

REPORT

AI in charity content

A study of over 70 charity content practitioners, exploring how generative AI is changing how content is made and how users find and engage with it.

Lauren Pope, in partnership with William Joseph

September 2026

Introduction

Charity content practitioners are experiencing two significant changes. The way users find and consume information online is shifting because of changes to search engines and adoption of generative AI tools. And the generative AI tools appearing in our workplaces have the potential to disrupt/are already disrupting the way content is planned, written, designed, produced and built.

Either one of these things alone would be a lot to absorb. But teams are dealing with both at once. And to compound it, the whole sector is under prolonged pressure: under-resourced, overstretched, and operating in a challenging political, social and economic climate.

Alongside the material ways AI is changing our work, it's adding a lot of noise. The relentless marketing telling us that generative AI is infallible and inevitable. The scramble to learn about and apply GEO.¹ The hot takes saying adopt now or get left behind. The public shaming for AI slop. The pressure from leaders to use it to create more content faster. There's also a double-edged moral pressure that's very specific to charities. If generative AI can help us work faster and do more, are we obliged to use it so that we can deliver more value for our donations and funding, and achieve more impact against our mission? (And it can be useful: I'm sure many of us will have felt a thrill when we enter a prompt and the output is something we could never have done without AI, or at least not as quickly.) But at the same time, AI contributes to the climate catastrophe, compromises rights and privacy, reinforces bias and stereotypes. Does using these tools compromise a charity's credibility and ability to work for the public good? And what about the compromise to our personal ethics?

I don't think we'll make good decisions about AI if we focus on that noise and on the technology itself. This might seem like an odd thing to say at the start of a report about AI. But if we focus the conversation on AI – what it can do, what everyone else is doing, what's the best tool – we're responding to the technology and the hype. What we should be responding to are the things that really matter: our mission, and the people and communities we serve. If we keep those two things in focus, the answers to the questions of if, how, and when to use AI are a lot clearer. That's the underlying idea that I've tried to hold on to throughout this report.

¹Generative engine optimisation (GEO) also called answer engine optimisation (AEO). For clarity and consistency, I'll use the term GEO throughout this report.

The aim of the report is to give content and comms leaders something that you can use to benchmark your team and organisation against: a sense of where the sector is right now. I also want to give you a clearer picture of what these tools can do and what their limits are, so you can make sharper decisions about capability and application. I also want to provide an honest view of the risks and compromises. Overall, I hope that this gets you closer to understanding what positive, strategic use of AI could look like for your content team.

A note on my personal perspective on AI

I'm a neo-Luddite. But that might not mean what you think it means.

The name 'Luddite' has become synonymous with being afraid of technology. But the Luddites weren't backwards or resisting progress out of fear. They were highly skilled artisan textile workers who knew how to use the weaving looms they vandalised. They protested as a strategic, collective bargaining tactic. They weren't fighting the looms, but the factory owners who were using technology to cut their wages and take away their agency. Their rebellion was about demanding a say in how technology was used and who could benefit from it.

So when I say I'm a neo-Luddite, I mean that I have a critical, sometimes resistant, perspective on generative AI. It's not rooted in fear, or a lack of understanding of the technology, and I'm not stuck in the 'resistance' half of the change curve. It's rooted in the ethics of the companies behind the products. The way they use people's work without credit or consent, the environmental cost, the damage to privacy and rights, and more. That perspective inevitably shapes this report, even though I have tried to write fairly and pragmatically.

Lauren

Findings summary

- **Charities are seeing a decline in organic search traffic that they attribute to generative AI.** This isn't just a traffic problem – it's a sign that the model (create content, get found, build a relationship) underpinning our digital strategies is breaking down.
- **AI overviews, AI summaries, and AI chat bots are creating misinformation about charities and the subject matter areas they deal in.** Sometimes this is down to hallucination, but sometimes it's down to AI surfacing outdated content from our own websites. This is creating problems for charity content teams, but also – and more importantly – for helplines and services. This is a risk that needs to be mitigated, and can be mitigated with content strategy best practice, but teams often don't have the capacity and backing they need.
- **Almost all charity content practitioners are using generative AI,** most commonly for things like drafting, checking, admin and synthesis. But most are also deliberately opting out of using AI for certain tasks, especially health and advice content.
- **There's a real split on using AI to write content.** One group – mostly individual, self-directed AI users with fewer organisational guardrails – use AI to help them get over intimidating blank pages. Another group sees the blank page as the place where the real thinking happens, and avoids AI for drafting as a result.
- **Most generative AI use is individual and informal – even though coordinated, team level applications are likely to have the biggest impact.** This isn't just a gap in terms of efficiency – even 'safe' individual uses of AI like summarising and transcription carry accuracy risks, and skipping the associated thinking that comes with them has its own cost.
- **It's not clear if generative AI is having a positive impact or not –** responses were mixed and no one shared concrete, measurable evidence of an AI application that was saving time or money, or creating better outcomes. Where people do feel that there is a positive impact, it tends to correlate with working somewhere that's provided training and support.
- **People feel extremely conflicted about using AI.** Ethical concerns are near universal, but this doesn't often translate into behaviour, and there's a lot of cognitive dissonance going on.
- **There's a lack of useful, meaningful policy, governance and training on AI.** Most respondents have not had formal training, and many do not feel supported by their organisation when it comes to AI.

- **Peer communities of practice are universally praised by those who have access to them.** They reduce anxiety, isolation and shame around AI use and provide valuable opportunities to learn.

1. AI, organic search traffic and content discovery

RESEARCH QUESTIONS

- Are charities seeing a fall in organic search traffic because people are finding the information they need from their search engine results page (for example Google's AI Overviews) or in chats with generative AI tools (like ChatGPT and Gemini)?
- Is there anything different about the traffic coming from AI sources?
- Are AI search and generative AI tools leading to misinformation?
- Are charities changing their content strategy as a result, and if so, how?

FINDINGS

Organic search traffic is declining and this is attributed to AI

- 39% of survey respondents have seen a decrease in organic website traffic over the last 12-18 months which they attribute to AI Overviews/AI search.
- Only 20% said that they had seen no decrease.
- Respondents reported declines of between 4 and 40% at the extremes, but most responses clustered between 10% and 25%.
- 80% see AI search as a risk to reach for their organisation (25% major, 41% moderate).

AI-driven traffic shows different patterns

- Respondents are seeing changes in user behaviour in the traffic coming to their websites from AI sources.
- The most common changes in user behaviour that people are seeing are:
 - lower engagement rates (25%)
 - users arriving with more specific or narrower queries (17%)
 - shorter time on page (17%)

36% have seen AI misinformation about their organisation, and it's creating operational issues

- 36% of respondents said that they had seen specific instances of AI tools producing inaccurate information about their organisation.
- The study also found a number of instances where AI tools were producing wrong or misleading information that was creating extra work and even conflict for helplines and services.
- Misinformation was a concern for all respondents – 46% said that they are very concerned about AI tools serving inaccurate information about their organisation or subject matter area.
- One respondent reported that the concern over AI misinformation is so great that their organisation has blocked AI from scraping its website.

Most charities haven't made changes to their content strategy (yet)

- Most respondents (65%) said that they have not made changes to their content strategy or approach in response to AI search. 38% said that they are actively planning changes. Only 27% have made some kind of change.
- Among those who have made a change, the most common steps people have taken are:
 - Monitoring AI search tools to see how our content is represented (54%)
 - Creating more structured or schema-marked-up content (44%)
 - Auditing or removing thin or outdated content (45%)

Capacity and leadership backing are barriers

- Respondents from smaller charities reported that a lack of capacity was holding them back from making changes and adapting their strategy.
- Some participants from larger organisations reported that leadership engagement or measurement uncertainty was their barrier.

