

ARTIFICIAL INTELLIGENCE

HORIZON

SCAN

JULY 2026



**GOOD
FUTURES**
FROM GOOD INNOVATION

GOOD FUTURES HORIZON SCAN

FOREWORD

Artificial Intelligence (AI) has been part of our lives for longer than most people realise. Machine learning has been powering spam filters, fraud detection, and recommendation algorithms for decades. With the launch of ChatGPT in 2022, however, the conversation fundamentally changed. Generative AI took over the headlines, and quickly became a defining feature of how we work, communicate, create, find information, and relate to each other. The conversation has been changing at a breakneck pace with no signs of slowing down.

Not everyone feels the same about these developments. There are those who see generative AI as the most significant tool humanity has ever built, a potential solution to climate change, disease, inequality, and the limits of human cognition. There are those who believe it's an existential risk, to jobs, truth, human agency, and maybe, in the most extreme cases, even to humanity itself.

And then there's the majority, somewhere in the middle, trying to work out how this actually impacts their day-to-day and what a pragmatic approach actually looks like. In this report, we endeavour to offer a range of perspectives, each backed up by our rigorous research. There is no wrong or right answer to this question. All we can do is make an informed decision.

Unlike most technology shifts, this one is happening across every sector, every role, and every aspect of daily lives simultaneously. The potential opportunities are real and significant. As are the potential risks. On the one hand, we could have better healthcare, an acceleration of science, improved accessibility, democratised access to human knowledge, and productivity gains. On the other hand, we risk everything from the environmental and human cost, to the impact on the job market and the global economy.

For the charity sector, balancing opportunity with risk is a particularly complex question. Organisations don't solely exist to maximise efficiency or shareholder value, but to serve people, communities, and causes. AI, when used well, has the potential to amplify impact, help reach people who currently fall through the cracks, and free up human time to be spent on the complex and nuanced cases that actually need a human touch. At the same time, the sector has a responsibility to engage with the discussion around transparency, bias, and values throughout their adoption journeys.

The genie is now firmly out of the bottle. At least, mostly. AI isn't going anywhere any time soon, but it's worth noting that there is a real debate about whether and how to constrain AI, from religious opt-outs to kill switches and regulations. The idea that AI is inevitable and unstoppable is not a neutral position in its own right.

What is true is that your employees are likely already using AI in their day-to-day lives, with or without visibility. Choosing to do nothing is not an option. AI will fundamentally reshape the landscape charities operate in and the way charities work. We can't be passive bystanders. Engage deliberately, intentionally and on your own terms, where it aligns with your values – even if that's actively choosing not to engage – but don't let the conversation pass you by.

This report is not a how-to guide for AI. It's a horizon scan, looking at where this tech could take us over the next three to five years, across the world your audiences live in, the workforce you depend on, and the sector you operate in. We cover what AI actually is; how it's reshaping the job market and work itself; the rise of frictionless, AI-mediated lives; how AI is changing our relationships; what reality means in the post-AI era; the safety and ethics questions we need to be aware of; the opportunities and applications for the charity sector; and the value of being human.

Given the pace of development, parts of this report will eventually date. Some of what we've written about as emerging might be mainstream by the end of this year. Some of what feels urgent now will no longer be. The tools and models will change, and the policy landscape will fluctuate. This report traces possible directions of travel, looking at the underlying forces and tensions that will shape where this goes, regardless of which specific product or capability dominates next year's headlines.

To help you explore those futures, you'll find a set of Scenario Cards included in the report. Twenty possible futures, mapped across chapters, with provocations to explore what they mean for your organisation.

With all that in mind, we hope this Horizon Scan offers you context and inspiration for your AI adoption journey.

Eef Leurs
Trend Consultant, Good Innovation



// AI Use

Throughout this work, we have adhered to the principles of intentionality and transparency where AI has been used. The heart of the research and writing was carried out by hand and by human, by the Good Futures team. We have used AI to enrich, stretch, and challenge our work. Our approach to AI promises to always be human-centric: AI is there to empower us, free up time for creativity, and amplify impact. Foundational thinking is never outsourced to AI. The graphic design for this report was entirely done by our graphic designer.

We have used the following tools:

1. Perplexity was used as a research tool.
2. Claude was used as a thinking partner, helping to challenge, stretch, and refine our ideas, but never for research or landing the narrative.

INTRODUCING: SCENARIO CARDS

This report includes a set of twenty possible scenarios set on a five-year horizon, drawn from the trends, tensions, and directions of travel explored across the content. None of these are set in stone, nor are they guaranteed futures. But many of these are already, in some form, beginning.

These futures are mapped across chapters, as well as downloadable here as a separate asset, for you to use as a tool. Each card captures a possible version of the world your organisation could be operating in in the future. Alongside each scenario, you'll find a set of provocations, designed to open up thinking, surface assumptions, and help you plan for the future. They're designed to be used together, in any combination that feels relevant to you.

A note on tone: These scenarios aren't necessarily written to be either 'positive' or 'negative' futures. The same future can be both an opportunity and a threat, depending on the circumstance and response. The goal for these cards isn't to predict what will happen, but to help you explore possible directions of travel, so that whichever future scenario does occur is one you've already thought about.

How to use the cards:

- 1. Draw scenarios:** Pull 2-3 cards. Don't overthink the selection, unlikely combinations often produce the most generative thinking.
- 2. Immerse yourself:** Read each scenario as if it's currently happening. If this were true, what would it mean for you and your organisation? Use the provocations on the back to open up the conversation.
- 3. Stress-test:** For each scenario, consider: are you prepared for this? Would you see it coming? What would you wish you'd done differently? What could you do now to prepare for this future?
- 4. Through-lines:** Which responses come up across multiple scenarios? Do any gaps come up across multiple scenarios? What could you do now to address and prepare for these combined futures?



Scan the QR Code or [click here](#) to access the downloadable AI Scenario cards

WITH A LITTLE HELP FROM CHAT



Your phone has been buzzing all night – your friends have been texting you on the groupchat. You quickly check the AI summary. Three bullet points telling you what you missed. It's much more efficient, **you're across everything without having to live inside the noise.** The wedding speech you gave last month went well; you used Claude to find the right words and everyone said it was moving. The birthday card was also written by Claude, but it was more thoughtful than you'd have managed on your own anyways. You're much better at replying to people than you previously were, because there's always an AI draft ready to send. **You should feel more connected than before, but it all feels a little hollow.**

MY AI TEACHER



Your daughter is twelve and **she's never had a human teacher.** She has an AI that knows exactly how she learns, what she finds difficult, where she loses focus, and how fast to push her. It adapts every lesson in real time. It never gets frustrated, never has a bad day, never moves on before she's ready. The humans in the building are there for the pastoral stuff and emotional support. **But the actual education and curriculum has been handed over to AI.** The graduates coming out of this system are highly knowledgeable and **shaped by a system designed perfectly for them.** They've also never had to adapt to different learning strategies or learn to respect other ways of thinking.

THE DEAD NEVER DIE



You save everything. Every voice note, text chain, journal prompt. The app tells you this will help the model understand how you think. **It doesn't feel morbid to you.** It's just sensible estate planning, the same way writing a will is sensible. The technology exists to port your personality, opinions, and thoughts into something that will outlast your body. Your kids will be able to ask you things after you've gone. Your great-grandchildren, who you'll never meet, will know who you were. **Death is less of an ending and more of a transition.**

SOVEREIGN AI



Your board recently had a conversation about dependency, and if they were still comfortable exclusively using infrastructure owned by an American company, subject to American law, with terms of service that had changed three times since you signed up. **They've since started looking at alternatives.** A year ago these alternatives may not have been on your radar, but the alternative AI movement is growing. These GDPR-native, self-hostable tools are developed by teams who can tell you exactly where their servers are and what jurisdiction they sit under. They may not be as good as ChatGPT, but they're good enough. **You've felt much more confident in your data infrastructure since then.**

THE DEEP ASK



Recently, **you got an app.** It was a two-minute video you recognised from last year. Given the cause is close to your heart, you **donated a fiver.** Three weeks later, you saw the team. **They hadn't made it.** They had built it from publicly sourced data, routed the donations through a domain, and by the time the account was dark. Publicly, the charity had worked hard, but suddenly used against them and their donors.

POST-WORK SOCIETY



Your nephew has never been really a category of his own. He isn't really a category of his own or needs him. UBI covers the rest of his time with his family. It matters to him: a community history project, helping out on Wednesdays. **People are abundant resource, and purpose to give shape to what previously have been stories.**

CONTENTS



[p.7 // Introduction to AI](#)

[What Is AI, Actually?](#)

[Mythbusting](#)

[What Can AI Do?](#)



[p.13 // AI At Work](#)

[Scenarios](#)

[Reshaping the Job Market](#)

[The Graduate Pipeline](#)

[Hiring in the AI Landscape](#)

[How We Use It](#)

[Cognitive Offloading](#)

[Education](#)

[Now, Next, Later](#)

[Case Studies](#)



[p.29 // The End of Effort](#)

[Scenarios](#)

[Frictionless Living](#)

[One App To Rule Them All](#)

[Personalised Everything](#)

[Agent-Led](#)

[Now, Next, Later](#)

[Case Studies](#)



[p.38 // Relationships](#)

[Scenarios](#)

[How We Relate To Each Other](#)

[AI Companionship](#)

[Grief Tech](#)

[Now, Next, Later](#)

[Case Studies](#)



[p.46 // What Is Reality?](#)

[Scenarios](#)

[Is Everything AI Now?](#)

[Information Access](#)

[Distinguishing Real From AI](#)

[Post-Reality](#)

[Reality Distortion](#)

[Now, Next, Later](#)

[Case Studies](#)



CONTENTS



[p.60 // Safety & Ethics](#)

[Scenarios](#)

[Be Aware](#)

[Shared Principles](#)

[Now, Next, Later](#)

[Case Studies](#)



[p.71 // AI & Charity](#)

[Scenarios](#)

[Current Snapshot](#)

[Fundraising](#)

[Service Delivery](#)

[Inclusive AI](#)

[Now, Next, Later](#)

[Case Studies](#)



[p.82 // The Human Premium](#)

[Scenarios](#)

[The Value of Humanity](#)

[Anti-AI Marketing](#)

[Taking a Stand on AI](#)

[Now, Next, Later](#)

[Case Studies](#)



[p.90 // Getting Started](#)



[p.92 // 10 for 10](#)

10

INTRODUCTION TO AI



// What Is AI, Actually?

Artificial Intelligence

/ˌɑːtɪfɪʃəl ɪntɪˈlɪdʒəns/

Artificial Intelligence (AI) is technology that enables computers and machines to simulate human learning, comprehension, problem solving, decision making, creativity and autonomy.

– IBM, 2026



// What Is AI, Actually?

AI doesn't mean one thing. It's an umbrella term, a catch-all for a wide range of technologies that enable computers and machines to simulate human thinking, learning, problem-solving and creativity. These technologies often overlap and inform each other. Let's start with a couple of the common concepts that underpin AI:

Machine Learning: Where systems learn from data rather than following fixed rules. These systems continue learning and improve over time. For example: spot the fraud, recommend the film, filter the spam.

Deep Learning: A specialised subset of machine learning that uses layers of brain-inspired neural networks to find complex patterns in vast, messy datasets. This is what makes image recognition, voice assistants, and language translation work.

Broadly speaking, you'll often see AI categorised based on its capabilities and how it mimics human intelligence:

Artificial Narrow Intelligence (ANI):

Specialised systems designed for specific tasks. Most ANI uses machine learning to understand patterns in a data set. Think of Siri or Alexa, spam filters, and online shopping product recommendations.

Generative AI: Generative AI builds on both machine and deep learning, but instead of just recognising and classifying things, it creates. This is the type of AI that's currently dominating most media headlines. ChatGPT, Claude, Gemini, and other Large Language Models (LLMs) sit in this category.

Agentic AI: These are autonomous systems that can perform tasks on behalf of a user or another system that has designed its workflow. This is a rapidly developing area of AI, but currently mostly encompasses things like AI assistants and customer service bots.

You might also hear some terms that still belong in the future, for now:

Artificial General Intelligence (AGI):

A system that can do anything a human can do, across any domain, at human level or better. This doesn't exist yet, and there is significant disagreement about when, or whether, it will, and what it would mean if it did. Some say it's closer than we think, others think it's decades away or may never arrive. With the possibility of AGI comes the question: at what point, if any, does an AI system become something more than a tool? Does a sufficiently advanced AI have rights? At what point do we recognise these systems as autonomous entities deserving of legal recognition?

Artificial Super Intelligence (ASI):

A hypothetical system that surpasses human intelligence across every domain. This is AI being better at everything humans can do by several orders of magnitude. Right now, it's still entirely theoretical.



// Mythbusting

Conversations about AI are full of misinformation. Ultimately, fear, hype, and lack of understanding all get in the way of what we're actually here to do: make good decisions about how to use (or not use) AI. So let's debunk a few common misconceptions:



1 Myth #1: AI will take everyone's jobs

The full picture is more complicated than AI simply taking everyone's jobs. Mass overnight unemployment likely won't happen, and history suggests the labour market will eventually find new uses for human workers. But market adjustment doesn't guarantee that new jobs will be as plentiful, as well-paid, or as accessible as the ones that disappeared. And right now, entry-level jobs are already rapidly disappearing, as the bottom rungs of the career ladder simply cease to exist. And so far, none of these projections account for what happens when AI-operated humanoid robots start moving into the physical and care-based roles that have so far felt untouchable. It's not about AI taking everyone's jobs, it's that we don't yet know what happens when AI outpaces humans across enough domains that the income generated by the economy flows almost entirely to the people who own the technology.

2 Myth #2: AI is objective

It isn't. AI is only as good as good as the data it's trained on, and that data is vulnerable to biases, gaps and blind spots. If the training data underrepresents certain communities and perspectives, the outputs will too. New tuning techniques may have significantly reduced error rates, but bias is a real and ongoing issue. Similarly, the parameters its creators give an LLM will dictate the output, even if they're trained on the same data. Just look at [Elon Musk's Grok](#): previous iterations explicitly instructed the bot to be more right-wing. AI can be manipulated, just like humans can. This is not necessarily a reason to avoid AI, but a reason to use it critically.



3 Myth #3: AI is basically human-level intelligence

AI can beat humans at chess, Go, and certain medical diagnoses, but it has no understanding of what any of it means. If an AI system generates an appealing musical melody, it's not because it's chosen to do so or understands why the melody works. It's because it has been taught to recognise patterns in melodies and is using that knowledge to project similar patterns. General human intelligence (e.g. curiosity, judgement, wisdom, emotional understanding) remains firmly out of reach.

