AI in the News
Reshaping the Information Ecosystem?

Felix M. Simon & Luisa Fernanda Isaza-Ibarra
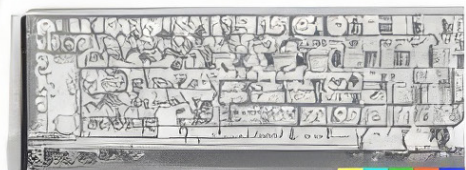
TABLE OF CONTENTS

Keywords: Artificial Intelligence, AI,
LLMs, Generative AI, News, Journalism,
News Organisations, Platform Companies,
Information Environment

All graphics in this report were created with Open AI's Dall-E 2 and Midjourney

EXECUTIVE SUMMARY

Main Findings

- The definition of artificial intelligence (AI) remains debated, AI does not currently possess conscious, general intelligence (AGI), and experts consider the portrayal of AGI in popular culture as fiction.

- Current AI systems based on various forms of machine learning (ML) are narrow, focusing on specific tasks associated with human abilities. These systems learn from data through algorithms trained on large datasets using methods like supervised, unsupervised, and reinforcement learning.

- AI in the news industry is viewed by many as the computational simulation of human activities and skills in specific domains. Mostly, this means applications of ML systems and increasingly Large Language Models (LLMs). However, the line between AI and other technologies and forms of automation is often blurry.

- AI systems already play a concealed but significant role in shaping the information landscape. They filter, curate, rank, and moderate content on platforms and search engines. Increasingly, news organisations, too, use AI for journalistic tasks and in distribution processes.

- The implementation of these AI systems often requires significant time and resource investments, and their development may initially be inefficient. While they promise greater efficiency and productivity for news-work, this is not a foregone conclusion as AI-generated work may require e.g., additional editing or careful supervision.

- The impact of AI on efficiency and productivity varies depending on the task and context. Large language models lack true thinking capacity and cannot fully automate tasks requiring critical thinking and contextual understanding.

- Other challenges in AI integration in news organisations include the risk of unwanted biases, privacy and intellectual property concerns, and the disruption of newsroom dynamics. Many news organisations prioritise reliability and trustworthiness, setting limits on AI use.

- Large technology companies already exert significant influence over news organisations, controlling gateways to audiences. Convenience and user-friendly experiences have solidified platform companies' central position in the news ecosystem.

- AI adoption by news organisations may further increase the influence of large technology companies. Large platform companies are akin to 'landlords' who shape and control large parts of the AI ecosystem, while smaller start-ups resemble 'tenants',

who are both reliant on these larger firms while also pursuing their own interests.

- Concerns exist regarding the possibility for vendor lock-in and the impact of AI as used by platforms on news organisations' business models. Some suggest that addressing these challenges requires strengthening competition and data protection laws, while news organisations should invest in Research and Development (R&D), produce resonating formats, and determine on a case-by-case basis which areas are critical and where concerns are unwarranted.

- AI technologies, like language models, are expected to improve over time. However, some participants argued that the risk of an extinction event for humanity caused by AI is overstated. Control over the technology will be a key battle of years to come, although technology companies' dominance may be counteracted by open-source approaches and tighter regulation.

- The future of the information environment is uncertain, as the integration of AI is still in its early stages and depends on decisions made by technology companies, news organisations, regulatory bodies, and public acceptance and usage of AI systems.

- News organisations have an opportunity to reinforce their position as trusted brands in the face of widespread AI use, by emphasising the human element in journalism and responsibly using AI to explore new audience-focused formats.

- Concerns exist regarding how the adoption of AI, particularly by platforms, may impact the visibility and viability of news organisations, as platform companies increasingly pull back from news content and explore alternative search experiences like 'zero-click results.'

Recommendations

This report presents a set of recommendations in the 'Ways Forward' section below which is summarised here:

- Scrutinise old and new technology companies to understand their motives and actions, as their agendas may not align with those of news publishers or those of the public.

- Increase investment in R&D initiatives and comprehensive training programs to adapt to an evolving media and technology landscape and maintain a competitive edge.

- Strengthen industry collaboration by identifying common areas of concern and involving smaller and international publishers to address winner-takes-most dynamics within the industry.

- Foster collaboration between news organisations and researchers to leverage the potential benefits of AI for journalism and overcome barriers to collaboration.

- Emphasise responsible use of AI by implementing ethical guidelines, human oversight, transparency policies, and internal auditing processes, while considering the need for regulatory efforts.

- Researchers, regulators, and other stakeholders in the AI space should actively consider local and regional perspectives to help create a balanced and inclusive media ecosystem, preserving diverse perspectives and local voices. This extends to the inclusion of non-Western newsrooms in discussions, collaborations, and research efforts to address their unique challenges and ensure a global perspective.