'I think internally the risk is that the senior team don't have the engagement or knowledge in this area to see it as a risk, or not one high enough to put some proper resource behind.'

DEEP DIVE

How AI misinformation is affecting charities – and how to mitigate against it

‘ChatGPT told a user that our organisation funds funerals... this seems to have been an AI hallucination.’

‘There’s anecdotal evidence from our advisers where they’re having people show up with printed out reams of responses from AI. That can be really hard to advise the clients on, because they trust the AI more than the person sitting in front of them, especially when the AI is more likely to tell them what they want to hear.’

These two examples from the research are representative of a common story shared in the survey and the interviews:

1. Users get incorrect information from AI while researching an issue
2. They contact a charity’s helpline or service and get different information or advice
3. They feel conflicted, surprised and don’t know what to trust
4. Calls, emails, chats and appointments are harder or more time-consuming for the charity to handle as a result

The stakes feel higher for some charities too, especially those in areas connected with health or serving vulnerable people. Here, misinformation can have more serious consequences.

[Shelter has carried out research on how people are using AI to find advice.](#) They found that:

- 79% of advisers said they are sure some clients ask AI for help before contacting Shelter
- 21% of advisers said AI gives clients incorrect information ‘a lot of the time’; 38% say ‘some of the time’
- 8% of advisers had clients tell them Shelter’s advice is wrong because AI said something different
- 58% of advisers felt clients using AI makes giving advice more difficult
- 57% of advisers said AI makes them less confident about the future

There are two potential reasons why this kind of misinformation might be happening. The first is that the AI could be ‘hallucinating’ – inventing details because the algorithm calculates that they are the right kind of words to use in the response, even though they have no basis in fact. The second is that AI could be finding outdated or inaccurate information on our websites or other sites, and using it in its response. This second scenario can be mitigated through established content strategy best practices. Things like deleting old content, improving thin pages, and adding structure. These things also have benefits for user experience, SEO, and GEO too.

ANALYSIS

Charity leaders need to be aware of AI disintermediation and changing user behaviours

Top ranking pages now get 58 clicks when they would have got 100 before – [Ahrefs study](#) comparing data from before and after the AI Overviews rollout

73% of the public have used AI in their day-to-day life in the past month, and 35% have used generative AI like Chat GPT – [UK government study](#)

5% of all ChatGPT messages globally are about healthcare – [OpenAI](#)

Declining organic search traffic and the increase in cases of misinformation are already a concern to charity content practitioners, and should probably be a bigger concern for charity leaders. But there’s a far more fundamental shift going on here that’s relevant to both content teams and charity leaders.

The level of AI adoption and usage means that the model we’ve been relying on and using as the cornerstone of digital strategy is breaking. For years, there’s been an unspoken agreement: we create content and make it freely available on our websites. In exchange, search engines and social platforms help people discover that content. Users click through and we can build a relationship, offer support, or ask for a donation. The search engines and social platforms sell ads. Everyone benefits.

AI is breaking this agreement through disintermediation. AI tools still use our content, they just don't send as many people to our websites. They're taking content from multiple sources, synthesising it, and delivering it directly to users. People don't need to visit the source. Or – as the study shows – the users eventually make it to the source, but with poor quality information or unrealistic expectations.

It's tempting to see this as a traffic problem. But that's missing the bigger picture: we're seeing a shift in how people find information online. Generative AI hasn't completely replaced search. What we are seeing is that it is now one of several parallel paths that people take when they're looking for information and advice. That poses three challenges for us as content practitioners:

1. A SEO, GEO and structured content challenge for visibility and accuracy in AI Overviews
2. A channel-strategy challenge now that a proportion of users are going straight to Gemini or Chat GPT, and expect to complete tasks in those platforms
3. A user insight challenge in understanding what our users will deploy and trust AI for, and therefore what's a threat to reach versus what's a more minor channel shift (plus this is likely to change a lot in the short to mid-term)

CASE STUDY

Scope

[Scope's content team carried out research with disabled people and their families to understand how they search for information](#). They found that behaviour varies a lot: some just use AI, some avoid it, and some sit in between. They also found that people changed the tools and methods they used depending on the task. For a task about PIP (Personal Independence Payment, a state benefit in the UK) some users avoided using AI on the grounds of trust. The complexity of the question also influenced behaviour – some participants were more likely to use a search engine if it was a simple question. However, if the topic was more complex or they wanted tailored information, they were more likely to use AI.

CASE STUDY

Patient Information Forum

[The Patient Information Forum has also been carrying out research in the health sector.](#) They found that:

- 9 in 10 people go online to seek health information
- 7 in 10 use the NHS website or app
- 2 in 10 use Google AI summaries and ChatGPT
- More than 2 in 10 use social media
- More than 6 in 10 ChatGPT users say they ask questions about their health and wellbeing.

CASE STUDY

The Department for Business and Trade

The team at the Department for Business and Trade found that AI models were often generating misinformation by scraping outdated or fragmented guidance from the Department's website, among other sources. In response to this, they have:

- Centralised information into signposting pages to give AI reliable, up-to-date, well-structured information
- Archived old, low-traffic and outdated pages to prevent the out-dated or incorrect information from appearing in AI tools
- Used a dedicated audit and tracking tool (CART – which is not an AI tool) to monitor what content needs to be updated or corrected
- Tested proactive content design processes to ensure that future policies are AI-ready, accurate, and consistently maintained

Source: [Department for Business and Trade](#)

TAKEAWAYS

Ways to address AI's impact on organic search traffic and content discovery

- **Make sure you know what's happening with your users.** If you haven't carried out research to find out more about how your users are engaging with AI, it's likely to be worth prioritising this. It could be as simple as a survey on

your website and to your mailing list, or you could run some interviews and user testing. The objective should be to learn more about whether your users are using AI, what they use it for, and if they trust it. This will give you a more solid foundation for making any tactical or strategic changes.

- **Make sure you know what's happening to your content.** If you're in a position of responsibility for content and you're not aware of how AI has affected your organic search traffic, make it a priority to find out. 30% of people in the survey did not know if their organic traffic had changed.
- **Understand how you show up.** Build up a picture of how your organisation and content are showing up in AI tools. 54% of respondents are monitoring AI search tools to see how their content is represented. Google now offers reports on generative AI performance, and there are also options like [Ahrefs Brand Radar](#) and [Microsoft Clarity](#). However, these tools are new, and getting an accurate picture of AI visibility is complex. The traditional keyword and page based SEO impressions and clicks don't translate to generative AI, so there's a big learning curve.
- **Brief your leaders.** If you feel like your leaders do not understand or appreciate the impact AI is having and could have in the future, brief them on:
 - How traffic patterns are changing as a result of AI and how this could affect the charity over time
 - What you know about how your users are using AI
 - Data on, and specific examples of, misinformation in AI and how this is affecting calls, emails, webchats and real life interaction with users
- **Get the basics right.** If you see evidence of misinformation, or you want to optimise for AI, focus on the established content, usability and readability best practices that are also beneficial for GEO. These kinds of activities are lower risk than focusing on some of the emerging tips for GEO (like adding FAQs), which sometimes contradict good practice for user-focused design. Things like:
 - Deleting out-of-date content
 - Strengthening content structure, meta data and schema
 - Improving old or thin content
 - Adding review dates or publishing dates
 - Producing more authoritative, original or expert-led content
- **Use the AI angle to do best practice work.** Consider if using the AI misinformation or GEO angles could help you get buy-in for core content strategy and content design work that you want to do, but have struggled to get support for in the past. These angles might also show the risk of having an under-resourced content team.