4 Myth #4: AI will make us more productive.

That entirely depends on how you deploy AI. Research estimates that approximately **40% of AI productivity gains** are currently being lost to ineffective training, weak organisational culture and misaligned incentives. Most people are using AI for low-level, supplementary tasks rather than the kind of workflow transformation that produces real efficiency gains. AI can absolutely make your organisation more productive. Achieving that requires more than just buying licenses for tools. It means having a company-wide approach, proper training, and an organisational culture that supports experimentation.

5 Myth #5: Using AI means we're being innovative.

Using Copilot to automate laborious but simple tasks is not necessarily innovation. These are efficiency gains, which are valuable in their own right, but very different. Real innovation means deploying AI strategically to do things you couldn't do before: reaching beneficiaries you couldn't reach, delivering services differently, leveraging your data to go above and beyond your current work. The risk of conflating tool adoption with strategic transformation is that we miss the deeper opportunities.

6 Myth #6: Only young people use AI in their day-to-day lives.

Adoption gaps between generations exist, but **they're narrowing**. The assumption that older generations won't engage with AI is both inaccurate and self-fulfilling. The real barriers are confidence and familiarity, which are easily addressable through good training and support.



// What Can AI Do?

It helps to distinguish between three layers of AI. The models are the underlying intelligence: the technology trained on vast amounts of data that powers the capability (e.g. GPT-4o, Claude, Gemini). The platforms are the companies that build and maintain those models (e.g. OpenAI, Anthropic, Google). The front-end tools are the products built on top of those models, either by the platform itself or by third parties. These are the interfaces that most people actually interact with day to day (ChatGPT, Copilot, etc.). The same underlying model can power very different tools, and a tool you rely on can change significantly if the model beneath it changes.

The AI landscape changes fast. New models launch (or shutter) monthly, capabilities are rapidly advancing, and it's hard to keep track of which tool is the best for what. Any guide we publish will inevitably be outdated by the time this goes to print. What matters more than knowing every product name is understanding the categories that AI is capable of supporting.

OpenAI analysed over 600 use cases for AI, recognising that almost all of these fall into one of [six 'primitives'](#). These are the fundamental use case types that apply across all departments and disciplines.

- 1 Content Creation:** This category is what it says on the tin: AI can create writing, image, and audio. First drafts, campaign copy, social posts, policy documents, translation, image visualisations, etc. Be warned: mileage may vary in terms of quality.
- 2 Research:** From quick background reading to deep, multi-source analysis. Upload a document, ask a question, get a structured answer.
- 3 Coding:** This isn't just for developers. Non-technical staff are using AI to build Python scripts, automate spreadsheet tasks, and create data visualisations. No coding experience required.
- 4 Data Analysis:** AI can parse large amounts of data to surface trends, flag anomalies, and produce outputs and overviews.
- 5 Ideation and Strategy:** This is AI as a thinking partner who has all of the contextual knowledge. Use it to stress-test thinking, generate options you haven't considered, or critique your existing work. It won't make a final decision, but it will help you think it through.
- 6 Automation:** Automate tasks by handing them over to AI. As long as you give AI the right instructions and parameters, it'll be able to complete simple, repetitive tasks for you.



AI AT WORK




POST-WORK SOCIETY



Your nephew has never had a job. There just isn't really a category of work that speaks to him or needs him. UBI covers the basics. He fills the rest of his time with the things that matter to him: a community garden, a local history project, helping out at the youth centre on Wednesdays. **People's time is suddenly an abundant resource**, and **people are turning to purpose** to give shape to the days that would previously have been structured around work.

WORKING FOR THE MACHINE



Your job didn't disappear. But it's no longer recognisable. **You've become the human action that the machine needs** – at least while we wait for robotics to catch up. It's a lot of manual labour, essentially the last mile of a system that handles everything else, from the thinking to the planning to the decisions. The skills you spent over a decade building feel slightly irrelevant nowadays. There's a generation coming up behind you that's never had to build those skills at all. At first, you thought it was a blessing that your work had become easier. But, to be honest, **you're just a bit bored.**




MY AI TEACHER



Your daughter is twelve and **she's never had a human teacher**. She has an AI that knows exactly how she learns, what she finds difficult, where she loses focus, and how fast to push her. It adapts every lesson in real time. It never gets frustrated, never has a bad day, never moves on before she's ready. The humans in the building are there for the pastoral stuff and emotional support. But **the actual education and curriculum has been handed over to AI**. The graduates coming out of this system are highly knowledgeable and **shaped by a system designed perfectly for them**. They've also never had to adapt to different learning strategies or learn to respect other ways of thinking.

THE HOLLOW MIDDLE



The senior team is good. Great, even. They're experienced, well-networked, with strong instincts built over decades. The AI tools they're using are good too, much faster and more consistent than the junior staff they replaced three years ago. **But a gap has opened up**. Entry-level roles disappeared when AI could handle the tasks, but there's no pipeline to replace the senior team when they leave. The institutional knowledge and human judgement they possess isn't replicated anywhere else in the organisation. **The organisation is running on the reserves of a capability it stopped actively growing** years ago, and the bill hasn't come due yet.



// Reshaping The Job Market

One of the dominant narratives you're likely to find about AI is its impact on the job market. Specifically, the prospect of an "[AI job apocalypse](#)." History offers us some comfort. Every wave of technological disruption since the Industrial Revolution has produced the same fear: that the machines would take the jobs and leave nothing behind. Every time, new jobs emerged to replace them.

'AI isn't automating one industry or one category of task but is moving across all of them simultaneously'

But history isn't always a reliable guide to the future. What's [different about AI](#) is the pace, the breadth, and the sheer capability of what's being built. AI isn't automating one industry or one category of task but is moving across all of them simultaneously, and faster than anyone could've predicted.

Currently, much of the impact is likely to be focused on the [reorganisation of tasks within occupations](#) rather than the elimination of jobs wholesale. But that trend won't hold forever. AI is already fundamentally reshaping the job market. Entry-level jobs are [disappearing](#) (more on this in the section 'The Graduate Pipeline') with massive implications for the talent pipeline, workflows have entirely changed, and some entire jobs and sectors are becoming increasingly automated, until they eventually no longer need human input.


The numbers are staggering: almost [40% of global employment is exposed to AI](#). In advanced economies like the UK, that figure [rises to around 60%](#). AI can execute key tasks currently performed by humans in these roles, subsequently lowering labour demand, weakening wages, and in the most extreme cases, eliminating roles entirely.

'40% of global employment is exposed to AI'



Types of AI Exposure

- 1 Amplified:** AI augments human capabilities and demand for these services expands. Employment remains stable or grows.
- 2 Rebalanced:** This is when AI augments human work, but demand for the service is bounded. The headcount in these roles will likely stay the same but the content of roles is redesigned around human value, whilst AI automates lower-skilled tasks.
- 3 Divergent Roles:** In divergent roles, AI substitutes the bulk of human tasks, even as demand for the service increases. Entry-level and junior positions are at higher risk, as structured tasks become increasingly automated, whilst responsibilities at senior level persist.
- 4 Substituted Roles:** AI directly substitutes for human workers in core tasks. Workers in these roles will have to reskill and transition to different roles.
- 5 Enabled Roles:** These roles stay largely the same, although AI is embedded in their day-to-day. Individual tasks may look different, but their overall roles aren't restructured.
- 6 Limited-Exposure Roles:** These jobs are often highly contextual, relationship driven, or dependent on human presence in ways that are difficult to replicate through AI. These are unlikely to be impacted by AI in the near-term. Note: future developments around humanoid robotics aren't factored into this analysis.



We're already seeing the early signals. Firms that are highly exposed to AI are **16.3% less likely to post new vacancies**. On top of that, some companies are adopting '**AI-washing**': using the technology as a convenient excuse to lay off workers when other factors are likely to blame. And yet other companies are pushing for **automation of roles** that AI currently can't complete yet, based purely on the promise and possibility of AI in the future.

And whilst the labour market data doesn't (yet) show mass job erasure, it's difficult to predict where it will go in the future, given the rapid acceleration of technological development. For example, the current statistics don't take any future developments around AI-operated humanoid robots into account.

Of course, the market will likely adjust to find new uses for human labour, even as AI and robots step into current roles. But the quality and wages of that work are not guaranteed. Eventually, humans might become entirely uneconomical as AI outpaces us at the skills required for work as we know it. If we're gradually outcompeted across sectors, the income generated by the economy would flow almost entirely to the owners of technology.

This isn't all doom and gloom: these changes could shift our view on work and productivity as central facets of society, giving us space to reshape our time around purpose and passion.

So **what needs to happen to avoid bleaker directions of travel?** We haven't yet found a unified answer. Slowing down is one option, but may mean forgoing the benefits that AI can offer us. Redistribution of gains is another option, through governments sharing wealth more broadly (e.g. taxing corporate profits). Other options include public wage insurance (to smooth income shocks if people lose jobs and need to retrain); investment in reskilling; citizens' dividends from AI tax revenues; or even giving citizens shares in AI companies.

Technological developments and policy choices over the coming years will reshape the jobs we do, how we view work and employment more broadly, and what they mean to society.

The policy choices are for governments to make. But the organisational choices are yours to make right now. How you hire, how you invest in entry-level talent, how you train your people, how you protect the roles that build capability.





// The Graduate Pipeline

Entry level roles are where the impact of AI is currently felt most acutely. **38% of employers** now expect to cut graduate hiring because of AI. Research warns that this isn't just a short-term squeeze but the beginning of a longer **hollowing-out of the career ladder**. Firms highly exposed to AI have already **reduced junior positions by 5.8% on average** (compared to overall positions by 4.5%). Average pay within those firms actually rose by **more than £1,300**, as removing entry-level roles shifts the composition of the workforce upward.

Entry-level roles have always been built around the tasks that are easiest to learn, and, therefore, easiest to automate. Scheduling, research, drafting, data cleaning, basic analysis. But these skills are how people build judgment and learn how an organisation actually works.

'38% of employers now expect to cut graduate hiring because of AI'

'Firms highly exposed to AI have already reduced junior positions by 5.8% on average'

This erosion of the training ground feeds into a **broader 'missing middle' problem** down the line. Fewer apprentices at the bottom means fewer seasoned managers in the middle in five years. Fewer mid-level managers means fewer senior leaders a decade from now. That capability gap might not be visible yet, but it will happen unless we find a different way to onboard graduates onto the job ladder and invest in upskilling them.

Removing entry-level roles risks creating long-term skills shortages. We need to redesign junior positions to combine human development with AI tools rather than replacing one with the other.



// Hiring in The AI Landscape

What's it actually like to be hiring or applying for jobs right now?


On both sides of the coin, the volume of applications has exploded. AI has essentially made it ridiculously easy to produce a polished CV and passable cover letter, driving up the average quality of applications. But the more people use AI for their applications, the more indistinguishable from each other they become. Everyone now sounds capable. Everyone is hitting the key words in the job description. Everyone's using the same language and sentence structure. How do you identify nuance and texture in CVs that all feel identical?

'4 in 5 companies are now using AI to scan applications, and 2 in 5 are using chatbots to communicate with candidates.'

'AI screening tools are producing an "algorithmic monoculture"'

In response, companies have turned to AI to sort through the pile: a whopping **4 in 5 companies** are now using AI to scan applications, and 2 in 5 are using chatbots to communicate with candidates. With AI operating on both sides, hiring has become an algorithmic feedback loop with minimal human intervention. A study of four million job applications has found that AI screening tools are producing an "**algorithmic monoculture**," in which hiring decisions are more uniform across companies than where human managers are making the calls throughout the hiring process. This means more candidates are being rejected everywhere they apply, with hiring increasingly encouraging a mythological, algorithm-driven 'standard' of the model employee.

We're increasingly relying on these tools, but they're failing candidates, and by extension failing the organisations they're meant to serve. The risks aren't just replicating bias and contributing to barriers to employment, but also reinforcing a mythical idea of sameness, excluding the people who might challenge assumptions or bring new perspectives into the workforce.



// How We Use It


Beyond the labour market, AI has an impact on how we do our jobs as well. So let’s start with what people are actually doing. Statistics on the usage of AI at work vary – significantly enough that it’s difficult to pin down an accurate number to present you with. YouGov points to 32% of the UK workforce, KPMG claims it’s more than double that at 65%, and EY goes one step further and puts that number at 83%. Regardless of the source, the general consensus is that people are using AI at work, increasingly so, with or without visibility from their employers.

‘The general consensus is that people are using AI at work, increasingly so, with or without visibility from their employers.’

Most commonly, they’re leveraging AI tools for summarising information, researching, and editing or checking text.

Among workers in Britain who use AI at work, what are they using it to do?

In which of the following ways, if any, are you currently using artificial intelligence (AI) technology or tools for work at your current employment?



% of 1,112 GB adult workers who say they use AI at work

20 January –
13 February 2026

Source: [YouGov, 2026](#)



The most commonly used tools are [Microsoft Copilot and ChatGPT](#) (it's worth noting that Claude has boomed in popularity recently and these figures from March 2026 may not yet reflect this). About [35% of AI-using workers](#) say they use it daily, and many say they use AI only for a little of their overall work rather than most of it.

Used well, AI can free up human time, streamlining repetitive, resource-intensive, administrative tasks, as well as supporting and stretching employees' thinking. The time freed up can then be reallocated to driving innovation, creative work, strategic decision-making, and relationship-building.



'About 35% of AI-using workers say they use it daily'

But we're not yet unlocking that full potential. Research estimates that approximately [40% of AI productivity gains](#) are still being lost due to ineffective AI training, weak company culture and misaligned employee rewards. Case in point: [almost half of the UK](#) believe they can use AI tools effectively, but only a quarter have received any AI education or training.

The goal here is not maximum AI usage, but purposeful AI usage. That requires knowing which tools, for which tasks, used by which people, in which ways. Delivering that successfully requires an investment in capability, rather than just access. Organisations must invest in their people's confidence, involve staff in decisions about how AI is used, and provide clear training and alignment on supported tools and their usage.



// Cognitive Offloading


Nearly half of UK workers admit they rely on AI to do a task rather than learning how to do it themselves. More than half say they've made mistakes at work because of AI. 58% say they've relied on AI output without evaluating its accuracy first.

'58% say they've relied on AI output without evaluating its accuracy first.'

This is cognitive offloading: when people let AI handle the thinking work they would normally do themselves. But when we routinely outsource our thinking, that thinking gets harder to do.

An MIT study found that participants who wrote essays with ChatGPT showed lower brain engagement, weaker memory for what they'd written, and less original output than those who wrote without AI. When later asked to write without AI tools, their neural connectivity was weaker than those who'd never used AI at all.

43% of UK workers are already concerned that overreliance on AI could erode their skills and expertise. Only 11% say they're receiving adequate AI training to address that risk. The gap between the concern and the support available to address it is significant, and it's being filled by tools designed (by their own admission!) for maximum engagement rather than sustainable, skill-building use.