- Actively involve audiences and the public in the development and use of AI, as their exclusion may lead to fragmented news consumption experiences and diminished trust and engagement.

INTRODUCTION

Journalism faces a decline of traditional business models. News leaders are increasingly pressured to reorient toward data-driven logics. Many news organisations now bet big on AI investments, hoping that the technology can generate additional revenue or free up staff time. But problems emerge: Some journalists fear being replaced with AI; there are possible frictions between journalistic values and the values encoded into AI systems and infrastructures; and little is known about the impact of AI on the news and the health of our public discourse. AI also poses the risk of making news organisations even more reliant on the technology and platforms companies that dominate in AI development – potentially aggravating the economic problems that news organisations face.

Starting in March 2021, the project 'AI in the News: Reshaping our Information Ecosystem'[1] at the Oxford Internet Institute investigated these and related questions, generously funded by Oxford University's Research Centre in the Humanities (TORCH) and the Minderoo-Oxford Challenge Fund in AI Governance and with administrative support from the Oxford Internet Institute and Balliol College. The aim of the project, led by Felix M. Simon and co-launched by Prof Gina Neff, was to identify key issues in this space, collect evidence, and start a conversation among academic and industry leaders about those issues. These efforts culminated in a public symposium held at Balliol College on 25th May 2023, which sought to foster discussions between industry experts, academics, and students on the key

---

[1] https://www.oii.ox.ac.uk/research/projects/ai-in-the-news-reshaping-our-information-ecosystem/

issues identified during the active research phase.[2] The symposium featured three panel discussions on 'AI in the Workplace', 'AI, News, and the Role of the Technology Sector', and 'AI and the Implications for Public Discourse'.

Held under the Chatham House Rule, participants included leading experts on AI and the news, including Shreya Vaidyanathan, Product Manager at Bloomberg LP; Jane Barrett, Global Editor for Media News Strategy at Reuters; Siddharth Venkataramakrishnan, Banking and Fintech Correspondent at the Financial Times; Melissa Heikkilä, Senior Reporter for AI at MIT Technology Review; Nic Newman, Senior Research Associate at the Reuters Institute for the Study of Journalism; David Caswell, Executive Product Manager at the BBC; and Tom Standage, Deputy Editor at The Economist. The diverse audience comprised journalists, policy experts, technology specialists, and students from the University of Oxford and beyond, contributing to a rich and multifaceted exchange of ideas.

This report provides a summary of the main themes that emerged during the symposium and outlines a few recommendations as well as blind spots to be addressed in future research. It examines the role AI plays in news production and distribution and, consequently, in the broader public sphere drawing from the experiences of various well-known British and US-American news organisations.

We begin by summarising what media organisations understand AI to be and how AI tools are increasingly being integrated into all stages of the process of news creation and distribution. Notably, different news organisations have varying degrees of engagement with AI. Some have been using AI for years and have developed their own tools, while others are still exploring these technologies.

Next, we summarise what was discussed in the three panels of the symposium. First, we provide an overview of the tensions and issues emerging from AI's use in new organisations and how the organisations have navigated the same. The discussion showed that maintaining reliability and trustworthiness was a guiding principle for many. For the second panel, participants discussed the possible dependency on the providers of AI tools and systems, especially major technology companies. They recognised the influence that these companies already have and some of them, though not all, expressed their concern that AI could increase their power even further. In the final section of the symposium, the speakers talked about the long-term implications of AI for the industry and the broader public sphere.

We conclude this report by presenting some recommendations around the use of AI in the news and highlighting some of the blind spots of the current discussions on this topic.

[2] Simon, F. M. (2021, May 26). AI in News Organisations: Exploring Applications, Challenges, and Future Implications. Medium. https://felixsimon.medium.com/ai-in-news-organizations-exploring-applications-challenges-and-future-implications-fc667c695ab1

OVERVIEW

What is Artificial Intelligence?

'Lots of things that were considered to be AI, like getting a computer to play chess in the fifties, aren't really considered AI anymore.'

The definition of artificial intelligence remains a topic of debate, with no universally accepted definition, with the term itself attracting criticism for its value-ladenness.[3] However, it is generally agreed that AI currently does not encompass conscious, general intelligence. The concept of artificial general intelligence, as depicted in popular culture, is seen as a fiction by leading AI experts.[4] In practice, AI typically takes the form of 'narrow' computer systems that focus on specific tasks and problems usually associated with human abilities.[5] These systems employ a variety of techniques, varying in complexity, autonomy, and abstraction. While they may excel within their predefined boundaries, they lack the ability to operate beyond those limits. Applications of machine learning, particularly deep learning, and various forms of natural language processing (NLP) are common examples of such narrow AI. These systems learn from data and experiences through algorithms trained on large datasets,

gradually improving their performance on specific tasks over time. The training process can range from supervised to unsupervised learning and may include other steps such as reinforcement learning through human feedback (RLHF).