- **Don't rush to act without evidence.** Don't feel rushed into acting if you don't have an evidence-based hypothesis. One interviewee – who leads a team that I see as a leading light in charity content design – said that their choice not to adapt their strategy or tactics was intentional, they felt that there's not enough data and insight available to make those choices yet. Watchful waiting and gathering evidence is a perfectly acceptable response if you're already getting the content basics and best practice right.

2. How charity content practitioners use AI

RESEARCH QUESTIONS

- Are charity content practitioners using AI?
- What tools are they using?
- What are practitioners using AI for, and is it effective?
- Is AI usage individual and informal or part of shared and defined processes?

FINDINGS

90% of charity content practitioners use AI

- 64% of respondents are using AI in their work at least once a week.
- 29% use it daily.
- 10% said that they do not use AI tools in any aspect of their content work.

Copilot is the most commonly used tool

- The most commonly used tool is Microsoft Copilot, which is used by 51% of respondents.
- Gemini is the next most common, at 39%.
- Grammarly or other writing assistants are used by 32%.
- Only one participant reported using a bespoke in house AI tool/platform.

AI is most commonly used for drafting, checking, admin and synthesis

- The most common tasks that respondents use AI for are:
 - Editing or improving existing copy (50%)
 - Proofreading, grammar or tone checking (48%)
 - Generating initial drafts or outlines (42%)
 - Summarising or condensing content (40%)
 - Internal tasks like emails, meeting notes, briefing documents (39%)

Some teams are building custom tools to draft or check content against their guidelines

- 3 participants use AI tools to check content against their guidelines, by building custom tools (for example using custom Gems, or adding guidelines to Grammarly) that can review their content against their style, voice, tone and/or writing guidelines.
- 7 participants specifically mentioned that they use a custom AI agent to draft content based on their voice, tone and/or writing guidelines.
- These participants reported different levels of success:
 - 'We have a custom Gem that has been trained on our style guide... it's really, really good'
 - 'We have tried to build a custom Gem that other teams can use to get feedback on their copy to make sure it's in the right tone and style, however, it's not that intelligent and does get things wrong or makes stuff up (for example it will call someone out for doing something against the style guide, when they haven't)'

People are asking AI rather than having conversations

- Some respondents reported that they now use AI for questions that they would previously have asked a manager, a peer, or an outside contact.
- Others said that they worry that people are doing this.

'Maybe I have some big overwhelming things I'm trying to get my head around... it's quicker and a bit less exposing to sometimes ask Gemini.'

'I hate that when you're stuck on something, the first answer from many people now is 'why don't you use AI to...' rather than trying to problem-solve yourself, or speak to other people.'

'I wonder how much the increase in hybrid working is kind of a perfect storm... we're just going to continually talk to robots rather than each other.'

70% deliberately do not use AI for certain tasks

- The most common things that this group said they do not use AI for are:
 - Full drafts and long-form writing (*'I would never use it for a first draft of anything long-form because I know what it will produce will be so bad and generic I'd just need to rewrite the whole thing anyway'*)
 - Imagery, photography and design (*'It's important we are authentic in our use of imagery and using AI tools to generate images feels in opposition to that principle'*)
 - Numbers, research and evidence (*'I don't use it for anything involving numbers because I've found it to be too unreliable'*)
- Health, clinical and advice content is an area where many practitioners are opting not to use AI. Several respondents reported having formal policies that prohibit using AI to generate this content, because it's high stakes, needs original thinking, and the kind of deep institutional expertise that AI can't provide.

'In my team, we wrote an AI and health content policy... It states that we don't use AI to write health content, it's authored by humans and reviewed by humans.'

'We do not use LLMs to generate health information... we do not trust LLMs to be accurate or to convey the correct supportive tone for our service users.'

Most AI use is individual and informal

- 52% of respondents described their AI use as 'mostly individual and informal'.
- 28% said their use is 'individual but within loose team norms'.
- Only 20% described their use as part of a defined team or organisational approach (10% team, 2% organisation-wide, 8% a mix).

It's not clear if AI is having a positive impact on work or not

- 35% – the most common response – said that the impact of AI is mixed, with some benefits, some drawbacks.

- 27% said the impact is broadly positive – it helps, though the benefits are modest.
- 14% said it is having a very positive impact, and significantly improving the quality or efficiency of their work.
- 8% said the impact is broadly negative, creating more problems than it solves.

CASE STUDY

Parkinson's UK

Parkinson's UK has a specific policy guiding how its content team uses generative AI for health content. The policy states that the team will not use generative AI to author health content. It also provides examples of how AI can be used in relation to health content, along with principles for responsible use. Thanks to Pia Dawson for sharing this example with me.

DEEP DIVE

The conflict over using AI to write content

The research uncovered some interesting ideas and some real conflict about using AI to write content. 42% said that they are using AI for generating initial drafts or outlines – a way of helping them get past the intimidation of a blank page.

'It helps me get over the initial fear on the blank sheet.'

'I use it lightly – often when I'm stuck or can't get started as I sometimes find it easier to edit than create.'

'AI tools have increased my efficiency and helped me more easily overcome frustrations like "writer's block".'

When I dug into the data, I found a pattern in this group that I'm calling 'drafters', (with a caveat that this is a small pool – 26 people). In the survey responses, drafting content with AI is associated with more frequent, self-directed, ungoverned AI use, and a positive attitude to AI. This group are daily users of tools that they pay for themselves, working for organisations without an AI policy, and with fewer concerns about the tech. This might be a sign of more agile organisations, or of a lower level of digital maturity.

For another group of respondents, staring at a blank page and having to generate the first version of something isn't a chore to be automated away – it's the heart of what they do as content practitioners. The blank page is where the thinking happens and where they work out how to solve for the user need. If AI fills that page, they are reacting to and editing the AI's framing rather than establishing their own – even if they rewrite every word afterwards.

'The blank page is the hardest because it is the most valuable.'

'First drafts I avoid letting AI do because I think that is giving too much steer to the AI, risks my own practice.'

'I was hired because I have the skills to write and design content, not because I have the skills to ask an AI the right questions to write and design content. Until AI surpasses my ability to write in the right voice in the right contexts with nuance and understanding I will continue to do those tasks myself.'

I dug into the data for this 'non-drafter' group too, and found that there are some non-drafters who are regular AI users, and feel good about using AI too. For this group, refusing to draft with AI isn't the same as being uncomfortable with using AI – they just don't think writing is a good use case for the technology. On the other hand, there are those who do not use AI because of their overall ethical concerns.

ANALYSIS

Even 'safe' individual uses of AI have risks – and coordinated, team-level applications are likely to have the biggest impact

What I can see emerging from the study is a sense that there are 'safe' uses of AI. Things like summarising, data-wrangling, transcription, capturing meeting notes, using an agent as a personal sounding board. These tend to be individual activities that people do alone as part of their own workflows, and that are seen as freeing up time and capacity for work that makes better use of their skills. The kind of work – like drafting content or making judgements – that would be riskier to use AI for.