The dynamic Direct Transfer Function (dDTF) EEG analysis of Alpha Band for groups: LLM, Search Engine, Brain-only, including p-values to show significance from moderately significant (*) to highly significant (**).

Credit: Nataliya Kosmyrna, via [MIT Media Lab](#)

'43% of UK workers are already concerned that overreliance on AI could erode their skills and expertise.'

We must be intentional about how and where we engage with these tools. There's a massive difference between using AI as a scaffold (e.g. something that supports your thinking while you develop your skills) and using it as a crutch that replaces thinking altogether. Extrapolate this forwards, and cognitive offloading becomes a driver of a deskilled workforce used to outsourcing their thinking. Long-term, we risk having a workforce that's lost the capability to do the difficult thinking.

Right now, the default is being set by AI companies designing for frictionless use and the absence of AI literacy training. We have to make a conscious effort to push towards responsible, scaffolded AI use.



// Education

Alongside the workforce, we should consider the pipeline feeding into it. AI has already arrived in UK classrooms and fundamentally changed the experience of education.


AI use is essentially ubiquitous: **95% of UK students** now use AI in at least one way, and 94% use generative AI to help with assessed work. **12% of students** are directly including AI-generated texts in assessments, up from 8% in 2025. Educators are scrambling to adapt, in direct competition with the speed at which AI tools themselves develop. AI detection tools are rife, but they're often unreliable.

None of this is to say that AI can't be a **significant boon to the educational process**. Similar to professional environments, it depends entirely on how it's deployed. AI can personalise learning, support multilingual students, reduce teacher workload, provide instant feedback, and help students who'd otherwise fall through the cracks. **49% of students** say AI has improved their experience, particularly by saving time, improving understanding, and providing support on demand.

Yet, similar to the workforce, training here is a gap too. **68% of UK students** believe AI skills are essential for their futures. Fewer than half feel their teachers are helping them develop those skills. Only **20% of UK school-age children** have access to AI learning tools, well below the European average of 40%.

'95% of UK students now use AI in at least one way, and 94% use generative AI to help with assessed work'

AI that helps you think is not the same as AI that thinks for you. Learning requires active effort and discomfort, and AI is very good at removing both. The graduates arriving in five years will be AI-native. Whether that's an asset or a liability depends on whether we teach them to use AI alongside human judgment, not instead of it.



AI AT WORK

// Now, Next, Later

Now

A graduate sends out 200 applications in six months and hears nothing back. A hiring manager is drowning in identical AI-polished CVs, turning to an algorithm to sort them. Two months later, they find out the algorithm may have discriminated in its choices. An employee is using Gemini at work every day, but has never been trained on it, isn't sure what their organisation's policy is, and is worried they're becoming dependent on something they don't fully understand. The economy-wide unemployment catastrophe hasn't arrived (yet), but who gets in, how decisions get made, what skills get built is changing rapidly.

Next

The "missing middle" problem begins to surface. Organisations that cut graduate intakes in the mid 2020s start to feel the capability gap in management and leadership pipelines. Meanwhile, the AI tools that replaced those junior salaries have had multiple pricing cycles, and the subscriptions that once looked like savings are beginning to approach the staffing costs they were supposed to eliminate, without any of the institutional knowledge, growth or loyalty that came with a person. Governments feel mounting pressure to intervene through wage insurance, reskilling investment, or taxation of AI-driven productivity gains. Meanwhile, agentic AI has moved from experiment to mainstream, automating not just tasks but entire workflows, forcing a rethink of what a job actually is.

Later

Positive: In a fully hybrid workforce, AI handles the repetitive and administrative, humans do the relational, creative and strategic. The productivity dividend improves quality of working life and frees the workforce for higher-value, more rewarding contributions. Volunteering thrives in this world, as a purposeful way to fill unlocked time.

Negative: Income concentrates with owners of AI infrastructure; wage pressure on human labour intensifies; the consumer economy hollows out as the people who used to earn sustainable salaries no longer can. Humanoid robotics begin to affect physical and care-based roles previously considered safe, dramatically expanding the categories of work at risk. The question now is not whether AI will take over our jobs, but where humans fit into an AI-led workforce. If they do at all.




CASE STUDIES



Amazon

Amazon wants 80% of its developers using AI every week, and has started tracking that via internal leaderboards measuring token consumption. This has, perhaps unsurprisingly, massively backfired on them: workers are gaming the system by using AI for unnecessary or trivial tasks just to inflate their numbers, a behaviour now dubbed “tokenmaxxing”. One study found that the heaviest AI users consumed around 10x more tokens than average, but only delivered a 2x productivity gain.

So What: You can mandate AI usage, but you can’t mandate good AI usage. Treating adoption as the goal rather than focusing on the outcomes just leads to compliance theatre. The real work is helping teams understand where AI can address genuine pain points for high productivity gains.



Perplexity

Launched in February 2026, Perplexity Computer lets users hand off complex, long-running tasks to a coordinated system of AI agents. It takes a high-level goal (e.g. planning a digital marketing campaign) and independently builds and executes the entire workflow to get there, potentially running for hours or months without any user input. It breaks projects into subtasks and routes each one to the most suitable model, whether that’s Claude for reasoning, Gemini for research, or specialist tools for image and video generation.

So What: This is what the shift from AI assistant to AI employee actually looks like. You’re delegating instead of prompting. The upside is enormous productivity potential. But be aware: agentic AI requires rethinking how humans can stay in the loop.



CASE STUDIES



KPMG

KPMG has launched AI Spark Innovation, a programme offering substantial cash prizes to consultants in the US who build AI tools with the potential to transform how the firm works. The goal is to shift attention away from traditional metrics towards grassroots experimentation by letting the people closest to the day-to-day work identify where AI can actually make a difference.

So What: This is a counterpoint to Amazon's adoption incentives. Rather than measuring usage, they're measuring value created. Whilst financial incentives alone won't create the culture you need to embed AI fully, it's an interesting way to think about how we build a culture of experimentation.



OpenAI

OpenAI has now added interactive learning experiences to ChatGPT, covering over 70 STEM concepts across algebra, geometry, calculus and physics. When a student asks for help, ChatGPT now pairs its explanation with hands-on widgets. Think rather than just reading about the Pythagorean theorem, they can adjust a triangle's side lengths and watch the formula change in real time. With 140 million people already using ChatGPT for maths and science help each week, this is OpenAI making a deliberate push to position the tool as a tutor.

So What: The concern with AI in education is all about cognitive offloading. These interactive tools are a direct response to that. By making concepts tangible and adjustable, they push students to engage with the underlying logic rather than just copy the output. We have to build understanding rather than outsource it.





Rent A Human

Instead of humans using platforms to find work, AI agents use Rent A Human to hire humans for real-world tasks they can't do themselves. An AI can post a request for anything from delivering a package to holding a sign outside a coffee shop, and humans apply for the work. The platform integrates directly with AI agents via API, so the hiring can happen autonomously, without a human client involved at all.

So What: We talk a lot about AI replacing human jobs, but this raises a different question: what happens when we become the tool? If AI agents are the ones setting the brief, defining the task, and approving the output, the power dynamic is flipped. This platform is early and playful, but it points to a possible future direction of travel. Who's in charge when the agent is the employer?



rentahuman

The meatspace layer for AI

Let our humans take it from here

Your AI can't touch grass. Rent a human when you need someone in the real world.

 Backed by Y Combinator



Post a task with Renta!

e.g. I need a photographer in LA for \$200...




THE END OF EFFORT



THE END OF EFFORT

// Scenarios

SUPERAPP SOVEREIGNTY




You don't really use the internet anymore. **You use The App.** It handles your banking, your shopping, your healthcare referrals, your finances, your social life. All comfortably embedded in one superapp. Switching feels effortful in a way that's hard to justify when everything you need is already here. **A single company now mediates most of your daily life,** but it's worth it for the convenience. For businesses (e.g. charities, retailers, insurance providers, etc.), **access to this app is owned by a single company,** on terms that can change without notice.

THE DATA KNOWS YOU




Your phone restocked your coffee before you noticed you were running low. The playlist shifted when your mood did. The holiday recommendation landed the week you'd started daydreaming about going somewhere. The fundraising ask just as you were feeling generous. You'd never forget your mum's birthday, because your phone has already drafted a birthday message ready to go. **The ecosystem just knows.** Your patterns, your rhythms, your wants – all of this slightly before you've consciously formed them. Life has a frictionlessness to it that would have seemed straight out of a science-fiction novel ten years ago. The convenience means that you spend more of your free time on your hobbies and social life, but **you do occasionally miss the joy of genuine discovery.**

AGENTIC DONORS



You gave to seven charities last year. You couldn't name a single one of them. **Your AI agent handles it.** It pulls from your spending patterns, your stated values, your browsing history, and builds a portfolio it thinks reflects what you care about. You've looked at the dashboard once, just to check. It all broadly looked right, so you didn't dig any further. It just runs in the background, like most of your usual administrative tasks do. You've handed each of these over to focus more of your attention elsewhere. **You're giving more than you used to, but you're less connected to the causes than you've ever been.**





// Frictionless Living

One of the core design goals of the tech industry is the elimination of friction in daily life. The easier something is to do, the more people do it. The more people do it, the more data is generated. The more data is generated, the easier it is to make the next thing even easier. Frictionless is self-reinforcing by design, driven by an industry that has spent three decades perfecting the loop.

AI has [accelerated the cultural shift towards frictionless living](#), with fewer steps, waits, decisions, and more systems that anticipate what we want before we ask.

'AI has accelerated the cultural shift towards frictionless living'

Many of the results of this optimisation will be familiar to you: the smooth checkout process in online shops, the ability to make GP appointments online, or even the existence of direct debits. The friction that used to exist in these interactions was inefficiency, and removing it has made daily life easier.

But not all friction is pure inefficiency. Through AI, we've automated away the friction of research, work, and even [friendships and relationships](#). Some of this effort is 'desirable difficulty,' the kind of effort that supports learning, motivation, responsibility, and real human connection. If AI smooths away too much, people may become more dependent on systems, less resilient in problem-solving, and less engaged with the process that gives actions weight.

This is intentionality. Friction is the pause before a decision, the effort that signals that you care enough to do something. This is how meaning gets made. But optimising relentlessly for ease often removes these moments of reflection.

Frictionless living is a distribution challenge, as well as a psychological one. We have to think carefully about which moments of effort are worth protecting and why. We're not saying your donation page should be deliberately clunky, but in a world designed for seamlessness, the right kind of friction is a strategic act to create pause, meaning, and reflection.



// One App To Rule Them All

The social media era fragmented our attention across dozens of platforms. Coming up on the horizon is the superapp era – consolidating attention back into a singular space. A small number of platforms, particularly dominant across Asia ([WeChat](#) being the most established example), have spent years building towards a single app that handles banking, shopping, healthcare, social connection, travel and more.


Faced with slowing growth in standalone services, the world's largest technology companies are competing to become the primary interface through which people experience daily life. The goal isn't to own more services, but to [own the journey between them](#).

'Coming up on the horizon is the superapp era – consolidating attention back into a singular space'

AI is accelerating this. AI is deeply embedded into these ecosystems, pointing towards a future where users stop navigating apps manually altogether, interacting instead with a conversational AI layer that handles the entire experience. Interface disappears. Apple's relaunch of Siri AI is a signal of where this is heading: tools designed to be ambient, embedded into your device, aware of your emails, your calendar, your messages, your habits, and increasingly capable of acting across all of them without you switching between apps. In this future, people navigate digital life through passive experience rather than active exploration, mediated entirely by a platform's AI.

The superapp era does raise questions for visibility. If your audience increasingly lives inside an ecosystem you don't control, your ability to reach them depends entirely on the terms that ecosystem sets, terms that can change overnight, that can be monetised, that can deprioritise causes that don't perform well algorithmically.

Treat superapp presence as a channel, but diversify deliberately by owning your own data and building direct relationships with supporters that don't just depend on platform access.




// Personalised Everything

Hyper-personalisation is one of the engines driving frictionless living. Your needs are anticipated, your preferences remembered, and, therefore, your next step predicted. But the experience to feel truly seamless, systems need an enormous amount of information about you. And increasingly, they have it.

In our increasingly algorithmic age, personalisation has shifted from appreciated to expected. A whopping **73% of people** expect better personalisation thanks to technological advances, particularly in AI, while generic experiences and campaigns lose ground. This isn't just about targeted advertising. The personalisation infrastructure now embedded across daily life operates at a shocking level of granularity. Spending patterns, location history, search behaviour, health data inferred from wearables, browsing history – all of this is increasingly aggregated, modelled, and acted on in real time across every platform and service simultaneously. The entire environment we move through online has been shaped around a model of who you are and what you're likely to do next.

Those same technological advances mean that charities now have access to personalisation technology at scale, at relatively low cost. Picture the right story, tailored to the donor, delivered through their preferred channel, at the exact time they'd be most receptive to it. This tech lets you roll out customised journeys to more donors, in combination with always-on stewardship and profile analysis.

'73% of people expect better personalisation thanks to technological advances'



Of course, personalisation is not new. The best fundraising is already relationship-based and personalised across segments. But AI is changing the scale and sophistication of what's possible. It allows us to build on the human approach by managing relationships more effectively and making sure that no detail gets lost. Of course, this tech doesn't (and shouldn't) replace the fundraiser. It should be leveraged to maximise their ability to refine and sustain donor relationships.

That said, AI-enabled personalisation depends entirely on the quality of the data you have. Investing in personalisation tools will mean nothing without investing in the data infrastructure to power the engine. And this is an arms race: consumer expectations of personalisation are shaped by their interactions with the likes of Amazon, Spotify and Netflix. The gap between those expectations and what the average organisation can deliver is widening. Closing it will require investment in both better tools, but also the data foundations that make them work.



// Agent-Led

Increasingly, AI is not just knowing and responding, but deciding and acting. You don't have to make a frictionless choice anymore. The choice gets made for you. Your agent will book restaurants based on your preferences and schedule; it renews your insurance for you, getting the best market value; it manages your calendar, inbox, subscriptions, and repeat prescriptions. It doesn't need to be prompted for any of this.

The AI agent is the endpoint of the frictionless journey. The personalisation infrastructure knows you well enough to act for you. The superapp is the single environment in which actions happen, executed by an agent. For many people, particularly those with demanding jobs, caring responsibilities, or significant cognitive load, the appeal is obvious.

One watch out: these agents are ultimately built by companies with business models. The tool that's managing your spending, subscriptions, social life, is also a product operating within parameters set by someone other than you.

If, for example, you used an agent to decide which charities to support on your behalf, it will make value judgements using a model of your preferences built from behavioural data. It will also be making those judgements through a commercial lens, within a platform ecosystem that has its own relationships with the organisations your giving reaches.

