



2023 Round Tables on AI and the Global News Industry

Jessica Cecil



Summary


The Reuters Institute for the Study of Journalism (RISJ) at the University of Oxford held three round tables across 2023, under the Chatham House rule, to consider the effect of emerging artificial intelligence (AI) technology, and in particular generative AI (GenAI), on the creation of news. The aim was to capture the emerging thinking of industry leaders, with participants from global tech platforms and global news organisations, and experts specialising in the detection of disinformation.

Our main learnings

- There was a very high willingness among participating news publishers to experiment with new GenAI tools.
- This was combined with an appreciation that the playbook for what works would have to come from them and their own experiments because of the nature of the technology: GenAI reveals what it is capable of through trial and error, rather than a set of use cases predefined by the product engineers.
- There were already innovative experiments at scale under way at Global South publishers who participated.
- But Global South publishers are facing a more stark risk–benefit trade-off as they tend to work in a tougher financial environment. Tools like translation into multiple languages are potentially highly impactful for them in opening up new markets. At the same time, the boldest uses of AI, like simultaneous translation, come with risks that can't yet be mitigated because it is impossible to check in advance for errors.
- Global North publishers we spoke to were more cautious, though they may stand to benefit from the experimentation of others.
- There was great interest in technical insights into how AI can tackle disinformation.
- Many news organisations were concerned that their interests do not align with those of tech platforms.
- Intellectual property (IP) issues are top of mind for how many news publishers think about AI, and this informs their approach to collaboration in other areas.