However, there is still risk in these kinds of AI applications. Firstly, AI makes a lot of mistakes, even with simple tasks like synthesis and transcription. [The 2026 AI Index Report from Stanford HAI](#) found that AI accuracy is still lower than human accuracy, and in a benchmark, hallucination rates across 26 top AI models are between 22% to 94%. One in five times, AI will get it wrong – that's the best case scenario.

'AI models can win a gold medal at the International Mathematical Olympiad but cannot reliably tell time' – [Stanford HAI](#)

There's also a much bigger, more insidious risk – skipping some of this work is skipping important thinking. Tasks like writing meeting notes might feel like drudgery, but they play an important role in cognition and understanding. The process of writing something down means we have to engage with it in a more active way than when we just read something. On a personal note, I was excited to use AI to create user and stakeholder interview summaries because of the time saving, but stopped doing it after realising that it was leading to poorer outputs because I didn't have the same level of understanding of the data.

'Even if you write down (or, let's face it, have an AI transcribe) a meeting or an interview into a document, knowledge does not automatically come along with that text. Knowledge is something that is produced in your brain as you transcribe the text, because transcription is an act of metacognition. It is synthesis. It is thinking.' – [Pavel Samsonov](#)

There's something complex and a bit painful to grapple with here. I've never worked with a charity content team that wasn't over capacity and under-resourced. The idea that AI can take away some of the work is appealing. It can also feel like we have a responsibility to use it, so that we can do more with our donations and funding. But we need to balance that moral imperative with risk mitigation.

I heard [Ellen Gofton](#) use the analogy of the Sorcerer's Apprentice to describe this. The young apprentice to a powerful sorcerer is tasked with fetching water. Left unsupervised, he decides the best way to do it is to magic himself more buckets and mops to carry it. Spoiler: it's not the best way. Chaos and a flood ensue. More people using more generative AI prompts isn't always the way to a better output, and it has potential to get messy.

One way to avoid some of the risks is to look for team-wide, visible applications for AI. The Stanford report found that productivity gains from AI are largest in structured, measurable work where the outputs are easy to monitor. And when all AI

usage is individual and hidden (even inadvertently) there's a risk of poor practice and inefficiency. As one content leader put it *'because the way that we work is so integrated, it's really important to know if someone is using AI. They need to be open about that'*. The survey data itself backs this up; the people who say AI is having a very positive impact are more likely to:

- work at organisations that have given them training and guidance on how to use AI
- say that they feel fairly or very well supported by their organisation in using AI.

In a talk at Lead With Tempo, Google's [Tim Hetland](#) described the concept of AI model behaviour design, which gives a good steer on what 'team wide and visible' looks like in practice. Solo AI usage looks like:

Prompt > AI > Output

Team-wide, visible, AI model behaviour design looks like:

Prompt > System instructions and examples > AI > Output > Rubrics and evaluations > Feedback loop to improve the system

In this approach, system instructions and examples mean that AI has a basis for what good looks like. Evaluations and rubrics mean the AI can self-correct and critique. And the feedback loop means the system gets better over time. Individual, 'shadow' AI use might skip these steps, or take them but not share the benefits among the team, and can compound poor quality as a result. I'd say there's a step before the prompt too, which is ensuring that the work you're doing is aligned to your strategy and user needs.

Examples like Tim's approach to model behavior design show the value of looking outside the charity sector for a view on what the future might look like. Product and tech sector content design has already ridden the AI rollercoaster. There have been waves of layoffs that have affected content folks specifically. But there are also a range of new AI-focused content roles, and in some places, a new appreciation of the skills that content practitioners have. Let's not forget that large language models are at the heart of generative AI. Understanding language, words and how they work is an increasingly valuable skill. As charities experiment more with building AI tools and interfaces for users, content teams should be at the heart of that development.

CASE STUDY

Citizen's Advice

At Citizen's Advice, the Innovation and AI team has created a custom Gem in Google Gemini to streamline the brief creation process. The content designer and subject expert have a 15-minute call using a predefined list of questions to capture necessary information. They download the transcript of the call and paste it into the custom Gem. The Gem summarises the transcript and formats it into a summary which is then pasted into Jira. This ensures the person picking up the ticket has full context and a clear understanding of what is expected before they begin drafting. This approach aims to reduce misunderstanding and back-and-forth communication between team members. Thanks to Phoebe Nightingale for sharing this example with me.

TAKEAWAYS

Ways to use and govern AI in your content team

- **Root your choices about AI in your mission and users.** Before you use AI, ask some questions about the compromise you're making, to understand if it's a good one, or just an easy one:
 - **Do you need more content/this content?** AI makes it easy to create more content faster. But more content doesn't always mean better or more effective. Will making more content help? Is there a true user need for the thing you're about to make?
 - **What's actually blocking your team, and can AI fix that?** Look at the things that are really affecting your team. If the blocker is a lack of strategy, AI can't fix that. If it's a lack of data and insight, AI can help you process it faster, but it can't fill a void or replace the judgement of knowing what matters. If it's poor collaboration, no AI tool can fix that.
 - **Where can AI free up time?** What work really doesn't need your time and attention, or needs less of your time and attention? Is compromising on this going to give you more time, energy and attention to work on something more impactful?
- **Create your own policy.** Create an AI policy or guidelines for your content team to provide guardrails on what AI should and should not be used for. This could be a good option if your organisation as a whole is not providing AI policy, or if you have sensitive, high-profile, or high-risk content that needs to be handled carefully

- **Consider the risks of using AI in ‘blank page’ scenarios.** By ‘blank page’, I mean situations where you don’t have previous material, in-house expertise, or precedent for how to solve the problem. In those scenarios, will you be able to effectively prompt, check any outputs, and maintain the results? If not, is using AI too high risk? This is a particular risk if you use AI to create interactive content.
- **Make sure you know what’s going on with AI in your team.** Set an expectation that people should declare when they use AI, but also create a sense of safety in sharing. In order to share freely, people need to know they will be met with support, curiosity and care, rather than judgement. (Giving people clear parameters on the things they shouldn’t use AI for is necessary to underpin this.)
- **Look at ways to use AI as a team.** Explore how AI could be beneficial as part of your established practices and process. That might look like:
 - Building a library of tried, tested and refined prompts
 - Developing your own agents, prompts, and skills
 - Using AI to support your existing processes.
- **Set the expectation that every single AI output should be checked.** As well as this, give people some ideas on how to go about it. For example, if you use AI to:
 - transcribe a call – read the transcript and correct it if necessary before you share it
 - summarise a document or write up notes – check that all quotes can be found in the materials, and check c.10% of notes against the originals to make sure that no errors are appearing
- **Don’t let asking AI become the default way you get an answer.** AI isn't the best option for every kind of question. Keep having human conversations – with managers, with peers, with external contacts – to expand your thinking.

Other ways that you could use AI in your content team

- **Creating formulaic, repetitive content at scale.** In product content design, teams are increasingly creating and managing large volumes of content with AI. Structured, formula-driven copy – things like button text, error messages, notifications – can be created by AI that’s been fed with system instructions and examples. Look at your content ecosystem – what repetitive content do you have, and how much time does it take? Could building an agent help speed up the process of creating event listings, service directory entries, or something else?

