years ago whether the average user needed strong data literacy to get value from AI, I would have said no. Generative AI has changed that.”

“ The skill is no longer in producing an answer. It is knowing how much confidence to place in it, what assumptions sit beneath it, and whether it is fit for purpose. ”

AI is only as useful as the data it can reliably draw on. While large language models are trained primarily on publicly available information, much of the data that matters most to charities is not public. It sits in internal systems such as supporter records, research, case notes, operational data and internal policies.

Public tools can draft text or summarise generic information, but they cannot offer meaningful organisational insight unless they are connected to internal data. Many charities begin this journey by uploading a small dataset or connecting a document library to an existing tool. As organisations mature, however, the opportunity expands. Linking multiple data sources, cleaning and structuring information, and storing it in reliable, reusable ways allows AI to operate across the organisation rather than in isolated pockets.

This is where data engineering and infrastructure become critical. At its core, this work ensures that data is accessible, well organised and fit for use, not only for AI, but for reporting, evaluation, learning and decision-making more broadly. Done well, it reduces friction, improves consistency and makes insight easier to share.

Strong data foundations also shape trust. AI systems built on incomplete, poorly understood or inconsistently managed data can produce fragmented or misleading outputs, quickly undermining confidence among team members and leaders. By contrast, well-designed data pipelines help ensure that AI outputs are reliable, explainable and appropriate for the contexts in which they are used.

Data engineering also plays a central role in governance. Decisions about what data AI systems can access, how consent is respected, how sensitive information is protected, and how outputs can be audited are embedded directly into data architectures and pipelines. In this sense, data engineering is a form of governance by design, shaping not only what AI can do, but what it is allowed to do.

Finally, strong data foundations determine whether AI can scale responsibly. Many pilots appear successful in isolation, but struggle when applied across teams or functions. Without robust data foundations, AI systems can amplify bias, surface misleading insights or expose sensitive information. In charity contexts, where trust, safeguarding and accountability are paramount, these risks carry real consequences.

Only at the upper end of this progression do applied AI skills come into play. These are the skills required to design, customise or significantly extend AI systems beyond off-the-shelf tools. Many charities will never need them. Others will, particularly those operating at scale, working with complex or sensitive data, or pursuing use cases where generic tools fall short.

Applied AI expertise is about shaping how AI behaves: deciding which problems it should solve, how models are adapted to specific contexts, how performance is evaluated, and how systems are monitored once deployed. This includes activities such as fine-tuning models, building bespoke decision-support systems, integrating AI into core workflows, and ensuring systems behave reliably under real-world conditions.

Applied AI skills become relevant only once organisations have clarity of purpose, sufficiently strong data foundations, and governance arrangements that allow AI systems to be deployed responsibly. Without these, advanced AI work risks increasing complexity, cost and risk without delivering meaningful benefit.

For charities that do reach this stage, however, applied AI skills can unlock significant value. They make it possible to move beyond generic assistance towards systems tailored to specific challenges, embedding AI in ways that support impact at scale while maintaining trust and accountability.

Rethinking skill in an AI-enabled charity

It is perhaps unsurprising that skills are cited as the number one barrier to AI adoption. As AI technologies mature, their value increasingly lies not with a small group of specialists, but in how effectively they are used across an organisation. What stood out in this research, however, was how different the skills people assumed they needed were from the skills that actually unlocked progress, and how important sequencing proved to be.

What emerged was a clear progression of capabilities that organisations develop over time. Progress begins with leaders who can frame problems well, set direction, and explain decisions clearly. It is then enabled by people who can translate organisational needs into practical use cases, assess outputs critically, and judge where AI fits and where it does not. More specialised data engineering and applied AI expertise matter, but usually later, and for far fewer organisations than is often assumed. In a sector defined by trust and accountability, this ordering matters.



Funding data and AI skills through the Growth and Skills Levy

In England, organisations with an annual pay bill of over £3 million pay 0.5% of that bill into the Growth and Skills Levy¹⁹. These funds are held by the government and can be used to pay for approved training and development, including apprenticeships.

Charities do not need to be levy payers themselves to benefit, as many employers donate surplus levy funds to other organisations. The Corndel Levy Pledge enables organisations to turn unspent Apprenticeship Levy into meaningful impact by gifting it to aligned charities.

Below are three Corndel programmes²⁰ charities are using to build their data and AI Skills.

Data Driven Professional

This programme helps teams use data more confidently in day-to-day decision-making, linking operational data to organisational goals and demonstrating how insight can support meaningful change. It introduces core AI literacy, including ethics, responsible use, and practical skills such as prompt engineering. Developed in partnership with Imperial College Executive Education, it builds the confidence and judgement that underpin effective AI use.

Business Impact with AI

Focused on making AI work inside real organisations, this programme develops the skills needed to identify use cases, design AI-enabled solutions, and lead change across teams. It supports those translating strategy into delivery and judging where AI fits and where it does not.

Applied AI Engineering

For organisations moving towards more advanced AI, this programme upskills people with technical aptitude to integrate, build and manage AI systems responsibly. It offers a pathway to develop in-house capability without relying solely on specialist recruitment.

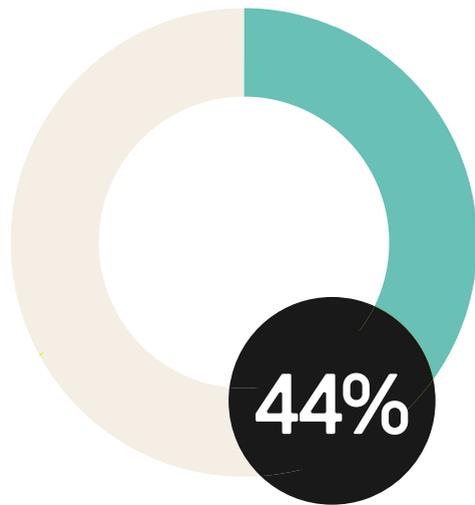
Corndel is the only national training provider rated double excellent by both learners and employers. It also has a dedicated charities team. Its programmes are developed in partnership with Imperial College and supported by Microsoft.



Time and capacity

Across the research, lack of time and capacity emerged as the third most cited barrier to engaging with AI, with **44% of leaders identifying it as a major challenge. Similarly, the 2025 Charity Digital Skills Report found that 63% of charities cite a lack of headspace and capacity as a significant barrier to digital and data work, ranking just behind squeezed finances and rising funding costs.**²¹

In that context, AI was often seen as “one more thing” rather than a solution.



44% of leaders identified a lack of time and capacity as a major challenge to engaging with AI

Working in survival mode

Several interviewees explained why the time barrier feels so immovable.

Paul Butcher, founder of Commonsensing AI and formerly Executive Director of Strategic Initiatives at Save the Children, described how funding shocks, cost pressures and rising demand have pushed the sector into a defensive posture. In that environment, innovation can feel like a luxury rather than a necessity. Investment in AI is rarely rejected on principle, but often deferred out of necessity, with the space to explore new tools shrinking quickly.

For Brian Rubenstein, President of the Rubenstein Impact Group and former leader at the advocacy affiliate of the American Cancer Society, this shortage is no longer an exceptional state for charities. It is the baseline. Most organisations are already operating at or beyond what feels sustainable. Introducing AI therefore carries an upfront time commitment that teams often do not feel they have the headspace for and oftentimes feel threatened by. The question is not whether AI could help, but where the time to learn, adapt and integrate it might realistically come from without putting themselves out of a job.

Daniel Gray, Director of Data, Digital and Technology at Crisis echoed this dynamic. He described how major infrastructure work can absorb almost all available organisational bandwidth. For Crisis, migrating core systems such as its CRM and data platform meant that capacity was effectively bunkered. Thinking time was devoted almost entirely to keeping critical enablers on track. In that context, experimentation with AI was not ignored, but postponed. Without these foundations, anything else would have been premature.

Together, these perspectives help explain why lack of time is so often cited. It is not reluctance, but triage.

The capacity paradox

Alongside this narrative sits a paradox: the organisations most constrained by time are often those most likely to benefit from even modest uses of AI. Many leaders argued that AI should not be framed as additional work layered onto already overstretched teams, but as one of the few tools capable of reducing workload itself.

Courtney Bugler, President and CEO of ZERO Prostate Cancer, captured this vividly. She rejected the idea that charity staff do not have time for AI, pointing out that many nonprofit professionals already feel they are barely keeping their heads above water. The promise of AI, in her view, is not to add new responsibilities, but to lower the waterline. Used well, it allows organisations to do work that previously required far more staff or external spend, freeing people to focus on the more human elements of their roles.

For Brittany Dunn, COO and Co-founder of Safe House Project²², the time and capacity question is existential. The scale of data required to identify trafficking victims cannot be processed by humans alone. Traditional hotline models require vast staffing to handle call volumes, yet still miss the majority of victims. AI, combined with a human-in-the-loop approach, is not an efficiency play so much as the only mechanism capable of expanding capacity to the level the problem demands. In this context, waiting longer is not neutral. It costs lives.

What happens once charities actually start using AI

This research offers an interesting glimpse into what happens once charities move beyond exploration and into use.

Among organisations already using AI, the most consistently reported benefit was saving time. 86% agreed that AI had made their organisation more efficient, with nearly a quarter strongly agreeing. For many, these efficiency gains were the first and most tangible outcome.

This matters because it speaks directly to concerns of those organisations where lack of time is the defining constraint.

Looking beyond efficiency, there is also little evidence of regret. Once an investment had been made, 60% of charities said AI had met or exceeded their

expectations. A further 32% took a neutral view, with interviews clarifying that most often this was because it was still too early to judge impact rather than due to dissatisfaction. Only 8% felt that AI had not met expectations.

The strongest signal, however, lies not in retrospective judgement but in future intent. All but one of the 275 respondents indicated an intention to increase their AI usage over the next 12 months, including all those who said that AI had not met their expectations. Even where outcomes were still emerging, there was virtually no appetite to pull back.

Finding time without stopping the organisation

One of the most practical challenges raised in interviews was not whether AI could help, but how organisations could realistically create space to learn and adapt without pausing day-to-day work.

Shawn N Olds, Co-Founder and Founding CEO of boodleAI, who began his career in the U.S. Army's 82nd Airborne Division, offered a useful analogy. His argument was not for intensity, but the opposite.

