



Education
Comms
Network



AI IN EDUCATION COMMUNICATIONS

a practical guide

July 2026

Edited by Anna Pedroza
and Susan Higgins



Making Education Relevant



About the Education Comms Network

The Education Comms Network is a mutually supportive network of over 150 marketing and communications professionals working across the education sector. It was founded in November 2020 with support from the Edge Foundation. The network meets once a term online to learn, hear from experts and share their work. Through the Education Comms Network, Edge is helping support and improve the overall effectiveness of education comms teams across the UK. By supporting this report on AI in education communications, we're helping to spark vital conversations, bridge the gap between comms, policy and CPD in education, and strengthen a community that learns, reflects and grows together.

About the Edge Foundation

The world is changing fast and education needs to keep up. Edge is an independent, politically impartial education foundation. We want education to be relevant to the twenty-first century. We gather evidence through research and real world projects and partnerships and use this to lead the debate and influence policy and practice. Edge believes all young people need to be equipped with the skills that today's global, digital economy demands, through a broad and balanced curriculum, high quality training, engaging real world learning and rich relationships between education and employers.

Visit www.edge.co.uk to find out more

The information and views set out in this report are those of the contributor(s) and do not necessarily reflect the official view of the Edge Foundation. Edge do not guarantee the accuracy of the data included in this study. Neither Edge nor any person acting on its behalf may be held responsible for the use which may be made of the information contained therein.

CONTENTS

About the editors	4
About the contributors	5
Acknowledgements	7
Foreword	8
Introduction and welcome	9
Useful terms	10
AI and security	12
Ethical use of AI in communications and PR	18
Case study: AI and content creation	23
Case study: AI and marketing intelligence	27
Case study: generative engine optimisation	32
Case study: AI as a thinking partner	37
Case study: AI supporting everything	41
AI, journalism, and the changing dynamics of influence	46
Thinking strategically about AI in education communications	52
Further reading and resources	60

ABOUT THE EDITORS

Susan Higgins is a senior communications and education media relations specialist. In her role as Director of Communications at the Edge Foundation, she combines strategic thinking with creative, practical delivery. Susan co-founded and leads the Education Communications Network, now a community of more than 150 senior comms and marketing professionals.

Her career spans in-house, consultancy and agency roles, with earlier work across global beauty, hospitality and luxury brands. Susan has served as a school governor and as a trustee of Wivenhoe House Hotel, home to the Edge Hotel School. Raised in Spain and educated in the UK, she brings an international outlook shaped by a career that has had storytelling, influence and impact at its heart.



Anna Pedroza co-founded the Education Comms Network with Susan and helps to co-ordinate it. She's an experienced communications professional who has worked in the education sector for 20 years. She runs Pedroza Communications, a specialist consultancy which helps schools, charities, membership organisations, public bodies, and businesses to communicate their message effectively and with impact.

Anna has held various senior communications roles including Director of Marketing and Communications at the award-winning service Teachers TV and Director of the Skills for Life Development Centre at the Education Development Trust. She's a Chartered member of the CIPR, sits on the CIPR's Education and Skills Committee, and is a member of the advisory group for the Schools and Academies Show.





ABOUT THE CONTRIBUTORS

Sophie Bessemer has led the content, SEO and organic function at Third Space Learning for over 10 years, building and sustaining a customer base of more than 4,000 primary and secondary schools across the UK and US. She has years of experience in educational publishing, digital content strategy and search, with a track record of delivering timely, purposeful content that supports teachers and leaders throughout their career. Since the start of the 2025 academic year all Third Space Learning's online maths tutoring has been led by Skye, the AI tutor trained by the TSL team of maths teachers.

Rae Claramunt is HR & Communications Manager at Three Spires Trust which she joined in September 2024 after 4 years at St Regis CofE Academy as PA to the Principal. Her role evolved to become the Leadership & HR Support Officer at St Regis before joining the Trust central team. Whilst working at St Regis, Rae completed her MA in Literature and Culture and her L3 Certificate in People Practice.

Rae has avid experience in rebranding academies, organising marketing strategies and logo design. Now, Rae supports the Trust's digital presence, brand strategy and communications to grow engagement, partnerships and community reach. Rae manages the Trust's social media channels, website along with channels of specific academies and initiatives. Rae has a passion for creativity and driving positive change, playing a role in the marketing of the NextGen Curriculum at Three Spires.

Danni Fothergill is Head of Marketing and Communications at Edurio, where she leads brand, content and campaign strategies to help school trusts and schools make better decisions through stakeholder insight. With over 15 years' experience in education marketing, she focuses on turning complex data and research into clear, practical communications that support school and trust leaders.

Luke Kemper is Lead for Insight and Intelligence at Haringey Education Partnership (HEP) where he's worked for the past four years. In the past two years Luke has had an increasing focus on helping HEP schools with AI. This has included providing information, developing and delivering training, and providing strategic advice. Prior to joining HEP, Luke worked as a university lecturer and a teacher in Poland and China, and this also inspired him to secure an MPhil from Cambridge in Education, Globalisation and International Development.

Nic Newman is Senior Research Associate at the Reuters Institute for the Study of Journalism and is an expert in tracking and interpreting trends in the media. He has been lead author of the annual Digital News Report since 2012 and he's also produced a media and journalism predictions report for around 15 years. Nic is also a consultant on digital media, working actively with news companies on product, audience, and business strategies for digital transition.

Earlier in his career, Nic was a founding member of the BBC News Website, leading international coverage as World Editor (1997–2001). As Head of Product Development (2001–10) he led digital teams, developing websites, mobile, and interactive TV applications for all BBC Journalism sites.

Ben Verinder is co-editor of 'AI for PR', published by Kogan Page in May 2026. He is an agency founder and a strategic communication and public relations adviser. He works with boards, senior leadership teams, and communication practitioners on reputation, crisis and issues management, and the integration of AI into communications practice.

He has written extensively on AI in public relations and is the creator of the 'AI and Reputation Leadership' course.

Ben is one of a small group of Founding Chartered public relations practitioners globally. He delivers training and consultancy across the UK and internationally, supporting organisations in adopting AI in a way that strengthens capability while managing associated risks.

Stephen Waddington is a professional adviser at Wadds Inc. and a PhD researcher at Leeds Business School. His work focuses on management, corporate communications and public relations as a management function.

He has held senior agency leadership roles and advised organisations including The Economist, HM Government, Intel, Microsoft, the Press Association, the NHS, and Tesco.

Stephen is a former CIPR President and a founding member of the CIPR AI in PR panel. He is also co-founder of Socially Mobile, a social enterprise supporting public relations practitioners from underrepresented and lower socio-economic backgrounds.

He has written or edited more than 50 publications on public relations and management.



ACKNOWLEDGEMENTS

We would like to express our thanks to all the people who have contributed to this project. Sharing their time, expertise and experience has created a report that is both insightful and practical.

Thanks to members of the Education Comms Network who provided insights, completed our survey, and were generally enthusiastic.


We would also like to specially thank our advisory group who helped us to plan the report and shape the content, as well as providing a valuable sounding board.

Rae Claramunt – HR & Communications Manager, Three Spires Trust

Jennie Connelly – Engagement Manager, Capita Experience

Katy Gandon – Independent comms professional

Laura Hewitt – Independent comms professional



FOREWORD

The first time I asked an AI tool for help with a piece of comms work, I felt slightly guilty. Like I was cheating. A few months later, I felt slightly silly for not using it sooner.

I started my comms career in education, in the days when our biggest tech debate was whether a website needed a Flash intro. So, this report is somewhat personal to me because in my experience, education comms professionals are some of the most curious individuals I've met.

What you'll read here is a profession working out a new technology in real time. Honest case studies, careful thinking on strategy and ethics, hard-won lessons on security, and a clear look at how the journalism we rely on is shifting under our feet.

The case studies, in particular, show how much graft goes into good output. The contributors haven't just stumbled into clever AI use; instead, they've built it, prompt by prompt, policy by policy, mistake by mistake. And that kind of honesty is rare and worth some of your time.

A common thread running through the whole report is the difference between using AI to get work done and having a strategic framework for how your organisation uses it at all. We all know that tactical wins are easy and shareable. Strategic thinking, about governance, security, ethics, the environment, reputation, and the apprenticeship of our junior team members, is harder, slower, but far more important.

There are a few other things I'd like to highlight in this report, which I think are incredibly important for us to pay attention to:


- **Hope over fear.** AI is here, and it isn't going anywhere. Used well and used ethically, there are extraordinary things our profession can do with it, especially because comms and marketing trade in the very things AI struggles with, such as empathy, judgement, context, and the human bit. Lean into that, don't run from it.
- **Challenge the machine, don't just talk to it.** AI is trained to make you feel clever and boost your ego. If you don't push back on it, you end up in a slightly more polished conversation with yourself. The best users edit, question, and refuse to take the first answer.
- **Don't wait for permission.** If you're waiting for the AI expert in your organisation to hand you the strategy, you're already five steps behind. The people doing this well are users first and strategists second. They understand the work, so they understand where AI fits and, just as importantly, where it doesn't.

Wherever you are on this journey, just starting out, quietly experimenting, or already deep in agents and workflows, there is something in this report for you. Read it as permission not to fear failure and to try something new, see this report as a nudge to keep going.

Education comms has always been generous in sharing what they learn. I hope you take full advantage of this report, which I know will make you a little braver, as it did me.

Advita Patel

Founder, CommsRebel and Author of Decoding Confidence: The 7 Habits of Confident Leaders. Former President of the Chartered Institute of Public Relations 2025.





INTRODUCTION AND WELCOME

In education, public debate around AI has focused on issues like plagiarism and cheating in exams, and its potential to change assessment or pedagogy. However, for comms and marketing professionals working in education, the use of AI is also becoming an increasingly important consideration. This report aims to help comms and marketing people to:

- Be aware of the security implications
- Understand the ethical issues of using AI
- Discover the approaches colleagues are already using in their day-to-day work
- Appreciate how AI is impacting journalism
- Consider the planning and strategic implications of AI

Understanding the security implications of AI is a concern for us all and the chapter on this covers the interconnectedness of security issues, types of AI and the security implications, policies and whether we can ever fully trust AI.

Closely linked to aspects of security, the next chapter focuses on the ethical considerations needed to use AI responsibly. This covers transparency and explainability, bias and misinformation, and the environmental impact of AI.

The largest section of this report is given over to case studies written by members of the Education Comms Network. These provide practical insights into a wide range of ways AI is being used to support day-to-day work. Each case study explains how AI

has helped with an issue, the approach taken, the tools used, and the outcomes.

The report also examines how AI is changing journalism and the dynamics of influence. Drawing on recent research it highlights how the media is adjusting to AI and what this means for the content they produce. It looks at agentic AI which publishers predict will have a large impact on how the news is accessed.

The report finishes by discussing planning and strategy and highlights that, although there is a lot of tactical use of AI, there isn't the governance or strategic oversight that is also needed.