- **Checking content quality against guidelines and standards.** You could build a custom agent that checks drafts against your content guidelines – voice, tone, style, accessibility, etc. Would this help colleagues outside the content team develop better drafts? Could it reduce the amount of time your team has to spend on checking content?
- **Converting, localising and translating content:** Using AI to convert content into easy read/audio formats or translating into different languages. Would this help you to deliver content to more users? Would it speed up the process so that staff only have to review and refine? Could it help you meet a user need you couldn't otherwise meet? Is the quality acceptable?
- **Building the things you don't have the budget or capability to build:** Using AI's coding ability to create content like interactive tools or visualisations. Would this enable something you couldn't otherwise afford to do? Can AI add skills your team lacks? Could it help you meet a user need you couldn't otherwise meet?
- **Helping people who struggle with writing to draft or contribute to content:** Speaking to an AI tool or answering a list of programmed questions can be easier and more accessible than writing for some people (especially for some neurodiverse and disabled people). Could AI help you get subject matter expertise more easily? Could it make it easier to get inputs from people with lived experience?

CASE STUDY

RNIB

The RNIB is using AI to convert content into accessible formats like Braille, audio and large print. It has an AI solution called Mailings that automatically extracts and converts content, before a member of staff reviews it for accuracy. It has cut conversion time for content from as long as two weeks to as little as 3 hours.

Source: [Royal National Institute of Blind People & Microsoft](#)

3. The ethics of AI for charities

RESEARCH QUESTIONS

- Do charity content practitioners have ethical concerns about AI?

- Does AI pose any unique ethical questions for the charity sector?
- Are charities thinking about the ethical implications of using AI and how it might interact with their mission or values?

FINDINGS

Ethical concerns are near-universal

- 98% of respondents said that they have ethical concerns about AI.
- The most common concerns – shared by almost all participants – are:
 - environmental impact (79%)
 - accuracy (79%)
 - ethical concerns about the AI industry itself (79%)
 - originality (79%)
- The average participant shared 7 ethical concerns.

People have complex, conflicting feelings about AI

- 51% said that they feel satisfaction – when AI works well. This was the most common response.
- 49% said they feel frustration – when outputs aren't good enough.
- 48% said they feel guilt or ambivalence about using it.
- 27% reported 'a sense of pressure or obligation to use it', and 33% 'concern about what it means for my role or profession'.

'I feel like I need to use it to stay relevant, but every time I do I feel guilty – even though I'm not using it that much at all!'

'It has made me lose the love I had for my job. Everything is about AI now – using it to be quicker and create more than humanly possible.'

DEEP DIVE

Exploring the depth of negative and conflicting feelings about AI – and the unique concerns for charities

The study revealed a depth of negative and conflicting feelings about using AI, and an angle which is unique to the charity sector.

There's a high level of concern about AI's energy, water and data-centre footprint. This was the top ethical issue independent of role, seniority or overall enthusiasm for AI. The free-text answers throughout the survey were also dominated by environmental language even when the question wasn't specifically about ethics. It's worth noting that the survey ran in July 2026 – when the climate and nature catastrophe became extremely visible in the UK through heat waves, drought and wildfires. On a personal level, using AI at this moment in time has felt incredibly jarring.

There's a reasonable chance that AI may become less resource intensive and environmentally damaging over time. Energy is a cost to AI companies, and they want to make as much money as possible, so they will look for ways to reduce this. But it would be naive to think that efficiency gains alone could solve the problem. [A UN report](#) found that by 2030, global data centre electricity use will be nearly triple the combined annual electricity use of Pakistan, Bangladesh, and Nigeria (home to more than 650 million people). And the associated water footprint will equal the annual domestic water needs of all 1.3 billion people in Sub-Saharan Africa. The researchers also found that as models become more efficient, they become cheaper and are used more frequently, so the resource reduction from efficiency is cancelled out by volume growth.

The environment was not the only concern. There were plenty more:

- **Intellectual property (IP) and copyright:** 54% cite 'copyright or intellectual property' as a concern they have about AI. Many respondents characterised AI companies as profiting from uncredited, unpaid use of others' creative and intellectual work. *'AI steals human-made content, art, music, (my brain cells) and tries to sell it back to us, creating a profit for the few off the backs of the many'*
- **Bias and reinforcing stereotypes:** This is a concern for 63% of respondents. *'It worries me when people think that it isn't biased because it's looking at everything on the internet, but everything on the internet is biased'*
- **Eroding judgement:** The responses also showed a fear that habitual AI use will erode people's thinking, writing skills or professional judgement over time. This was frequently expressed as a concern for others (juniors, students, the next generation) rather than for the participants themselves. *'It's making the job a little too easy. I may have already outsourced my brain! I used to be a really good writer... who knows now!'*

Most interestingly, 80% of survey respondents think that using AI in charity content work raises concerns that are specific to the sector – the charity context specifically

sharpens their unease, on top of whatever concerns they'd have about AI regardless.

'We work on a lot of [campaigns about]... big tech and how it affects human rights, democracy and the environment... we should really be practicing what we preach.'

'We publish research about the consequences of AI, and campaign against the use of AI in mass-surveillance, processing applications for social care and other state operations. As such, our natural position is to not use it at all ourselves. However, there are conversations about identifying ways AI can be used to be more efficient, while also confidently avoiding the risks we call out in our external campaigning.'

'Charities exist to improve things in some way. How can charities balance that with using a tool that could be accelerating environmental changes in a negative way?'

For some respondents, the issue is that using a technology associated with environmental harm, IP extraction or corporate power clashes with what their organisation exists to do. This is different from general environmental or IP-theft concerns covered elsewhere in the study – it's specifically about internal consistency between mission and practice. For most, this is a personal concern rather than a formal position for the organisation. But at least two organisations in the interview sample have built mission/values considerations explicitly into how they govern AI.

Several respondents framed the AI risk for charities in terms of trust. Charities' ability to fundraise, deliver services and make change depends on public trust, and they see AI as a threat to that.

'Trust and authenticity are central to donor giving – AI content can be in danger of eroding all that.'

'It would only take one scandal about misuse or a data leak via AI to make a lot of people switch off from supporting charities.'

'Our organisation's whole selling point is that we are independent and trustworthy. If AI is seen as fake or downright lying to people, then using it hurts us.'

It seems like content teams might be feeling this more keenly than their leaders. Only 15% say that they think the ethical implications of AI have been explicitly considered and addressed; 28% think it hasn't been considered at all.

ANALYSIS

Why aren't ethical concerns leading to behaviour change?

The survey responses showed a high level of all-round concern about AI ethics, but not a lot of action happening as a result. Or to put it another way: huge amounts of cognitive dissonance.

People (including me) are using these tools even though we know they clash with our values, personally and organisationally. We feel excited about the potential, while knowing that the outputs are often poor quality and cannot be used unedited and unquestioned. The group of people who do not use AI because of their moral objections is a very small one. It's increasingly hard to resist the allure, the pressure and the hype.

One diary study respondent flagged it as a concern without an impact on their behaviour: *'It doesn't impact my day to day decisions, it's more of a background concern.'* No participants mentioned using more ethical AI options, like privacy-focused tools or less energy-intensive self-hosted AI options. It's particularly interesting that this is true even of respondents from the organisations that acknowledged a specific tension between AI and their mission, or those that had built considerations about their mission and values into how they govern AI.

This moral disengagement from the ethics and impact of AI feels like a risk. It has the potential to damage the morale of practitioners and even erode the culture of organisations if it's not addressed.