This approach speaks directly to concerns about time. It recognises that charities cannot pause day-to-day operations to learn AI, while challenging the assumption that waiting for a quieter moment is the safer option. In this framing, capacity is not something that must be fully in place before adoption, but something that is built through use.

Olds also offered a broader strategic lens for leaders tempted to delay. Drawing on a principle articulated by Cassie Kozyrkov, former Chief Decision Scientist at Google, that the hallmark of an AI-first organisation is paradoxically being AI-last. Leaders should begin with the problem, not the tool. If a challenge is generic and likely to be solved by the market, waiting may make sense. If it is specific to your mission, waiting simply delays progress. The problem will still be there next year.

The “Airborne School” approach to learning AI

Shawn N Olds compares organisational AI learning to Airborne School training. “Airborne School is three weeks long,” he explained. “The joke is it’s a five-day school shoved into three weeks, because you just do it again and again.”

The lesson, he argues, is that forcing AI learning into a short, intensive burst rarely works for charities. No organisation can realistically take four weeks off for an entire team to learn AI.

“With AI,” Olds told me, “I shove four weeks’ worth of training into six months. Every week I teach them a little more, and every week that allows them to do a little more.”

He likens the process to the boiling the frog analogy. Change happens gradually enough to feel manageable, but cumulatively it becomes transformative. “If you take people from nothing and try to do it all in one week, it’s caustic,” he said. “When you do one day a week, by month two they’re excited to come and show me how they’re using it. That’s what you want: people who feel empowered to explore it on their own.”

“ That’s what you want: people who feel empowered to explore it on their own.”

The risk of waiting

None of this implies that organisations should move without caution. There is, however, also risk in prolonged hesitation.

Capacity constraints rarely resolve themselves. In a sector facing rising demand, waiting can quietly harden into inertia. The risk is not that charities fail to adopt AI perfectly, but that they fail to engage with it at all, missing opportunities to relieve pressure where it is most acute.

Yet the very pressures that make AI feel out of reach are often those that make engagement most worthwhile. While lack of time is a defining constraint for many charities, initial engagement with AI frequently delivers returns through reduced manual effort, fewer bottlenecks and more resilient processes. Once charities cross the initial threshold and begin using AI in a considered way, the overwhelming majority choose to continue.

For many charities, the question is no longer whether they have time for AI, but whether they can afford not to.



Budget and funding

Financial constraints emerged as the fourth most commonly cited barrier to doing more with AI, with 43% of survey respondents identifying limited budget or funding as a significant obstacle. For many leaders, interest in AI quickly meets a familiar constraint: a lack of discretionary funding to support another initiative.

Similar findings are referenced in the 2025 Charity Digital Skills Report, where 69% of charities cited financial constraints as their biggest barrier to digital progress²³, and 60% have not accessed any specific funding for digital costs in the last year. Add to that, nearly a third report that funders will only cover limited core costs, with over a quarter struggling to find any funder willing to support a digital project²⁴.

This section explores why funding is so frequently cited as a barrier to AI adoption, and why that perception, while understandable, can become more constraining than necessary. It examines three related dynamics: sector-wide financial pressure, how AI investment is framed and justified, and the funding routes charities are using in response.

When funding pressure meets investment decisions

For many charities, the difficulty is not simply whether AI is affordable, but whether it is justifiable.

AI is often discussed in terms of transformation, scale and ambition. That framing can make it feel inherently expensive, experimental or risky. In organisations already making difficult trade-offs between frontline services and rising costs, it is understandable that AI can feel like an indulgence rather than a priority.

There is also a genuine opportunity cost to consider. Money spent on technology is money not spent elsewhere. Leaders are rightly cautious about committing funds to initiatives whose benefits may not be immediate or guaranteed.

This framing can, however, obscure an important distinction. AI does not need to be transformational to be valuable. As earlier sections of this book have shown, many

effective uses of AI involve incremental improvements using commercially available tools. These are often relatively low cost, and in some cases deliver net financial benefits through improved productivity, reduced manual effort or avoided spend. The more useful question, then, is not whether charities can afford AI, but where modest, well-judged use of it could relieve pressure and protect impact over time.

Funding the unglamorous essentials

For many organisations starting out on their AI journey, what hinders progress is actually the quality of their underlying infrastructure. While this is explored in more depth in the next section, poor data quality, fragmented systems and outdated platforms often make even modest AI use difficult.

This is a foundational issue. The IBM 2025 Chief Data Officer Study identifies data trapped in silos as the “Achilles’ heel” of AI initiatives, noting that without strong data foundations, every AI project becomes a prolonged and costly clean-up exercise. Using innovation funding to address data quality, integration or governance is not a diversion of funds, but a prerequisite for meaningful AI use.

Against this backdrop, several interviewees described the difficulty of securing funding for what they called the “boring but essential” parts of technology.

Amy Sample Ward, CEO of nonprofit community NTEN and Author of ‘The Tech That Comes Next’, noted that many charities struggle to secure funding for foundational work such as CRM migrations, data cleaning or website overhauls. These activities are critical enablers of future impact, but they are often labelled as overhead and deprioritised by funders. In that environment, AI can sometimes act as a proxy through which organisations are able to access investment for work they have needed for years.

Jim Fruchterman, veteran technology for social good advocate, explained that while many funders are not explicitly demanding AI, some are increasingly offering grants framed around AI or innovation. His advice to leaders encountering such opportunities, particularly those without strong data foundations or a desire to pursue AI for its own sake, was to design projects where the majority of the work focuses on data infrastructure, governance and internal capability, with a smaller portion allocated to AI itself. Even if the AI component proves less valuable than hoped, the organisation is left with stronger foundations.

Seen this way, AI funding can act as a route into long-overdue investment in digital basics rather than a distraction from them.

From cost to impact

A second, often overlooked, factor shaping funding decisions is what charity leaders believe their donors will tolerate.

As part of this research, 275 organisations at different stages of their AI journeys were asked how positively they believed their donors would view the use of AI. Very few expected active opposition. However, among organisations yet to adopt AI, only 23% believed their donors would view its use positively.

By contrast, this figure was roughly double among organisations that had already begun adopting AI. This may indicate that concerns about donor perception diminish once charities gain direct experience of the technology. Alternatively, it may reflect a selection effect, where organisations that believe their donors will respond positively are more likely to proceed with adoption in the first place.

Either way, uncertainty about donor sentiment appears to influence decision-making long before any conversation takes place. Several interviewees argued that this uncertainty reflects outdated assumptions rather than real donor resistance.

Joshua Tripp, who leads on innovation at World Central Kitchen and has extensive experience in nonprofit finance, argued that whether AI spend counts as overhead or programme cost depends entirely on application. Automating accounts payable may reasonably be considered overhead. Using AI to analyse drone footage to route food to hungry communities faster is programme delivery.

As AI becomes more embedded in service delivery, that boundary increasingly blurs. Tripp's point was that sophisticated donors are often willing to fund infrastructure when the connection to impact is explicit. When charities can articulate how data and technology enable a faster, smarter or more community-led response, the conversation shifts from software to outcomes.

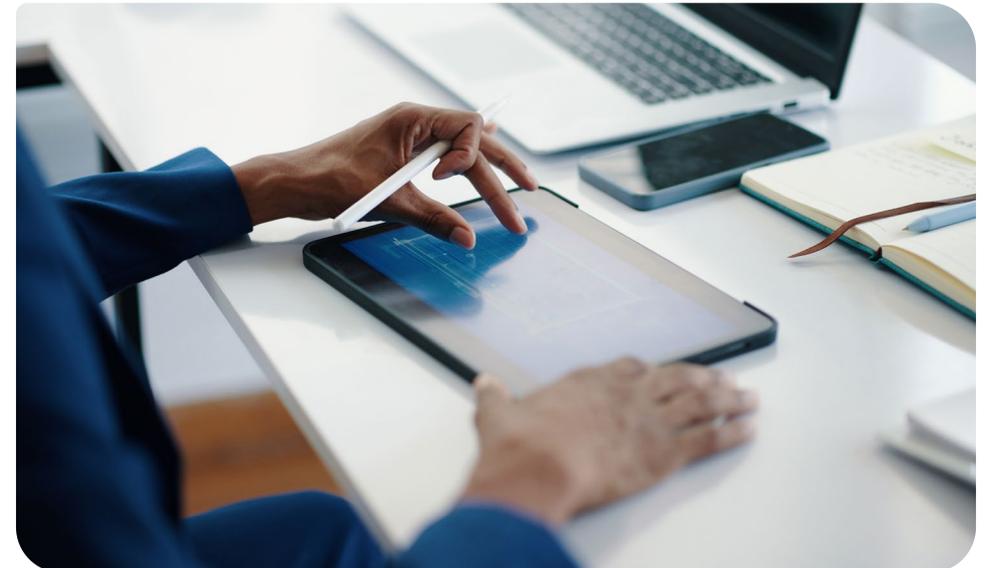
Paul Butcher framed this more fundamentally, drawing on the widely used concept of 'Impact Return on Investment'. A charity might appear efficient by spending very little on employees or systems, yet still deliver minimal overall impact. Conversely, an organisation that invests more in people and infrastructure may deliver significantly greater change. The relevant question, he suggested, is not how little a charity spends, but how much impact a donation creates.

This view is supported by wider research. Studies cited by organisations such as Vanguard Charitable²⁶ and GiveWell²⁷ have shown that having low overheads is by itself a poor proxy for effectiveness. In some cases, charities with slightly higher

overheads perform better precisely because they invest in people, systems and learning. While much of this research predates the current wave of generative AI, the insight remains highly relevant: robust infrastructure enables better outcomes.

Shawn N Olds, Co-Founder and Founding CEO of boodleAI, was more direct. He argued that the idea that "our donors won't tolerate this" is misplaced, because most donors are business-savvy and understand ROI (both return on investment and return on impact). He pointed out that if a nonprofit delivering meals to vulnerable people told donors that it did so using horse and cart rather than vehicles, it would not be praised for avoiding technology. Not only would it cause spoiled food and slower delivery, donors would rightly question the inefficiency. Olds argued that AI sits in the same category. Donors are not paying for effort. They are paying for results.