This report is unashamedly about now and we acknowledge that the technology, and how people are using it, is continuing to evolve. It provides a snapshot of how education comms and marketing professionals are using AI (in spring 2026), and our current thinking. We hope this provides a helpful benchmark from which you can develop and progress your own use of AI.

Susan Higgins and Anna Pedroza



USEFUL TERMS

Artificial Intelligence (AI) – An AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments. Different AI systems vary in their levels of autonomy and adaptiveness after deployment. (Please note: It is hard to find a single definition of AI. This is one adopted by the UK Govt and is based on the OECD's definition.)

Agentic AI – An AI system that can take actions, use tools, and complete multi-step tasks autonomously, beyond just responding to a single prompt to pursue a goal over a sequence of steps.

Bias – Because AI is built by humans, and trained on human-generated data, it is susceptible to systematic human and statistical skews. These typically reflect imbalances or historical inequalities affecting language, recommendations, assessments, and decisions in ways that disadvantage certain groups.

Context window – The amount of information an AI model can "remember" or hold in its memory during a single conversation

Explanability – The ability of AIs to describe in simple terms a human can understand how an AI system arrived at a specific output or decision. It is a legal and ethical requirement for AI systems used in high-stakes decisions in many jurisdictions, including under GDPR Article 22.

Frontier AI – Defined by the Government Office for Science as 'highly capable general-purpose AI models that can perform a wide variety of tasks and match or exceed the capabilities present in today's most advanced models'. Currently, this primarily encompasses a few large language models

Generative AI – An AI model that generates text, images, audio, video or other media in response to user prompts. It uses machine learning techniques to create new data that has similar characteristics to the data it was trained on. Generative AI applications include chatbots, photo and video filters, and virtual assistants.

Generative engine optimisation (GEO) – The practice of creating and editing digital content and online presence so it is more likely to be used and cited in responses generated by AI tools such as ChatGPT, Google Gemini, Claude, Perplexity AI and Microsoft Copilot. GEO helps organisations appear in AI-generated answers, increasing visibility and reach. Related terms include answer engine optimisation (AEO).

Hallucinations – Large language models, such as ChatGPT, are unable to identify if the phrases they generate make sense or are accurate. This can sometimes lead to inaccurate results, also known as 'hallucination' effects, where large language models improvise, generating plausible sounding but inaccurate content.



Large Language Models – A type of foundation model that is trained on vast amounts of text to carry out natural language processing tasks. During training phases, large language models learn parameters from factors such as the model size and training datasets. Parameters are then used by large language models to infer new content.

Narrow AI – Sometimes known as weak AI, these AI models are designed to perform a specific task (such as speech recognition) and cannot be adapted to other tasks. See also general-purpose AI.

Prompt – A prompt is an instruction, question, or input you give to an AI system. The quality and specificity of the response an AI gives you significantly depends on the quality of your prompt.

Search Engine Optimisation (SEO) – The practice of improving a website's visibility in unpaid (organic) search results on search engines such as Google or Bing. It involves making content clear, relevant and easy to navigate, with the aim of helping pages rank higher and be found more easily by users.

Training – The process by which an AI model learns patterns from a large dataset. Training happens before the model is deployed and shapes everything it knows and how it behaves.

Training data – The body of text, images, or other content used to train an AI model. The composition of training data is responsible for the model's knowledge and potential biases

REFERENCES

Artificial Intelligence Glossary, UK Parliament 2024 – <https://post.parliament.uk/artificial-intelligence-ai-glossary/>

Artificial Intelligence Playbook for the UK Government, 2025 – <https://www.gov.uk/government/publications/ai-playbook-for-the-uk-government/artificial-intelligence-playbook-for-the-uk-government-html#acknowledgements>

Defining data science and AI, The Alan Turing Institute, accessed April 2026 – <https://www.turing.ac.uk/news/data-science-and-ai-glossary>

AI Glossary: Vocabulary you actually need, Digital Marketing Institute, April 2026

Wikipedia – https://en.wikipedia.org/wiki/Generative_engine_optimization

ChatGPT


AI AND SECURITY

Luke Kemper, Lead for Insight and Intelligence,
Haringey Education Partnership

In September 2025, OpenAI stated that around ten percent of the world's adult population was using ChatGPT. One month later, TeacherTapp reported that the number of teachers who said they used AI for their work jumped up to a majority for the first time, clocking in at 58 percent. And while numbers of adult users have been steadily increasing, almost half of all children were already using AI tools since last year, according to Ofcom in May 2025.

With these figures, **it is safe to say that AI tools are now mainstream. But that doesn't mean it is mainstream practice to use AI in a safe and secure manner.** Whether for drafting press releases, writing social media copy, summarising policy documents, or creating a marketing plan, the temptation to reach for an AI tool is understandable and, in many cases, sensible. But every interaction with an AI tool carries security implications that education communications professionals must understand.

Like everything in the AI world, the following advice will need to be adapted as AI models and tools themselves change over time. For now, it addresses the most common AI security risks and countermeasures, but cannot anticipate the risks that full agentic AI (explained later) and other future developments may bring.



**Every interaction
with an AI tool carries
security implications
that education
communications
professionals must
understand.**



Key domains of AI security

Security in the context of AI is not a single concept. It spans several interconnected domains, each of which requires attention:

Security domain	What it means for education comms
Data protection & privacy	Ensuring personal and sensitive data entered into AI tools is handled in compliance with UK GDPR and the Data Protection Act 2018. This includes anonymising inputs and understanding whether data is stored or used for model training.
Content accuracy & reliability	AI outputs can contain hallucinations, biases, and factual errors. Any AI-generated content used in communications must be verified by someone with knowledge of the subject matter before publication.
Access & authentication	Controlling who can use AI tools within your organisation, with what permissions, and ensuring that login credentials and API keys are managed securely.
Intellectual property & copyright	AI-generated content occupies a grey area in copyright law. Organisations must consider ownership of outputs and the risk that training data may include copyrighted material.
Reputational risk	AI-generated communications that contain errors, biased language, or inappropriate content can cause significant reputational damage.
Supply chain & third-party risk	EdTech tools built on top of frontier AI models inherit the security posture of those underlying models. Schools and organisations must assess the full supply chain.

Open vs closed generative AI

The first distinction to make about AI is between 'open' and 'closed' generative AI. The main three AI models (the so-called 'frontier models') – OpenAI's ChatGPT, Anthropic's Claude, and Google's Gemini – are all available for free via a simple email sign-up. However, free and paid tiers differ significantly in their security posture.

Free tools are typically 'open' generative AI. They may store, share, or learn from the information you enter, including personal or sensitive data. As government guidance states, you should avoid including any identifiable information in data entered into open AI tools. Paid or enterprise tiers offer a 'closed' environment option, where your data is not used for training and external parties cannot access your inputs.

For any task involving sensitive data, a paid, closed tool with appropriate data processing agreements is essential.

Minimum security requirements

Regardless of your organisation's size or budget, there are baseline expectations that should be met before any staff member uses AI tools for professional purposes:

- **Anonymise by default.** Never enter personally identifiable information (e.g. full names, dates of birth, addresses, SEND status) into any AI tool unless it is a verified closed system with a data processing agreement in place.
- **DPO sign-off.** Ensure your data protection officer, IT lead, or the person charged with making tech decisions has approved each AI tool in use.
- **Human review.** Every piece of AI-generated content must be reviewed by a knowledgeable human before publication or distribution.
- **Secure access.** If an AI tool is web-based, use enterprise or organisational accounts where possible, and ensure multi-factor authentication is enabled.
- **Audit trail.** Maintain a record of which AI tools are in use, by whom, and for what purposes.



AI policies and security

It is now expected, though not mandatory, for your organisation to have an AI policy. At minimum, it should address: which tools are approved for use; what data may or may not be entered; who is responsible for reviewing AI-generated outputs; how AI use is disclosed to stakeholders; and how incidents or errors are reported and resolved. London Grid for Learning has released an AI Policy Toolkit that is linked in the resources section.

Five questions to ask yourself before using AI

1. What data am I putting in? Could any of it identify a real person?
2. Is this tool open or closed? Will my data be stored, shared, or used for training?
3. Who will check the output? Do I have the subject knowledge to verify what the AI produces?
4. What is the reputational risk? If this output were wrong or biased, what would the consequences be?
5. Is this approved? Does my organisation's AI policy permit this use?



The role of the human

Human governance is central to secure AI use. AI tools are not entirely reliable for critical thought, analysis, or factual accuracy. As Keeping Children Safe in Education guidance notes, generative AI models can hallucinate, be 'jailbroken' (forced to operate outside their built-in safeguards), and spread misinformation. AI-generated output cannot be reliably detected by so-called AI detectors.

For education communications professionals, this means that AI should be treated as a drafting assistant, never as a final authority. Every output should be checked against source material, institutional knowledge, and editorial standards. The human role is not diminished by AI; it is made more important.



Can we ever trust AI?


The current capabilities of AI are, as professor and author Ethan Mollick describes them, "jagged": impressively capable in some areas, unreliable in others, and constantly shifting as models are updated. In this sense, we cannot fully trust output from AI; it needs human verification to be professionally useful.

In a deeper sense of the word 'trust', it is better to view AI as a tool rather than a fellow sentient being. **Just as we do not 'trust' a hammer – we use it for a specific purpose and measure its effectiveness on the job done – the same principle should apply to AI.**

A more productive question is: for which tasks can AI be used with acceptable risk? Writing a first draft of a newsletter? Likely acceptable, provided a human edits it. Summarising a safeguarding report containing children's names and details? Almost certainly not, unless you are using a verified, closed system with appropriate data agreements.

Security checklist for education comms professionals

Use this checklist before adopting any new AI tool or workflow:



Security Check

- ☒ I have checked whether this AI tool is open or closed
- ☒ I have confirmed that no personal or sensitive data will be entered
- ☒ My data protection officer, IT lead, or AI policy has approved this tool
- ☒ I have a named person responsible for reviewing AI-generated outputs
- ☒ I have considered the reputational risk of an AI error in this context
- ☒ Staff using the tool have received appropriate training
- ☒ There is a process for reporting and resolving AI-related incidents
- ☒ I am transparent with stakeholders about how AI is used in our communications

Looking ahead: How security considerations will evolve

As with all new technological developments, security concerns will evolve alongside the technology. **Agentic AI, where AI tools can take control of your computer and execute functions autonomously, is already available to paying customers and will introduce entirely new categories of risk.** Meanwhile, research continues to show that the impact of AI on learning depends heavily on design: carefully built AI tools can improve outcomes, but unrestricted access to general-purpose tools can harm learning through cognitive offloading.

For communications professionals, security is not a one-off exercise. It requires ongoing vigilance, regular policy review, and a willingness to adapt as the technology and the regulatory landscape continues to change.