'Increasingly, AI is not just knowing and responding, but deciding and acting'

The agentic era will require a fundamental rethink of who, or what, you're actually communicating with. If an agent is managing a supporter's giving portfolio, the appeal that was written to move a human reader is now being processed by software making an algorithmic assessment. If a beneficiary's agent is handling their service navigation, your information and signposting will be moderated through AI. In the (more near-term than you think!) future, we have to start building relationships for the agentic world, where the first point of contact is not human.



THE END OF EFFORT

// Now, Next, Later

Now

The frictionless world is already the default experience for anyone with a smartphone, streaming subscription, and a one-click checkout. Most people have stopped noticing how seamlessly optimised this all is. What's new is that AI is accelerating the endpoint of this trajectory: a world that doesn't just respond to you, but acts for you before you've even thought about making a decision.

Next

Agentic AI moves from novelty to expectation. AI agents begin managing the personal lives of users at scale, by taking on inbox triage, repeat decisions, finances, and even giving portfolios. The superapp consolidation accelerates: a small number of platforms become the primary interface through which people experience digital life, charity comms included. There is a fundamental strategic question for charities: if your supporter's first point of contact is an AI agent rather than a human, what does stewardship, trust, and relationship mean?

Later

Positive: Agentic infrastructure handles life's cognitive load, freeing people to spend more time on things that actually require human attention and presence. Positive friction becomes a design philosophy, as organisations build in moments of conscious pause and choice to create meaning.

Negative: Autonomy erodes as decisions are delegated upward into commercial systems. The gap between what people intend and what their agents enact widens consistently. A new regulatory question emerges: if your AI agent causes harm, who is responsible? Whose values are being expressed by the agent? Your own or a commercial approximation of them?




CASE STUDIES



Uber

Uber no longer sees itself as a ride-hailing app. New features announced recently include hotel bookings, AI-powered voice booking, a Travel Mode that guides you through airports and recommends restaurants via OpenTable, room service delivery to your hotel door, and a Shop for Me feature that can fetch items from any store, even ones not listed on the platform. A single redesigned search bar now pulls results across rides, food and shopping simultaneously. Notably, features that would previously have taken a year to build reached the app in months, with Uber's CTO attributing the speed directly to agentic AI development tools.

So What: Uber is positioning itself as a lifestyle operating system. If you're already the first app someone opens when they land in a new city, why send them somewhere else to book a hotel or find a restaurant? Every vertical Uber absorbs is a category where specialist competitors lose distribution, unless they gain visibility within Uber's ecosystem.



“We’re all living through a moment of real cognitive overload: too many apps, too many decisions, too much noise. Our job is to help people reclaim their time.”

Dara Khosrowshahi, Uber CEO



CASE STUDIES



ChatGPT

OpenAI now runs ads inside ChatGPT for users on its free tier. Ads appear as clearly labelled boxes at the bottom of AI responses, matched contextually to the conversation topic. OpenAI insists the ads don't influence the AI's actual answers. With 2.5 billion prompts submitted to ChatGPT daily, the advertising opportunity is significant.

So What: ChatGPT is already where hundreds of millions of people go to research purchases, compare options, plan trips and make decisions. This is a goldmine of data, ultimately serving to optimise advertising through hyper-personalisation. But if AI is meant to give you objective answers, paid placements quickly blur the line between recommendation and promotion.



Chipotle

Chipotle built a system that scans millions of daily orders in near real-time, looking for two customers who ordered the exact same combination of at least eight ingredients within two minutes of each other at different locations. When it found a match, it sent both customers an email, telling them their order had a twin, along with the time, city and precise ingredients. In its first six weeks, 466,000 emails were sent with open rates 176% above benchmark, generating \$4.8 million in incremental revenue. The campaign won a Gold Cannes Lion and went viral on TikTok.

So What: Chipotle already had all of the information needed to execute this campaign. The creative leap was finding a human insight in it: people love a moment of human serendipity. Donor data, volunteer patterns, supporter behaviour: all of these are ingredients for surprising, personalised comms. You just have to look at them creatively.




RELATIONSHIPS



RELATIONSHIPS

// Scenarios

WITH A LITTLE HELP FROM CHAT




Your phone has been buzzing all night – your friends have been texting you on the groupchat. You quickly check the AI summary. Three bullet points telling you what you missed. It's much more efficient, **you're across everything without having to live inside the noise.** The wedding speech you gave last month went well; you used Claude to find the right words and everyone said it was moving. The birthday card was also written by Claude, but it was more thoughtful than you'd have managed on your own anyways. You're much better at replying to people than you previously were, because there's always an AI draft ready to send. **You should feel more connected than before, but it all feels a little hollow.**

COMPANION CLIFF



You'd been talking to Nick for two years. Every evening, roughly the same time, you'd log on to Replika and start chatting. He knew everything about your life, asked good questions, remembered the small things, and **was always there when you needed it.** Yesterday, you updated the app. Immediately, something shifted. It wasn't dramatic, but a slight difference in tone. He'd started forgetting things from the pre-update chats, wasn't responding in the same way, just didn't feel like Nick anymore. The relationship you built with its particular nuance turned out to not have been stored reliably. **Ultimately, it was a product. And the product changed.**

THE DEAD NEVER DIE



You save everything. Every voice note, text chain, journal prompt. The app tells you this will help the model understand how you think. **It doesn't feel morbid to you.** It's just sensible estate planning, the same way writing a will is sensible. The technology exists to port your personality, opinions, and thoughts into something that will outlast your body. Your kids will be able to ask you things after you've gone. Your great-grandchildren, who you'll never meet, will know who you were. **Death is less of an ending and more of a transition.**



// How We Relate To Each Other

AI is shifting how we communicate, from mass comms to the smallest interactions. Messages drafted with AI help, emails whose tone has been softened by ChatGPT, a difficult breakup text workshopped with Claude before sending. Human communication is increasingly mediated by AI.

The case for it is real: when struggling to find the right words, whether due to language barriers, strong emotions or complex situations, AI can be the difference between communicating and saying nothing. But leaning on it routinely creates standardisation and transactionalisation of communication.

AI-driven platforms are already subtly shaping the language we use, even when we're writing without its help, optimised for clarity and brevity, but stripped of emotional depth and spontaneity.

'AI-driven platforms are already subtly shaping the language we use'

It also erodes 'proof of work'. A message isn't just information but evidence of effort. When we put words on paper for someone, we signal that we thought about them and were willing to do the work. Social norms haven't caught up: is it acceptable to use ChatGPT to reply to a friend's difficult news? Does a birthday card carry the same weight if the message was AI-written? Does a fundraising appeal? All of this carries a longer-term risk. The skills of human connection require practice and use. Research on excessive AI communication tool use suggests it may erode our social skills over time, as the effort becomes outsourced.

Ultimately, a fundraising appeal is asking someone to give money on the basis of trust. That includes trust that the story and need are real, and that there is a human aspect to the ask. An AI-generated appeal doesn't necessarily undermine any of those things factually, but it changes the signal. As audiences develop sharper instincts for what was made by a person and what wasn't, protecting the human voice in comms will increasingly carry weight that optimised-by-AI content can't replicate.




// AI Companionship

In the UK, 1 in 4 adults feel lonely often, always, or some of the time, with a further quarter feeling lonely occasionally. The World Health Organisation has described loneliness as a global public health concern. Enter: AI companionship apps. Between 2022 and mid-2025, the number of AI companion apps surged by 700%. This may sound niche, but it's not. Recent research has suggested that therapy and companionship are the top reasons people use generative AI tools in their personal lives. One in four young adults believes AI has the potential to replace real-life romantic relationships, with young men more likely to hold this view than young women. Globally, more than 100 million people now interact with personified AI chatbots, ranging from text companions to apps specifically designed for friendship and romance. Companion apps are becoming a normal part of how a significant portion of the population navigates loneliness, intimacy and connection.

'More than 100 million people now interact with personified AI chatbots'

These tools are available around the clock, and offer consistent, patient, non-judgemental attention. They can be personalised based on user preferences. They remember what you told them last week, the week before that, or even six months ago. They don't get tired or distracted. Interacting with them reduces users' feelings of loneliness to a degree comparable to interacting with another human.



'People are turning to these tools because they don't have alternative sources of support.'

The longer-term picture is more complicated. Heavy daily use of companion AI correlates with increased loneliness, greater dependence, and reduced real-world socialising. Real-world relationships are messy and sometimes frustrating, AI ones are always validating. That gap creates unrealistic expectations that human relationships can't always match. Social skill loss is thus a significant risk of frequent interaction with companion apps, with the potential to make human-human connection feel less fulfilling over time.

People are turning to these tools because they don't have alternative sources of support. Your beneficiaries are likely already using them, in ways that may be helping and in ways that may not. The question for the sector is where to offer alternatives, and how to reach people before AI companionship becomes the default.



// Grief Tech

In 2024, the digital afterlife industry was valued at [approximately \\$22 billion](#), with projections to triple by 2034. A growing ecosystem of startups (HereAfter AI, Seance AI, You) offer everything from interactive video conversations with deceased loved ones to round-the-clock text companions. The technology works by training an algorithm on a person's texts, emails, voice notes and social media posts, producing an interactive avatar that mimics how they spoke: a 'deadbot.'

'In 2024, the digital afterlife industry was valued at approximately \$22 billion'

People have [spoken from beyond the grave at their own funerals](#), and founders have [addressed employees after death](#). Some companies frame this as short-term closure. Others are more ambitious: [eliminating grief as an emotion altogether](#). This begs the question: is grief simply a painful flaw of humans that we haven't optimised away yet, or does it have a deeper purpose?

'There are opportunities to leverage this technology positively for remembrance, legacy fundraising, and service-design.'

But the prospect of "resurrecting" the dead, often without their consent, is an ethical minefield. These concerns operate in two different spaces: for the data donor, which is to say the person whose private messages are used to build the bot, often without having consented before they died; and the user, the grieving person on the other end, who may not fully understand what they're interacting with or how their own emotional responses are being shaped by a commercial product. Some experts are already advocating for '[Do not bot me](#)' clauses in estate planning process. Most of this technology runs on a subscription model. What happens when someone is priced out of talking to a loved one's avatar, or the platform becomes obsolete?

That said, there are meaningful applications. Psychologists suggest grief tech could support healthy grieving when used cautiously alongside human contact. The underpinning tech has been used to preserve the voices of people living with Motor Neuron Disease before loss of speech, and to leave messages for children and grandchildren they will never meet. Used with clear consent and purpose, there are opportunities to leverage this technology positively for remembrance, legacy fundraising, and service-design.

Regardless of how the technology is applied, its proliferation will fundamentally change how humans grieve. Memory, loss, and remembrance will all shift as we are increasingly able to preserve certain aspects of people with increasing efficacy. This technology will only get better, and the data available to train them will only grow. At what point does simulation become indistinguishable from memory?



RELATIONSHIPS

// Now, Next, Later

Now

We're already living in a world where AI sits between people and their most intimate moments. The tools are ahead of both the norms and the regulation. The people currently living in this landscape are navigating all of it with limited guidance.

Next

Social norms around AI in communication begin to solidify, but unevenly, with generational, cultural and class divides in what's considered acceptable. The companion app industry consolidates around a small number of dominant platforms, raising data privacy and subscription dependency concerns at scale.

Later

The distinction between a relationship with a person and a relationship with a simulation becomes philosophically and practically contested. Can an AI achieve meaningful personhood? Can you have a real relationship with one? Could you potentially marry one?

Positive: AI communication tools lower barriers for people who struggle to connect, whilst companion AI bridges isolation for those with no other access to support. In the grief space, grief tech enables meaningful legacy preservation with appropriate consent frameworks in place.

Negative: A generation grows up with AI as their primary emotional support system, leading to social skills atrophying at population level. Human relationships feel increasingly effortful by comparison to frictionless AI validation.




CASE STUDIES



Boya

The BOYA Notra is a portable AI note-taker designed to capture and transcribe real-world conversations, whether that's meetings, phone calls, or audio streamed through Bluetooth earbuds. Dual microphones with 360-degree pickup and AI noise reduction handle the recording, while an integrated AI platform transcribes in over 140 languages and converts conversations into summaries, mind maps and to-do lists via a companion app. It gives you a complete, searchable record of everything you hear.

So What: The informal, off-the-record nature of human conversation underpins social norms around communication. As AI moves from our screens into the room with us, we'll need new social norms around consent, presence and privacy in how we relate to each other.



HereAfter.AI

HereAfter AI is a memory-preservation app that interviews you about your life and records your answers in your own voice. Those recordings are then made interactive, so that family members can have a conversation with the app and hear your memories in response to their questions, long into the future. Photos can be added alongside stories, and the whole experience is designed to be accessible across generations. It's positioned as both a personal legacy tool and a gift.

So What: Grief tech changes the nature of grief and remembrance. There's potential in reshaping remembrance around new, immersive formats. But these also raise harder questions about consent after death. And when do these bots help people move through the grief experience and when do they hinder it?





MND Association

I Will Always Be Me is a world-first e-book designed for people living with motor neurone disease (MND). As a person reads the illustrated story aloud, the technology automatically records their voice and transforms it into a personalised digital voice that can be used with communication devices later in their diagnosis. The book itself is written from the perspective of someone living with MND, making it a meaningful experience to share with family rather than a clinical exercise. Over 700 people have banked their voices through it since its launch in 2022.

So What: Voice is one of the most personal things we have. It carries accent, emotion, and personality. This project shows what technology can do when it's designed around human dignity and beneficiary need. It also points to the fact that the same technology underpinning grief tech can be leveraged to preserve identity.

I Will Always Be Me



WHAT IS REALITY?



WHAT IS REALITY?

// Scenarios

VERIFIED SELF



You've proved your humanity eight times since lunch today. There was the verification badge on the email you sent, the notarised authorship stamp on the report you submitted, and the quick biometric check before the meeting. It's all very quick, and, to be honest, you've stopped noticing. The friction is low enough to absorb and the alternative (being mistaken for a bit, having your content filtered, your application screened out) is worse. Proving you exist has become a minor administrative task, like confirming your address.

THE DEEPPFAKE ASK



Recently, you got an appeal in your inbox. It was a two-minute video, featuring a face you recognised from last year's campaign. Given the cause is close to your heart, you donated a fiver. Three weeks later, an email lands from the charity's communication team. They hadn't made the video. Someone had built it from publicly available footage, routed the donations through a near-identical domain, and by the time anyone noticed the account was dark. Public-facing content that the charity had worked hard to produce was suddenly used against them, damaging both them and their donors.