As Courtney Bugler, President and CEO of ZERO Prostate Cancer, told me, if AI enables a team to do more with the same resources, then using it can be one of the most respectful ways to treat a donation.



Budget cycles

Even where funding exists, another barrier often intervenes: the mechanics of nonprofit budgeting.

Several leaders described a mismatch between the pace of AI development and the rigidity of traditional financial planning. Multi-year strategic plans and annual budget cycles are designed to promote stability and accountability, but they can struggle to accommodate rapidly evolving technologies.

Sandra Diaz, Director of Impact and Evaluation at Restore NYC, described how generative AI emerged midway through a three-year strategic cycle. While the organisation recognised its potential, it was effectively locked into a plan designed for a pre-AI world also being impacted by shifts in federal funding. By the time AI could be formally incorporated into strategy and budgets, the organisation risked being significantly behind. So they launched an initiative to map requirements, develop a prototype, and estimate costs of full implementation to pursue fundraising.

Joshua Tripp of World Central Kitchen was blunt about the implications. “I think the world of three-year budget cycles is long gone.” Rather than allocating funds to specific tools years in advance, Tripp sees a growing trend among charities working in fast-moving contexts, such as humanitarian response, to set aside flexible investment pools that can be deployed as needs arise, with rapid board engagement when decisions are required. This approach allows leaders to remain accountable for risk and spend, while recognising that agility itself is a form of risk management in fast-moving environments.

“ Several leaders described a mismatch between the pace of AI development and the rigidity of traditional financial planning.

Choosing how to invest: buy, configure, build

For most charities, funding decisions around AI quickly come down to a practical question: how should technology be acquired and deployed?

Over coffee in New York, Steven Francisco, VP of Technology at LEE Foundation for Civic Leadership, offered a useful framework: buy, configure or build:

BUY

By purchasing off the shelf tools, organisations gain immediate utility, yet no competitive advantage, because every other charity can do the same.

CONFIGURE

Platforms such as Salesforce can be adapted to organisational needs, delivering tailored value without the cost and risk of full product development.

BUILD

High risk and high reward. It is expensive and difficult, but can deliver genuinely unique capability where the market is unlikely to solve the problem.

Francisco's warning was that leaders must be ruthless about where they play. If an organisation cannot realistically win at “build”, it is better to wait until the capability becomes configurable or purchasable.

Yet for Jim Fruchterman, charities shouldn't be in a hurry to move into the build phase. “99.9% of nonprofits should not be trying to build an AI product.” Building is expensive, risky, and demands capabilities most organisations do not have. He described it as a “step five” ambition when most charities are still on “step one”.

A nonprofit, he argued, should no more have a software engineering team than a local restaurant or golf club should. The golf club buys its booking system. It does not build one. In this view, administrative and support functions are best served by affordable tools where the technology is already powerful and the cost of failure is low.

Funding AI Responsibly

Budget and funding are real constraints. They reflect trade-offs and accountability to donors and beneficiaries. But this research suggests they are more flexible than they often appear.

As with time and capacity, funding tends to feel most restrictive when AI is framed as transformational, bespoke or experimental. When it is reframed as incremental, enabling and focused on foundations, the picture changes. Many AI-enabled improvements are affordable. Some reduce costs directly. Others increase impact without requiring proportional increases in spend.

The challenge for leaders, then, is not to find unlimited funds for AI, but to ask better questions about return on impact, opportunity cost and long-term viability. In a sector under sustained financial pressure, investing in tools that help organisations use scarce resources more effectively is not a departure from charitable values. It is one way of safeguarding them.

This has important implications for how AI is funded. Funding AI responsibly depends on understanding what money is actually buying, and where value is being created. That, in turn, rests on data quality, systems and organisational readiness. Those foundations, and the risks of neglecting them, are the focus of the next section.

“ The challenge for leaders, then, is not to find unlimited funds for AI, but to ask better questions about return on impact, opportunity cost and long-term viability. ”

Technology and data

In the survey, 35% of respondents pointed to data quality or availability as a barrier to progressing with AI, and 28% said their systems were not set up for it. Although this places data and systems below skills, ethics, time and funding, what matters is that these constraints behave differently from others.

Unlike cultural or leadership concerns, which often soften as organisations gain experience, data and systems constraints tend to persist. In this research, among charities yet to adopt AI strategically, 36% cited data problems and 28% cited systems not being set up for AI.

Among those that had adopted AI strategically, the figures were still 31% and 27% respectively. This suggests that data and systems are not simply “entry requirements” for adoption. They remain live constraints that determine whether early experimentation can become dependable, scaled practice. Confidence may grow, but underlying systems do not improve simply because AI has been introduced.

This matters because AI does not operate independently of organisational systems. It depends on them. It draws directly from the data, structures and permissions that already exist. Where data is fragmented, systems are outdated, or access is poorly controlled, AI does not compensate for those weaknesses. It exposes and amplifies them.

It is also important to be clear about what this does not mean. Charities do not need perfect data, fully modernised systems, or enterprise-grade infrastructure before they can benefit from AI. For many organisations, meaningful progress comes from working with what they already have, provided expectations are realistic and risks are understood.

Why data and systems matter for AI

AI works by identifying patterns in information and using them to support decisions, predictions or actions. Whether drafting communications, forecasting demand or prioritising cases, these systems depend on information that is consistent, timely and interpretable.

This is why technologists repeatedly returned to the same point in interviews: AI is only as reliable as the systems feeding it. When data is incomplete, inconsistent or poorly governed, AI does not simply become less effective. It becomes unpredictable.

For example, a predictive model flagging risk relies on definitions that remain stable over time. If one team records a “case closed” as service delivery completed, while another uses it to mean beneficiary disengagement, the system cannot reliably learn from that data. In these situations, the limiting factor is rarely the model itself. It is the condition of the data and systems it connects to.

Connecting generative AI systems such as Copilot or ChatGPT to internal records, operational systems or decision-making processes does not introduce new weaknesses so much as reveal existing ones. Long-standing issues with data quality, system design and governance become visible once AI starts relying on them. In practice, this means that the limits of AI are often set less by the sophistication of the model than by the reliability of the systems it connects to. Where those foundations are strong, AI can amplify capability. Where they are weak, it can amplify inconsistency, risk and uncertainty just as quickly.

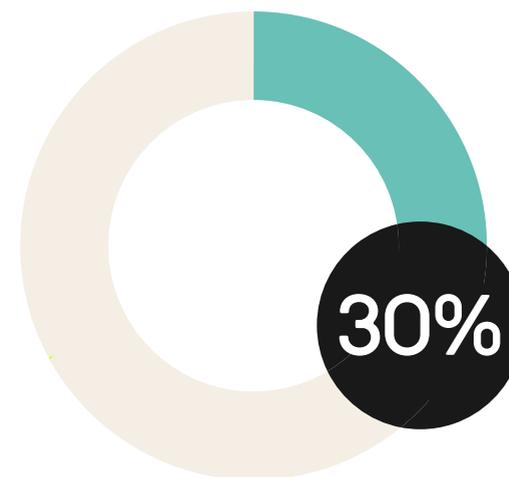
Fragmented data

A consistent theme across interviews was fragmentation. Data spread across multiple systems, recorded in different formats, governed inconsistently, or simply not trusted by the people expected to use it.

This fragmentation usually develops gradually. International charities, for example, may find that the same dataset contains both UK and US date formats. Daniel Gray, Director of Data, Digital and Technology at Crisis, described fragmentation as a legacy issue for the sector. Years of organisational growth, system changes and under-investment has left data for many charities in a poor state. When data cannot be relied upon, leaders become cautious about acting on it.

The Charity Digital Skills Report 2025 reflected this picture, finding that 30% of charities rate their systems and databases as poor or non-existent, and nearly a quarter say the same of their IT support and hardware. These are not edge cases. They shape everyday judgements about what organisations feel able to attempt.

Yet for those who are daunted by the prospect of solving this problem, Alice Kershaw, Head of Digital Transformation at The Wildlife Trusts, had a word of encouragement. Rather than attempting to perfect everything, Kershaw suggests that charities should focus on whether a dataset is good enough to support a specific decision. She suggested distinguishing between research-quality data which many charities will have experience producing, and decision-quality data. Her point was that organisations sometimes hold themselves to an unrealistic standard before taking action. Trying to fix all data at once can exhaust limited capacity. Prioritising the data that matters most allows progress without boiling the ocean.



30% of charities rate their systems and databases as poor or non-existent, and nearly a quarter say the same of their IT support and hardware.

Legacy Systems

Alongside data quality, legacy systems were repeatedly cited as a limiting factor. Many charities rely on systems that were never designed to work together, or that were implemented without the training and resources needed to use them fully. Information may be stored differently across platforms, or team members working across multiple systems may find no obvious way to transfer data between them. To keep work moving, teams create manual workarounds.

The most common workarounds involve employees using spreadsheets to bridge gaps between systems. The reality for many charities is that so much of the knowledge about how things really work lives with individuals rather than in the systems themselves. Over time, a small number of people become an informal translation layer, connecting data, correcting errors and compensating for limitations in tools or skills. It's the 'way we've always done things' knowledge that every new employee learns.

This human translation layer allows organisations to function, but it comes at a cost. When key individuals leave, much of that knowledge leaves with them. AI systems also have a low tolerance for such informal fixes. When someone extracts data manually and adjusts it, the AI system does not see the reasoning behind those changes. It only sees the underlying inconsistency. Over time, the logic that explains how a process really works exists only in people's heads, rather than in the system itself.

As a result, processes that rely on human interpretation become much harder to automate. Where work depends on moving data between legacy systems, there is often no simple or obvious technical fix, particularly when systems are old or poorly maintained. In these cases, the difficulty is not identifying where AI might help, but translating human workarounds into something a system can reliably act on.

For resource-constrained charities, this creates a real tension. Organisations are often drawn to AI as a way to increase impact with limited headcount and funding. Yet addressing legacy systems and embedded workarounds can feel like a prerequisite that consumes time, money and attention before any visible benefit appears. Research by McKinsey²⁸ shows that even in leading organisations, data readiness remains one of the biggest barriers to scaling AI. Without addressing these foundations, AI initiatives risk becoming prolonged data-cleaning exercises rather than genuine improvements to service delivery.