TAKEAWAYS

Ethical ways to approach AI usage

- **Know the issues.** Think about the limitations of generative AI and the impact that it has before you use it:
 - Environmental impact – generative AI tools use substantial natural resources
 - Reinforcing bias and stereotypes – AI can reproduce biases and stereotypes in its training data
 - Data privacy and security – AI tools rely on gathering large amounts of data and may threaten user privacy
 - Human exploitation – some AI companies have exploited workers as part of developing their models
 - Other AI applications – some AI companies do more than provide generative AI tools for the public, they might also provide surveillance tools, military intelligence, and other applications.
 - Originality and intellectual property – AI content is built on other people's work, who will not be credited or compensated
 - Accuracy – AI might produce incorrect information
- **Use your influence.** If you're reading this report, you're probably not the person in charge of IT procurement for your organisation. But you can still use your influence by sharing your ethical concerns, and getting conversations started about more ethical generative AI options.
- **Explore more ethical tools.** There are generative AI tools on the market that are aiming to provide a more ethical and sustainable experience,* for example:
 - [Mistral Vibe](#): EU based, GDPR compliant
 - [Lumo by Proton](#): EU based, privacy focused, open source
 - [Thaura](#): EU based, renewable energy, independently owned
 - [Eurix](#): Switzerland-based, privacy focused, renewable energy, GDPR compliant
 - [Green PT](#): EU based, privacy focused, renewable energy, GDPR compliant, open source
 - [BrowserOS](#): open source, privacy focused
 - [Worthwhile Chat](#): built and governed by non-profits (CAST and The Developer Society), UK and EU-based, privacy focused. This is in beta as of September 2026 and not publicly available.
 - **Self hosting**: You can self-host a large language model via a tool like [Ollama](#) or [Open WebUI](#), which gives a high level of privacy, control and ownership, as well as lower energy consumption.

- **Only use AI when there's a real reason to.** Use AI mindfully, not as a default. GOV.UK's [Using AI ethically and sustainably](#) puts it well: 'There are many ways to do a task, and you may not need to use AI. Consider if a simpler tool could do the job (for example, a spreadsheet or search engine). If you do need AI, if your model allows it, pick the smallest one which will do what you need.'
- **Consider the cost and whether it is worth it.** [Vu Le](#) has written a very helpful list of questions to help the sector think through the cost of using AI. How are we:
 - harming marginalised communities?
 - traumatising people, especially women of color in poorer countries?
 - supporting fascism without realising it?
 - contributing to the entrenchment of racism and white supremacy?
 - destroying the livelihoods of artists?
 - creating a more egotistical, sycophantic, narcissistic society?
 - enshittifying ourselves and our world?
 - perpetuating the injustice that we are trying to fight?
- **If you feel guilt, shame or conflict, talk about it.** Carrying something like this alone is heavy. If you feel safe to do so, talk to your team and leaders. They may be feeling the same, and it may turn into something productive.

*Ethics can mean a lot of different things. Here, I'm thinking about things like environmental sustainability, independent ownership, GDPR compliance, transparency, not being US-based and therefore not subject to the CLOUD Act, and I've tried to explain why the tools have been listed.

4. Organisation-wide policy, leadership, and support on AI

RESEARCH QUESTIONS

- Do charities have AI policies, and what approach do they take?
- Are they providing training and guidance?
- Do content practitioners feel supported in using AI?

FINDINGS

Over half do not have a formal AI policy

- Only 38% of respondents said their organisation has a formal written policy on AI.
- A majority (56%) have either no policy or one that's still in progress:
 - 13% have informal guidance only
 - 18% said a policy is in development
 - 18% said there's no policy but conversations are happening
 - 7% said it hasn't come up at all

87% have not had any formal training on AI

- Only 13% of respondents said that their organisation has provided formal training on how to use AI tools.
- 42% said that there has been some informal guidance or resources.

Half do not feel supported when it comes to AI

- 49% feel somewhat or very unsupported by their organisation in navigating AI in content work.
- Only 5% said that they feel very supported by their organisation.

DEEP DIVE

Peer communities of practice are an excellent source of support and learning

A small number of participants mentioned that they are part of communities of practice groups focused on AI. These vary from small, informal peer-led groups, to formal organisation-wide cohorts. Where they exist, communities of practice were almost universally-praised. People credited them with:

- surfacing what's actually working when it comes to AI
- creating a space where fragmented individual practice can evolve into shared team ways of working
- reducing anxiety, shame and isolation around AI use

For example, one participant shared that they are part of a peer-led community of practice in their organisation's marketing and communications department. The members of the community meet regularly to share learnings, prompts and use cases. They have found the group to be a useful forum for hearing different perspectives on AI, as well as comparing their own learnings.

Another participant who is part of two different AI communities of practice in their organisation, said that they value these spaces because they allow people to discuss the benefits, risks, and concerns regarding AI, helping to avoid 'spiraling about doom and gloom' or adopting an overly optimistic, 'tech bro' mindset.

[CAST's AI survey 2026](#) also found that '86% of respondents wanted access to peers within the sector to discuss AI' and '62% of those who responded... wanted to see coordination between charities'.

ANALYSIS

Governance often feels tick-box or disconnected from practice

AI governance maturity seems to vary widely across the sector. In the study, respondents reported everything from having no formal policy at all, to generic guidance that's rarely referred to in practice, to bespoke, consultative policies developed with staff input. There were even instances where, within a single organisation, the team-level policy was more developed than the organisation-wide one it sits under.

The study echoes the findings of the [Charity Digital Skills Report](#):

- 44% of charities are not taking action to manage risks and progress with AI
- 51% of charities do not have an AI policy
- 19% are reviewing data protection, security and safeguarding in order to progress with AI and manage risks
- 34% are making strengthening digital, data and AI governance and decision making a priority

One sub-theme was that, for some participants, their organisation's AI policy and governance feels very disconnected from the day to day ways AI is being used, and more like a box checking exercise than something that's usable.

'The policy we have was written without involvement from our digital teams... it feels more like a tick-box exercise.'

'The policy about AI use is very short and seems more concerned with security risks of using AI tools, rather than anything else.'

The study also found a few instances where people are using personally paid-for or otherwise unofficial AI tools and workflows because the official options haven't caught up with what they want to do. 18% of respondents say that they personally pay for AI tools they use at work (15% regularly, 3% occasionally), and a further 7% have done so in the past. This is notable given that 85% say their organisation provides access to at least one official AI tool (53% one tool, 32% a few). This 'shadow' use isn't because there's no tool in place, it's more about tool preference or capability gaps.

'I'll be honest, I cheat and I... pay for Claude... I can't believe I'm the only person who's using their own version.'

When there's a gap for governance, training and support on AI, individual actions and judgement fill that gap. People bring in their own tools, come up with their own rules, and set their own ethical boundaries. This isn't always a bad thing – there's plenty of strong individual practice around AI – but it comes with risks. The risks falls into two main categories:

1. **Exposure for the organisation:** there could be breaches of data and privacy, increased risk of sharing misinformation publicly, embarrassing values and mission clashes.

2. **Burden on individuals:** having to absorb tool costs (or not being able to afford tools that others can), concerns over making decisions that are too high level, cognitive surrender and deskilling.

Thoughtful, considered, practical governance is an enabler. It doesn't just help to mitigate the organisation against the risks, it gives individuals the foundations they need to build good practice on.

It's important to remember that we're still in the early days of generative AI. Governing AI well requires organisational capacity: time to evaluate tools, legal and IT input, and confidence about a technology that's moving fast and being marketed aggressively. Leaders are managing this alongside other competing pressures and priorities, and without some of the unique first-hand exposure that content practitioners have.