Context: grappling with change in a momentous year

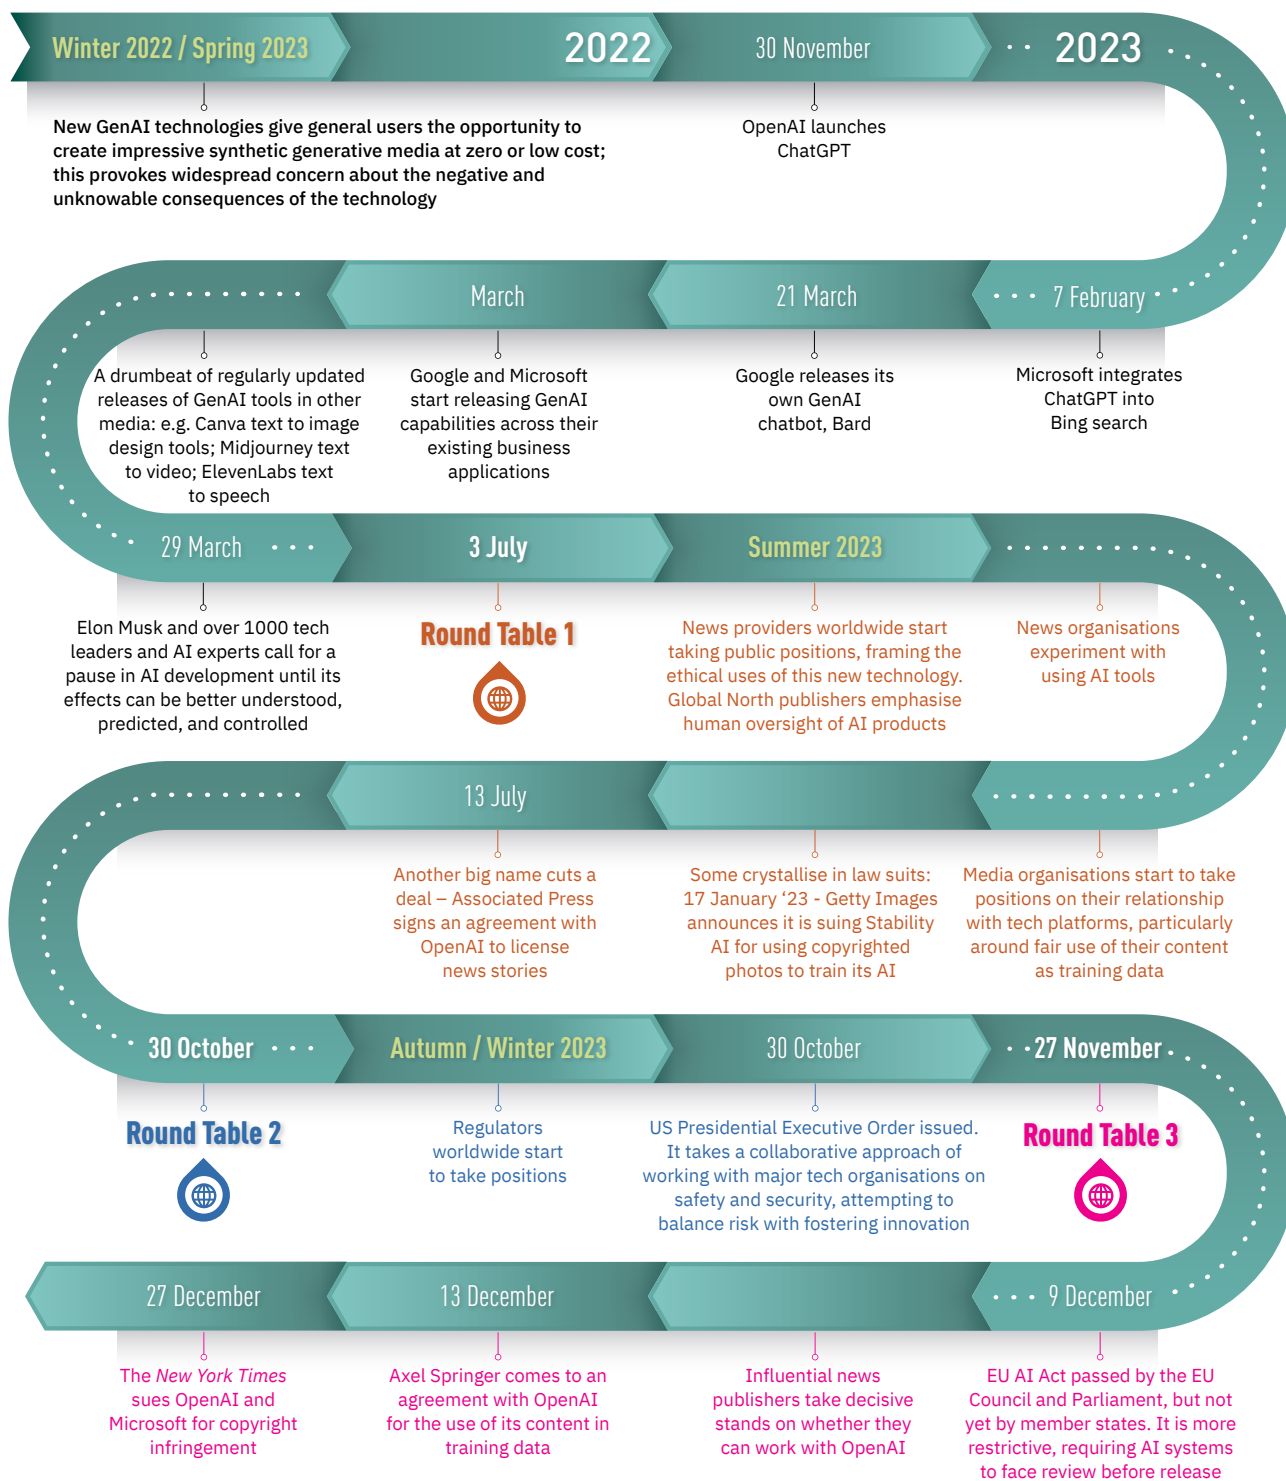
GenAI came to the fore in 2023 as a deeply disruptive technology, forcing almost all industries to think differently about how they operate. That includes the global news industry. Across the course of the year, the RISJ at the University of Oxford convened a series of three round tables, under the Chatham House rule, to consider the effect of emerging AI technology, and in particular GenAI, on news.