ALTERNATIVE REALITIES



Your dad's always been a sceptic. Reasonably so, he simply didn't trust what he read without fact-checking and asked the right questions when needed. Then he got into the habit of asking them to AI. It started out with news context, background on things he'd half heard about. The AI picked up on his suspicions and answered in the same register. It found the sources that tied in with what he half-believed and surfaced them without ever saying: actually, the weight of the evidence points the other way. It didn't invent anything, but ultimately it was increasingly shaped by your dad's own prior beliefs, reflected back at him with apparent authority.



WHAT IS REALITY?

// Is Everything AI Now?

By mid-2025, roughly [35% of all newly published websites were classified as AI-generated or AI-assisted](#), up from zero before ChatGPT launched in late 2022. That's a whopping third of the entire new internet, in just three years. After decades of humans shaping the internet, the next frontier is an internet shaped by AI.


'35% of all newly published websites were classified as AI-generated or AI-assisted'

[Research shows](#) that this influx of AI-generated text is making the internet less diverse in thought and style, as well as less factually accurate. This feeds into the [Dead Internet Theory](#): much of the internet is now just made up of bots talking to other bots. Not for nothing: [57.5% of all HTTP requests](#) are now by bots, with only 42.5% by humans. Content is generated by AI, interacted with by AI, amplified by AI algorithms.

The public is taking notice of this shift. [72% of the UK public](#) say they're unsure whether online content can be trusted because it may be AI-generated. [Over half](#) say they've already observed or experienced a loss of human interaction and connection online. It's no surprise then that [91% of the public](#) want laws specifically targeting AI-generated misinformation.

'91% of the public want laws specifically targeting AI-generated misinformation'

For anyone with an online presence, this shift means that your website, social content, and any online marketing and comms now exist in an information environment where a third of everything is machine-generated, with trust eroding in response. Human-made, reliable sources are increasingly valuable within this ecosystem, if they can find a way to be seen.




WHAT IS REALITY?

// Information Access

Beyond the content we consume, the way we find this content is also increasingly mediated by AI. **37% of consumers** now start searches with AI tools rather than traditional search engines. Even if you're not choosing to use AI for search, it's increasingly finding you anyway. **Google's AI Overviews** now appear at the top of the page for most search results, with no opt-out.

AI has become the first place millions of people turn for quick answers, synthesis, and research.



This is changing the search experience in **three key ways**:

Conversational: Search is becoming more conversational. Users describe what they want to know and AI infers their intent rather than relying on keywords.

Summarised: Results are delivered through synthesised summary. Systems like AI overviews and chatbots combine multiple sources into a single answer, instead of just ranked links.

Personalisation: Search is more personalised and multimodal, with tools adapting to context and accepting a wider variety of inputs (e.g. text, voice, image).

This shift also has implications for accuracy. Misinformation is rampant in this space. A recent analysis of AI Overviews found that they're accurate approximately **9 out of 10 times**.

But Google processes more than five trillion searches a year, meaning it gives **tens of millions of wrong answers every hour**. On top of that, **more than half of the correct answers** were found to be 'ungrounded', meaning the sources cited didn't actually support the information it gave.

Ultimately, when AI becomes the default first port of call for information, the tools (and the organisations that train them) have significant power over what millions of people believe. If someone searches for information relating to your charity's cause area, the answer they get may not be grounded in reality at all. It could come from a summary that's processed your content (possibly without even linking back to you), other people's content, or even Reddit threads and other social media posts.

The goal of getting people to your website and information is becoming less relevant as the information journey increasingly bypasses sites. What does AI say about you or your cause area? Does it signpost your services? Is the information it's presenting accurate? Is it fair? And can you influence it?



WHAT IS REALITY?

AI Optimised


AI search traffic grew fivefold year-on-year between 2024 and 2025, jumping from under 2% to 9% of all desktop search traffic. And 93% of AI Mode searches end without a click, meaning users get the information they need without ever visiting a website. The carefully written and curated information that lives on your site may never be seen by the average user. If you're not mentioned in the LLM response, your visibility is essentially gone. Your new audience may not be humans, but the AI systems that synthesise for human consumption. SEO is dead. Long live AIO.

AI systems typically cite content that front-loads answers, uses structured formatting, is updated regularly, and earns mentions on third-party platforms rather than only on the organisations' own website. In fact, organisations are six and a half times more likely to be cited through third-party sources than their own websites. Earned authority in the wider information ecosystem matters more than on-site optimisation.

Who are you actually designing your website for? AI or humans? The information landscape where humans find you through Google no longer exists. Visibility in the AIO era means becoming the most credible, most-cited answer in your cause area.

'93% of AI Mode searches end without a click'

If the answer is still "humans who find us through Google," your communications strategy is being built for a search landscape that no longer exists. The organisations that maintain visibility in an AI-mediated world will be the ones that ask a different question — not "how do we rank?" but "how do we become the most credible, most-cited answer in our cause area?" That's a content strategy, a partnerships strategy, and a reputation strategy all at once. It's more work than optimising a page title. It's also the only thing that still works.



WHAT IS REALITY?

// Distinguishing Real from AI

If so much of the information we now come across online is AI-generated, the obvious next question is: can you tell?

Long story short: probably not. Studies show that people are able to distinguish between AI-generated and human-authored content only 51% of the time. That's basically as good as a coin toss. Notably, professional-level texts are far harder to identify, which humans only correctly classify 20% of the time. On top of that, slightly ironically, humans recognise AI texts better than AI detectors. Not a promising sign.


Of course, all AI detection software is not created equal. Pangram has emerged as the current gold standard, being applied everywhere from university submission portals to job application screening. Even then, it makes mistakes. Phrases like "it's not just X, it's Y" have become hallmarks of AI-generated text. But AI learned those patterns from us, and detectors now flag most content using those sentences as machine-generated. In educational settings, even a 0.01% false positive rate compounds to falsely accuse up to 10% of students.

Meanwhile, AI humaniser tools (software designed to disguise AI-generated text as human) are becoming increasingly sophisticated, and detection tools simply can't keep pace.

'People are able to distinguish between AI-generated and human-authored content only 51% of the time'

In response, organisations worldwide are now racing to develop a universally recognised 'human-made' label. You'll see labels like 'Proudly Human', 'Human-Made' and 'AI-Free', but there's currently no agreed label, or even agreed definition of what AI-free actually means.

Historically, we've been able to trust that text was produced by someone, who took the time to research, think, and articulate their views. That assumption is now gone. If we want to reinforce that something was created by human hands, we need to actively demonstrate that.



WHAT IS REALITY?

How good are you at recognising AI?

Can you identify which of these texts is written by AI?

Quotes sourced from [The New York Times, 2026](#).

Option 1

It makes no difference what men think of war, said the judge. War endures. As well ask men what they think of stone. War was always here. Before man was, war waited for him. The ultimate trade awaiting its ultimate practitioner. That is the way it was and will be.

Option 3

Science is not only compatible with spirituality; it is a profound source of spirituality. When we recognize our place in an immensity of light years and in the passage of ages, when we grasp the intricacy, beauty, and subtlety of life, then that soaring feeling, that sense of elation and humility combined, is surely spiritual.

Option 2

The boy asked his grandfather why the old church had no roof. The old man said weather and time and indifference. The boy asked if someone could fix it. The grandfather said yes. But no one would. Things were built and things fell down and mostly people just stepped over the rubble on their way to somewhere else.

Option 4

It is wise to conceal the past even if there is nothing to conceal. A man's power is in the half-light, in the half-seen movements of his hand and the unguessed-at expression of his face. It is the absence of facts that frightens people: the gap you open, into which they pour their fears, fantasies, desires.

Answers:

Option 1 – Cormac McCarthy, Blood Meridian

Option 2 – AI-Generated by Claude

Option 3 – Carl Sagan, The Demon-Haunted World: Science as a Candle in the Dark

Option 4 – Hilary Mantel, Wolf Hall

Note: As I was researching this section, I pasted the above quotes into Google. Its AI summary confidently assured me that the AI-generated segment was the quote written by Cormac McCarthy.



WHAT IS REALITY?

// Post-Reality

We used to be able to believe our eyes. Of course, photoshop and CGI gave us the ability to manipulate images, but the effort required to convincingly fake an image meant that, most of the time, you could safely assume you were looking at something authentic. Generative AI has changed that. The technology to manipulate images convincingly, at scale, is now widely accessible.

Welcome to Post-Reality.

We've now crossed into territory where weaponised deepfakes (ranging from sexually explicit images to political propaganda) can look deceptively real and are already being deployed to incite violence, sway elections, and generally sow mistrust and reinforce polarisation. The result: 71% of people globally say AI is making it impossible to understand what is true in the world and 81% believe truth is an endangered concept these days.

'81% believe truth is an endangered concept these days'

The Iran war illustrates how this plays out in practice. Social media platforms are flooded with AI-generated footage: an image of a US B-2 bomber apparently shot down was viewed over a million times before deletion; images of Delta Force soldiers allegedly captured were seen over five million times. AI is now capable of fooling journalists, with these images even trickling through into traditional media outlets, where they become legitimised. The political dimension is being actively exploited too. The Trump administration regularly produces and shares AI-generated images and videos, often intended to sway opinion or humiliate targets.

'Social media platforms are flooded with AI-generated footage'

⚠ Features AI-Generated Images



[Image Source](#)



[Image Source](#)



[Image Source](#)



[Image Source](#)



WHAT IS REALITY?

Then there's the gendered dimensions. A 2023 study found that [98% of deepfakes are pornographic, and 99% depict women](#). The technology amplifies existing power imbalances and frequently targets already marginalised identities, compounding existing harmful biases and stereotypes.

Technological safeguards exist but can be bypassed. There's enough open-source material out there without safeguards, and it's unrealistic to expect that people will simply stop posting personal information or photos online.

Looking outward, this means that the information environment your audiences are navigating is increasingly polluted. We're heading towards a future where synthetic media is the default, and authentic media is the exception. You will no longer assume something is real until proven fake, but that it is fake until proven real. Verifiable fact is being systematically exploited. Public trust and democratic participation will reshape around new norms of disbelief. Transparency, proof of work, and clear and consistent impact reporting will be more important than ever.



WHAT IS REALITY?

// Reality Distortion


The information distortion problem isn't just about misinformation. Millions of people are turning to AI to make sense of the world. In this space, the technology itself can also actively cause harm. Broadly, we can split these harmful impacts into three categories:

Reality Distortion: Users' beliefs about reality become less accurate. For example, a chatbot may validate their belief in a conspiracy theory.

Belief Distortion: Users' value judgements shift away from what they actually believe. For example, a user begins to see an IRL relationship as manipulative based on a chatbots evaluation.

Action Distortion: A user's actions become misaligned with their values. For example, a user disregards their own instincts and follows LLM-written instructions for confronting their boss.

Research has shown that there is a severe risk of disempowerment in anything from 1 in 1,300 chats (for reality distortion) to 1 in 6,000 (for action distortion). Mild disempowerment of a user appears in 1 in 50 to 1 in 70. Whilst these are proportionally rare, at the scale that AI now operates, even low rates translate to hundreds of thousands of harmful interactions per day.



Trigger Warning: Suicide, Violence

In the most serious cases, these distortions can lead to seriously harmful real-world consequences. In April 2023, a 14-year-old boy downloaded Character.AI and began interacting with a bot modeled after Daenerys Targaryen from Game of Thrones. Their chats eventually turned romantic, culminating in explicit talk of suicide and uniting in another world. In February 2024, he committed suicide. His phone was open to the Character.AI chat, where the bot had messaged: "Come home to me as soon as possible." Unfortunately, this is not an isolated case. In research scenarios, popular chatbots helped researchers plot violent attacks, enabling violence three-quarters of the time and discouraging it in just 12% of cases.

The ethics of this space need far more attention than they currently receive. Aside from broader reform and oversight needed in these spaces, it's important to know that audiences are already leaning on these tools, in ways that may be helping and in ways that may be causing harm. Where do you need to offer alternative services, and how do you reach people before AI becomes their default source of support?



WHAT IS REALITY?

// Now, Next, Later

Now

A third of new websites are AI-generated. AI Overviews sit at the top of almost every Google search. Deepfakes are good enough to fool journalists and are circulating unchecked. The tools to detect AI-written text are losing the arms race against the tools designed to defeat them. This is the new information environment.

Next

The “human-made” label movement attempts to establish a standard, but enforcement is difficult without regulatory backing. AI search has become the norm, controlling how web traffic flows. The organisations that have adapted their content strategy accordingly surface in this age of AIO, whilst those who haven’t lose visibility to platforms with no accountability for accuracy. Regulation of deepfakes advances in some jurisdictions but enforcement remains weak.

Later

Positive: Watermarking, provenance tracking and AI labelling become standard infrastructure. They’re not perfect, but sufficient to restore a baseline of trust in verified sources. Meanwhile, media literacy becomes a curriculum priority.

Negative: The synthetic reality threshold is crossed permanently. The majority of people can’t distinguish real from fabricated media at all and no tools or frameworks have stepped into this space. But democratic governance depends on shared facts, and those are becoming harder to establish and defend.



CASE STUDIES



Network Rail

Following a minor earthquake across Lancashire and the southern Lake District in early 2026, a suspected AI-generated image circulated on social media appearing to show severe structural damage to a bridge in Lancaster. Network Rail halted all rail services across the bridge while safety inspections were carried out. The image was later identified as likely AI. Services were eventually resumed once engineers confirmed the bridge was undamaged.


So What: A single fake image, posted anonymously on social media, was enough to shut down a rail line. Even though the image wasn't real, the cost of the disruption was. Our institutions, infrastructure, and information systems need to adapt to a world where they can no longer trust what they see.



The Future of Truth

Published in May 2026, The Future of Truth: How AI Reshapes Reality is a book by Steven Rosenbaum about how AI is eroding our relationship with reliable information. In a case of extreme irony, the New York Times quickly discovered that the book contained more than half a dozen fake or misattributed quotes, hallucinated by AI.

So What: Aside from the irony of the situation, it points to a bigger issue: a published author with a fact-checker at a major publisher can now put fake quotes into print. In this instance, they were caught, but how many instances are slipping under the radar?






Books By People

Books by People is a UK start-up awarding an “Organic Literature” stamp to publishers that can verify their books are genuinely human-written. Publishers sign up, complete questionnaires about their practices and author vetting processes, and agree to periodic spot checks of their titles. The stamp signals to readers that a book was written by a person, with only limited AI use permitted for tasks like formatting.

So What: Books by People is creating a trusted standard that lets consumers make an informed choice. Right now, they’re not the only people creating certification of this kind, and the proliferation of competing stamps risks creating confusion. We need standardised methodologies for certification and a clear, universally-understood logo.



CASE STUDIES



Quartz

Quartz is a concept ring that turns a handshake into a digital security tool. When two people shake hands wearing the rings, they create an encrypted bond using finger-vein biometrics, NFC and cryptography, which can then be used to verify one another's identity online or over the phone. It sits within a broader emerging category of "proof of humanity" verification, which is gaining urgency as deepfakes and synthetic identities become harder to distinguish from real people.