Cybersecurity

AI also intensifies a risk charities already face: cybersecurity.

Charities hold sensitive and valuable information, including financial details, personal histories, location data and case notes. This makes them attractive targets. Yet many organisations lack the resources to defend themselves at the same pace as attackers evolve.

The Okta Businesses at Work 2025 report shows that nonprofits are now the second most targeted sector globally for cyber attacks, behind only energy, mining, oil and gas²⁹. Other studies indicate that phishing and account compromise affect nearly every charity, and that nearly half of all charities have experienced digital financial fraud. More concerning still is the rise of sophisticated surveillance tools used against organisations working in human rights and humanitarian contexts.

Paul DePonte, Executive Director of the National Crime Prevention Council, illustrated this risk starkly. His own organisation, dedicated to crime prevention, was itself the victim of a cyber attack involving financial fraud. His conclusion was blunt. If an organisation whose mission is to prevent crime can be compromised, anyone can.

Criminals are early adopters. They are not constrained by ethics, procurement cycles or governance. Charities, by contrast, operate within tight budgets and high accountability. This creates a structural disadvantage. There is, however, an important counterbalance to these risks.

“ Nonprofits are now the second most targeted sector globally for cyber attacks. ”

The benefit of cloud platforms

One of the most significant advantages charities now have is that they no longer need to manage much of their data infrastructure alone. Moving data into cloud-based platforms allows organisations to outsource part of the burden of storage, resilience and security to providers with vastly greater resources and a strong commercial incentive to invest in protection.

Most charities no longer maintain their own systems. They rent them through monthly or annual subscriptions. This means security updates, patches and improvements are pushed automatically, rather than relying on scarce internal capacity to identify and fix issues.

These platforms are no longer just storage solutions. Charities may store documents in tools such as SharePoint or Google Drive, but they also gain access to increasingly powerful capabilities layered on top. Microsoft, for example, is embedding AI across its ecosystem through Copilot, enabling tasks such as drafting documents, summarising emails and analysing files within familiar tools. Platforms such as Microsoft Foundry extend this further, allowing organisations to connect modern AI models directly to their data, often without that data leaving the Microsoft environment. Google offers similar capabilities through Vertex AI and its integrations with Google Workspace.

For many charities, AI is thus being introduced inside environments they already use and understand, rather than through unfamiliar standalone tools. This lowers the barrier to experimentation and reduces some security risks associated with moving data between systems.

It does, however, come with trade-offs.

Vendor lock-in and cost pressure

When charities rely heavily on a small number of digital platforms, switching becomes difficult. Large technical migrations are costly, disruptive, and place significant pressure on teams already managing multiple change programmes. In practice, this means that price increases are more likely to be absorbed than avoided.

Several interviewees for example, shared that in 2025 Microsoft made changes to its nonprofit licensing model, including ending the availability of free Business Premium licences for many organisations and moving to discounted paid plans instead. These changes were presented as part of a broader shift towards long-term sustainability and greater parity across customer groups, rather than a withdrawal of support for the sector. Even so, they illustrate how pricing can evolve as platforms expand their capabilities and embed new functionality.

Similar dynamics exist in the CRM market, where many interviewees described AI features being introduced primarily as paid add-ons, sometimes at a significant cost. For organisations already locked into a particular platform, these changes can feel difficult to challenge, even when budgets are under strain.

During our interview, Ha Cole, EMEA Chief AI Officer for Microsoft Elevate, which focuses on supporting the company's nonprofit and social impact work, emphasised that these developments need to be understood in the context of how large technology providers balance sustained investment with affordability.

She noted that Microsoft continues to subsidise the nonprofit sector globally through a combination of grants, discounted software and services, and matched employee donations, amounting to several billion dollars each year. At the same time, pricing adjustments often reflect years of cumulative innovation, with new AI capabilities increasingly integrated into core products rather than offered separately. As a result, charities may be receiving substantially expanded functionality, including enterprise-grade security and embedded AI tools, within licences they already hold, even where headline prices increase. A recurring challenge, she observed, is that many organisations are not fully aware of what is already included, and may assume that adopting AI necessarily requires new expenditure when this is not always the case.

The risk, however, is losing sight of the foundations in the rush to adopt visible AI tools. In a sector already under pressure, neglecting data quality, security, and systems resilience in favour of surface-level innovation is a false economy.

Charities are already doing extraordinary work under intense constraints. AI does not require them to solve every technical problem at once. It does require them to recognise which problems matter most, and to treat data, systems, and security as central to delivering impact safely and sustainably.

When AI exposes the foundations

Technology and data readiness are rarely the most visible barriers to AI adoption, but they are among the most consequential. Skills can be developed, confidence can grow, and use cases can evolve. Fragmented data, legacy systems and unmanaged risk persist until they are addressed deliberately.

This does not require charities to modernise everything at once. It does require clarity about which foundations matter most, where risk is accumulating, and which weaknesses AI is likely to expose rather than resolve. Treating data, systems and security as secondary concerns may allow experimentation to begin, but it rarely supports safe or sustainable progress. AI can extend what organisations are capable of doing, but only when the systems beneath it are strong enough to carry that weight.



Responsibility in high-trust environments

Charities operate in environments where trust is essential. They rely on the confidence of donors, partners, regulators and the communities they serve, often in contexts where mistakes can carry serious consequences. This makes decisions about AI different in kind from those faced by many commercial organisations.

Ethical concerns and questions of mission fit were the second most frequently cited barrier to AI adoption in this research. When leaders were asked what those concerns actually meant in practice, two issues were consistently brought together: environmental impact and wider social risks, such as algorithmic bias. This section explores what responsibility looks like when AI is introduced into high-trust settings. It examines ethical concerns, transparency, environmental impact and governance, and how charities are responding in practice.

Environmental Impact

The environmental impact of artificial intelligence has attracted a great deal of attention in recent years, and rightly so. The rapid growth of large language models has coincided with a wider shift in how resources are consumed by digital infrastructure. After a long period of relative stability, electricity demand from data centres is now rising sharply, driven in large part by AI workloads. Against this backdrop, it is unsurprising that many charities identified environmental impact as one of their most pressing barriers to using AI.

This section examines those concerns more closely. It looks first at what the available data tells us about the environmental footprint of AI, before asking to what extent is avoidance a prudent or effective response for charities deciding whether to use AI?

Across interviews, leaders acknowledged a wide range of environmental issues, yet

two were cited most frequently: electricity consumption and water use. The International Energy Agency projects that by 2030, annual electricity consumption from data centres could reach between 945 and 2,200 terawatt-hours, roughly equivalent to the current annual electricity consumption of India³¹. It is often assumed that this footprint is driven primarily by the training of large models. In practice, the balance is quite different. Between 80 and 90%³² of energy demand comes from the everyday use of AI systems once they are deployed. This is known as 'inference'.

Water use presents a related, and in some respects more localised, concern. Data centres require cooling to prevent overheating, and many rely on water-intensive systems. Global AI-related water demand is projected to reach between 4.2 and 6.6 billion cubic metres by 2027³³, a volume exceeding the entire annual water use of Denmark³⁴ and approaching half that of the United Kingdom.

The two key indicators here are water withdrawn and water consumed. In many data centres, a proportion of the water used for cooling is returned to the local system, while the remainder is lost through evaporation and therefore unavailable for reuse. It is this consumptive use, rather than total withdrawal, that determines the majority of the local impact on water availability, and importantly from a societal view, this impact of this is not evenly shared.

A single AI prompt processed in a hot, water-stressed region such as Arizona can consume around 30 millilitres of water, compared to under 8 millilitres in cooler regions³⁵. This raises the risk of placing a disproportionate burden on areas that are already resource constrained. Given the urgency of climate change and growing pressures on water systems, it is therefore understandable that charities are cautious about adding to these pressures.

At the same time, environmental charities I spoke to were keen to stress that environmental impact is rarely assessed in isolation. Charities do not operate outside the systems they are part of. They consume energy, rely on digital infrastructure, heat offices, travel to meet funders, and use tools whose environmental costs are rarely zero. None of these choices are neutral. They are accepted because the value they enable is judged to outweigh the cost.



How many prompts are in a burger?

It was against this backdrop that several experts argued environmental concerns around AI risk being taken out of context. During my conversation with nonprofit marketing expert, George Weiner, I asked whether avoiding AI use was an important step. He offered a simple comparison: “You know... don’t eat a burger tomorrow, and you’re good for ChatGPT for a year.”

The comparison initially felt implausible. Yet when set against the available data, the comparison appears to understate the difference rather than exaggerate it. A recent study by Google³⁶, which measured the end-to-end environmental impact of AI inference in a live production environment, found that a median AI text prompt consumes around five drops of water, approximately 0.26 millilitres. This figure reflects water consumption (that not returned to the water system) rather than total water withdrawn³⁷.

By contrast, producing a quarter-pound beef burger is typically estimated to require between 1,500 and 3,000 litres of water. This figure reflects consumptive water use across the full supply chain, including animal feed, irrigation, processing and transport. Even using the lower estimate, skipping a single burger would equate to around 5 million AI text prompts, enough for years of everyday use, even at heavy prompting rates.³⁸

During our interview, Ha Cole, EMEA Chief AI Officer at Microsoft Elevate, emphasised the importance of charities considering the environmental impact of their use of AI, without allowing those concerns to become a barrier to using the technology where it can deliver the greatest benefit. To illustrate the scale involved, she pointed to research by Hannah Ritchie, which suggests that, in carbon terms, a return flight such as London–Copenhagen can be on the order of hundreds of thousands of prompts, depending on assumptions.

As Cole put it, the challenge is “to prioritise low-risk uses that minimise harm to beneficiaries, colleagues, and organisations, and to learn through doing rather than remaining stuck in theory or abstract fear”. While those concerns are real, she warned that failing to break them down into practical decisions can have a paralysing effect.

Ha Cole was keen to stress that more compute-intensive uses, such as image and video generation, carry substantially higher environmental impact than text-based prompts alone. Charities therefore need to consider their likely patterns of use before drawing conclusions from any single comparison. Even with this caveat, several experts cautioned that charities risk over-weighting the environmental

impact of everyday AI use while under-examining other, far more water- and energy-intensive choices already embedded in daily life.