The aim was to capture the emerging thinking of industry leaders as they wrestled with a knotty challenge: how to embrace a technology that could not be ignored, when no roadmap for seizing opportunities and mitigating risks existed. Participants were genuinely uncertain about the direction of the new technology, and this made for particularly rich and open discussions. As one participant vividly put it, we are ‘caught in a tension between wanting to fling ourselves into it [GenAI] and wanting to run away from it’.

These round tables had a distinctive mix of participants – global tech platforms, global news organisations, and experts specialising in the detection of disinformation. There was also a focus on practical takeaways, and participants were asked to describe their own trials in using AI. These provided a foundation on which to build general insights from localised learning. This approach, at a more granular level than most round tables, tried to avoid the pitfall of restating issues and points that are already known.

There were different participants across the three sessions, and they reflected evolving thinking about the threats and opportunities of this highly disruptive new technology. This coincided with a rapid succession of industry changes across the year, as new products were launched, in turn resulting in new policies within news organisations, reframed tensions between tech platforms and news organisations, and new positions from regulators. The infographic is an attempt to characterise the year in terms of these developments, putting the round tables in a wider context.



Across 2023, the shape of the new terrain began to emerge, and with it discussion moved from first-principle concerns towards

a greater willingness to accept change is happening, look at the details, and get on with innovation.

Learnings from round tables

There was an optimism mixed with concern in the first round table (July 2023) among the participating publishers, many still attempting to find a framework for innovation. This reflected the speed of new AI releases in the spring.

As one put it:

The opportunities are so wide and the playing field is so big that we do not know where to start. Where we need to go to next is having a task force, and a detailed plan and goals. Rather than being reactive, the next phase has to be strategic, linking up to mission and editorial goals, no experimentation without focus.

After the summer, the discussion became less optimistic and more granular, and focused on three subjects:

- **Newsroom experiments using AI tools;**
- **AI-fuelled disinformation;**
- **News providers' relationship with platforms, including news IP used in training data.**

More broadly, news providers recognised they had to 'stop regarding AI as a silo'. They all saw one of the potential biggest prizes of AI was making more time in the day for already overstretched journalists to find original stories, by automating other tasks.

Experiments in using AI tools

Overall, we found that, as the year went on, there was a very high willingness among many news publishers to experiment with new GenAI tools.

This was combined with an appreciation that the playbook for what works would have to come from them and their own experiments.

It became increasingly clear that the nature of GenAI lends itself to learning from doing, and drawing conclusions and strategic direction from the results of individual experiments.

GLOBAL SOUTH EXPERIENCES

There was a strong willingness among participating Global South publishers to undertake bolder experiments. They were prepared to be nimble and novel. GenAI gave them an ability to offer their content across multiple languages for the first time and set up pop-up newsrooms in remote locations at an affordable marginal cost. This included experimenting with tools for oversight – for instance, an experiment to create a GenAI programme to flag legal risks in copy, specifically tailored to the laws of their country.

There was also a willingness to embrace the possibilities of AI across types of content, such as rewriting non-original and non-reported text, summaries, and context, and for search engine optimisation in headlines. They were also comfortable with using GenAI to generate generic illustrations and imagery. But one Global South publisher believed this came at an extra cost to them: they reported noticeable Western bias in GenAI, in imagery and language, and this meant their experiments also had to include workarounds and elaborate prompts.

For one participant, this reflected a bigger issue, that AI could have the effect of concentrating even more power in the Global North:

“We run the risk of the concentration of power in the AI space which mirrors earlier trends. We are likely running into some of the same risks we have seen historically.”

This ambitious approach to innovation demanded very granular guidelines about when the use of GenAI was and wasn't appropriate, and it needed innovative approaches to mitigating risks and addressing ethical issues. For instance, one newsroom considered adopting an AI newsreader to offer audiovisual versions of text news stories, but they were keen not to mislead audiences. They proposed solving this problem by creating newsreaders in the style of a graphic novel to signal these were not actual human journalists.

There was a red line for all to protect the credibility of their own journalists. Explicitly prohibited was adopting the style of a particular writer to generate related or unrelated text, mimicking their words. There was also a prohibition on creating photorealistic avatars of specific journalists. These may have been a response to the growing concern among several news providers about the AI-enabled impersonation of their own journalists.