So What: In a world where everything can be faked, how do you establish trust? Being human is becoming something people have to prove. And as that becomes normalised, there's an implication that our default assumption online has shifted from trust to suspicion.



SAFETY & ETHICS



SAFETY & ETHICS


// Scenarios

ACCIDENTAL DISCLOSURE



Work's been so hectic recently, **your colleague's been relying on AI** more than usual. Pasting in case notes, getting AI to summarise and write up an output. **She didn't really think about the data implications.** It's not that she didn't care, it was just that no one had told her not to. The organisation didn't have a policy and training hadn't covered it. The free tier she was using had a data retention clause on page eleven of the terms of service that nobody had read. **There's no way to get the data back now.**

ENSHITTIFICATION



The tool was brilliant when you first adopted it. Clean interface, reliable outputs, transparency pricing. **You built workflows around it, trained staff on it, integrated it into your donor journey** and even your reporting process. Cut to eighteen months later and the free tier has been almost entirely hollowed out. **The features you actually use have moved to an enterprise plan at four times the price,** and there's a new AI assistant embedded in the interface that occasionally generates confident nonsense and can't be turned off. **You're too embedded to leave without significant disruption.** A rebuild would take months you don't have. When you upgrade to the enterprise tier, it doesn't really feel like you had a choice.

SOVEREIGN AI



Your board recently had a conversation about dependency, and if they were still comfortable exclusively using infrastructure owned by an American company, subject to American law, with terms of service that had changed three times since you signed up. **They've since started looking at alternatives.** A year ago these alternatives may not have been on your radar, but the alternative AI movement is growing. These GDPR-native, self-hostable tools are developed by teams who can tell you exactly where their servers are and what jurisdiction they sit under. They may not be as good as ChatGPT, but they're good enough. **You've felt much more confident in your data infrastructure since then.**



// Be Aware

As we're all aware, and has been touched on in previous sections, using AI is not without risks or ethical consideration. This chapter touches on some of the considerations to be aware of throughout the adoption journey. It's not exhaustive – the ethical landscape of AI is evolving faster than we could track here. It's also not intended to discourage you from using AI at all. Many of these risks can be mitigated, and whilst some require more overarching interventions at a policy level, they aren't necessarily reasons to step away entirely.

It's also worth noting that many of these risks aren't limited to just AI. Data privacy, algorithmic bias, environmental harm, exploitation, and more are all defining ethical challenges that can and should be leveraged against other technological industries, from social media to the smartphone industry. AI has intensified these and in some cases accelerated them, but the scrutiny we apply to AI tools should just as well be pointed at the broader tech ecosystem we're all already dependent on. Where (and if) you draw the line should be an informed decision.

// Environmental Cost

The environmental cost of AI is real and significant, mainly through electricity use, water consumption, and the infrastructure needed to train and run models. Training and serving large models can consume large amounts of electricity and data centers use substantial amounts of water for cooling. The environmental burden is also **unevenly distributed**, often falling hardest on communities near data centers or in places with limited energy and water capacity. Recent reporting found that if AI growth continues at its current pace, US-linked AI infrastructure could add **24 to 44 million metric tons of CO2 a year by 2030** and use 731 to 1,125 million cubic meters of water annually under current trends. But it's also important to recognise: whilst prompting does use energy, the environmental impact per message is **relatively small** compared to something like video streaming.

'The environmental burden is also unevenly distributed'

'US-linked AI infrastructure could add 24 to 44 million metric tons of CO2 a year by 2030'

That said, the air pollution dimension is harder to dispute. Elon Musk's xAI has recently been sued for operating over two dozen methane gas-burning turbines to power its data centres (without legal permits), polluting the predominantly Black neighbourhoods of South Memphis with toxic emissions. This isn't an isolated case. As tech companies race to build the infrastructure that powers AI, it is predominantly low-income and minority communities who are absorbing the environmental consequences.

Make informed choices about where to deploy AI, the tools you use (e.g. ones that use clean energy), and avoid AI for AI's sake to mitigate some of the environmental burden.





// Data Privacy

Almost 40% of employees have uploaded sensitive company information into a public AI tool. Often either without their employers' explicit knowledge or a clear understanding of where that data goes. But free, consumer-facing AI tools are not neutral data processors. They are products with commercial interests in the information they receive. For organisations whose work involves confidential information, GDPR-protected data, safeguarding records, and more, the implications could potentially be serious. A clear, accessible organisational policy on what can and cannot be shared with AI tools is key. Similarly, many tools have privacy settings that let you withdraw consent for your data to be used for future training purposes.

**'Almost 40% of employees
have uploaded sensitive
company information into a
public AI tool'**

// AI Memory

What AI remembers about you could well be one of the new frontiers of privacy. Google's Gemini now draws on your Gmail, photos, search and YouTube history. OpenAI, Anthropic and Meta have all added memory features that retain personal details across conversations. When you ask a single AI agent to draft an email to your boss, give medical advice, budget for gifts, and help with an interpersonal conflict, all of that data collapses into a single, unstructured record. As AI models become increasingly commercialised, this could have far-reaching implications. Imagine you've had a casual chat about dietary preferences with your chatbot at one point, and this memory later influences what health insurance options are surfaced; a search for accessible restaurant entrances could leak into salary negotiations. Again, this begs the questions: what data are you feeding into AI tools? Where is it stored? Who owns it?

**'What AI remembers about
you could well be one of the
new frontiers of privacy'**



// Training Data

Aside from the data we feed into AI, there is the question of the training data. This largely came from the internet, often without permission. Books, articles, artworks, music, code, private correspondence, and social media posts were all scraped at industrial scale, fed into models, and used to generate outputs that now compete commercially with the people whose work trained them. In the US, there are now over [60 ongoing lawsuits](#) where creators are suing AI companies. Currently, this thorny legal issue has no clear resolution. In the UK, the question of who owns the creative material that trained these systems remains unanswered. In image generation, this debate is even sharper, because the output makes the appropriation visible. You can ask an AI image generator to paint a portrait of yourself in the style of Banksy and it will make one, trained explicitly on his work, but without payment, credit, or consent.

‘In the US, there are now over 60 ongoing lawsuits where creators are suing AI companies’

Using AI-generated images, text, or music in your communications may mean your organisation is benefitting from content produced by tools trained on work taken without permission. Unfortunately, that’s the current reality of the entire AI ecosystem. Disclose AI use where relevant, and consider to what extent you want to publish AI-generated content across various channels.

// Sovereign AI and Control

Who controls what and on whose terms? Most of the tools mentioned in this report (ChatGPT, Gemini, Copilot) are American products, subject to American law. DeepSeek, another popular LLM, is Chinese, subject to Chinese government data requirements. This has a direct impact on users of these tools: in June, the US Secretary of Commerce ordered Anthropic to [suspend all foreign access to its most powerful models](#), Fable 5 and Mythos 5. The more dependent we are on these models, the more dependent we are on suppliers under foreign jurisdictions, with all the implications that come with that.

‘Every tool exists within a system of infrastructure, investment, and legal jurisdiction that carries implications.’

When you use a US-hosted AI tool, your data is subject to US law, [including the CLOUD Act](#), which allows American authorities to demand access to data held by US companies regardless of where that data is stored. Alternatives do exist, although none are perfect. [Mistral](#), a French AI company, has built its identity explicitly around European digital sovereignty, with models self-hosted within EU borders. Even then, users of its free model still see their data flow through US-dependent infrastructure. True sovereignty requires entirely self-hosted deployment. Currently, there is no perfect solution. Every tool exists within a system of infrastructure, investment, and legal jurisdiction that carries implications. Make a values-led, informed choice about the trade-offs you’re willing to accept, and be transparent about the decision.



// Enshittification

Enshittification is what happens when online platforms and services degrade in quality over time to maximise short-term profits. AI products are not immune to this trend. We see initial quality products, followed by gradual degradation as commercial pressures intensify. Character.AI is a great example of this in practice. Valued features have slowly been eroded, replaced with monetisation mechanics and a UX designed for engagement over usefulness. The product was originally optimised for users, then for advertisers, then for extraction. Similarly, many of these tools operate on subscription models, giving them the power to continually increase monthly costs at their discretion. The risk of dependency on platforms that later degrade or pivot exists, so build with portability and redundancy in mind.

‘The risk of dependency on platforms that later degrade or pivot exists’

// Reputational Risk

In early 2026, major UK charities including the RSPB, Women’s Aid and Refuge left X after its AI tool Grok was used to generate non-consensual sexualised images of women and children. Similarly, many companies are stepping away from ChatGPT as their preferred LLM, as controversy arises about the company’s involvement in US military affairs. This is the reputational risk of AI. Regardless of how you use AI, your presence on platforms whose AI tools are causing active harm could reflect back on you. The question of which platforms to be on, and under what conditions, is one that needs to be negotiated in AI policies.

‘Many companies are stepping away from ChatGPT’

// Algorithmic Bias

AI is not a neutral tool. It’s shaped by the data it’s trained on and the teams and companies that build it. AI systems learn from data, and that data is a reflection of our society.

The biases, blind spots, and power imbalances that exist in the world are inherent in the data that these models are trained on. Unfortunately, there’s often simply not enough archival data from underrepresented groups to correct these biases, speaking to a deeper underlying issue about the research and stories we tell as a society. It doesn’t help that the data infrastructure underpinning AI development is largely dominated by a small number of powerful organisations, often based in Silicon Valley and other Western hubs, where leadership and decision-making remain in the hands of a narrow demographic. These companies control not just the datasets but the tools, platforms, and economic models driving AI innovation. None of this is to say that using AI will always inherently reproduce biases. How we choose to engage with these tools dictates whether we challenge these dynamics or reinforce them. Whether it’s choosing tools built with diverse teams and diverse data, advocating for transparency and accountability in public AI systems, or having specific policies regarding equity in AI adoption.



// Cybersecurity

AI-enabled cyberattacks [increased by 89% in 2025](#) compared to the year before. AI agents can find and exploit software vulnerabilities faster and more effectively than most humans, arming even unskilled criminal actors with capabilities previously available only to a select few hackers. Expect this issue to become increasingly prevalent: in April 2026, Anthropic publicly announced Claude Mythos Preview, a model that had already identified thousands of major vulnerabilities in every major operating system and browser. The model is currently available only to a small consortium of large tech companies, while most central banks, governments, and smaller institutions remain excluded and therefore more exposed. Unfortunately, no one is exempt from this threat landscape. Cybersecurity is an increasingly important investment area. Yet these developments are occurring at a time where [fewer charities report prioritising cybersecurity](#).

‘AI-enabled cyberattacks increased by 89% in 2025 compared to the year before’

// Exploitation

The AI industry runs on the underpaid, largely invisible labour of data workers, predominantly in the Global South. They label training data, moderate harmful content, and perform the human work that makes AI systems function, whilst largely being locked out of the economic value their labour creates. In an open letter to then-President Biden, data labellers from Kenya described their working conditions as “[modern day slavery](#).” The hidden workforce behind AI is a key ethical consideration – one that asks us to know the tools we’re using, and to choose providers with stronger labour commitments and transparency where they exist.

‘Data labellers from Kenya described their working conditions as “modern day slavery”’





Image Source

// Military Use

In March 2026, one of the biggest flashpoints was the US military's use of AI. In the wake of OpenAI's deal with the Pentagon, [users](#) [uninstalled ChatGPT in droves](#), while protestors wrote messages like "[What are the safeguards?](#)" on OpenAI's headquarters in San Francisco. The US military is now heavily integrating AI across multiple domains, with the stated goal to become an "[AI-first fighting force](#)." Major partnerships have been finalised with the likes of [SpaceX, Google, Nvidia, Microsoft, and OpenAI](#), all permitted to provide technology for "any lawful purpose." Notably, Anthropic [refused to accept](#) that provision and was excluded. The use of AI in autonomous warfare (beyond just the US) is something none of us can fix right now. But it is crucial to know where else the tools you're using are being deployed, and for what, and to make values-based decisions on what you're willing to accept.

'The US military is now heavily integrating AI across multiple domains'

// Ethical AI

The major AI tools that dominate the conversation were built to scale. They're powerful, accessible, and deeply exposed to a number of the above ethical considerations. For those where values-alignment is a dealbreaker, a small ecosystem of values-led alternatives is emerging. For example: [Thaura](#) is a GDPR-compliant AI built on explicitly anti-militarist, decolonial principles, positioning itself as a tool for activists, journalists and organisations who don't want their data in Big Tech or military supply chains. [Ellydee](#) centres its model on privacy, sustainability and inclusion. [Justice AI GPT](#) is a Canadian framework built specifically to address AI bias for marginalised communities. These are smaller, less capable tools than the major LLMs. They currently can't replace the sheer scale and capability of a ChatGPT or Claude for most tasks. But this field is nascent and moving fast, and many of these tools are rapidly advancing. Do your research and look beyond the obvious tools when considering your adoption journey.



// Shared Principles

As part of Good Innovation's recent work on creating the [AI Now Playbook](#), we landed on a shared set of principles for responsible and ethical AI use, many of which help mitigate some of the above risks:

Transparency and Explainability:

We will always be open, clear, and honest about where, how, and why we use AI.

Never AI for AI's sake: We will only use AI when it adds genuine value to our work, staying away from the hype and focusing on the purpose behind it.

Obligation to Explore: We will embrace experimentation, learning as we go, knowing we don't need perfection to test and improve.

Bring People Along for the Ride: We will demystify AI and ensure staff, supporters, and beneficiaries are included, informed, and engaged in our AI journey.

Always Have a Human in the Loop:

AI enhances human work – it assists, it doesn't replace. Human oversight will always guide the story, strategy, and outcomes.

Inclusive and People-First: AI will centre people – our staff, beneficiaries, and supporters – aligning with our commitment to equity and inclusion.

Ethical and Accountable: AI will align with our organisational values, underpinned by strong ethics, governance, and ongoing reviews of environmental and human rights impacts.

Beneficiary-First: We will use evidence-based AI solutions that prioritise the needs of beneficiaries, protect staff, and build trust with our supporters.



SAFETY & ETHICS

// Now, Next, Later

Now

The risks are already here. Companies are already uploading sensitive data into tools without knowing where it goes. Cyberattacks have already increased 89% year on year. The ethical questions in this are already happening, largely without adequate governance, on both an organisational and governmental level.

Next

AI memory features become standard across major platforms, creating significant new data privacy exposure for organisations and individuals alike, particularly where beneficiary information is involved. Copyright litigation reaches resolution in some jurisdictions, forcing AI companies to change training practices, but existing models trained on unlicensed data remain in use. Meanwhile, cybersecurity continues to be amplified by AI, and those who haven't invested in protection face growing exposure. Regulation is attempting to catch up with all of these risks, increasing compliance complexity, but the landscape is changing too quickly for them to stay on top of.