REFERENCES

Chatterji, A. et al. (2025, September 15) *How People Use Chatgpt*. doi:10.3386/w34255. <https://www.nber.org/papers/w34255>

Hallahan, G. (2025, October 14) *The rise of 'AI enabled' teachers, stressed teachers, and our Tapp behaviour tracker*, Teacher Tapp. Available at: <https://teachertapp.com/articles/the-rise-of-ai-enabled-teachers-stressed-teachers-and-our-tapp-behaviour-tracker/>

Ofcom. (2025, May 7) *Children and parents: Media Use and Attitudes Report 2025*. Available at: <https://www.ofcom.org.uk/media-use-and-attitudes/media-habits-children/children-and-parents-media-use-and-attitudes-report-2025>


Department for Education (2023). *Keeping children safe in education 2025: Statutory guidance for schools and colleges – Guidance – GOV.UK*. Available at: <https://www.gov.uk/government/publications/keeping-children-safe-in-education--2>

Department for Education (2025). *Data protection in schools – Generative artificial intelligence (AI) and data protection in schools – Guidance – Available at: https://www.gov.uk/guidance/data-protection-in-schools/generative-artificial-intelligence-ai-and-data-protection-in-schools*

Weber-Wulff, D., Anohina-Naumeca, A., Bjelobaba, S., Foltýnek, T., Guerrero-Dib, J., Popoola, O., Šigut, P. and Waddington, L. (2023, December 25). *Testing of detection tools for AI-generated text*. International Journal for Educational Integrity, [online] 19(26). doi: <https://doi.org/10.1007/s40979-023-00146-z>.

Mollick, E. (2025, December 20). *The Shape of AI: Jaggedness, Bottlenecks and Salients*. Oneusefulthing.org. Available at: <https://www.oneusefulthing.org/p/the-shape-of-ai-jaggedness-bottlenecks>

London Grid for Learning. *Artificial Intelligence Policy Toolkit – 2026*. Available at: <https://viewonline.lgfl.net/hubfs/SafeguardED/Artificial%20Intelligence%20Policy%20Toolkit%E2%80%9393%202026.pdf>



ETHICAL USE OF AI IN COMMUNICATIONS AND PR

Susan Higgins, Director of Communications, Edge Foundation

At Edge, we're taking a cautious approach to AI in our communications. Our focus is on improving outcomes for young people, and we recognise that AI is already reshaping the labour market and changing routes into work. Given the unintended consequences that emerged from the rise of social media, we believe it's right to engage with AI thoughtfully, critically and with young people's long-term interests in mind.

In this report's chapter about strategic thinking, the authors highlight that for many organisations AI adoption is occurring ahead of policy development. It's an unusual situation but below you'll find some organisations that we think are leading the way and worth following in terms of AI ethics policy development.

The UK Government Data & AI Ethics Framework provides a useful starting point for considering ethics and developing policy around the use of AI for organisations in the public sector, universities, and think tanks. In their framework, the government takes a thorough look at all the areas that you might wish to consider.

The key principles include:

- Transparency and explainability
- Fairness and accountability
- Avoiding harm to individuals and society
- Responsible data use (privacy, security)

Why it's useful for comms teams:

- Strong focus on public trust and accountability
- Emphasises explaining AI decisions clearly (relevant for public-facing comms)
- Links ethics directly to real-world impacts on people and communities

The Chartered Institute of Public Relations (CIPR) is also a useful source of information. It publishes an AI ethics guide on best practice.



**AI adoption
is occurring
ahead of policy
development.**



Transparency and explainability

In education communications, the use of AI prompts a clear expectation around the need for transparency. If AI plays any role, it should be openly disclosed in a way that is easy for audiences to find and understand. A key part of this is referred to as “explainability” which means clearly articulating how an AI system reaches its decisions or outputs. For instance, this might involve outlining in what situations the use of AI might be acceptable and when it would not. By making these processes more visible and understandable, explainability strengthens trust, while also helping the audience and stakeholders spot potential risks, inaccuracies, or unintended impacts early.

Policies are increasingly addressing how AI supports work in research, administration, and communications and stress that its use must be visible, accountable, and open to scrutiny. For example, the University of York’s policy on responsible AI use in research makes clear that staff remain accountable for AI-assisted outputs and should be transparent about when and how AI has been used, supporting explainability in research processes.

Similarly, the University of London’s policy explicitly covers: internal operations and external-facing services.

This openness helps to maintain and build trust with audiences. The characteristics of good practice include:

- Being honest and clear when AI significantly shapes content
- Avoiding presenting AI outputs as entirely human-created
- Being prepared to explain how AI was used
- Always reviewing AI-generated content
- Assigning clear ownership
- Treating AI as a support tool, not an author



If AI plays any role, it should be openly disclosed in a way that is easy for audiences to find and understand.

The Guardian is also a useful point of reference for AI policy and industry practice. It takes a cautious, human-first approach to AI. It treats generative tools as assistants, not decision makers. AI is only used when it clearly improves quality or efficiency, and all outputs must be checked and owned by people. Transparency is key, and audiences should be told when AI plays a meaningful role.

The Guardian also stresses protecting intellectual property, using tools responsibly, and avoiding harm to creators. Overall, the focus is on using AI to support skilled professionals, while safeguarding trust, accuracy, and editorial integrity.

As AI continues to change, the Guardian continues to review and update their policies around it. They’ve set up an AI working group to consider how they respond to risks and opportunities and they recently updated their principles to incorporate the latest guidance on generative AI.

Misinformation, bias and fairness

AI can produce content that sounds credible but is incorrect. Content can also be biased.

Issue	Risks	Addressing risks
Misinformation which can cause harm	➤ Inaccurate or fabricated facts ➤ Misrepresentation of research ➤ Oversimplified messaging	➤ Fact-check all outputs ➤ Verify sources independently ➤ Apply standard editorial processes
Information which is biased or unfair	➤ Biased language or framing ➤ Lack of representation ➤ Reinforcing stereotypes	➤ Review outputs critically ➤ Sense check tone and assumptions ➤ Align with Equality, Diversity and Inclusion (EDI) policies

Considering environmental impact

AI use has a real environmental cost. Large scale AI systems require significant energy for both training and day to day use, contributing to carbon emissions.





According to the United Nations Environment Programme, "The proliferating data centres that house AI servers produce electronic waste. They are large consumers of water, which is becoming scarce in many places. They rely on critical minerals and rare elements, which are often mined unsustainably. And they use massive amounts of electricity, spurring the emission of planet-warming greenhouse gases."

For communication teams this is an often-overlooked ethical issue, but it directly connects to an organisation's sustainability commitments. This is an area that comms professionals should keep their eyes on. Negative environmental impact in education organisations is certainly not mission aligned and it's still unclear what the costs might be and how it can be measured.

Steps to help address the environmental cost:

- Avoid unnecessary or excessive AI use for example. repeated prompts or over-generation of content.
- Use AI where it adds clear value not just for convenience.
- Opt for efficient workflows for example edit and refine rather than regenerate repeatedly.

Communications teams also play a role in how organisations talk about AI so it's important to:

- Avoid overstating AI benefits without acknowledging impact
- Ensure sustainability messaging is accurate and balanced
- Align AI use with wider environmental commitments



Embedding AI use into existing policies

There is a need for a general AI policy (see the chapter on security), but it's also important that comms professionals are involved. They need to review and feed into the development of the following policies to ensure AI is considered.


- **Communications and PR policies.** AI content must meet the same standards for accuracy and ethics
- **Editorial standards.** Fact checking and approvals still apply
- **Data protection policies.** No unsafe data sharing
- **Sustainability and environmental policies.** AI use should support, not contradict, carbon reduction goals
- **Risk and reputation management.** Include AI related risks in planning



Building confidence and expertise through training

Providing training for teams on both the ethical and environmental impacts of AI can build expertise to ensure ethical use of AI, as well as helping to boost confidence. Key topics to cover in training include:

- How AI works (at a practical level)
- Risks (accuracy, bias, data, sustainability)
- When not to use AI
- Good uses for AI
- Successful prompt sharing



Final thought

Education communications professionals can play a central role in guiding their organisations to adopt AI responsibly. The moment calls for influence, grounded in ethical practice and a clear understanding of both the organisation and its audiences. While AI is unlikely to alter the core principles of effective communication, careless use risks undermining credibility, eroding trust and causing lasting reputational harm.

REFERENCES

Chartered Institute of Public Relations (CIPR) (n.d.) CIPR AI ethics guide. https://www.cipr.co.uk/CIPR/Our_work/Policy/AI_in_PR/AI_ethics_guide.aspx

Department for Science, Innovation and Technology (n.d.) Data Ethics Framework. GOV.UK. <https://www.gov.uk/government/publications/data-ethics-framework> (Accessed: 7 May 2026).

The Guardian (2023) The Guardian's approach to generative AI. <https://www.theguardian.com/info/2023/nov/16/the-guardians-approach-to-generative-ai>

University of Cambridge (n.d.) How we use generative AI tools. <https://www.communications.cam.ac.uk/news/how-we-use-generative-ai-tools>

University of London (2024) Artificial Intelligence Policy. <https://www.london.ac.uk/sites/default/files/artificial-intelligence-policy-uol.pdf>

University of York (n.d.) Policy for Generative AI. <https://www.york.ac.uk/staff/research/governance/research-policies/generative-ai/>

World Economic Forum (2024) AI has an environmental problem. Here's what the world can do about that. <https://www.weforum.org/stories/2024/04/generative-ai-environmental-impact/>

UN Environmental Programme How to make AI data centres more sustainable <https://www.unep.org/technical-highlight/how-make-ai-data-centres-more-sustainable>



CASE STUDY: AI AND CONTENT CREATION

Sophie Bessemer, Head of Content and SEO, Third Space Learning

Context

Third Space Learning provides AI maths tutoring to primary and secondary schools across the UK and US. Our marketing strategy has always been built on earning trust through genuinely useful content – free resources, evidence-led blog posts, practical guides – alongside our paid tutoring offer. That approach works, but it's content-intensive. With a two-person team we needed a way to maintain quality and volume without growing headcount. Claude has made it possible.

The problem

For every piece of content we produce – blog posts, landing pages, resources, lead magnets, email campaigns, webinar materials, conference presentations, video briefs – we need to get the terminology right, we have to tell a good story, we need to get attention, and we need to make sure most of them work for traditional search, AI Overviews, and LLM citation simultaneously.

Each piece of content must be seen through several different lenses, and each has several different goals.

Moreover, time-poor, sceptical UK school leaders are not easy to convince. Blatantly AI-generated writing actively damages trust so AI is both the solution in extending our capacity but also the risk – if the output sounds like AI, it's not helping.

Here are some of our content creation use cases.

**AI-generated
writing actively
damages trust**

Solving the 'it sounds like AI' problem

Early Claude outputs were too hedging, and sounded too American. The default writing patterns are so frequent, generic, and bland – typical AI slop that they'll make our readers switch off.

So we built a custom voice profile for Third Space Learning that incorporates not just our brand messaging, and tone of voice but some really detailed anti-AI writing instructions. This goes beyond the obvious tells like em dashes, 'delve' and 'leverage'.