As George Weiner reflected near the end of our interview, “I think putting more good words in the world at that cost is better value than, frankly, ignoring a tool that can save you that kind of time and output potential.”

Environmental impacts of everyday activities

A median AI text prompt consumes approximately 0.24 watt-hours of electricity, roughly equivalent to watching nine seconds of television³⁹. An hour of phone charging equates to around 50 to 70 prompts, while an hour of laptop use equates to between 125 and 300 prompts. At a rate of ten prompts per day, it would take roughly 148 years to use the same amount of water as a single cup of coffee.⁴⁰

Between 2024 and 2025, the median AI prompt underwent a 33-fold reduction in energy use and a 44-fold reduction in carbon intensity⁴¹ (Elsworth et al., 2025). Not only is the energy use of everyday text prompts modest compared with many common activities, it is also falling rapidly over time.

Beyond individual use cases, organisations must account for the wider impacts of AI as a system, not just as a set of tools. These impacts arise at different points in the lifecycle of the infrastructure that supports AI, from how it is built and cooled to how data is moved and stored.

At the physical level, the rapid pace of hardware upgrades in data centres generates significant electronic waste, which can release toxic substances and add to already stretched landfill capacity. The scale of modern data centres also brings land-use consequences, sometimes displacing habitats and affecting local biodiversity.

Operational impacts extend beyond hardware. Water use is not confined to cooling a single task or system, but accumulates across many data centres drawing on the same local water supplies. In regions already facing scarcity, this cumulative demand can place additional strain on shared resources. Upstream supply chains add further pressures, from the mining and processing of metals used in servers and storage devices to the emissions associated with global manufacturing and transport.

Even once systems are in place, impact does not stop. The movement of data itself carries an energy cost, as storing and retrieving information across cloud networks

requires electricity. Cooling systems introduce another layer of risk, relying on refrigerant gases that, despite being designed as closed systems, can leak over time and contribute to greenhouse gas emissions.

Yet it is important to remember that many of these environmental costs are already being incurred by charities through their existing use of cloud storage, digital collaboration tools and online services. Choosing not to use a tool such as ChatGPT does not mean avoiding this infrastructure altogether. The question is not whether charities participate in these systems, but how deliberately they do so, and whether the value created justifies the additional burden imposed.



The choice for leaders to make

In this light, many leaders reframed the question they were asking. The issue was not whether AI has an environmental impact (it does), but whether that impact is proportionate to the benefits it might unlock. Several warned that focusing too narrowly on AI's costs risked obscuring the value these technologies could deliver for beneficiaries, particularly in a sector under increasing pressure to do more with limited resources.

For environmental organisations, this type of judgement is not new. Alice Kershaw, Head of Digital Transformation at The Wildlife Trusts, offered a helpful analogy drawn from land management. From a narrow emissions perspective, livestock could be considered some of the more carbon intensive assets across the federation. Yet they are retained because they deliver outcomes that outweigh their cost. As Kershaw said, cows are part of land management not because they are low impact, but because they are high value.

The same logic now shapes how The Wildlife Trusts approach AI. Certain uses, such as AI-generated video, are currently ruled out because the environmental cost is high and the added value marginal. At the same time, machine learning is used where it meaningfully improves conservation outcomes. The decision is not whether to use AI, but where its use can be justified.

This principle extends beyond environmental charities. Michael Sheldrick of Global Citizen argued that the comparative advantage of the charity sector does not lie in marginal reductions in electricity use achieved through abstention from digital tools. It lies in using those tools to drive policy change, mobilise public support and accelerate transitions at a systemic level.

Taken together, these perspectives point to a shared conclusion. Environmental responsibility is not about avoiding every activity with a measurable footprint. It is about understanding trade-offs and making conscious, proportionate choices. Each charity will arrive at its own view, shaped by mission, context and values. Environmental impact is a real and legitimate concern, but for charities, abstaining from AI is rarely a proportionate or effective environmental response. Responsible use, governance and value-based trade-offs matter more than avoidance.

As I reached the end of my research, one line shared by a leader stayed with me: "If all charities were to 10X their AI use overnight, it would have only a negligible effect on energy and water consumption, yet a huge positive impact on the lives of their beneficiaries."

Ethics

If environmental concerns prompt charity leaders to ask whether AI is a responsible use of shared resources, ethical concerns raise a deeper question: does its use increase risk for employees, beneficiaries or supporters?



Bias, inequality and the limits of neutrality

Few ethical issues generate as much anxiety for charity leaders as algorithmic bias. Charity leaders are rightly concerned that AI systems trained on historical data may replicate or amplify existing inequalities, particularly when used with people who already experience discrimination or exclusion. When the stakes involve access to support, safety or dignity, even small distortions can have serious consequences.

As Sarah Wyrer, a PhD researcher in AI ethics, explained to me, much of the training data used in early large language models was drawn from uncurated portions of the internet, including forums, comment threads and social platforms.

“These models were trained on vast amounts of online text, much of it scraped from spaces that were never designed to represent people fairly or accurately. The result is not a neutral picture of the world, but an accumulation of biases, stereotypes and distortions that reflect how people behave online, not how they necessarily live or treat one another offline.”

As models trained on this material are increasingly used to summarise information, guide decisions or mediate access to services, there is a risk that these distortions become normalised and perpetuated at scale.

Research from organisations including the Alan Turing Institute and the Centre for Data Ethics and Innovation⁴² reinforces these concerns. Their work shows how biased data, poorly chosen proxies and unexamined assumptions can lead to discriminatory outcomes in areas such as healthcare, recruitment, policing and welfare decision-making.

Proxies matter because AI systems often cannot measure what we care about directly. Instead, they rely on substitutes to stand in for more complex realities. A postcode may stand in for deprivation, previous arrests for risk, or historical service use for need. These shortcuts are rarely neutral. They can encode social inequalities into technical systems, even when no sensitive characteristics are explicitly included. Similarly, unexamined assumptions about what counts as success, risk or eligibility can quietly shape outcomes in ways that disadvantage certain groups.

For charities, this creates a distinct tension. Many organisations are seeking to use AI precisely to reach people who have historically been overlooked or underserved. Yet those same people may be poorly represented in data, misclassified by proxies, or rendered effectively invisible to automated systems. In trying to fix one set of inequalities, charities risk inheriting another.

In response, some organisations draw clear ethical boundaries. Berenice Levenez, Strategy, Technology and Transformation Director at Girlguiding UK, described the organisation's decision never to use AI-generated images. Generative image models often reproduce deeply embedded gender stereotypes, and Girlguiding judged that the risk of reinforcing those patterns outweighed any operational benefit.

Yet at the same time, several leaders were quick to caution against treating human judgement as a neutral alternative. Human decision-making is itself shaped by bias, habit and limited perspective. In some contexts, carefully governed AI may reduce harm rather than increase it. Brittany Dunn, COO and co-founder of the Safe House Project, described how human gatekeepers, such as police dispatchers, can carry implicit biases that contribute to the under-identification of victims from marginalised groups. In her view, AI systems that analyse patterns across large volumes of reports can sometimes surface cases that would otherwise be missed.

The ethical question, then, is not simply whether AI is biased, but under what conditions it may be more or less biased than the human systems it supports, and whether it can still serve a charity's purpose despite those limitations. That judgement cannot be made in the abstract. It depends on context, data quality, oversight, and what is at stake when decisions go wrong.

Vulnerability, data and unintended harm

Ethical risk increases most sharply when AI is applied to sensitive data about people. Many charities work with survivors of violence, people in crisis, children, refugees and others whose safety depends on discretion, care and trust. In these settings, decisions about technology can shape safety, dignity and willingness to seek help.

Discussions of AI risk are often divided between external threats, such as hacking or malicious misuse, and internal risks arising from design choices, everyday use and unintended consequences. Both matter, and for charities they are closely intertwined. As explored earlier in this book, nonprofits are now the second most targeted sector globally for cyber attacks. Charities hold valuable information, often lack extensive defensive resources, and operate in contexts where exposure can cause serious harm. AI intensifies this landscape by increasing the volume of accessible data, the number of systems interacting with it, and the speed at which misuse can occur.

External threats and the changing risk landscape

External threats are evolving rapidly. Techniques such as prompt injection allow users to manipulate AI systems into revealing information they were not intended to share or to bypass safeguards built into a tool. Poorly designed conversational interfaces can allow sensitive data to be inferred or extracted through repeated questioning, even when no single response appears problematic on its own. In beneficiary-facing systems, well-meaning users may inadvertently disclose information about themselves or others. These risks now form a significant part of the threat environment charities must consider when deploying AI.

At the same time, ethical risk does not come only from outside. Many of the most serious harms described by interviewees arose from internal practices that shifted subtly after AI was introduced. As tasks that once required considerable time become quick, automatic and routine, it becomes easier to capture more detail, retain it for longer, search across it and reuse it in new contexts. In doing so, AI changes the risk profile of ordinary work.

Sandra Diaz, Director of Impact and Evaluation at Restore NYC, highlighted a particularly subtle example involving AI note-taking and transcription. In trauma-informed practice, it is common for survivors of human trafficking to express guilt or self-blame. If those statements are captured verbatim by an AI system and stored digitally, they may later be subpoenaed and used against the individual in court. The ethical issue here is durability. Information that would once have existed only within a fleeting conversation, or as handwritten notes filed away, becomes persistent, searchable and extractable. Accompanying the efficiency gain is a new form of exposure.

Similar issues arise when AI is used to improve access to information inside organisations. Daniel Gray, Director of Data, Digital and Technology at Crisis, described how AI-powered search tools can surface sensitive material that was previously hard to find. AI assistants with access to search across email, shared drives and document libraries can reveal information from files someone was once given access to, such as historic HR folders or beneficiary records, even when those materials were never intended to be routinely visible. Without careful data governance, improving access can unintentionally broaden who can see what. In these cases, AI does not introduce new sensitivity, but removes the practical barriers that once limited exposure. This makes long-standing data practices newly consequential and links internal choices more directly to external risk.

Common AI risks charities must consider

Prompt injection and manipulation

Users may be able to manipulate an AI system into ignoring safeguards or revealing information it was not designed to share.