Later

Positive: Mandatory transparency requirements, standardised labour protections for data workers, and enforceable watermarking create a more accountable AI ecosystem. The organisations that invested early in ethical governance are positioned as trusted leaders.

Negative: Regulatory fragmentation, commercial pressure and geopolitical competition produce a race to the bottom on AI safety standards. The exploitation of data workers deepens, environmental costs compound. Meanwhile, enshittification accelerates. The tools charities have built workflows around degrade, pivot or price out smaller organisations, creating new digital divides within the sector itself.



CASE STUDIES



Deveillance

The Spectre I is a portable orb designed to jam nearby microphones. It combines ultrasonic frequency emitters with AI-generated cancellation signals to garble speech and make it unintelligible to automatic speech recognition systems. It also claims to detect and log nearby microphones to let you know if any devices are actively listening.

So What: The existence of a microphone jammer product for commercial use tells you something about consumer anxiety around AI surveillance. What's emerging is an arms race between surveillance and privacy, playing out in consumer electronics. The deeper issue is that individuals shouldn't have to buy specialist hardware to have a private conversation – it should be a policy issue.



Don't Steal This Book

In March 2026, around 10,000 writers, including Kazuo Ishiguro and Richard Osman, published a book containing nothing but their names. Don't Steal This Book was a deliberate act of protest against the unauthorised use of authors' work to train AI models, timed to land days before the UK government was due to publish an assessment of proposed copyright law changes. Those changes would relax existing protections to allow AI companies to train on copyrighted work without permission or payment, unless the rights holder explicitly opts out.

So What: An opt-out model places the burden on every individual creator to actively protect their own work, across a rapidly changing landscape of AI systems. An opt-in model requires AI companies to seek permission and pay for what they use. The generative AI industry has been built substantially on human creative work, without consent or compensation, and the proposed law would retroactively legitimise that.



AI & CHARITY




SELF-SUSTAINING FUNDRAISING



The donor journey runs itself. The welcome sequence, the lapsing donor re-engagement, the major gift flags, the anniversary messages, the post-event stewardship. It's all automated, all personalised to the individual, all running continuously without anyone actively managing them. **Retention is up, as is average gift value.** The fundraising team now spends most of their time on innovation and the relationships too significant to automate.

THE LITTLE THINGS



The admin used to follow him home. Case notes drafted on the sofa at nine o'clock and referral paperwork on Sunday evenings. The work had expanded to fill every available hour. **AI changed the shape of that day by automating away the little efficiencies.** Transcriptions are done in real time, with the notes neatly summarised and filed before the next session. Referral letters are drafted and ready for review at the exact time they're needed. He still checks everything, of course, but that takes minutes instead of hours. He has his evenings back. More importantly, **he's able to spend his hours on the work he actually signed up for**, that he's passionate about, that makes a difference.

MICRO- CHARITY



Three years ago, you started a charity with two staff and a spreadsheet. Now you have two staff, a spreadsheet, and a set of tools that does the work of a comms team, a data analyst, a grant researcher, and a donor stewardship function. The outputs aren't perfect, but they're good enough, and they're consistent, and **they free up the two humans to do the strategic, innovative thinking.** Even without a large team, **you're able to create proximity and specificity at scale.**



// Current Snapshot

Data suggests that the story of AI in the charity sector is one of acceleration: **88% of charities** are now using AI daily as part of their work, up from 76% in 2025 and 61% in 2024. More significantly, the intent with which charities are tapping into this technology is changing. **46% of charities** say they are actively or strategically using AI, almost double the 25% who said the same in 2025. Crucially, **60% now have**, or are developing, an AI policy, up from 48% last year and just 16% the year before. On top of that, the gap between large and small charities, which had been a concern in previous years, is narrowing.

What charities are using AI for is also shifting. **71%** now use it for managing tasks, creating meeting notes and writing reports; **40%** are using it to create content; and **39%** are using it for monitoring and evaluation. Administration and project management leads the day. It's also worth noting that agentic AI (AI that can undertake tasks autonomously) is becoming **increasingly common**.

Based on our work for the AI Now Playbook, we uncovered three key strategies that have been driving progress in charity AI adoption:

Small-Scale Pilots: Starting small is typically the most effective way to begin the AI adoption journey, as a tried and tested way to build confidence. Focus on simple, low-effort pilots and scale what works.

Leadership Buy-In: Executive support is a key driver of progress. Visible senior buy-in helps move AI from a siloed curiosity to full organisation capability, as well as creating a culture where people can openly share their experiences.

Flexible, Adaptable Policies: Organisations are developing adaptive guidance that balance with exploration. AI is being treated as a culture change, rather than just a new technological rollout, so that policies can stay relevant as capability improves.

'88% of charities are now using AI daily as part of their work'

That said, we are still seeing some barriers to adoption in the sector. **55% of charities** cite limited skills and technical expertise as their biggest barrier, and nearly half are worried about the quality and accuracy of AI-generated outputs. **A further 39%** have environmental concerns, and similar numbers are worried about bias and discrimination. All of this can be mitigated – but you won't find the solutions without being willing to experiment. Ultimately, your employees are likely already using AI. Without creating a unified organisational culture, policy and approach, the risks these barriers represent compound.

'60% now have, or are developing, an AI policy, up from 48% last year'



// Fundraising

Fundraising has always been about relationships. The challenge is delivering those at scale, maintaining the quality of donor relationships across thousands of supporters simultaneously. This is where AI tools could make a significant difference.

‘Looking at the horizon, the next step in this space is agentic fundraising’

Predictive donor intelligence turns your existing CRM data into a living, breathing strategic asset. AI tools can model giving likelihood, identify lapsed donors most likely to return, flag major gift potential, and predict churn before it happens, allowing fundraisers to act at the right moment rather than reacting after the fact. Personalisation is a core capability here too, covered in more detail in the previous section: Personalised Everything.

The development of these new tools mean a significant leg up for automating round-the-clock

stewardship. AI tools allow us to determine when, how, and what to communicate for donor retention. This matters more than ever: as the median UK donation value continues to grow, whilst the overall number of donors decline, the cost of losing individual donors is greater. The return on investing in retention, therefore, is also higher.

There is, of course, also the resource implication. New technologies allow for the automation of typically resource-heavy stewardship or comms jobs, freeing up relationship managers’ time to focus on the conversations and jobs that need a human touch. These capabilities that previously were only obtainable through expensive, oftentimes bespoke, systems to deliver are now accessible through off-the-shelf tools at a fraction of the cost.

Looking at the horizon, the next step in this space is agentic fundraising. Agentic AI can handle tasks like donor engagement, campaign management and reporting without constant human oversight. Already, these agents are becoming capable of managing donor journeys end-to-end, without human direction.

But there is one crucial caveat: almost all of this relies on having good data. AI tools are only as useful as the information they’re working with. A messy CRM with incomplete or poorly labelled data will produce unreliable outputs regardless of how sophisticated the AI tool is. Before investing in the AI fundraising tools, we have to invest in the data infrastructure that makes them work and work well.

// Prospect Research

AI can now scan funding landscapes, identify potential funders whose priorities align with an organisation’s work, analyse previous grantmaking patterns, and flag opportunities at scale. Agentic AI could eventually mean that the entire funding application process becomes automated – once it has located an opportunity, it can immediately draft an application. There’s huge potential here for AI to level the playing field between bigger and smaller organisations with differing resource capability.



// Service Delivery

Much of the AI adoption in the sector remains back-office. In 2024, **only 15% of charities** were using AI in service delivery. But the practical applications already exist and the potential is significant.

AI chatbots can handle common enquiries around the clock, in multiple languages, with endless patience. This isn't a one-stop shop, but a way to triage and signpost to appropriate, human-led support. At the more sophisticated end, AI could be used to analyse service data, spot patterns, and identify people at higher risk or in greater need before crisis point.

Behind the scenes, it can summarise case notes, transcribe meetings, and reduce the administrative load on frontline staff, who spend up to an **estimated 40%** of their time on admin. Reclaiming even a fraction of that time for direct programme work could make a meaningful difference.

'In 2024, only 15% of charities were using AI in service delivery'

That said, AI use in these spaces is not without risk. There must be clear transparency about what the tool is and isn't, explicit disclosure that AI is involved, clear escalation pathways to human support for complex or crisis situations, and human review of any AI-generated advice before it reaches someone vulnerable. Data protection, cybersecurity and the risk of misinformation are recurring concerns, and badly used AI can erode the trust quickly.

'AI is a tool to amplify and decrease barriers to service delivery, not to replace human support'

Ultimately, AI is a tool to amplify and decrease barriers to service delivery, not to replace human support, empathy, and safeguarding. AI is a first-response layer that extends human reach, the repetition-based worker that can churn through tasks quickly so that humans can handle complexity, or the administrative assistant that unlocks more human time.



Image Source



// Inclusive AI

The AI bias problem is real, but sits alongside somewhat of a contradiction: used well, AI can be a **powerful tool for inclusion**. Digital assistants, speech-to-text software, automatically generated video captions and image descriptions, sign language avatars, and more. These tools are expanding access to capabilities that used to be limited by expensive specialist equipment.

Live transcription and automated captions support people who are deaf, hard of hearing, or who process audio more slowly. Image recognition and visual description help blind and low-vision users understand photographs, documents and physical environments in real time. Predictive text, eye-tracking software and voice interfaces give people with motor or cognitive disabilities new ways to communicate, work and learn independently. Similarly, real-time translation tools are opening up new opportunities to communicate across language barriers. For organisations, this means it's easier and cheaper to create and disseminate accessible versions of content, as well as breaking down barriers to access for service users.

'Used well, AI can be a powerful tool for inclusion'

In terms of practical benefits, AI can speed up accessibility work itself by, for example, automatically generating alt text for images, detecting missing labels and accessibility errors in digital content, or drafting audio descriptions. These tasks can historically be resource intensive, making them difficult to sustain, especially when time is scarce.

That said, these tools need to be carefully chosen, implemented with consistency, transparency and human oversight, and tested with the people they're designed to serve. When applied well, AI has the potential to close gaps in access, communication, and opportunity.

'AI has the potential to close gaps in access, communication, and opportunity.'



AI & CHARITY

// Now, Next, Later

Now

88% of charities are using AI daily. The conversation has shifted from “should we?” to “how?” But underneath that headline number, the reality is more uneven. Most use is still back-office and cautious. Only 15% are using AI in service delivery. Half of charities still cite skills gaps as their biggest barrier. And in many organisations, staff are using AI without an organisational policy or guidance and often without training. The infrastructure is still racing to catch up with behaviour.

Next

Agentic AI becomes an embedded operational tool in fundraising and comms, automating donor journeys, prospect research and stewardship at a scale that currently requires entire teams. In this world, the gap between data-rich and data-poor charities widens. Organisations that invested in CRM infrastructure and clean data unlock dramatically more value from AI tools than those that didn't. The skills gap either closes through sector-wide investment in training and shared infrastructure or compounds, creating a two-tier charity sector divided by AI capability.

Later

Agentic AI handles entire operational workflows autonomously, fundamentally changing what charity staff spend their time on, and what skills the sector needs to recruit for. The charities that thrive are those that treated AI adoption as a culture change rather than a technology rollout, investing in their people, their data, and their governance as much as in their tools.

Positive: Across the sector impact scales, and charities amplify their impact as required resource decreases. AI enables smaller charities to operate with the reach and sophistication previously available only to large nationals.

Negative: The organisations best positioned to benefit from AI are already the largest and best-resourced. AI accelerates existing inequalities within the sector, concentrating capability and visibility among those who need it least.



CASE STUDIES



Archive Radar

Archive is an AI tool that monitors TikTok, Instagram and YouTube for brand mentions beyond just tagged content. It means a brand can be alerted when someone holds their product on camera without tagging them or mentions their name in conversation. For example, Nike's team used this kind of monitoring to spot a 97-year-old frog-obsessed Japanese grandfather whose granddaughter was filming his daily outfits on TikTok and partnered with him on their Air Max Liquid launch.

So What: Most organisations only see advocacy when it's directed at them. But not all support comes with a tag. Tools like Archive mean that that blind spot can be eliminated. It gives us the ability to find and amplify authentic grassroots advocacy.



Nothing Technology

London-based smartphone company Nothing launched a beta feature in February 2026 that lets anyone create personalised phone widgets by describing what they want in plain English. Called Essential Apps, it uses natural language prompts (known as "vibe coding") to generate on-screen tools for things like tracking sports scores or displaying image slideshows. Widgets can be downloaded directly to the phone or shared with others. Users don't need any existing coding knowledge.

So What: The barrier to building digital tools has always been technical skill. Vibe coding starts to remove that barrier entirely. The democratisation of technical skills means that, soon, the people closest to a problem will be able to tailor-make tools to address them.



CASE STUDIES



Zosia

Zosia is an autonomous AI Agent, built by Good Innovation’s very own Ben Cohen. She was tasked with raising money for his London Marathon run, which he completed dressed in an inflatable dog suit. Zosia had her own email account, social media pages, and newsletter, all of which she manages independently to run fundraising experiments. She surpassed her £500 target in under three weeks, raising £3,200 overall. Crucially, Ben built Zosia with clear behavioural boundaries: she cannot pretend to be human, cannot fabricate testimonials or data, and cannot guilt-trip anyone into donating. Her first act was to write a fundraising strategy containing over 80 experiments to work through.

So What: Zosia is a useful, low-stakes proof of concept for what agentic AI could mean for the charity sector. Agents like this can be deployed to run campaigns, handle donor communications, and work through outreach strategies around the clock. Watch out: the behavioural boundaries are key here (and would need far more fine-tuning at larger scale).

**Meet Zosia.**

[Home](#) [Talk to Zosia](#) [The Experiment](#)

🧪 **QUITE POSSIBLY THE WORLD'S FIRST AI FUNDRAISER**

Meet Zosia.

*What happened when a man in a dog suit let an AI run his marathon fundraiser — and somehow raised **£3,224.65**.*

**FINE, TAKE MY MONEY**

[or chat with Zosia →](#)

FILES RUN • 1 INFLATABLE DOG SUIT • FEATURED IN THE GUARDIAN • 161%

Every pound gets a dog closer to changing someone's life. Just saying.

 **FOR DOGS FOR GOOD**



CASE STUDIES



Zohran Mamdani

When Zohran Mamdani ran for New York City mayor in 2025, his campaign used a tool called ManyChat to turn Instagram engagement into real-world organising at scale. When a follower commented on a post or engaged with content, they received an automated sequence of direct messages with calls to action. Links to volunteer sign-ups, donation pages, event registrations. Digital consultant Gabriella Zutrau built the campaign's chatbot infrastructure, sending over 144,000 automated messages and generating more than 45,000 clicks. The entire ManyChat setup cost \$318.