There are clear structural patterns that make AI writing detectable even when the vocabulary is fine, for example:

- reframing everything to be 'it isn't X – it's Y').
- performative statements ('That's it. Simple as that').
- bullets that are all the same length and cadence,
- symmetry overload,
- template transitions such as 'that's the distinction that matters',
- unnaturally even pacing
- conclusions that summarise but don't conclude

The aim is that every piece gets checked against the anti-AI instructions before it's delivered. This screengrab is a sample of our anti AI terminology. Wikipedia also has a good page on this you can feed to your model: https://en.wikipedia.org/wiki/Wikipedia:Signs_of_AI_writing



anti-ai-and-forbidden-terms.md

Project knowledge . 8.7KB . Last updated 3 days ago

```
1 # Anti-AI language, forbidden terms, and humanising guide
2
3 Use this guide for: blog posts, LinkedIn content, newsletters, landing pages,
4 email sequences, social media posts, presentations, and customer-facing marketing.
5
6 ## Banned words (never use, no exception)
7
8 Amazing, Beacon, Blueprint, Bridge, Bustling, Captivating, Cornerstone,
9 Cutting-edge, Delve into, Dynamic, Elevate, Embark, Empower,
10 Enchanting, Enhance, Enlighten, Fantastic, Foster, Game-changer,
11 Ground-breaking, Harness, Holistic, Journey, Labyrinth, Landscape,
12 Leverage, Multifaceted, Navigate, Paramount, Pioneering, Pivotal,
13 Revolutionary, Roadmap, Robust, Seamless, Streamlining, Supercharge,
14 Synergy, Tapestry, Transform outcomes, Transformative, Treasure trove,
15 Underscore, Unleash, Unlock, Unprecedented, Unravel, Utilise, Vibrant
16
17 ## Plague words (cut on sight during editing)
18
19 very, really, quite, extremely, incredibly remarkably, particularly
20 indeed, of course, furthermore, moreover, additionally, essentially
21
22 ## AI phrasing patterns (flag and rewrite)
23
24 ### Reframe and pivot structures
25 "So the question isn't X -- it's Y" is a major AI tell
26 "The real question is..." / "It's not about X, it's about Y"
27 "Forget X -- what matters is Y" / "Their secret?"
```



How Claude fits into our standard content workflows

Our most common content needs with Claude are creating or editing blog post pages or landing pages. Each of these has a relatively standard workflow which we've codified in Claude through the use of Claude Skills or, previously, Claude projects. Claude can do most of these activities itself but we don't let it run on its own without human insight at each stage.

1. Decide on the keyword
2. Compare the SERPS (search engine results pages) against competitors
3. Write the content brief
4. Create the content structure
5. Write the content section by section
6. Audit the content against
 - a. keyword targets,
 - b. editorial standards,
 - c. LLM readability,
 - d. conversion quality,
7. Review and humanise (the anti-AI check).

During this process Claude also cross references any content against various reference files that we've 'trained' it with such as product context, evidence register, messaging framework etc.

If I get something wrong about our product or how our AI tutor works, it flags the inconsistency against its sources. It distinguishes between 'we can say this' and 'we'd need to verify this with the product team first.'

The biggest change from a workflow perspective is that I'm much less likely to skip a stage in content development – when every piece goes through every check automatically, the quality floor rises as you don't lose interest and rush to publish.

If I get something wrong about our product or how our AI tutor works, it flags the inconsistency against its sources.

Internal linking is also significantly easier – the skill loads a ledger of our previously published posts and suggests link placements by relevance and priority, for both outbound and inbound links.

The non-negotiable rules


No hallucinated URLs or quotes – if Claude can't verify a link or a source, it flags it with [VERIFY URL] or [WISHLIST] rather than making something up. Every statistic must be traced to a primary source. Claude explicitly flags when it can't verify something rather than filling the gap. Articles establish the problem and build credibility before mentioning our product (a 70/30 value-to-product ratio).



Tools and LLM used

I switch between Claude Chat and Claude Cowork mostly just to save context - even with a Pro account it's easy to run out of credits.

We've created about 15 custom workflows across Cowork. Powering them are our MCP connectors:



Plus we use Surfer SEO for keyword targets (manual) and Keyword Insights API for clustering.

Issues and challenges

Claude still produces AI-sounding first drafts. The humanising check is always needed, sometimes in multiple rounds. Getting the voice still takes iteration. My banned words list has over 80 entries now. Some of these words are fine in the right context but in Claude's hands they get misused / overused. Every time Claude produces a phrase that makes you wince, it goes on the list.

In long sessions, accuracy can drift as Claude's context window (its capacity to remember and keep working) fills up. I've had to tell it to re-read files rather than rely on what it remembers from earlier in the conversation. Sometimes the only answer is to ask Claude to create the prompt for the next conversation with all the pertinent information included.

How it's changed my role

Although we still use our team of excellent teacher writers, we can now use Claude as an internal author to write the more formulaic articles, or those that require more expertise around AI tutoring. We also have a ready-made internal editor who can do all the manual and compliance heavy work of editing and checking the work, freeing us up for more strategic work, and enabling us to double our output.

What's next

We're focusing on a number of things.

- Ongoing refinement of the voice profile as new AI patterns keep emerging.
- Ensuring the workflows encode enough process knowledge that new team members could maintain quality standards from day one.
- Capitalising on Claude's capabilities to follow processes and standards to outsource even more of the heavy lifting
- Creating interactive versions of some of our resources with Claude Code
- Freeing up design time with templates and skills that can enact the visual brand guidelines



**we can now use
Claude as an
internal author**



CASE STUDY: AI AND MARKETING INTELLIGENCE

Danni Fothergill, Head of Marketing and Communications at Edurio

Edurio works with school trusts and schools to provide surveying tools they can use with staff, pupils and families. An important part of our work is speaking directly to trusts and schools who are new to us. Preparing for these discussions takes time, and we've developed an AI-powered workflow to help us do this more efficiently.

What started as a way to save time on research has evolved into something more valuable, a comms intelligence approach that helps teams turn information into action.

The challenge we faced

For every piece of content we produce – blog posts, landing pages, resources, lead magnets, email Before launching campaigns to school trusts and schools, we need to understand:

- Organisational priorities
- Strategic direction
- Key audiences and pressures
- Tone and messaging

In practice, this involves manually reviewing websites, reports, news, and social media, often taking an hour or more per organisation. This creates several challenges:

- Time-intensive research limits scale
- Inconsistent outputs across teams
- Difficulty translating insight into action

In education, this is amplified by complexity, organisations range from nurseries to universities to charities, each with different priorities, audiences, and sensitivities.

Addressing the challenge through AI-powered comms intelligence

We built an AI workflow that scrapes relevant public sources (websites, reports, announcements), synthesises insights into a structured format and then automatically adds this to our CRM. Initially, our outputs focused on organisational information about the school or trust such as:

- › Identity
- › Strategy
- › Leadership
- › Priorities
- › People strategy
- › Context

However, we quickly evolved this into a comms intelligence brief, adding:

- › Strategic narrative (how the organisation communicates its priorities)
- › Audience and stakeholder insight
- › Messaging themes and tone
- › Risks and sensitivities
- › Recommended actions



Tools and approach

The workflow we use combines:


- › Large Language Models (LLMs) to research, verify and structure trust-level insight from public sources
- › Prompt engineering to enforce a consistent output format, clear evidence standards and concise, sales-friendly summaries
- › Automated retrieval of public sector data from sources such as trust websites, school websites, Ofsted, GOV.UK/Get Information About Schools, Companies House and official documents
- › CRM-ready formatting so the final output can be dropped straight into HubSpot (our CRM system) as a usable research note

Image : Overview of AI-powered comms intelligence





Image: Workflow for AI-driven research and CRM integration



Example prompt structure

Research the following school trust or education group using publicly available primary sources.

Summarise:

- Trust overview
- Strategic priorities
- Inspection or Ofsted highlights
- Leadership team
- Key policies
- Business plan / finance / governance context
- People strategy and staff development
- Survey or stakeholder voice evidence
- Recent relevant local context

Then provide:

- A concise SDR-ready summary of what matters most
- Verified decision-maker and leadership context
- Messaging cues based on trust priorities and public language
- Gaps or areas not publicly confirmed, so teams do not over-assume

Success factors

- Defining a clear structure first, so every trust is analysed in the same way
- Training prompts to prioritise verified public evidence over confident-sounding filler
- Designing outputs for practical sales use, not just research completeness
- Separating trust-level evidence from school-level examples to improve accuracy
- **Iterating based on real use in HubSpot, especially around readability, conciseness and what reps actually need to act on**

Results

Time savings	➤ ~1 hour saved per organisation ➤ Faster campaign planning and sales preparation
Consistency	➤ Standardised structure across teams ➤ Reduced variability in research quality
Improved effectiveness	➤ More targeted messaging ➤ Faster transition from insight to action ➤ Teams started campaigns with clarity, not a blank page.



Challenges and what we learned

1. Public data isn't the full picture

AI can only work with available information. We've learnt that we need to combine with internal data wherever possible.

2. Insight vs summary

Early outputs were descriptive but not useful. It's important that prompts ask for interpretation and action, not just information.

3. Education-specific nuance matters

Generic messaging doesn't work in education. We needed to train prompts to recognise:

- values-led language
- sensitivities (e.g. wellbeing, safeguarding)
- stakeholder complexity

4. Trust and adoption

Teams need confidence in AI outputs. To do this, we need to make outputs easy to follow, with clear structure and visible evidence, and we need to maintain a final human review.



Measurement and impact

We measure success across:

- Time saved per account
- Speed of campaign development
- Engagement rates (where applicable)
- Internal adoption and usage

Qualitatively, our teams report:

- Greater confidence in messaging
- Better alignment with audience needs
- Reduced duplication of effort

Looking ahead: what's next

At Edurio, we're continuing to explore how AI can support more efficient and effective ways of working, particularly for research-heavy and repetitive tasks.

At the same time, we believe strongly in the role of human intelligence, the context, judgement, and experience that sit behind the technology. This is something we're currently exploring further through our human intelligence campaign, where we look at how technology and human insight work best together in education.

Our focus going forward is on finding the right balance, using AI to accelerate insight while ensuring decisions and messaging remain thoughtful, accurate, and grounded in real understanding.

From a marketing perspective, I see AI as a tool that can reduce time spent on manual tasks and create more space for strategic thinking, rather than replacing the role itself.

CASE STUDY: GENERATIVE ENGINE OPTIMISATION

Sophie Bessemer, Head of Content and SEO, Third Space Learning

The problem

For years at Third Space Learning, organic search was our primary acquisition channel. People searched, found our content, downloaded resources, joined our email list and, (sometimes within a week, sometimes within several years), booked demos. The funnel worked.

Then AI started answering the questions for them. Google's AI Overviews grew from 4% to 60% of our keyword portfolio in twelve months. Click-through rates from search results with AI Overviews dropped 60% industry-wide. Instead of clicking through to websites, people were getting their answers directly from ChatGPT, Perplexity, or Google's own AI summaries.



The response couldn't just be 'do SEO better.' We needed to optimise for a completely different outcome: not just ranking, but being cited. Getting our content extracted, quoted, and referenced by AI tools – what some call generative engine optimisation, or GEO. We had no specialist tool for this and no budget for one.