Unintended disclosure through conversational tools

Sensitive information can be revealed gradually through interaction, even when no single response appears problematic.

Data leakage through logs and outputs

Prompts, transcripts and model responses may be stored or reused, creating exposure beyond the original use

Over-collection of data for context

AI systems often perform better with more information, encouraging the capture of detail that may be unnecessary or inappropriate.

Expanded visibility through search and summarisation

AI can surface documents or insights that were previously difficult to locate, increasing the risk of inappropriate access.

Retention or reuse by external services

Data processed by AI tools may be retained or handled in ways that are not immediately visible to the organisation.

These risks span technical, ethical and organisational boundaries and must be addressed together rather than in isolation.

Privacy, autonomy and restraint

Beneficiaries often engage with services because they trust their information will be handled with care. AI challenges this by opening the door to practices that slowly weaken privacy. Systems designed to improve performance can also reward the collection of additional context, on the assumption that more information leads to better results. Over time, this can undermine principles such as data minimisation and purpose limitation. Information gathered for one purpose can later be used for others.

Using third-party AI services adds additional complexity. This data is often processed outside the charity, retained for quality assurance, or handled in ways that are difficult to see or control. For charities working with sensitive personal information, these questions go to the heart of trust. Ethical use of AI in such settings often involves collecting less data, retaining it for shorter periods, and resisting the pull to centralise everything simply because technology makes it possible.

Questions of autonomy and consent run through these concerns. Charities often work with people who have few alternatives. Individuals may depend on services for safety, housing or essential support. Introducing AI into these interactions can change how choice is experienced. In principle, people may consent to AI-mediated services. In practice, alternatives may be unavailable. Ethical use of AI requires attention to whether consent is meaningful, and whether opting out carries an implicit cost.

These issues are particularly acute for children, people with additional support needs, and those in acute distress. Capacity varies, and the duty of care should increase accordingly. AI can assist employee judgement in these contexts, but should not dilute responsibility or reduce opportunities for human engagement. This is where human oversight becomes essential.

Reliability, access and fairness

Trust also depends on reliability. Charities are trusted intermediaries, and people often act on the information they provide. Human error is not new. Leaflets sometimes contain mistakes. Advice given on helplines can occasionally be outdated. Yet these errors are usually limited in scale and scope. They tend to arise from misunderstanding or ambiguity and affect relatively small numbers of people.

AI changes this pattern. Errors can be repeated consistently, delivered with confidence and distributed widely. The types of mistakes also differ. Generative systems may fabricate information, blur distinctions between sources or present uncertain outputs with unwarranted authority. In contexts such as medical, or safeguarding, these failure modes carry serious risk. Ethical practice therefore depends on guardrails: clear limits on what systems are allowed to do, reliable escalation to human support, and the ability to trace outputs back to trusted sources. Without these, scale and automation amplify risk rather than reducing it.

Public trust and ethical practice

Research by the Alan Turing Institute and the Ada Lovelace Institute⁴³ shows strong public support for AI in areas such as cancer detection and climate modelling, where it augments expert judgement. At the same time, there is significant concern about AI used in areas such as welfare assessment, policing and care. Support for AI is consistently strongest where systems support professionals, and weakest where they appear to replace human discretion in high-stakes decisions.

For charities, trust is shaped not only by outcomes, but by how decisions are made, reviewed and challenged. Where AI is used in contexts affecting people's safety, dignity or access to support, expectations of accountability are high, and tolerance for error is low.

Ethical use of AI in charities is therefore not about arriving at a single, fixed answer. It is about understanding how responsibility, organisational practice and external expectations interact. As efficiency increases, exposure often does too. Benefits and risks scale together, particularly when systems are applied across large numbers of people or embedded into frontline decision-making.

The challenge, then, is how charities move from ethical concern to ethical practice. That shift depends less on intent than on the structures that shape how decisions are made, monitored and corrected. This brings us to governance.

AI Governance

The environmental and ethical concerns explored in the previous sections are not reasons for charities to avoid artificial intelligence altogether. They are signals that AI cannot be treated as just another tool.

What distinguishes AI from earlier waves of digital change is not its novelty, but its reach into decision-making. AI systems can shape decisions, prioritise attention, surface information and, in some cases, act autonomously. When things go wrong, it can be unclear who made the decision or where responsibility sits. In high-trust environments, that ambiguity is itself a risk.

For charity leaders, responsibility therefore becomes a question of clarity: what AI is used for, where boundaries sit, how risks are identified early, and how accountability is maintained when decisions are partly automated. Relying on individual judgement, however well intentioned, is no longer enough. Ethics without governance leaves leaders exposed.

Governance, done well, allows ethical intent to be acted on consistently and safely over time.

In this research, regulatory uncertainty was cited as a barrier by 31% of respondents. Alongside skills, it was the only barrier that increased in line with AI use. This is perhaps unsurprising. As charities move from experimentation to more sustained use of AI, informal arrangements that were sufficient to get started begin to feel inadequate, and the need for more formal structures becomes clearer.

For organisations seeking to progress with AI, effective governance must therefore be designed into systems, documented in advance, and reviewed over time. It cannot rely on individuals making the right judgement in that moment.

Responsibility for AI use cannot be delegated to individual employees, outsourced to technology vendors, or deferred to informal assurances that issues will be addressed if and when they arise. As one governance expert put it, "We'll know when we see it, is longer good enough".

While day-to-day oversight may sit with executives and operational teams, ultimate responsibility rests with the trustees.

Trustee duties and AI

UK charity law already provides a strong foundation for governing the use of AI. While artificial intelligence is new, trustees' duties are not. Under Charity Commission guidance, trustees are required to:

- Ensure activities further the charity's purposes for public benefit
- Comply with the law, including data protection requirements
- Act in the charity's best interests
- Manage resources responsibly
- Act with reasonable care and skill
- Ensure appropriate accountability

Each of these duties applies directly to governing AI.

It is for each board to decide how best to discharge these responsibilities in practice. None of this means trustees must understand how AI models work at a technical level. It does, however, mean they should be able to answer basic governance questions: why a tool is being used, what risks it introduces, who is accountable if something goes wrong, and how they would know.

In practice, many charities are still some distance from this position. One early indicator is the absence of basic policies to guide employee use of AI. This research found that only 25% of charities surveyed had an AI policy in place, with a further 29% stating that one was in development and the rest admitting that they had none. Similar patterns appear in the Charity Digital Skills Report. In many cases, AI use is emerging faster than the governance designed to support it.

This gap is most visible at board level. If a charity is running an AI pilot or deploying AI in live services, that activity should appear as a standing item on the board agenda. Yet many boards lack the confidence or skills to engage with these discussions. The Charity Digital Skills Report found that 44% of charities now rate their boards' AI skills as poor, with only 3% rating this as excellent.

Why trustee voices matter

The research behind this book helps explain why trustee capability matters so directly. In examining how charity leaders decide whether to prioritise AI, this work draws on the UTAUT framework, a well-established model for understanding why organisations adopt new technologies and what determines whether interest turns into action. One finding stood out clearly. Social influence emerged as the strongest predictor of whether organisations intended to adopt AI.

In this context, social influence reflects whose judgement leaders rely on when decisions feel uncertain or carry risk. It is shaped by the examples leaders see from peer organisations, the views expressed by their own executives and trustees, and whether AI use feels accepted and defensible within the sector. When these influences align, leaders are more confident moving forward. When they do not, hesitation is more likely.

In charities, trustees are among the most influential voices in that process. Boards play a decisive role in determining whether AI is seen as a responsible and justifiable choice, or as something that is premature or hard to defend. Where trustees have the confidence to ask informed questions, set clear boundaries and express support, AI use is far more likely to move from tentative interest to considered action. Where that confidence is lacking, even well-founded proposals can stall.

Questions leaders should feel comfortable asking

If these questions cannot be answered, governance is not yet sufficient.

- Where are we currently using AI?
- Who is accountable for each system?
- How do we test for harm or bias?
- What would trigger a decision to stop?

Digital Trustees

Lack of expertise is not, in itself, a failing. Yet it is an issue to be addressed. One response adopted by many charities has been the appointment of a digital trustee: someone with sufficient understanding of the digital landscape to support informed discussion, provide challenge, and help boards interpret emerging risks and opportunities. This approach is increasingly common, with 30% of charities having appointed at least one digital trustee⁴⁴, and 12% appointing more than one. The role is not to make decisions on behalf of the board, but to strengthen the board's collective ability to do so.

Regulatory change reinforces the importance of this capability. While most charity leaders are familiar with UK GDPR and the Data Protection Act 2018, newer UK data legislation, including the Data (Use and Access) Act 2025, will shape how organisations can access, share and reuse data when deploying AI systems.

This complexity is heightened for charities operating across borders. As Paul Butcher, founder of Commonsensing AI and formerly Executive Director of Strategic Initiatives at Save the Children, highlighted to me, determining which rules apply is not always straightforward. "If you're in the UK and you're using ChatGPT for a donor in France," he noted, "what are the rules?" As both the technology and the regulatory landscape continue to evolve, questions like this are likely to become more frequent and more complex.

For leaders and trustees, the implication is straightforward. If a charity uses AI in ways that affect beneficiaries, donors or employees, it should be able to explain those decisions in plain language, demonstrate how risks are being managed, and retain the ability to intervene when needed. That is not an additional burden, but the practical expression of trustees' existing responsibilities in an AI-enabled context.

As the pace of change accelerates, both technologically and regulatorily, some boards may find that existing digital expertise is no longer sufficient on its own. Just as the rise of digital trustees reflected a shift in how charities operate, the growing role of AI in decision-making raises a further question: whether boards will increasingly seek more explicit AI-focused governance capability.

Whether this takes the form of expanded digital roles, deeper collective understanding, or the emergence of a more clearly defined AI trustee role remains to be seen. What is clear is that governance structures will need to evolve in step with the responsibilities AI introduces, rather than attempting to catch up after the fact.

The building blocks of AI governance

One reason AI governance feels intimidating is that terminology is often muddled. Leaders are often told they need strategies, policies, frameworks and roadmaps, yet are sometimes not clear on what each is for.