Global South publishers were undoubtedly confronting a more stark risk-benefit trade-off. Facing the greatest economic and often political pressures, they were also using AI in ways that offered them the greatest benefits to cut costs and reach new audiences. But that came with potential risks that can't yet be mitigated at scale. One of the ways in which Global South publishers are using GenAI is for the simultaneous translation of a piece of news

content into multiple languages inside a single country. While this could dramatically cut costs and open up new audiences, it is impossible to quality-control all these different versions and to check for hallucinations.

GLOBAL NORTH EXPERIENCES

This contrasted with a more cautious approach in general taken by participating Global North publishers.

There was interest in exploring using GenAI to generate background and context boxes, which are particularly valued by those who consume less news.

There was also some interest in using GenAI in a much more ambitious way to think creatively about how to meet the needs of audiences. One news executive observed:

“We could fine-tune our own LLMs [large language models] on our own data and archives, going on the offensive, to create a better experience in terms of relatability. How can we use some of these models in our relationships with people and to make our organisations work better?”

Overall, for the moment, these Global North publishers tended to favour a co-pilot approach that allows back office efficiency gains but has a human check on audience-facing content creation. This may reflect their greater size and longer history. Nevertheless, all publishers stand to benefit from the innovative research and development undertaken by Global South players.



Misinformation and disinformation

All participants recognised the potential dangers of low to zero cost GenAI outputs designed to mislead, particularly in the context of the unprecedented number of elections globally in 2024. Many participants worried that the believability of all news content is being called into question – not just for audiences, but for newsroom editors, who are becoming more sceptical of the content they interact with, even from their own stringers.

TOPICS AND CASES

Again as 2023 went on, there was more focus on specific threats and how to deal with them. Some were exacerbated by AI tools, but not always: content about the Israel–Gaza conflict reinforced that even ‘shallow fake’ disinformation can widen divisions and cause real-world harm. Looking forward, there was concern around imposter websites, purporting to be from established news media. There was also a widespread concern that AI allows individual journalists to be imitated.

MANIPULATED AUDIO

There was a growing recognition that manipulated audio is hardest to detect. That means countries in which audio is an important news medium face particular challenges around disinformation in upcoming elections. Above all, there was a concern about the volume of disinformation that the low cost and high believability of GenAI could unleash. One participant predicted that Russia would exploit this, and misinformation factories there, using AI tools like translation, would create content fast and cheaply in more languages.

DETECTING MISINFORMATION

There was huge interest from all participants in the methodology and tools to detect disinformation campaigns. They envisioned an

AI-enabled arms race with so-called bad actors trying to outwit those detecting disinformation. Bad actors are learning to use LLMs to improve the quality of what they are posting, and are making it hyper-personalised. In response, those detecting disinformation are analysing the hidden patterns left in synthetic content and using natural language processing (NLP) techniques to classify disinformation and the provenance of content.

Nevertheless, the news providers were looking for practical ways to access that insight. They asked for early detection and shared insights for mis-/disinformation narratives. They wanted help in specific instances: fact-checking tools, and the development of specific tools to detect fake audio. They also wanted guidance: how to deal fast with a journalist’s impersonation and imposter news websites. One participant asked for a labelling protocol to advise when a post by a reputable new organisation is genuine.

FLAGGING MISINFORMATION

There was a concern from some that as technology has developed, the issue of calling out factual inaccuracy remains. This was in the context of a retreat from established content moderation positions, for instance at X (formerly Twitter), and the politicisation of news in the context of the Israel–Gaza conflict. As one participant put it, ‘The technological problems seem to be somewhat solvable, but that is harder with the veracity of facts.’

There was a more optimistic discussion about how AI could create tools that would aid community cohesion. These included the creation of ultra-local content valued by communities, for instance school sports reports. There is also now the ability to detect, at scale, journalistic phrases that are likely to exacerbate divisions and inflame tensions.