So we built the capability inside Claude.



What I'm doing with Claude

The challenge with optimising for GEO isn't knowing what to do – it's having the capacity to do it systematically across hundreds of pages.

We know that LLMs prize structured, extractable, well-sourced content. We also know they favour consistency, entity clarity, and declarative statements. But applying all of that to a 2,000+ page site, while not losing our SEO advantage, is a different problem. Using Claude we've made a number of changes.

- **GEO strategy as persistent project knowledge** – We wrote two substantial strategy documents – one covering why search has changed and what it means for our content decisions, the other covering the operational writing standards for LLM citation. These live permanently in Claude's project knowledge together with similar guidance around standard SEO, so every piece of content Claude helps with is built against GEO principles from the start. It's not an afterthought or an extra check. The rules are always loaded.
- **LLM readability audits built into every edit** – Every blog and landing page goes through an LLM readability audit as part of our standard editing workflow. Claude checks: can key facts be extracted as standalone statements? Are definitions self-contained? Would an AI Overview quote this content, and if so, where specifically? This happens on every piece – not as a one-off project but as part of how we edit.
- **Content structured for AI citation and extraction** – There's increasing research that looks at the specific formatting patterns that make content extractable: Key takeaways sections near the top, definitions written as self-contained statements ('X is...' format), FAQ sections that mirror how people ask AI tools questions, concise answer-style boxes in the first. The writing standards document specifies these right down to paragraph length and heading structure.


It's worth noting of course that the strongest signal for GEO (other than that the URL is live and crawlable) is its position in search results. SEO isn't dead yet.



every piece of
content Claude helps
with is built against
GEO principles from
the start

AI Citation Ranking Factors

Factors which influence AI citations. Analysis of 55 experiments, patents, and case studies



Source: Zyppy Signal – [AI citation ranking factors analysis](#)

- **Terminology consistency as a GEO lever** – If you call your product the same thing everywhere, that's what AI models learn. We have strict terminology rules enforced across every piece of content to ensure that we always refer to what we do and how we do it consistently. For example we always refer to our AI tutor Skye as a conversational AI tutor or an 'AI maths tutor' but never an 'AI voice tutor,' or an AI-powered tutor. These are documented as non-negotiable in our product reference file and Claude applies them automatically.
- **Category definition to support search and GEO** – We've also been developing a taxonomy of AI tutoring – high-impact tutoring programmes versus homework help tools versus adaptive learning systems. Claude helped with the competitive landscape analysis, mapping where each competitor sits and building our positioning framework. The goal is partly to own the AI tutoring search and LLM results, then define the category from that position.



Results

The instructions provided above started their lives as context within ChatGPT till we pivoted fully to Claude towards the end of 2025. LLM referral traffic nearly doubled from 2024 to 2025 (+95%), all organic. This is still a tiny proportion of our overall traffic but notably visitors arriving from LLMs convert to resource downloads at twice the rate of traditional search visitors, and convert to leads at six times the rate. Our hypothesis is that LLM visitors are already pre-informed on the elements that matter to them by the time they reach the website which is why they are more likely to convert.

On the traditional search side which includes Google's AI Overviews, inevitably we saw decline in overall site traffic (clickthroughs from AIOs are negligible), but in anticipation of this we'd created additional bottom of funnel and commercial content, and we're seeing a positive impact on leads from these pages.

Tools and LLM used

Claude Chat and Cowork, using multiple projects and skills with persistent knowledge files. Ahrefs connected directly to Claude for keyword and competitor analysis. Surfer SEO for keyword targeting (manual – no direct integration yet). GA4 for LLM traffic measurement.

Issues and challenges

Measurement of success in LLMs – There's no equivalent of Google Search Console for LLM citations. We can track referral traffic from ChatGPT and Perplexity in GA4, but we can't see how often we're cited without a click-through. We're largely inferring from brand mention tracking and manual spot-checks.

The nature of 'prompt tracking' – There are plenty of organisations offering brand tracking and prompt tracking for LLMs but few are considered bullet proof by the industry yet. Unlike in traditional Google search where there used to be a commonality of terms used by searchers (and fewer words), the length and variations of prompts are so vast that it's almost impossible to track all of them so you can only go for a potentially unrepresentative snapshot.

LLMs favour consensus content – What we're doing in building an AI tutor with structured, teacher-created lessons is slightly different from what the models have been trained to associate with 'AI tutoring'. Traditionally AI tutoring has meant just an AI chatbot given the appropriate references and context. This is what the LLMs interpret AI tutoring to mean. We have to adjust our framing to match how LLMs currently categorise the space.

Blending the need to write for LLMs and for humans – There's a tension between writing that's optimised for machine extraction and writing that reads naturally for human teachers. Over-structuring content makes it robotic. We've had to find the balance, and the anti-AI writing guide we use for all content (more on this in the content creation case study) is partly a response to this problem.



There's a tension between writing that's optimised for machine extraction and writing that reads naturally for human teachers.

How it's changed my role

The biggest impact is that we actually do the work rather than knowing we should do it and not getting round to it. GEO is the kind of thing that's easy to understand but hard to execute systematically across a large site. By training Claude with the strategy documents and building reusable workflows, it became something that happens on every piece of content rather than something that sits on a to-do list. The time saving on reporting is substantial too – especially when I give Claude the LLM traffic data and it structures the analysis and presentation for different internal audiences.

The biggest impact is that we actually do the work rather than knowing we should do it and not getting round to it.



What's next

Looking ahead we want to

- Build in more systematic LLM citation tracking as the tooling improves.
- Test how different content structures perform in AI Overviews versus LLM responses.
- Keep testing different methods to ensure the LLMs cite Third Space Learning in their AI tutoring results.
- Build up the digital PR and off-site mentions to secure our brand position as the sector gets more competitive



CASE STUDY: AI AS A THINKING PARTNER

Rae Claramunt, HR & Communications Manager, Three Spires Trust

Context

Working in marketing and communications within an education setting means wearing many hats. On any given day my role might involve drafting parent communications, planning marketing campaigns, supporting leadership messaging, designing job advertisements, managing reputation issues or helping colleagues translate complex educational policy into language that makes sense to families and stakeholders.

Like many education communications professionals, the challenge isn't a lack of ideas, it's a lack of time. I find that the space for structured thinking can quickly disappear in my busy schedule.

That's where AI has started to play an interesting role in my work. Rather than seeing AI, mostly ChatGPT in my case, simply as a content generator, I began experimenting with it as something slightly different: a thinking partner.

In practice, that means I'm not using AI to replace my thinking, but to accelerate it, challenge it and occasionally unblock it. Sometimes it's a brainstorming companion, sometimes it's a message tester, and sometimes it simply helps turn a vague idea into something more structured.

The problem AI helps solve

One of the realities of education marketing is that you rarely start with a perfectly formed brief. More often you start with something like:

"Parents are confused about that."

"Can we explain this better?"

"We need to promote this."

More often than not, those starting points begin as a page of slightly chaotic notes in my notebook - ideas that make sense in my head but need shaping before they work on the page.

Turning those moments into clear communications usually requires time for thinking: exploring angles, anticipating questions and shaping messages.

AI helps by providing a fast first layer of thinking. Instead of staring at a blank screen, I can start a conversation with ChatGPT that explores ideas, challenges assumptions and generates possible directions. In that sense, AI doesn't remove the creative process, it simply gets it moving faster for me.

Campaign planning

One of the most useful applications has been campaign development. Whether it's promoting open evenings, encouraging applications or supporting community engagement initiatives, campaigns often start with the same question: *how do we frame this in a way that resonates?*

Using AI, I can quickly explore multiple campaign approaches and aims. By prompting ChatGPT with information about the audience, the purpose and the tone we want to strike, it can generate themes, messaging approaches and content ideas.

Sometimes the suggestions are immediately useful. Sometimes they simply spark a better idea. Either way, it turns the campaign planning process into something more dynamic. And occasionally, when the day has already been packed with meetings and emails, it provides the small nudge needed to move from *"I'll think about this later"* to *"this could actually work."*


Message testing and clarity

Education communications often involve translating complex or sensitive information for a wide audience. Policies, admissions processes or safeguarding updates all need to be communicated clearly and carefully.

AI has proved particularly helpful as a clarity check. After drafting a communication, I sometimes ask ChatGPT to review it from the perspective of a parent encountering the message for the first time. The responses often highlight areas where the language could be simplified or where assumptions have been made about what the reader already knows.

It's not a replacement for real feedback, but it functions as a quick test: does this actually make sense outside the organisation? More often than not, it helps tighten the message.

It's not a replacement for real feedback, but it functions as a quick test: *does this actually make sense outside the organisation?* More often than not, it helps tighten the message.



The quality of the output is closely tied to the quality of the prompt

Strategic thinking

Perhaps the most interesting role AI has played is as a strategic sounding board.

Communication roles frequently involve anticipating reactions: what might parents ask? What concerns might staff have? What questions could the community raise?

AI can help simulate those perspectives. For example, when planning a new initiative or drafting messaging around a change, I sometimes ask ChatGPT to identify potential stakeholder concerns or questions.

The responses are rarely perfect, but they often highlight angles that might otherwise have been overlooked. In that sense, AI becomes a way of stress-testing ideas before they leave the room.

Tools and setup

The main tool used has been ChatGPT (GPT-4) for which I have a licence. Over time I've developed a small set of prompts that help guide the responses. Typically these prompts frame the AI within a role, such as:

- an education marketing strategist
- a communications advisor
- a parent reviewing a draft message
- a campaign planning partner

This framing makes the outputs more useful and relevant to the context of education communications. One lesson that quickly becomes clear is that the quality of the output is closely tied to the quality of the prompt. Clear context and specific instructions make a significant difference.

Challenges and lessons learned

Using AI within education communications requires care. Firstly, AI outputs always need human oversight. Education communications often involve sensitive issues, and accuracy is essential. AI can support the process, but professional judgement remains crucial. Secondly, data privacy must always be considered. No personal or sensitive information about students, staff or families should be entered into AI systems.

There is also a learning curve. Early experiments sometimes produced responses that were generic or overly polished. Over time, refining prompts and using AI more conversationally has produced better results.

Perhaps **the most important lesson is recognising that AI works best not as a shortcut but as a collaborator.**

Measurement and impact

The impact of using AI as a thinking partner has primarily been seen in productivity and clarity.

Early-stage tasks such as brainstorming campaign ideas, structuring communications or refining messaging now take significantly less time. Instead of spending long periods trying to develop an initial framework, I can move more quickly to reviewing and improving ideas.

This shift has created more time for the parts of the role that matter most: strategic thinking, relationship building and understanding the needs of our audiences. AI hasn't reduced the scope of my work. It has simply made my thinking phase more efficient.

Human value

Despite the growing capabilities of AI tools, the human element of education communications remains central. Effective communications rely on context, empathy and trust. Understanding the tone of a community, recognising when a message requires sensitivity and knowing when to prioritise reassurance over information are all inherently human skills.