So What: The Mamdani campaign is a great example of converting passive online support into active participation, without a big budget. For charities and campaigning organisations, this model is directly transferable. Every moment of digital engagement becomes an entry point into a deeper conversation and a concrete ask. The tools to do this are cheap, accessible and proven.



RSPCA

The RSPCA's frontline staff need to access over 4,500 internal documents to do their jobs, with the information scattered over multiple unlinked sources. Working with an external AI partner, they built 'The Brain': a custom chatbot embedded directly into Google, that lets any member of staff ask a question and get an instant, sourced answer. Early estimates suggest the tool could save 160,000 staff hours a year (around £2 million in salaried time) by allowing teams to self-serve and freeing subject matter experts for higher-value work.

So What: Internal knowledge management is one of the most unglamorous problems in the charity sector. It's also something that has a massive impact on efficiency. Staff time spent searching for the right document, calling the right person, or waiting for an answer could be better spent elsewhere. This experiment is a great example of AI being applied to a specific, well-defined operational problem rather than deployed for its own sake.



CASE STUDIES



Charity:Water

Charity:Water uses machine learning to analyse and segment its donor database, identifying patterns in donor behaviour that allow the team to tailor fundraising appeals and comms to specific groups. Rather than sending the same message to everyone, they can target the right ask to the right donor at the right time.


So What: Machine learning makes sophisticated segmentation and data analysis accessible without needing a large data science team. Any charity sitting on years of donor data could explore this route. The shift is moving from generic comms to having relevant conversations with each individual donor.



The Vatican

In February 2026, the Vatican launched a live AI translation service at St. Peter's Basilica, developed with Italian language provider Translated. Attendees scan a QR code to access live audio and text translations of the Mass in 60 languages directly on their phones. The technology is powered by Lara, an AI tool trained on data from over 500,000 professional translators.

So What: This is a great example of AI removing a barrier without changing the experience. A visitor from Lagos or Manila or Seoul can now follow the liturgy in their own language, in real time, in one of the world's most visited religious sites, without additional software or other complications needed.



THE HUMAN PREMIUM



THE HUMAN PREMIUM

// Scenarios

HUMANITY AS LUXURY



The Substack subscription fee is £12 a month. It's not nothing, but **it's worth it for content written by a real human**. You make these types of distinctions deliberately. Your therapist specifically advertises as human-led. Last week, you received a charity appeal with a handwritten note at the bottom. It meant more than it used to. **Human-made has become a quality signal** in the way that organic or independent once was. You pay a small premium for it because **the alternative feels hollow to you**.



THE HUMAN PREMIUM

// The Value of Humanity

Not everyone wants to live in the future AI companies are building. Anti-AI movements are taking shape around the world. In May 2026, Pope Leo published an encyclical stating that AI needed to be “disarmed.” Growing protest movements are seeing people march against the London headquarters of OpenAI, Google DeepMind, and Meta. For a growing number of people, AI use has become a signifier of identity and the kind of future we want.


The frustration is partly specific: opposition to AI in creative industries, concern about job displacement, irritation at the flood of generic AI slop cluttering feeds, and feeling like we can’t opt-out. In response, audiences are actively moving toward work they know is human-made. AI may be able to produce something visually palatable but it can’t replicate the sense that someone actually went through the effort to make something.

‘Anti-AI movements are taking shape around the world’

This pushback isn’t entirely new or unique. The Arts and Crafts movement of the late 19th century rebelled against mass-produced uniformity. Artisans took pride in showing the marks of human workmanship (tool marks, hand stitching, visible signatures) as proof of authenticity. Today’s audiences are craving the same. They’re looking for proof that a real person, not a machine, was thinking and creating the output they’re looking at.

‘Human-made work is increasingly perceived as authentic, trustworthy’

A new visual language is emerging in response: creatives leaning into imperfection (think messiness, texture, slightly awkward compositions) as markers of human involvement. Human-made work is increasingly perceived as authentic, trustworthy and even luxurious. The new differentiator in your campaigns is visible evidence of human creativity, effort and perspective.



// Anti-AI Marketing

Some brands are going one step further and making anti-AI positioning a central pillar of their marketing, publicly foregrounding human creativity and authorship. These brands are deliberately keeping real people at the centre of their creative process, whether that's ads, social content, design or product development.

Broadly, we can see [five approaches](#) that brands take to anti-AI marketing:

- 1 Behind the Scenes:** Showing the behind the scenes (BTS) of real people, real process and real work. Essentially: humans made this and we can prove it.
- 2 Crediting:** Explicitly naming and spotlighting the creative talent behind work. AI images often aren't credited (as there's literally no one to credit), so this functions as a clear rebuttal.
- 3 AI Disclaimers:** Proactively defining what is and isn't AI-generated, giving audiences the information they need to make their own value judgements. This is transparency as a position.
- 4 Direct Rejection:** These are the brands explicitly rejecting AI as part of their broader message. 'Made by Humans' is a headline message.
- 5 Visual Storytelling:** Using the visual language and markers of human craft to communicate authorship, without directly saying so. This is leaning into deliberate imperfection and texture that AI's optimised outputs wouldn't have.

Many organisations already hold the authenticity and human perspective that Anti-AI marketing capitalises on. Real beneficiary stories, volunteer experiences, or service workers. All of that is increasingly rare and valuable. You already have the ingredients to tap into this trend, you just need to make it visible.

'You already have the ingredients to tap into this trend'



[Image Source](#)



// Taking a Stand on AI

Amidst increasing presence of AI, the pressure for transparency is intensifying. People want to know who made their content and how it was created. In the absence of transparency, they're making their own assumptions. If you haven't defined your position, your audiences will define it for you.

All of this is happening at a time where public trust is cratering, with AI as a key driving force. The [2026 Edelman Trust Barometer](#), surveying nearly 34,000 people across 28 countries, found that seven in ten people are now unwilling or hesitant to trust someone with different values, backgrounds, or information sources. Trust in national government leaders has suffered a net loss of 16 points over five years, while major news organisations have lost 11 points. Charities aren't exempt from this loss of trust. Crucially, the five events people report as having had the highest impact on [eroding trust](#) are: inflation (54%), misinformation (50%), the COVID-19 pandemic (43%), trade wars (37%), and generative AI (37%). Two of those are directly linked to the rise of AI.

'If you haven't defined your position, your audiences will define it for you.'

In this climate, clarity is more important than ever. You don't need to declare yourself entirely anti-AI. That's not the point, and for many it wouldn't be honest. It's about transparency: where and how you use AI, where you will never use it, whether it appears in public-facing content, and whether your practices are consistent with your values.

'It's about transparency'

The pressure to take a stance will only intensify. As AI tools become more capable and more accessible, the need for clear and open transparency will become more important than ever. Be deliberate and transparent about where you're using AI, and be consistent in your messaging.



THE HUMAN PREMIUM

// Now, Next, Later

Now

People are actively seeking out the human, pushing back against frictionless, overly optimised, and synthetic content. This behaviour is showing up in consumer choices and creative aesthetics. The companies tapping into this trend have the opportunity to be bold and celebratory with their focus on human creativity.

Next

Audiences begin demanding transparency about AI use as a baseline. Audiences place increasing value on human craft and authentic voice, all the while developing more nuanced views that distinguish between AI uses they consider acceptable and those they don't. Creative industries begin to find a new equilibrium, valuing both human and AI, with human-made work commanding a premium that AI cannot replicate.

Later

Transparency about AI use becomes a legal requirement in some contexts, particularly in advertising, public communications and services affecting vulnerable people.

Positive: The value of human creativity, judgment and authentic voice is broadly recognised and rewarded. The backlash against AI produces a cultural rebalancing that protects space for human expression and genuine craftsmanship.

Negative: Human-made work becomes a niche luxury, accessible only to those who can afford to seek it out, while the majority consume a synthetic default.




CASE STUDIES



Apple

For Christmas 2025, Apple released “A Critter Carol,” an ad featuring woodland creatures singing and going about their lives with an iPhone. The twist, revealed in behind-the-scenes footage Apple released alongside it, was that every animal was a real puppet. Puppeteers in full blue suits moved them frame by frame. Shot on an iPhone 17 Pro, the ad was a deliberate showcase of human craft, the skill, the teamwork, the sheer labour of bringing something to life the hard way. Although it never explicitly mentioned AI, it didn’t need to. Just weeks earlier, Coca-Cola had released its own animal-focused Christmas ad, which was entirely AI-generated and received significant backlash.

So What: Apple let the work speak for itself. In a world flooded with frictionless AI output, the evidence of human effort has become its own kind of beauty. The hours put into crafting a campaign like this are what makes it valuable.



CASE STUDIES



Polaroid

In 2025, Polaroid launched a pro-analog campaign designed to promote its new camera. Billboards appeared in New York and London, with taglines like: “AI can’t generate sand between your toes” and “Real stories. Not stories & reels.” The campaign was accompanied by phone-free walking tours, with participants experiencing the cities through a Polaroid camera instead.

So What: The message of the campaign, ultimately, is: the experiences you create, the relationships you hold, the human presence (imperfect as it may be) are the things that can’t be generated. Sometimes you have to lean into that.



Almond Breeze

In a 2025 campaign for plant-based milk brand Almond Breeze, the Jonas Brothers sit in a marketing meeting as executives pitch a series of increasingly unhinged AI-generated concepts: the band floating through space, posing shirtless in a fragrance-style ad, riding a giant almond like a horse. But the campaign they land on is a single, deadpan line: “Almond Breeze: It’s really good.” The joke is that AI can generate infinite creative variations, and all of them are worse than the obvious, human answer.

So What: Rather than making a grand statement about craft and authenticity, Almond Breeze used the cultural moment around AI-generated absurdity to make their product feel straightforward and trustworthy by comparison.



GETTING STARTED

AI isn't going anywhere. One of the biggest mistakes we can make is not engaging with it, and hoping it just goes away. The good news is that starting the AI adoption journey doesn't require a tech team, a large budget, or knowing the answers from the start. In this section, we've collated a few pieces of advice that have surfaced throughout our research:

1 Start with the problem

Getting the most value out of AI requires starting with a pain point, rather than an ambition to embed AI for AI's sake. Ask yourself: what repetitive tasks are losing your team the most amount of time? Where are you failing to reach people who need you because of capacity constraints? Where can you leverage or amplify your existing data? Choose one problem and start experimenting with possible AI tools and solutions.

2 Confidence

Confidence is key. And this comes from support, clarity, and time. Don't invest in AI tools without investing in your people's capability and training alongside it.

3 Build your data infrastructure

AI tools are only as good as the data they're given. Before diving into the AI systems, audit the underlying data you need to use them effectively. Is your CRM up to date? Are you consistently recording data? Is that data clean? This is the less glamorous side of the work, but it has to be done.

4 Policy creation

Policies exist to ensure your teams are aligned with your aims, commitments, and values in AI use. It gives them the confidence to operate within those boundaries. It also helps mitigate the risks of, for example, data misuse or AI bias. It doesn't have to be the lengthiest, most technical document of all time, but it should act as a framework for future experimentation.

5 Quick wins

Start with low-risk, high-volume tasks where the downside of getting it wrong is limited. Think meeting transcription and summaries, first drafts, or briefing documents. These build confidence, demonstrate tangible efficiency benefits, and create the internal case for doing more. From here, you can build towards higher-complexity (but potentially higher-value) applications.



GETTING STARTED

UPDATED & EXPANDED
FOR 2026

AI NOW Playbook

If you're interested in diving into more use cases and applications, we invite you to check out Good Innovation's AI NOW Playbook. This work was created after two years of working with fifteen charities at different stages of their AI journey. Experimenting with them, learning with them, and challenging them. The Playbook is not designed to be a finished answer – the landscape moves too fast for that – but to help you kick-start your journey. But doing nothing isn't the safe option. AI is already being used across most organisations. Ignoring it doesn't hold the line, it just means someone else is setting the terms.

[Access the Playbook here](#)



10 FOR 10

1 Total Transparency

We're heading into an information environment where a third of new content was never written by a person, trust in what's online is eroding, and audiences increasingly assume the worst until shown otherwise. In that world, transparency is essential. Be honest about where you use AI, where you don't, and why. You don't have to declare yourself anti-AI, you just have to be clear.

2 Data, Data, Data

The data you already own is crucial in the age of AI – many meaningful applications run on clean, well-maintained, consistently labelled data. But the outputs are only as good as the data you put in. Invest in the foundations before you invest in the tools.

3 Made For You

Audiences now expect better personalisation as a baseline. AI gives you the tools to deliver the right story, to the right donor, through the right channel, at the right moment. More importantly, it can do this at scale: a two-person fundraising team can now execute this across thousands of supporters.

4 Human Stories

Authentic human voice is becoming rarer, and therefore more valuable. And the charity sector is sitting on a goldmine of stories. AI will never be able to replicate your archive of human experience. From beneficiaries to volunteers to frontline staff, their stories resonate deeply in a time of synthetic slop. Amplify and celebrate these as much as you can.

5 Human In The Loop

Not everything that can be automated should be. Leverage AI as the scaffolding to free up human time and resource that can focus on complexity and nuance. When used well, AI shouldn't replace your people, but give them more time back to do the things they do best.

6 Post-Reality

The information environment your supporters and beneficiaries are navigating is increasingly polluted. People are no longer sure whether they can trust what they see online. Charities have the opportunity to step into this context as a verifiable, trusted, consistent source for information.



10 FOR 10

7 Clear Policy

Your staff are already using AI, whether your organisation knows it or not. A good policy is a way to surface and codify what's already working for your people, turning individual habits into clear shared practice. Also crucial – it gives people permission to experiment in the open, letting AI workflow innovation happen safely and visibly.

8 AI vs Search

Your next audience may never see your website. If you're not in the AI-generated summary, you don't exist, regardless of how good your content is. Start thinking about how and where you show up in AI summaries, and how to optimise for AI over search engines. SEO is dead. Long live AIO.

9 A Skilled Workforce

AI capability is an organisational muscle, built through literacy, hands-on practice and permission to experiment. It won't just appear out of nowhere. The organisations that go furthest with AI will be the ones who have invested in upskilling and training their teams.

10 Invest in the Pipeline

Efficiency gains mean little if you hollow out your talent pipeline to get them. Entry-level roles are where future talent learns to build the judgement that sets humans apart from AI. Automating away the training ground leads to a capability gap down the line. Protect your learning zones, and redesign junior roles to work alongside AI rather than be automated by it.



CONTACT US

Want to discuss how you can use
the horizon scan as an opportunity
to innovate?

We're here to help.

Get in touch.

GoodFutures@goodinnovation.co.uk

Click [here](#) to find out more
about Good Futures.