AI STRATEGY

The high-level "why". How AI supports your mission, where you will and will not use it, and what level of risk is acceptable to the organisation.

Audience: Board and senior leadership.

AI POLICY

The ethical and legal guardrails. Sets expectations on fairness, transparency, accountability and data protection obligations.

Audience: All employees and volunteers.

AI ROADMAP

The "how and when". Pilots, training, review points and scaling decisions.

Audience: Leadership and operational teams.

ACCEPTABLE USE POLICY (AUP)

Day-to-day rules. Which tools are approved, what data must never be entered, and when human review is required.

Audience: Everyone using AI.

AI REGISTER (OR INVENTORY)

A living record of where AI is used, what type of tool it is, what data it touches, who owns it, and the level of risk it carries.

Audience: Leadership, governance and assurance roles.

Charities do not need all of these in place on day one. But as AI use grows, clarity at each level becomes essential.

Automation does not absolve accountability

One area that often causes confusion is how AI, and especially AI that automates workflows interact with existing legal duties. A helpful way to think about this is that AI does not create a new category of legal responsibility. It works within existing rules.

A legal expert I spoke to offered a simple example. UK law contains clear protections against discrimination in recruitment and hiring, and those protections apply regardless of how decisions are made.

Imagine a longstanding HR vendor introduces an automated CV shortlisting feature, promising to reduce the time needed to filter candidates. Keen to be seen to adopt AI, a team within the charity begins using it without subjecting it to proper testing or oversight. If the resulting hiring outcome is discriminatory, and the system cannot clearly explain how decisions were reached, responsibility does not shift to the technology or the vendor. It remains with the charity that chose to use it. Discrimination is still discrimination, whether it is carried out by a person or mediated through a machine.

Responsibility in the age of AI

The strongest case for governance does not come from law or regulation. It comes from the people charities exist to serve.

Many reading this will work with children, trauma survivors or people in crisis. In these contexts, errors are not abstract. An automated triage system that discriminates, a chatbot that responds inappropriately, or a data breach exposing sensitive histories can cause real harm. This is why the technology sector's mantra of "move fast and break things" does not translate to a charity context.

Responsible governance does not require perfection. It requires humility. Keeping humans meaningfully involved. Being willing to pause, adapt or stop when risks become clear. Above all, it requires recognising that AI does not reduce responsibility.

Taken together, the environmental, ethical and governance questions explored in this chapter point to a clear conclusion. Responsibility in the AI age is not about

saying no. It is about being clear-eyed about trade-offs, deliberate about choices, and accountable for outcomes.

Charity leaders are not expected to be technologists. They are expected to exercise judgement: to set boundaries, to ask difficult questions, and to ensure that innovation serves mission rather than distorting it, and that trust is strengthened rather than quietly eroded.

AI will continue to advance whether individual charities engage with it or not. The real choice leaders face is not whether to participate, but how. Governance is what turns that choice from something reactive into something intentional.

In a sector built on trust, responsibility is the foundation on which everything else stands.

Conclusion

My journey began with a simple observation. Something has shifted in the charity sector.

Not in the hardware on desks, nor even the software licences held, but in the kinds of questions leaders are now asking. Questions about scale, judgement and responsibility. About whether long-standing ways of working are still sufficient in an environment shaped by rising demand, tighter funding, and increasing expectations of transparency and impact. Artificial intelligence sits inside that wider change.

Part 1: The Promise of AI explored what this moment makes possible. Across fundraising, operations and mission delivery, the examples in this book show charities beginning to use AI to extend capacity, reduce friction, and in some cases reach people who were previously invisible to their systems. In fundraising, AI is reshaping how organisations understand supporters and communicate relevance in a crowded, competitive environment. In operations, often away from public view, it is quietly reducing key-person risk and freeing up time that can be redirected toward mission. In mission delivery, organisations are using AI to improve their support, widen access, and act earlier than previously possible.

What is striking about these examples is not that they are spectacular, but that they are practical. They are not driven by novelty or a desire to appear innovative, but by people who have spent years, often decades, understanding a problem in depth, and who are applying new technologies in service of needs they know intimately. They show that AI, applied thoughtfully, can become part of the everyday work of charities rather than a separate or special activity.

That promise, however, is inseparable from constraint.

Part 2: 'What gets in the way' examined why progress often feels slower or harder than the opportunity suggests. Skills gaps, time pressure, funding uncertainty, data quality and leadership caution all surfaced repeatedly. So did concerns about ethics, trust and mission fit. These are not abstract worries. They reflect the reality that charities operate in high-trust environments, where mistakes carry consequences.

One of the most important findings of this research is not that charities are enthusiastic or sceptical about AI, but that they are weighing it carefully against other responsibilities. Many leaders are already under pressure to deliver more

with less, to respond to public scrutiny, and to steward resources in line with their charitable purposes. AI enters that mix as another variable that must be judged, prioritised and governed.

The sections on responsibility, ethics and governance are therefore central to this book. They are not a brake on innovation, but a condition for it. Environmental impact, bias, safeguarding, transparency and accountability are not side issues to be dealt with later. They shape where AI can be used legitimately, and where it cannot. The work of governance, setting boundaries, documenting decisions, and retaining human oversight, is what turns technological possibility into something that can be sustained over time.

Over the course of this research, I met fundraisers, operations leaders, technologists, policymakers and charity CEOs each grappling with these questions. Some are cautiously experimenting. Some are holding back. Others are already deploying AI at scale. What they shared was not a single position, but a seriousness about the choices before them and a desire to get them right.

I came away convinced of two things. First, AI will continue to shape the context in which charities operate. Second, that the charity sector has particular strengths at this moment. It brings not only deep familiarity with the problems it seeks to solve, but also experience in measuring impact, a culture of ethical reflection, and a long history of making decisions where the stakes are measured in human lives and wellbeing.

My hope is that this book helps clarify that landscape. Not by offering certainty, but by exploring the questions that matter. Not by prescribing action, but by supporting the exercise of judgement that this moment demands.

AI acts as a force multiplier. It amplifies whatever it is placed upon. For charities, this means it will not only strengthen or weaken systems, but also magnify leadership itself. The values and assumptions you bring to AI will shape what it becomes. Used with care, it can support and extend human judgement and capacity. Used without, it can scale harm just as efficiently.

The question for you, as a charity leader, therefore is not simply whether to adopt AI, but what you are choosing to multiply. In that, I wish you the very best of luck.

Appendix

Acknowledgements

This research would not have been possible without the generosity, trust and support of a great many people.

I am deeply grateful to Sean Cosgrove, Tom Neal and Rachael Gordon at Corndel, whose encouragement and flexibility made it possible for me to undertake this work alongside my day job. My thanks also go to Craig and Maeve in the Corndel design team, whose work has greatly strengthened this book.

I am equally indebted to the Churchill Fellowship, which for more than sixty years has enabled people across the United Kingdom to research areas of benefit to national life. It has been a privilege to join that tradition.

I am grateful to the 275 charity leaders across the UK and the United States who took part in this study, as well as to the many advisers, collaborators and supporters who guided, challenged and encouraged me along the way. The insights in this book are grounded in their experience, candour and commitment to the organisations and communities they serve.

To name just a few:

Alison Badget, Courtney Bugler, Paul Butcher, Garlen Capita, Mark Casper, Nathan Chappell, Ha Cole, Bill Cockerill, Alison Court, Paul DelPonte, Sandra Diaz, Jonathan Douglas, Brittany Dunn, Nat Dyer, Charles Eales, Ali Feldhausen, Steven Francisco, Jim Fruchterman, Micah Goodman, Daniel Gray, Dr Mike Hoffman, Yves Jaques, Alice Kershaw, Dan Kershaw, Berenice Levenz, Tim Lockie, Chris Lyons, Gavin Maggs, Nilli Martinez, Shannon McCracken, Sara Mierke, Dr Sven Mikolon, Michelle Mitchell, Cliff O’Gorman, Alisa O’Malley, Nadia Niesterowicz-Newstead, Shawn Olds, Mary O’Shaughnessy, Gemma Peters, Helen Reed, Scott Rosenkrans, Brian Rubenstein, Albert Scerbo, Sofia Schmidt, Michael Sheldrick, Claire Stacey, Amy Sample Ward, Jonathan Townsend, Joshua Tripp, Lauren Tyger, Julie Uva, George Weiner, Jenna Winnocur, Sarah Wyer, and many, many more.

Finally, my greatest thanks go to my wife, Tarhana, and my daughter, Evie, who lived alongside this project through more than seven hundred hours of work, through long evenings of writing, and by my side as we travelled across the United States during the research. Any clarity in these pages owes much to their patience.

About the Author

Jake O’Gorman leads Data and AI Strategy at Corndel, where he works with enterprise organisations to build capability in AI, leadership and emerging technologies. His work focuses on supporting senior leaders as they navigate the opportunities and risks of AI adoption, particularly in complex, regulated environments where people, processes and technology intersect.

In 2025, Jake was awarded a Churchill Fellowship, the UK’s national memorial to Sir Winston Churchill. His research explores responsible AI adoption within UK and US nonprofits. This book shares the findings of that work.

Jake’s connection to the charity sector is longstanding. In 1988, his family founded Children with Cancer UK following the loss of two children in the same year. The charity has since raised over £330 million to support research, treatment and care, shaping Jake’s enduring interest in how organisations can use technology responsibly to increase impact.

Alongside his professional work, Jake is a Freeman of the City of London, and the Worshipful Company of Information Technologists, where he supports AI4C, a learning exchange for charities working with AI and machine learning. He is a fellow of the Royal Statistical Society, and holds a professional certificate in AI Governance as well as a postgraduate certificate in Corporate and Commercial law, focusing on AI technology transfer.

Jake’s career spans AI governance, commercial strategy and hands-on technical work. He has twice represented England at the World Memory Championships and won the Mind Sports Olympiad Memory World Cup. He remains fascinated by how people learn, think and solve complex problems.



Jake and family

Research Methodology

The quantitative component of this research was designed to provide a structured, comparative backdrop to the qualitative interviews, rather than to act as a representative survey of the charity sector as a whole. An online survey was conducted in mid-2025 with senior leaders from 275 nonprofit organisations in the UK and the United States (UK n = 146; US n = 129), typically holding roles with strategic or operational responsibility for technology, data, or organisational performance.