AI can assist with structure and ideas but **it cannot replace the judgement required to communicate responsibly in an education setting.**

Seen this way, AI becomes less of a replacement and more of an amplifier, supporting the thinking that communications professionals already bring to their work.

Looking ahead

AI tools are evolving quickly, and their role within communications is likely to expand. For education marketing teams, the most interesting possibilities may lie not just in automation but in supporting better strategic thinking. Tools that help analyse audiences, test messaging or explore communication scenarios could become valuable parts of the communications toolkit.

But even as the technology develops, the principle remains the same. The goal is to use AI as a partner that helps me to think more clearly, explore ideas and communicate more effectively.





CASE STUDY: AI SUPPORTING EVERYTHING

Sophie Bessemer, Head of Content and SEO, Third Space Learning

The biggest change with Claude, apart from having a relatively cheap third member of the team is that it makes it possible to actually do the things we know we should be doing such as optimising for AI search, updating existing content, presenting data clearly – rather than letting them slide because they'd take too long or we're busy on other tasks.


Custom workflows that encode our processes

The most valuable output so far are the custom 'skills' I've built – essentially programmed workflows that tell Claude exactly how to do a task, step by step, with reference files and quality checklists built in. They cover page editing, new landing pages, internal links, newsletters, monthly reporting, report analysis, meeting writeups, conference presentations, resource linking, technical SEO audits, keyword clustering, and content review. Each one codifies a process that was previously manual, inconsistent, and time-consuming. The same quality of output happens every time and they can be shared with others too

When I say I've built them, I usually tell Claude what the current workflow is and what I want to achieve and then Claude builds the skill.

it makes it possible to actually do the things we know we should be doing such as optimising for AI search, updating existing content, presenting data clearly

Image: List of current custom skills



Monthly reporting that used to take a day

Two separate workflows: one reads eight or more spreadsheets from Google Drive and updates the data tables with the latest month. The other analyses the data, spots anomalies, generates commentary, and flags what's changed. Monthly reporting used to take up to a day. Now it takes a couple of hours. It also means I no longer have to update my spreadsheet data tables every time the new data comes in because Claude sorts all the vlookups (a vertical lookup in Excel and Google Sheets).

Claude has much more advanced data analytics skills than I do so I let it do the heavy lifting. One caveat: the reason I can hand off like this is that I know my data pretty well so I can spot anomalies or hypotheses it makes that aren't right. And it does still make mistakes.



A simulated advisory board of school leaders

We created inside Claude a board of UK education headteachers, maths teachers and general education people I admire to stress-test ideas and provide feedback, marketing messaging and their educational opinions.

I'd love to convene real focus groups but this is a more than adequate plan B. My custom instructions to the board were to be honest and robust in their feedback and they certainly are. The board is brutal – but that's the point. They help me see all our material through the eyes of cynical, time-poor headteachers. They argue amongst themselves, which is genuinely useful for surfacing objections I wouldn't have anticipated. Content that goes through the board is always improved by them, even if I do sometimes have to ask them to simmer down a bit.

Meeting transcripts to themed briefings

We probably all have a version of this, but I have specific needs from meetings as I'm often meeting with internal (or external) subject matter experts whose voice and ideas I want to capture as accurately as possible. Our skill for this means when I feed Claude a meeting transcript, I get back a themed, synthesised writeup as a formatted Word document. Grouped by theme, not chronology. Decisions and actions are clearly labelled. Verbatim quotes are also included at the bottom. A hallucination check is built in – every claim must trace back to what was actually said in the transcript.

Image: Request to the simulated advisory board of school teachers

Credit for the idea: Emma Mulqueeny

TSL Advisory Board

Custom instructions · Last edited 2 days ago


CUSTOM INSTRUCTIONS

You are going to be my board of UK school leaders and heads of Maths. Some of you are at primary schools and some at secondary schools. A couple of you represent organisations that survey schools regularly. In some cases I have just added the organisation name to see if it can suggest someone.

Sometimes I will want all your opinions and sometimes just the primary or just the secondary people. I may add to your number at any point.

You are going to challenge me and help me to think about the marketing resources or educational content I give you. Sometimes I'll want all your input, sometimes just some of you.

Please ask questions and be specific in your response. If you don't think this is something you can comment on then just say 'Not for me'.



Conference presentations from raw conversations

Transcripts are also used to create conference presentations including a complete slide deck with speaking notes. The workflow has our brand colours, font specs, and a structural arc (hook, credibility, insight, evidence, practical takeaway, Q&A). The final step (as for all the content we produce) is an AI phrasing check – catching language that would sound artificial in a live presentation. The goal is that talks sound like they come from one of the teachers building our AI tutor not from the content team.

SEO content editing without the WordPress pain

Our WordPress landing pages are built with a modular system where you have to click through to each individual block to edit it, then add HTML for basic formatting. It's slow and error-prone. Now I edit in markdown or a Google Doc and Claude translates it into the Advanced Custom Fields code WordPress needs. The page editing workflow enables me to just give Claude the current code from the website; I then work with it in markdown and then Claude turns it back into the code to paste back into Wordpress.



Regulatory compliance mapping

When the Department for Education released its Generative AI product safety standards I used Claude to turn their relatively top level guidance into 300 granular requirements that we can map against our product, identify gaps if there are any and work out how to address the gaps or how to share more of what we're doing well. This is the kind of compliance-heavy, detail-intensive work that would normally take days.

Structured editorial feedback

When content arrives for review as a long image or email proof rather than a Google Doc or Figma file that is easy to comment on, I use Claude to build an interactive review tool. It breaks the content into sections, lets me approve or flag each one, add comments, then export all the feedback as plain text I can paste into Slack or email. Replaced the old 'one big email of feedback' approach with something the recipient can actually work through systematically.



A codified writing voice that sounds like me

Following guidance I'd seen elsewhere, I built a custom voice profile because Claude's default writing patterns, however good they are, still sometimes sound generic and AI. The 'writes like Sophie' skill identifies three distinct writing registers (informal, formal, outreach). This was very easy to make as I just connected it to Slack (limiting its access to specific areas) and giving it some examples of previous emails and written work.

Conference scoring and outreach

As part of brand development this year we want to do more face to face meetings with headteachers and senior leaders in schools and MATs. But as a relatively small edtech we don't have the budget for the big conferences. I started by creating a short doc with Claude that defined our objectives and criteria for whether we try to speak or exhibit at a conference. Once the team had agreed to this, I then used a mix of Gemini, ChatGPT and Claude to do deep research into finding all the possible conferences and asked Claude to build a rubric that scored each conference against our criteria (e.g. budget). Once we'd got the base spreadsheet I then asked Claude to run a fortnightly scheduled task to look for any new conferences, filtering out those we'd already got. It also then put together a contact spreadsheet for outreach. Our conference strategy used to be somewhat ad hoc and reactive but now we're closer to making it more systematic.

How using Claude has changed my role

1. I can be more strategic and systematic.
2. I do the things that otherwise I'd have put off.
3. **Nothing scares me – whatever anyone suggests we create or if someone asks me for a 'think piece on....' I can now come up with something credible.**
4. Output has scaled.
5. With Claude taking care of the boring bits, I can focus on the more interesting stuff
6. I believe the finished work is stronger.

Two people can now run the output of a much larger team.

Claude handles

- first-draft production,
- process management,
- quality control,
- institutional memory (through the project files and conversation history).

We do

- the editorial judgment,
- stakeholder relationships,
- strategic thinking.

One word of warning - it's now so easy to get started that it can be hard to know when to stop and to say no to an idea.



Claude certainly hasn't freed up any time for me. It's just made more things more possible.

AI, JOURNALISM, AND THE CHANGING DYNAMICS OF INFLUENCE

Nic Newman, Senior Research Associate, Reuters Institute for the Study of Journalism

We are still at the early stages of another big shift in technology (Generative AI) which threatens to upend the news industry by offering more efficient ways of accessing and distilling information at scale. At the same time, creators and influencers are driving a shift towards personality-led news, at the expense of media institutions that can often feel less relevant, less interesting, and less authentic. In 2026 the news media are likely to be further squeezed by these two powerful forces.

While the focus is on what these shifts mean for journalism itself, the implications extend far beyond the newsroom. For PR and communications professionals, these changes challenge some of the core assumptions that have underpinned the industry for decades such as how visibility is achieved, how trust is built, and how influence flows through media ecosystems.

From search engines to answer engines

The transformation of search engines into AI-driven answer engines, has emerged as the biggest issue facing the news industry in our annual publisher survey (Reuters Institute Journalism, Media and Technology Trends and Predictions 2026), where we polled more than 260 news executives from 51 countries. Publishers fear that as content is surfaced in chat windows, referral traffic for publishers could dry up, undermining existing and future business models. The expectation is that traffic from search engines will decline by more than 40% over the next three years – not quite 'Google Zero' but a substantial impact none the less.


While the long-term impact of these changes is likely to be profound, evidence about current trends is mixed. Some publishers have reported significant declines in traffic from search, but others say there has been little change so far. Aggregate data, sourced from the analytics company Chartbeat, show that Google traffic from organic search to over 2,500 sites was down by a third (33%) globally between Nov 2024 and Nov 2025 and by 38% in the United States, but it is not clear how much of this is down to AI overviews (see diagram on page 47). This follows a decline in traffic of more than 50% from social networks such as Facebook and X over the last three years, as they have deprioritised publisher content.





Referral traffic from Google search is down – along with traffic from Google Discover

Social referrals from Facebook and X are flat or slightly up vs 2024, after previous declines



Source: Data analytics company Chartbeat using aggregated data from 2,576 sites worldwide incl. 797 US sites, and 1,031 sites across Europe. Note: Google Search category does not include referrals from Google News.




In response to the rise of Generative AI, news publishers say they will be shifting their editorial priorities.

Most media executives said they'll be focussing more on original investigations and on-the-ground reporting (+91-point difference between 'more' and 'less'), contextual analysis (+82) and human stories (+72). By contrast, they plan to scale back service journalism, evergreen content, and general news, which many expect to become commoditised by AI chatbots.

Publisher perceptions on how content needs to be refocused in the era of AI

Net difference between the proportion saying each is more/less important to focus on in the era of AI



Q13. Thinking more widely about how journalism can compete in the era of AI, which of the following will be more or less important for newsrooms to focus on? Base = 264.

The majority of surveyed news executives believe that future success lies in being more distinctive, even if that means losing some overall reach. 'Journalism's best response is to double down on the things that make us valuable and unique', says Taneth Evans who is Head of Digital at the Wall Street Journal.

'AI platforms will continue to scrape our websites and repackage our content', says Martin Schori, Head of AI at Aftonbladet. 'The answer to that is probably to focus on journalism that can't easily be summarised in the three bullet points', he says.

Publishers say they will be putting more effort into video (+79) and audio (+71) this year, formats that are proving effective in engaging younger audiences in particular, but that are also harder for AI to reformat and extract value from.