The gap between content practitioners' concerns and leadership engagement isn't a sign that leaders don't care. It reflects that this is hard to get right, and that most organisations are still working out what good governance looks like.

TAKEAWAYS

How to fill a governance gap around AI

- **Use your influence.** If your organisation doesn't have an AI policy in place, use your influence and speak to the person who's responsible for this. Explain the risks of not having a robust policy and the opportunities you can see in developing one.
- **Develop a policy or guidelines just for content.** You could create a policy or guidelines just for content. Even if your organisation has a policy in place, you might find that you need something that deals with content-specific scenarios, use cases and risks. [William Joseph's James Gadsby Peet has a very helpful presentation on developing an AI policy.](#)
- **Start or join a community of practice:** Start or join a community of practice for AI in your organisation, department or team. These groups can be a great informal governance mechanism, as well as a source of valuable, bolstering peer support. You could also look beyond your organisation for groups in the wider sector.

Starting an AI community of practice

A community of practice is a group of people connected by a shared interest in something they actually do – and who get better at it by learning from each other over time.

Three things make a community of practice distinct from a working group, a mailing list, or a team meeting:

1. **A shared domain.** Something everyone in the group has in common, whether that's a skill, a responsibility, or an area of work. In this case: AI.
2. **Genuine community.** Members who connect, share, and build relationships over time. A channel in Teams or a list of people with similar job titles is not enough on its own. The interaction is what creates the community.
3. **Practice.** Members are doing the thing, not just interested in it. Over time, they build up a shared repertoire: resources, ways of working, stories, tools, things that have and haven't worked. That collective knowledge is what makes the community of practice useful.

It's also worth being clear about what a community of practice is not. It is not a project working group. It doesn't have deliverables or targets in the conventional sense. It's not a consultation exercise, and it shouldn't replace existing team structures.

Questions to help you develop your community of practice

- **What is the domain?** What specifically is the shared practice you're bringing people together around? 'AI' is broad – is it all AI usage, or are you starting with a narrower focus, like AI and content, AI and comms, or AI in a specific directorate or project? Being clearer about the domain early helps people understand whether the community is for them.
- **Who is the community for?** Who across the organisation is using AI? Which teams, roles, or individuals are most likely to benefit from a space like this? Is there anyone who is already doing interesting things with AI that you could build around? Who do you most need to reach, and who might champion it within their own teams?
- **What problem does it solve for members?** Why would someone choose to come? What would they get out of it that they can't get elsewhere? The community needs to be genuinely useful to members. What's the honest answer to 'why should I make time for this?'


- **What does participation look like?** How often would the community meet? What format – a call, an in-person session, something async? What happens between meetings? Is there a channel, a shared space? How formal or informal should the tone be? Would all members attend the same thing, or are there different tiers of involvement? How do you track and share learnings?
- **Who facilitates, and how?** Who will lead? What does facilitation involve – preparing sessions, keeping a record, following up, growing membership? How do you avoid this becoming a task that's easy to deprioritise when things get busy? And how does facilitation gradually shift towards members over time?
- **What does success look like, and how will you know?** It's worth agreeing early on how you'll track whether the community is working – not through vanity metrics like headcount, but through signs of genuine value. Are members showing up consistently? Are they contributing, not just attending? Are they telling you it's useful? Are things from the community finding their way into people's practice? You don't need a formal measurement framework at the start, but having a rough sense of what 'working' looks like will help you make good decisions about the community as it develops.

Methodology

About the participants


- 74 total participants:
 - 63 survey respondents
 - 8 one-to-one interview participants
 - 3 diary study participants
- Participants were recruited via my email newsletter, posts on my social media accounts, LinkedIn ads to reach outside my network, and emails and social posts from William Joseph and other contacts
- All participants work for nonprofits, except for a small number (7 respondents) working for agencies/as consultants in the sector

How many people work in your organisation in total?




How many people work in your organisation in total?

How many people have content as part of their role in your organisation (including yourself)?



What is your organisation's annual income?



How I used (and did not use) AI in this project

It feels important to be transparent and disclose when and why I chose to use, or not use, AI in on this project:

- I used AI (Claude) to help me write the full survey questionnaire. I did this by entering a very detailed prompt, containing a list of the questions I wanted to ask and the topics I wanted to explore. AI's main role was to help me refine the wording, and generate the options for multiple choice answers. I reviewed the AI output word-by-word.
- I did not use AI to build the survey. I have tried this feature in Tally, my survey tool, in the past but have found that it makes mistakes and also that it's hard to understand if the output is reliable and the logic is correct until it's too late.
- I used AI (Google Gemini) to produce interview transcripts. I did a quick scan to review each transcript for overall accuracy. I did not use Gemini's interview summaries at any point, as in my experience they lack nuance.
- I did not use AI to do the initial analysis of the survey results. It felt like an important step in understanding to turn the raw data into stats and to read through all the free text responses. I did use Claude to do some more detailed analysis and data comparisons later in the process to save time. I checked each stat that Claude produced against the original data.
- I used AI (Claude) to speed up the coding of qualitative data. By coding, I mean assigning descriptive tags to text responses and interview transcripts so that I could get a picture of the key themes. I ran a spot check on c.10% of tags to check that they matched my understanding. I found Claude to be accurate based on this sample. I only felt comfortable with taking this step because I had carried out the interviews myself and read all the survey responses, which meant that I already understood what it contained.
- I used AI (Claude) to augment my own draft notes on what the trends in the data set were. I asked Claude to review the qualitative and quantitative data and identify the key trends. I then compared these to my own conclusions. This step added value, as Claude spotted a couple of patterns that I had not noticed. However, it also came up with a number of suggestions that I disagreed with and felt did not hold up to scrutiny.
- I did not use AI to conduct research, identify extra source materials, or do any of the background reading for me. Everything you see in the bibliography was something I found through my own research or something that someone else shared with me. I have not found AI useful as a research tool, and tend to find it easier to use a search engine and communities to find material. I do not tend to use AI to summarise reading material, as I think the summaries can be inaccurate and I need to read the material myself to understand it.
- I wrote the report myself. I did not use AI to draft, write, reword, or edit it. AI-generated wording may be present in small snippets:
 - Quotes – the transcripts were generated by AI
 - Survey question wording – AI was used to generate this wording

- Findings – I copied some wording from Claude’s trend analysis into my first draft, but this has largely been replaced and reworded
- I got feedback on my drafts from human beings. I made the edits based on this feedback myself.
- I used AI as a sounding board when I wrote the introduction. I yapped at Claude for about 10 minutes to get all my thoughts out, and asked it to turn that into a bullet point list. I used this as a guide when I was writing the introduction. I would rather have done this with a human collaborator.
- A human editor – Ruth Oliver – proofread and copyedited the report.