To explore differences in attitudes and experiences, respondents were segmented by stage of AI adoption. Organisations reporting no AI use or only informal experimentation were classified as pre-adoption (n = 197), while those that had made a deliberate leadership decision to adopt AI strategically were classified as post-adoption (n = 78). The two groups were asked partially overlapping, tailored question sets, enabling comparison between current perceptions among non-adopters and recalled perceptions at the point of adoption among adopters, alongside questions on post-adoption outcomes and challenges.

Perceptions of AI were measured using constructs adapted from the Unified Theory of Acceptance and Use of Technology (UTAUT): performance expectancy, effort expectancy, social influence, and facilitating conditions. Responses were captured on five-point Likert scales, with acceptable internal consistency. Relationships between these constructs and behavioural intention were examined using exploratory regression analysis and structural equation modelling (SEM), allowing the coherence of the framework to be tested as an integrated system.

Given the sample size, SEM was used cautiously as a complementary method rather than a standalone basis for inference. Results were interpreted alongside regression and descriptive analyses and were not used in isolation. The modelling aimed to assess the plausibility and internal consistency of the theoretical framework, rather than to produce precise estimates or causal claims.

Overall, the quantitative findings identify patterns and decision dynamics that help explain how nonprofit organisations move from emergent to strategic use of AI. As with all self-reported survey data, findings should be interpreted with appropriate caution, particularly where respondents were asked to recall prior beliefs. The analysis was intended to inform and be interrogated by the qualitative interviews that followed, rather than to replace them.

References

¹ Charity Commission for England and Wales (2025) New data points to growing social and economic impact of the charity sector amidst a challenging financial environment. Press release, 10 July.

² Ruiz, D., Ruben, M. and Benton, J. (2025) Searches up, traffic down: How the AI search cliff impacted 17 nonprofits. M+R Lab, <https://www.mrss.com/lab/searches-up-traffic-down-how-the-ai-search-cliff-impacted-17-nonprofits/>

³ <https://www.mrss.com/lab/searches-up-traffic-down-how-the-ai-search-cliff-impacted-17-nonprofits/>

⁴ <https://www.nonprofitpro.com/article/charity-site-traffic-growth-slows-by-2-5-as-ai-search-rolls-out/> - Charity websites were the second least affected sector in terms of organic traffic, which has grown by 12.5% over the past year, compared to 15% growth the year before.

⁵ https://www.cjr.org/tow_center/we-compared-eight-ai-search-engines-theyre-all-bad-at-citing-news.php

⁶ <https://www.furniturebank.org/postcard/>

⁷ <https://techmatters.org/should-i-be-using-ai-for-this/>

⁸ <https://www.microsoft.com/en/customers/story/24729-british-heart-foundation-power-apps>

⁹ <https://www.microsoft.com/en/customers/story/1772353582649569700-ocgoodwill-azure-ai-services-nonprofit-en-united-states>

¹⁰ Charity Digital Skills Report 2025 (Page 70)

¹¹ <https://www.adalovelaceinstitute.org/wp-content/uploads/2025/03/How-do-people-feel-about-AI-Wave-2-EMBARGOED.pdf>

¹² <https://springact.org/sophia-chatbot/>

¹³ <https://www.sobot.io/article/ai-chatbots-vs-human-agents-customer-service/>

¹⁴ <https://restorenyc.org/wp-content/uploads/2026/01/Restore-Mental-Health-2026-2.pdf>

¹⁵ (p < 0.001)

¹⁶ Amar, Z., & Ramsay, N. (2025). Charity Digital Skills Report 2025. AI leadership skills, p. 75.

¹⁷ Amar, Z., & Ramsay, N. (2025). Charity Digital Skills Report 2025. Artificial intelligence (AI), p. 76.

¹⁸ Amar, Z., & Ramsay, N. (2025). Charity Digital Skills Report 2025. Artificial intelligence (AI), p. 76.

¹⁹ <https://find-employer-schemes.education.gov.uk/interim/growth-and-skills-levy>

²⁰ <https://www.cornel.com/ai-academy>

²¹ Amar, Z., & Ramsay, N. (2025). Charity Digital Skills Report 2025. "Barriers to digital progress" section on Page 55.

²² <https://www.safehouseproject.org/eradicating-human-trafficking/>

²³ Amar, Z., & Ramsay, N. (2025). Charity Digital Skills Report 2025. Key findings, p. 9.

²⁴ Amar, Z., & Ramsay, N. (2025). Charity Digital Skills Report 2025. Funding, p. 16.

²⁵ IBM Institute for Business Value. (2025). The 2025 Chief Data Officer Study: The AI multiplier effect. Page 2

²⁶ Vanguard Charitable. (2017, updated July 29, 2025). Debunking the nonprofit overhead myth "The idea that low nonprofit overhead equals high impact has been a longstanding, persistent myth

²⁷ BBB Wise Giving Alliance, GuideStar, & Charity Navigator. (2013). The Overhead Myth: To the Donors of America "In 2011, the charities which GiveWell reviewed and recommended had higher overhead than the charities they review and didn't recommend, 11.5% versus 10.8%"

²⁸ Caserta, J., Harreis, H., Rowshankish, K., Srinidhi, N., & Tavakoli, A. (2023). The data dividend: Fueling generative AI. McKinsey & Company.

²⁹ Isola, L. (2025). Securing good: How strategic tech support is protecting nonprofit missions. Okta.

³⁰ ANNOUNCEMENT: Microsoft 365 Business Premium & Office 365 E1 grant discontinuation | Microsoft Community Hub

³¹ <https://www.iea.org/reports/electricity-mid-year-update-2025/demand-global-electricity-use-to-grow-strongly-in-2025-and-2026>

³² O'Donnell, J., & Crownhart, C. (2025, May 20). We did the math on AI's energy footprint. MIT Technology Review. <https://www.technologyreview.com/2025/05/20/1116327/ai-energy-usage-climate-footprint-big-tech/>

³³ Li et al. (2025) via The Hidden Footprint of AI

³⁴ Li, P., Yang, J., Islam, M. A., & Ren, S. (2025). Making AI less "thirsty": Uncovering and addressing the secret water footprint of AI models <https://arxiv.org/abs/2304.03271>

³⁵ Li, P., Yang, J., Islam, M. A., & Ren, S. (2025). Making AI less "thirsty": Uncovering and addressing the secret water footprint of AI models <https://arxiv.org/abs/2304.03271>

³⁶ Elsworth, C. et al (2025). Measuring the environmental impact of delivering AI at Google scale. <https://arxiv.org/pdf/2508.15734>

³⁷ There is an important debate about how far "returned" water is immediately usable in practice. In many cases, discharged cooling water is warmer, more saline, or released into systems where it cannot be readily recovered for local use. As a result, even non-consumptive withdrawal can carry environmental and social consequences.

³⁸ A median AI text prompt consumes approximately 0.26 millilitres of water, based on a comprehensive, full-stack measurement of AI inference and serving across Google's production infrastructure. Producing a quarter-pound beef burger typically requires between 1,500 and 3,000 litres of water. Converting litres to millilitres (1,500 litres = 1,500,000 mL) and dividing by per-prompt water use gives approximately 5.77 million AI prompts per burger at the lower estimate ($1,500,000 \div 0.26 \approx 5,769,231$). Using more typical mid-range assumptions, a burger requiring around 2,500 litres of water (2,500,000 mL) and a prompt consuming approximately 0.5 mL would equate to around 5 million prompts ($2,500,000 \div 0.5 = 5,000,000$). Even under more pessimistic assumptions, the comparison remains in the order of millions of prompts per burger. The comparison is not perfectly symmetrical. Burger estimates include water used across agricultural production, feed, processing and supply chains, while the AI figure reflects operational water use associated with inference and serving, rather than model training or hardware manufacture. Making AI less "thirsty": Uncovering and addressing the secret water footprint of AI models <https://arxiv.org/abs/2304.03271>

³⁹ Miller, C. (2025). The Real Environmental Footprint of Generative AI: What 2025 Data Tell Us. Online Learning Consortium.

⁴⁰ A median AI text prompt consumes approximately 0.26 millilitres of water, based on a full-stack measurement of AI inference and serving across Google's production infrastructure. At a usage rate of ten prompts per day, this equates to 2.6 millilitres per day, or approximately 949 millilitres per year ($10 \times 0.26 \text{ mL} \times 365 \text{ days}$). A single cup of coffee is commonly estimated to require around 140 litres of water when agricultural production, processing and supply chains are taken into account. Converting litres to millilitres (140 litres = 140,000 mL) and dividing by daily AI water use gives approximately 53,846 days, or around 148 years, of AI use at ten prompts per day to equal the water footprint of one cup of coffee ($140,000 \div 2.6 \approx 53,846 \text{ days}$).

⁴¹ Measuring the environmental impact of AI inference | Google Cloud Blog

⁴² Centre for Data Ethics and Innovation (2020). Review into bias in algorithmic decision-making.

⁴³ Modhvadia, R., Sippy, T., Field Reid, O., & Margetts, H. (2025). How do people feel about AI?

⁴⁴ Charity-Digital-Skills-Report-2025.pdf

Artificial intelligence is reshaping how charities raise funds, run operations and deliver services. Yet for many leaders, it remains uncertain territory, full of promise, risk and difficult questions.

Drawing on research with 275 charity leaders across the UK and United States, alongside in-depth interviews and case studies from organisations including UNICEF, Save the Children and Oxfam, Jake O’Gorman examines where AI is genuinely strengthening impact and why so many organisations still struggle to use it well.

This is not a book about technology. It is about leadership. It argues that the question is no longer whether charities will use AI, but how deliberately they will do so. Used with judgement, AI can deepen relationships, improve decision-making and extend reach. Used poorly, it can undermine trust and distract from mission.

AI and the Future of Doing Good offers a practical, evidence-based framework to help leaders make informed choices and ensure that technology serves the values that define their organisations.

“ Jake O’Gorman has created a must read, current view, extremely well researched book on the state of AI in the nonprofit world. A comprehensive, concise and critical piece of research every nonprofit leader should read.

CHRIS LYONS, GROUP PRESIDENT, NONPROFIT PRO