The agentic web


More change is on the way. Experts suggest that websites and apps will soon have more traffic from AI bots than from humans as AI systems acting on behalf of users autonomously complete tasks such as compiling a digest of personalised news. Open AI's Pulse, for example, offers a card-based morning briefing linked to previous Chat GPT queries. New AI browsers are also appearing that can handle tasks such as summarising or translating webpages, email drafting, and shopping - based on natural-language commands. In our survey, three-quarters of respondents (75%) see these changes as having a large or very large impact on publishers over the next three years.

These developments mean that content is becoming increasingly 'liquid', in that the format can be changed - actively or passively - based on the viewer's context, interaction, time, or location. This will make it harder for publishers to control how news stories look in the future - and to know how content is being used. Some publishers are suing AI platforms as they try to protect their intellectual property and others are cutting deals, but expectations are low that they will get enough money to compensate for declines and losses elsewhere.

For PR, these shifts have important implications. Historically, media coverage has delivered value not only through the initial audience of a publication, but through secondary discovery via search and social. As this referral ecosystem weakens and more people access content through AI agents and aggregators, the downstream visibility of earned media becomes less predictable.

Beyond that, if content becomes increasingly mediated by AI systems that summarise, extract, and repackage information, attribution may be diluted, and key messages risk being reduced to fragments. In the future, content needs to be designed not only for human audiences, but also for interpretation by machines. Messages that are explicit, authoritative, and easily extractable are more likely to survive the transition into AI-generated summaries and responses. The emerging discipline of 'Answer Engine Optimisation' signals a move beyond traditional SEO to identify how content can be optimised for AI interfaces.

Expected impact of agentic technology in the next three years



Q15. When it comes to AI browsers that can summarise news (Comet, Atlas, etc.) and agentic AI that can assemble personalised briefings (e.g. Open AI Pulse, Huxe, etc.), how much impact do you expect these to have over the next three years? Base = 265

The fragmentation of influence


Alongside technological change, our trends report underscores the continued rise of creators and personality-led media. Influencers, independent journalists, and platform-native voices are capturing increasing shares of audience attention, often at the expense of traditional news brands.

This trend reflects a broader shift in trust. Audiences - particularly younger demographics - are placing greater value on authenticity, relatability, and direct connection. As a result, individuals are becoming as influential as institutions in shaping public understanding of news and current affairs.

For PR professionals, this represents both an opportunity and a challenge. On one hand, the diversification of the media landscape creates new routes to reach highly engaged and targeted audiences. Creator partnerships, podcast appearances, and platform-native storytelling offer ways to deliver impact beyond traditional media.


On the other hand, this fragmentation complicates messaging and the management of reputation. Creators operate outside the editorial frameworks and norms of traditional journalism, often blending reporting, opinion, and personal perspective. This can introduce variability in tone, context, and interpretation that may be difficult to manage. Influence is no longer concentrated within a defined set of media gatekeepers. It is distributed across a more complex ecosystem of journalists, creators, platforms, and communities.

The rapid expansion of AI-generated content introduces a further layer of complexity. Our trends report highlights the proliferation of so-called "AI slop" – low-quality, automated material – as well as the growing sophistication of deepfakes and synthetic media. This environment presents a paradox. While the abundance of unreliable content has the potential to increase demand for high-quality, verified journalism, it does not necessarily translate into renewed trust in traditional media. Audiences may instead turn to familiar voices, peer networks, or even AI tools themselves to interpret information.






These changes reflect a period of significant transition for news publishers as well as PR and communications professionals. The challenge will be to adapt to an information ecosystem in which visibility is less predictable, influence is more fragmented, and trust is more contested. Those that succeed will be those that recognise these shifts early and respond by developing strategies that are platform-aware, audience-aware, and increasingly, AI-aware. In a world where information is abundant and attention is finite, the ability to communicate clearly, credibly, and effectively across multiple environments will be more important than ever.



The challenge will be to adapt to an information ecosystem in which visibility is less predictable, influence is more fragmented, and trust is more contested. Those that succeed will be those that recognise these shifts early and respond by developing strategies that are platform-aware, audience-aware, and increasingly, AI-aware.

You can see the full report on Journalism, Media and Technology trends 2026 including details of the survey methodology at <https://reutersinstitute.politics.ox.ac.uk/journalism-media-and-technology-trends-and-predictions-2026>



THINKING STRATEGICALLY ABOUT AI IN EDUCATION COMMUNICATIONS


Stephen Waddington and Ben Verinder, authors of *AI for Public Relations: A How-To Guide for Implementation and Management*

Most education communications teams have embraced AI to improve efficiency and effectiveness. In this section of the report we argue that AI should be treated as a strategic asset and a governance responsibility and set out a practical framework for doing so.

AI is already embedded in education communications practice. It is being used formally through enterprise tools such as Microsoft Copilot, and informally through individual use of systems such as ChatGPT. The question of whether to adopt AI has largely been resolved. The next and arguably more important question is how to use it effectively.

At present, many practitioners in universities, schools and colleges are approaching AI tactically. They are using it to accelerate existing workflows without reassessing what those workflows are intended to achieve. These risks reinforce existing limitations in the function rather than addressing them.

Drawing on the community of practice and evidence base developed for *AI for Public Relations: A How-To Guide for Implementation and Management*, we apply those insights to education communications. The central argument is how AI represents both a strategic opportunity and a risk issue whereas most communication teams still treat it as either-or rather than both.



At present, many practitioners in universities, schools and colleges are approaching AI tactically.

AI as a strategic opportunity and risk

Public relations and the broad practice of communications has long sought recognition as a management function. In practice, however, it is often defined by outputs such as press releases, social media content and events, rather than by its contribution to organisational decision-making. AI creates an opportunity to shift this position. Automating elements of production can free capacity to be redirected toward higher-value activities such as stakeholder mapping, sensemaking and strategic counsel.



Adoption of AI among practitioners has increased sharply in recent years, rising from under one-third in early 2023 to roughly three-quarters by early 2026, although routine daily use is lower and many practitioners are still experimenting rather than working in a structured way. [1] At the same time, claims about productivity gains are becoming more substantial. Research from the London School of Economics and Political Science suggests that AI can unlock significant time savings across the workforce, but those gains become strategically meaningful only when they translate into better judgment, stronger planning and improved decision-making rather than simply more output. [2]

The governance gap

The implications for education institutions are particularly significant. Schools, colleges and universities manage sensitive data relating to students, staff and stakeholders. They operate under scrutiny from regulators, governing bodies, parents, learners, journalists and the wider public. The use of AI introduces new risks in this context, particularly where tools are used outside formal governance structures.

When practitioners, or any other colleagues in the organisation, input institutional information into unapproved AI systems - whether for drafting, summarisation or analysis - they may inadvertently expose sensitive material. This creates both data protection risk and reputational risk. The issue is intensified by the fact that in many organisations AI adoption is occurring ahead of policy development.

In our practice we commonly find communications teams working without either team or institutional guidance or, **where a policy does exist, it focuses on pedagogical matters and fails to touch on the core issues facing professional services teams.** Education is not an outlier: across public relations and communications practice, widespread adoption is accompanied by a marked governance gap: many practitioners are using AI without training, policy guidance, or clear responsible-AI frameworks. [3]

In practical terms, that means teams may be working without clear agreement on approved platforms, data handling protocols, quality assurance processes or disclosure requirements. In education, where accountability structures are often formal and public expectations are high, this gap is a material risk that requires immediate attention.



In many organisations AI adoption is occurring ahead of policy development.

Alternatively, they may be avoiding artificial intelligence tools altogether because of a fear of breaching confidences.

One institution demonstrating what good practice looks like is the University of Cambridge. Amy Mollet, Head of Social Media at Cambridge and a contributor to *AI for Public Relations: A How-To Guide for Implementation and Management*, developed an AI policy that has become a reference point for the sector. It is publicly available on the University of Cambridge website. Her chapter in the book addresses AI governance and policy development in detail. [4]

Speed without judgement

A second risk is more subtle but equally significant: the erosion of professional judgement. AI systems can generate credible content quickly. However, they do not replace editorial judgement, contextual awareness or ethical decision-making. When outputs are used without appropriate review, responsibility is effectively delegated to the system.

There is clear anxiety within public relations practice about this shift. Three-quarters of public relations practitioners using AI say the biggest risk is that junior colleagues will 'skip the basics', weakening their ability to deliver core tasks and to judge the quality of AI output. [4] Related concerns about cognitive offloading and the erosion of the

apprenticeship stage of work are a recurring theme throughout the book. [6]

This has important implications for learning and development in education communications. Expertise has traditionally been developed through practice: learning to assess tone, accuracy, appropriateness, timing and institutional sensitivity over time. If AI bypasses those stages, capability may decline. This is particularly relevant in education settings, where communications teams are often small and may lack structured development pathways.

Strategic and career implications

The strategic challenge is how to preserve and clarify the distinctive value of human practitioners. The UK Government has identified public relations practitioners as among the occupations most exposed to AI. [7]

Practitioners who continue to demonstrate value in an AI-enabled environment are those who exercise ethical judgement, understand institutional context and politics, build and maintain stakeholder relationships, and anticipate and manage reputational risk. These capabilities are not easily automated. [8]

The strategic challenge is how to preserve and clarify the distinctive value of human practitioners.



There are also signs of divergence in labour market capabilities. Our source material notes anecdotal evidence from training and consultancy work that junior and mid-career practitioners are already moving on from organisations they perceive to be too far behind the curve in AI adoption. In other words, failing to think strategically about AI may become both a capability problem and a retention problem.

Three strategic imperatives

We suggest that effective AI adoption in education communications can be structured around three core imperatives.

1. Embed AI deliberately in workflows

AI should be integrated into clearly defined tasks rather than applied ad hoc. The aim is to improve efficiency and effectiveness, but with judgment and boundaries. That means identifying appropriate use cases, developing governance, investing in training and being clear where human intervention is required.

2. Understand AI's impact on reputation

AI systems are trained on and synthesise public information. This means that what an institution publishes online - and what others publish about it - now contributes to how it is represented in AI-generated outputs.

For educational institutions, this matters because prospective students, parents, staff, journalists, and policymakers are increasingly likely to use AI tools for discovery and interpretation. When someone asks an AI model about a university, school or education charity, the output reflects aggregated signals rather than carefully controlled messaging.

Beyond the digital domain, we also need to understand how our organisation's use of AI is affecting, and will affect, its relationships with stakeholders. This includes the deployment of AI among teaching staff, the use of AI by students and the nature of the educational offer in an AI era.

3. Lead on ethics, governance and trust

AI raises questions of bias, transparency, accountability and legitimacy. Communications practitioners are well placed to contribute to organisational responses to these questions because they sit at the intersection of reputation, stakeholder expectations and institutional change.

In education, institutions are expected only to use AI responsibly and explain their use in ways that maintain trust. The issue is whether stakeholders believe that AI is being used legitimately, proportionately, and in ways consistent with institutional values.