Reading list and resources

[AI and Accessibility: abdicating engagement?, UNESCO](#)

[AI and British Sign Language, RNID](#)

[AI and user behaviour: 8 months of meetups, Stephanie Coulshed, Scope](#)

[AI Cheerleading, AI Abstention and AI Redirection, Vanessa Andreotti](#)

[AI in Action: Insights from our AI Essentials charity bootcamps, Media Trust](#)

[AI Is Your Bridge to Smarter Content Operations, Content Science](#)

[AI Knowledge Hub, GOV.UK](#)

[AI Skills for Life and Work: General Public Survey Findings, GOV.UK](#)

[AI as a Healthcare Ally, OpenAI](#)

[AI-driven productivity gains enable more CO₂ emissions than they avoid in a global energy–economy model, Will Alpine, Nathan Geldner, Holly Alpine and Maksym G. Chepeliev, npj Climate Action, npj Climate Action](#)

[Australian influencer Lily Jay's tangled web of AI manipulation, ABC News](#)

[Can AI replace your content team? Hear from someone who tried \(and failed\), Dan Onken](#)

[CAST's AI survey 2026: All the results — and the support available, CAST](#)

[Charity Digital Skills Report, Zoe Amar, Nissa Ramsay](#)

[Claude Skills for content design: inside Intuit's AI build, Sarah Moes, UX Content Collective](#)

[EasyMaker: Easy Read Software to Make Documents, Photosymbols](#)

[Enabled emissions: How AI helps to supercharge oil and gas production, Global Witness](#)

[Exploring the role of AI in our design process at Co-op, Matt Tyas](#)

[Getting started is not getting it right, Nicole Alexandra Michaelis](#)

[How AI-ready content can boost policy effectiveness, Department for Business and Trade](#)

[How NN/G Uses AI in Its Editorial Process, Raluca Budiu, NN/G](#)

[How we stopped outsourcing decision-making to AI, Eva Ratcliffe, UX Content Collective](#)

[How we're preventing AI misinformation at DBT, Giorgio Di Tunno, Neil Starr, GOV.UK](#)

[Idea Exchange: Our AI solution helps improve independence of people with learning disabilities, Local Government Chronicle](#)

[Library of AI experiments, CAST](#)

[Plausible but Not Valid: A Psychometric Audit of LLMs as Synthetic Survey Respondents, Mantas Lukauskas, Viktorija Šarkauskaitė](#)

[Practical actions for the AI search era – how to drive traffic to your website, Madeleine Sugden](#)

[Royal National Institute of Blind People makes the UK more accessible with Microsoft Azure AI, Microsoft](#)

[Supporting your organisation with AI: A free self-serve course, CAST](#)

[The 2026 AI Index Report, Stanford HAI](#)

[The Invisible Charity – how the top 15 charities in the UK are doing in AI search, Platypus Digital](#)

[The Web Is Being Made Accessible for AI, Not People, TechPolicy.Press](#)

[Thinking—Fast, Slow, and Artificial: How AI is Reshaping Human Reasoning and the Rise of Cognitive Surrender, Steven D Shaw and Gideon Nave, The Wharton School \(SSRN\)](#)

[Top Ten: 10 ways the sector is experimenting with AI, Ellen Smyth, CAST](#)

[Trauma-informed content and AI: Amplifying empathy at scale, Adrie van der Luijt, Button](#)

[Unlawful by design: Exposing the human rights costs of generative AI, Amnesty International](#)

[UN Scientists: AI Is Threatening Natural Resources for Billions, UNRIC](#)

[Update: AI Overviews Reduce Clicks by 58%, Ahrefs](#)

[Using AI ethically and sustainably – AI Knowledge Hub, GOV.UK](#)

[Walton Family Foundation Gen Z Research Hub, Gallup](#)

[We need to have much more serious conversations about AI and the nonprofit/philanthropic sector, Vu Le, Non-profit AF, Vu Le, Non-profit AF](#)

[What Makes Content Operations Successful? Full Report 202%, Content Science](#)

[Why your charity website needs to be a museum, not just a library, James Gadsby Peet, William Joseph](#)

Reading list and resources

[AI and Accessibility: abdicating engagement?, UNESCO](#)

[AI and British Sign Language, RNID](#)

[AI and user behaviour: 8 months of meetups, Stephanie Coulshed, Scope](#)

[AI Cheerleading, AI Abstention and AI Redirection, Vanessa Andreotti](#)

[AI in Action: Insights from our AI Essentials charity bootcamps, Media Trust](#)

[AI Is Your Bridge to Smarter Content Operations, Content Science](#)

[AI Knowledge Hub, GOV.UK](#)

[AI Skills for Life and Work: General Public Survey Findings, GOV.UK](#)

[AI as a Healthcare Ally, OpenAI](#)

[AI-driven productivity gains enable more CO₂ emissions than they avoid in a global energy–economy model, Will Alpine, Nathan Geldner, Holly Alpine and Maksym G. Chepeliev, npj Climate Action, npj Climate Action](#)

[Australian influencer Lily Jay's tangled web of AI manipulation, ABC News](#)

[Can AI replace your content team? Hear from someone who tried \(and failed\), Dan Onken](#)

[CAST's AI survey 2026: All the results — and the support available, CAST](#)

[Charity Digital Skills Report, Zoe Amar, Nissa Ramsay](#)

[Claude Skills for content design: inside Intuit's AI build, Sarah Moes, UX Content Collective](#)

[EasyMaker: Easy Read Software to Make Documents, Photosymbols](#)

[Enabled emissions: How AI helps to supercharge oil and gas production, Global Witness](#)

[Exploring the role of AI in our design process at Co-op, Matt Tyas](#)

[Getting started is not getting it right, Nicole Alexandra Michaelis](#)

[How AI-ready content can boost policy effectiveness, Department for Business and Trade](#)

[How NN/G Uses AI in Its Editorial Process, Raluca Budiu, NN/G](#)

[How we stopped outsourcing decision-making to AI, Eva Ratcliffe, UX Content Collective](#)

[How we're preventing AI misinformation at DBT, Giorgio Di Tunno, Neil Starr, GOV.UK](#)

[Idea Exchange: Our AI solution helps improve independence of people with learning disabilities, Local Government Chronicle](#)

[Library of AI experiments, CAST](#)

[Plausible but Not Valid: A Psychometric Audit of LLMs as Synthetic Survey Respondents, Mantas Lukauskas, Viktorija Šarkauskaitė](#)

[Practical actions for the AI search era – how to drive traffic to your website, Madeleine Sugden](#)

[Royal National Institute of Blind People makes the UK more accessible with Microsoft Azure AI, Microsoft](#)

[Supporting your organisation with AI: A free self-serve course, CAST](#)

[The 2026 AI Index Report, Stanford HAI](#)

[The Invisible Charity – how the top 15 charities in the UK are doing in AI search, Platypus Digital](#)

[The Web Is Being Made Accessible for AI, Not People, TechPolicy.Press](#)

[Thinking—Fast, Slow, and Artificial: How AI is Reshaping Human Reasoning and the Rise of Cognitive Surrender, Steven D Shaw and Gideon Nave, The Wharton School \(SSRN\)](#)

[Top Ten: 10 ways the sector is experimenting with AI, Ellen Smyth, CAST](#)

[Trauma-informed content and AI: Amplifying empathy at scale, Adrie van der Luijt, Button](#)

[Unlawful by design: Exposing the human rights costs of generative AI, Amnesty International](#)

[UN Scientists: AI Is Threatening Natural Resources for Billions, UNRIC](#)

[Update: AI Overviews Reduce Clicks by 58%, Ahrefs](#)

[Using AI ethically and sustainably – AI Knowledge Hub, GOV.UK](#)

[Walton Family Foundation Gen Z Research Hub, Gallup](#)

[We need to have much more serious conversations about AI and the nonprofit/philanthropic sector, Vu Le, Non-profit AF, Vu Le, Non-profit AF](#)

[What Makes Content Operations Successful? Full Report 202%, Content Science](#)

[Why your charity website needs to be a museum, not just a library, James Gadsby Peet, William Joseph](#)