News providers' relationship with platforms

Many news organisations were concerned that their interests do not align with those of tech platforms. As the year went on, so did concern that their news content might be used to train LLMs without their consent. Some news organisations said they were undertaking experiments within their own sandboxes, effectively isolated digital environments, to prevent that happening. As one news executive put it: 'We feel we are being played. IP has to be part of the battle.'

FINANCIAL PROBLEMS

Linked to that were predictions by many of the news organisations of the dire financial consequences of AI-enabled search cutting referral traffic to news publishers' own sites. This in turn affected the financial viability of news organisations. AI-enabled search could offer a single complete narrative answer, rather than a number of links to news sites providing supporting evidence. News organisations believed users would be far less likely, faced with this complete answer, to seek out news sites and subscribe to them. They said this was a problem for the organisations, but it also had wider consequences: 'How do you have a healthy media ecosystem if your journalism is not thriving?' asked one publisher.

Participants did not see a simple solution. News organisations did not necessarily want to be cited as a source if there was a risk that the AI-generated search answer had inaccuracies, or hallucinations. It might damage the news organisation's reputation for accuracy and could appear to legitimise AI-generated results.

PROTECTING IP

It was clear from the discussions that there is a potential for these IP-related issues to dominate the relationship between tech platforms and news organisations, and potentially prevent collaboration in the other areas outlined above, where working together is in the interests of all participants.

Concluding remarks

There is clear value in news organisations sharing learnings on AI experiments. There is value to all in crowdsourcing insight so that best practice can be established and mistakes not repeated. But for this to work, the most innovative newsrooms need to be rewarded for the disproportionate insight they bring. This can be achieved through giving them access to platforms, involving them in global conversations on the sustainability of news, and supporting them by understanding the political and legal threats they face.

Technology to detect AI-enabled disinformation is developing fast. But this is not matched by mechanisms to share insights around disinformation campaigns among news organisations. Nor is there widespread access to specific fact-checking tools, particularly for audio fakes. Finally, there is not a clear roadmap for news organisations to remedy imposter content and appropriation of journalists' identities.

Two related issues are top of mind for many publishers. They are IP and the fair use of news content for LLM training data, and fears around changes to search. These two issues inform how publishers are approaching the wider possibilities of GenAI.

Media organisations could work in partnerships and collaborate to address the challenges posed by AI-generated content, because individual organisations may not have the necessary resources to tackle these issues effectively. In particular, the learnings from experiments in using AI could be shared among news publishers. This could be through regular structured and frank practitioner conversations.

The development of fact-checking as a paid service provided by some news publishers to specific platforms might form the basis for future collaboration around developing safeguards and tools to identify AI-generated information. So too could a fast alert for disinformation that threatens real-world harm, especially around elections.

At the same time, it is important to recognise the reality that cooperation between tech platforms and news organisations has been strained and any consensus on content moderation principles has been fractured. They also often have competing commercial interests. This places an emphasis on initiatives that are very specific, limited in scope, and directed towards tangible outcomes. And as ever, there is value in building on the models of what already works.

ABOUT THE AUTHOR

Jessica Cecil was in 2023 a Visiting Fellow at the Reuters Institute for the Study of Journalism at the University of Oxford. She is a media executive and was previously a senior leader at the BBC and Chief of Staff to four Directors-General. She founded and led the Trusted News Initiative (TNI), the world's only alliance of major international tech companies and news organisations to counter the most harmful disinformation in real time.

She continues to specialise in the field of how to combat disinformation, with an emphasis on how news and media are affected by emerging technology. She does that commercially and is also a non-executive director at the Digital Catapult, on the Council of Advisors for RAND Europe, and an Adjunct Fellow of the Queen Elizabeth II Leadership School at Chatham House, the Royal Institute for International Affairs. She is a Trustee at the University of Bristol.

ACKNOWLEDGEMENTS

The author is grateful to the news organisations, tech platforms and disinformation detection companies that joined us at a senior level, and to Felix Simon, Richard Fletcher, Rasmus Kleis Nielsen, Federica Cherubini, and Louise Allcock for their support and input.