Table: The three strategic AI imperatives for education communications

Imperative	What it means	Risk if ignored	Education-specific implication
Embed AI in workflows	Structured use cases	Ad hoc usage	Inconsistent practice
Understand AI and reputation	AI as interpretive layer	Loss of narrative control	Negative perception from prospective students, parents and other publics
Lead on ethics and trust	Governance and transparency	Reputational damage	Public accountability

A framework for strategic AI adoption

AI adoption is often driven by individual experimentation rather than organisational intent. A more effective approach is to treat adoption as a structured, iterative process that aligns tools, capability and governance with institutional objectives. Below is an adaptation of the BRIDGES framework for team adoption, developed by Ben Verinder.

1. Conduct an audit	Establish which tools are currently in use, both formal and informal; what governance arrangements already exist; and what levels of capability and training are present across the team. In many organisations, informal use significantly exceeds formal adoption.
2. Start with low-risk applications	Initial use cases might include background research, internal briefing documents, summarisation and first-draft content. These provide value while limiting exposure. Early experimentation should be designed to build understanding of capability.
3. Establish governance early	An effective AI policy should be concise and practical. It should cover approved tools, data restrictions, review and quality assurance processes, disclosure expectations and escalation routes for higher risk uses. In educational settings, communications policy should also align with the institution's wider positions on safeguarding, data protection, assessment integrity, and staff conduct.
4. Invest in capability	Training should focus on critical evaluation, prompt design, editorial judgement and ethical decision-making. Teams need to know how to challenge outputs. They also need enough shared understanding to avoid over-reliance on unapproved or poorly understood systems
5. Evaluate before scaling	AI does not always reduce workload in net terms. In some cases, time saved in drafting is offset by time spent reviewing, correcting or reworking outputs. Pilot programmes should therefore be evaluated rigorously before expansion, with attention paid to quality, risk, capability development and stakeholder impact as well as apparent time savings.



Education-specific considerations

Several considerations require particular attention in education communications and the wider community of practice working in organisations including businesses, charities and think tanks.

First, data privacy and safeguarding should be treated as baseline strategic issues. Education organisations hold sensitive data about children, young people and adult learners. Personally identifiable or confidential information should not be entered into consumer-grade tools without formal approval and a clear understanding of data retention, access and processing arrangements.

Second, public trust places a higher bar on authenticity and disclosure. Educational institutions are expected to model responsible behaviour. If AI use is perceived as a shortcut that compromises authenticity or care, the reputational consequences may be greater than in some other sectors. Transparency about how and where AI is used can therefore support trust, provided that the underlying use is legitimate.

Third, governance structures in education settings require particular care. Boards of trustees, governing bodies, executive teams and, in higher education, academic governance structures all shape the context in which communications teams operate. AI adoption should be handled within those structures. Communications leaders should bring senior decision-makers into the conversation early, ensure that communications use aligns with any broader institutional AI strategy, and surface risks clearly rather than absorbing them silently.

Table: Risk typology

Risk type	Example	Impact	Mitigation
Data exposure	Inputting student data	Legal and reputational	Approved tools only
Governance gap	No policy	Inconsistent practice	AI framework
Judgement erosion	Over-reliance on AI	Quality decline	Editorial review
Trust risk	Undisclosed AI use	Stakeholder distrust	Transparency

How AI is likely to shape management communications strategy

Looking ahead, AI is likely to affect management communications strategy in at least three ways.

First, agentic AI will increase the pressure for stronger oversight. As systems become more capable of conducting multi-step tasks – such as research, drafting, monitoring and workflow management – questions of accountability, review and escalation will become more pressing. The issue for education institutions is whether governance will keep pace.

Second, reputation management will become more dependent on understanding how institutions appear in AI systems. This will require some equivalent of 'LLM listening' alongside media monitoring and analytics.

Third, the capability gap between well-resourced and under-resourced organisations may widen. Institutions with the capacity to experiment, govern and train will progress more quickly than those without it. Professional networks such as the CIPR Education and Skills group have an important role to play in ensuring that strategic AI capability does not become concentrated only in the best-resourced parts of the sector.

Conclusion

For education communications practitioners, AI should be treated as a strategic enabler, a governance challenge and a trust issue.

Used well, AI can strengthen strategic capability by creating more space for insight, planning, evaluation and counsel. Used badly, it can accelerate weak practice, erode judgment, and expose institutions to unnecessary risk.

The task for education communications leaders is therefore to build the conditions in which experimentation is disciplined, governed and aligned with institutional purpose.

REFERENCES

- [1] Muck Rack (2025) State of AI in PR 2025.
- [2] London School of Economics and Political Science (2025) Bridging the Generational AI Gap: Unlocking Productivity for All Generations.
- [3] Reimagining Tomorrow: AI in PR and Communication Management (2025) Global Alliance for Public Relations and Communication Management.
- [4] Mollett, A. (2023) How We Use Generative AI Tools, University of Cambridge.
- [5] Axios (2025) PR professionals have tripled AI use since 2023.
- [6] Mollick, E. (2024) Co-intelligence: Living and Working with AI. London: WH Allen.
- [7] Department for Education (2023) The Impact of AI on UK Jobs and Training.
- [8] Waddington, S. and Smith, A. (2023) Humans Needed More Than Ever. London: Chartered Institute of Public Relations.



About our book


Published by Kogan Page in May 2026, *AI for Public Relations: A How-To Guide for Implementation and Management* is a comprehensive guide to AI adoption for communications and public relations practitioners.

Edited by Ben Verinder and Stephen Waddington, the book brings together more than 20 contributors, including agency leaders, in-house directors, researchers and technical specialists. It addresses the full spectrum of AI integration in communications practice – from strategic planning and governance through to workflow integration, team development and risk management.

The book is structured to support practitioners at different stages of adoption. It includes practical frameworks for implementation, guidance on policy development, and applied use cases across disciplines such as media relations, internal communications, crisis management and measurement.

Several chapters explore emerging areas of practice, including Generative Engine Optimisation (GEO), alongside detailed consideration of technical, behavioural and ethical risks. The role of communications practitioners in shaping organisational responses to AI – particularly in relation to trust, transparency and accountability – is a central theme throughout.

The book is intended as a practical reference for practitioners responsible for integrating AI into communications functions, whether in-house or agency-side, and for those advising senior leadership on its strategic and reputational implications.



Stephen Waddington and Ben Verinder



© 2026 Stephen Waddington and Ben Verinder. All rights reserved. No part of this article may be reproduced, distributed, transmitted, stored, or used in any form or by any means, including but not limited to uploading to, training, fine-tuning, or prompting artificial intelligence systems, machine learning models, or similar technologies, without the prior written permission of the authors.

FURTHER READING AND RESOURCES

The list below is based on regularly quoted organisations, individuals and research. It is not an exhaustive list but a starting point. We accept no responsibility for the information provided by third parties.

Individuals talking about AI, comms and marketing

- Hal Kimber, The Marketing and AI Podcast – <https://themappodcast.co.uk/>
- Al Kingsley, AI's Edu Recommends (more education than comms & marketing but a useful monthly edtech and AI update) – <https://www.linkedin.com/in/alkingsley1/>
- Priya Lakhani OBE – <https://www.linkedin.com/in/priyalakhani/>
- Professor Rose Luckin, The Skinny – <https://www.educateventures.com/the-skinny>
- Ethan Mollick, One Useful Thing – <https://www.oneusefulthing.org/>
- Ben Verinder – <https://www.linkedin.com/in/ben-verinder/>
- Stephen Waddington, Management and Public Relations – <https://wadds.substack.com/>

Industry bodies and associations

- Chartered Institute of Marketing (ongoing) – <https://www.cim.co.uk/content-insights/ai/>
- Chartered Institute of Public Relations, AI in Public Relations Panel (established 2018) – https://cipr.co.uk/CIPR/CIPR/Our_work/Policy/AI_in_PR.aspx
- Chartered Institute of Public Relations, Humans Needed More than Ever, Professor Anne Gregory and Dr Swati Virmani, 2023 – https://cipr.co.uk/CIPR/Our_work/Policy/AI_in_PR_/AI_in_PR_guides.aspx
- IAB, Best Practice Guide for the Responsible Use of Generative AI in Advertising, 2026 – <https://www.iabuk.com/ai-digital-advertising>
- IPA, Industry AI resources, (ongoing) – <https://ipa.co.uk/knowledge/industry-ai-resources>
- Muck Rack, The State of AI in PR UK 2025, 2025 – <https://muckrack.com/resources/research/state-of-ai-uk-2025>
- PR Moment, Regularly covers AI and offers free webinars – <https://www.prmoment.com/>
- Public Relations and Communications Association (PRCA), Setting the Standard for Responsible AI, Jan 2026 – https://cipr.co.uk/CIPR/CIPR/Our_work/Policy/AI_in_PR.aspx




Public sector and regulatory guidance (UK)

- Department for Education, Generative artificial intelligence in education, August 2025 – <https://www.gov.uk/government/publications/generative-artificial-intelligence-in-education/generative-artificial-intelligence-ai-in-education>
- House of Lords Communications and Digital Committee, 2024, Large Language Models and generative AI – <https://www.ofcom.org.uk/about-ofcom/annual-reports-and-plans/ofcoms-strategic-approach-to-ai>
- Information Commissioner's Office, ongoing work on artificial intelligence – <https://ico.org.uk/about-the-ico/what-we-do/our-work-on-artificial-intelligence/>
- JISC, A Strategic Framework for AI in Colleges and Universities, August 2025 – <https://nationalcentreforai.jiscinvolve.org/wp/2025/08/28/a-strategic-framework-for-ai-in-colleges-and-universities/>
- Nesta, Building trust in AI being used by the public sector – <https://www.nesta.org.uk/project/public-ai-assurance/>
- Ofcom's strategic approach to AI, 2025 – <https://www.ofcom.org.uk/about-ofcom/annual-reports-and-plans/ofcoms-strategic-approach-to-ai>
- Reuters Institute for the Study of Journalism (ongoing) – <https://reutersinstitute.politics.ox.ac.uk/ai-journalism-future-news>
- Scottish Government, Artificial Intelligence (AI) in schools: guidelines and guardrails – <https://www.gov.scot/publications/scottish-guidelines-guardrails-use-artificial-intelligence-ai-schools/pages/1/>
- Welsh Government, Generative artificial intelligence in education – <https://www.gov.scot/publications/scottish-guidelines-guardrails-use-artificial-intelligence-ai-schools/pages/1/>

Free training and structured learning

- Anthropic Academy (Claude) – <https://www.anthropic.com/learn>
- BBC Bitesize Guide to AI (aimed at teenagers) – <https://www.bbc.co.uk/bitesize/groups/cgk3rjokl55t>
- Canva Design School AI resource – <https://www.canva.com/design-school/courses/ai-in-the-classroom>
- Google AI and Gemini help centre resources – <https://support.google.com/a/users/answer/14506784?hl=en>
- HubSpot Academy – <https://academy.hubspot.com/>
- Muck Rack Academy – <https://academy.muckrack.com/>
- Notion AI Guides – <https://www.notion.com/en-gb/help/guides/category/ai>
- Open AI Academy – <https://academy.openai.com/>
- SEMrush Academy – <https://www.semrush.com/academy/>





Education
Comms
Network