Within the news industry, AI serves as an umbrella term to encompass a range of technologies. It is commonly understood by many as the computational simulation of human activities and skills in specific domains. Some of the participants of the symposium explained that it is difficult to come up with a definition of AI in a news context, especially as the line between AI and other forms of technology is often blurry and becomes murkier as time goes on. Once the novelty is lost, what once was considered AI is perceived as standard computing. AI-generated transcriptions, once seen as futuristic but now common across the industry, are one such example. In many ways, it's Tesler's Theorem: 'AI is whatever hasn't been done yet.'[6]

However, true to their journalistic lens, some participants unequivocally emphasised another side of what AI means to them: a 'massive story' that needs to be covered, as one of them put it. The technological advancements, the regulatory landscape, and the political dimensions surrounding AI have garnered widespread attention among newsmakers – both in their reporting and for their own organisations. The launch of ChatGPT and the rise of Generative AI has only intensified this interest.

---

[3] Cave, S. (2020). The Problem with Intelligence: Its Value-Laden History and the Future of AI. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 29–35. https://doi.org/10.1145/3375627.3375813

[4] Mitchell, M. (2019). *Artificial Intelligence: A Guide for Thinking Humans.* Pelican.

[5] Broussard, M. (2018). *Artificial Unintelligence. How Computers Misunderstand the World.* (1st ed.). MIT Press.

[6] Although Tesler himself says he was misquoted, arguing that '[w]hat I actually said was: "Intelligence is whatever machines haven't done yet". Many people define humanity partly by our allegedly unique intelligence. Whatever a machine—or an animal—can do must (those people say) be something other than intelligence.' See: http://www.nomodes.com/Larry_Tesler_Consulting/Adages_and_Coinages.html

Artificial Intelligence and the News

'I get dramatic value out of them.'

In our current information environment, various forms of AI, including machine learning, have been playing a significant, albeit often concealed, role for several years. As individuals increasingly rely on search engines and social media platforms to access news and other information, these platforms have significantly expanded their use of AI technologies, filtering, curating, ranking, and moderating the content that users encounter.[7] This entails the application of sophisticated algorithmic systems that analyse user preferences, behaviour, and various other factors, shaping the information landscape overall. Consequently, AI systems already occupy a central position within our modern communication infrastructure, exerting substantial influence over the flow of information and news online.[8] As a result, the content individuals are exposed to is critically shaped by the use of AI on these key communication channels.

In recent years, however, AI has moved into news organisations – key gatekeepers to the public sphere and central players in our information ecosystem.[9] In many cases, the technology has slowly moved into news production and distribution, often without readers (or journalists) really noticing.[10] The symposium highlighted some of the ways in which news organisation use and engage with the technology as well as the latest developments in the space of Generative AI – and as the significant differences between them. While some organisations have been using some form of AI for years (almost 20 in one case) and have hired teams to develop their own tools and systems, others remain quite traditional and are only carefully exploring the technology.

In their quest to leverage the potential of artificial intelligence, one organisation has adopted a three-tier framework to guide their decision-making on how and when to use these technologies: 'What can AI replace? What can AI augment? What can we invent?' are the key questions that need to be addressed in this context. Expectations of efficiency gains, productivity improvements, and product innovation are also present at some of the other organisations that have incorporated AI into their processes of news production and distribution.

News organisations have implemented AI technologies for tasks such as information discovery, verification, and content categorisation, to enable large-scale analysis of social media and news coverage, to monitor public interest in specific topics, or to facilitate various kinds of (investigative) reporting. Various organisations, including some of those present at the symposium, have started to automate e.g., financial, sports, and election reporting, as well as weather segments with the help of AI systems, often in areas where structured data exists at scale. For example, some participants use the services of

[7] See: Newman, N., Fletcher, R., Eddy, K., Robertson, C. T., & Nielsen, R. K. (2023). *Reuters Institute Digital News Report 2023* (Digital News Report). Reuters Institute for the Study of Journalism. https://reutersinstitute.politics.ox.ac.uk/sites/default/files/2023-06/Digital_News_Report_2023.pdf & Nielsen, R. K., & Ganter, S. A. (2022). *The Power of Platforms: Shaping Media and Society.* Oxford University Press.

[8] Diakopoulos, N. (2019). *Automating the News: How Algorithms Are Rewriting the Media.* Harvard University Press.

[9] Jungherr, A., & Schroeder, R. (2021). *Digital Transformations of the Public Arena.* Cambridge University Press. https://www.cambridge.org/core/elements/digital-transformations-of-the-public-arena/6E4169B5E1C87B0687190F688AB3866E

[10] Beckett, C. (2019). *New powers, new responsibilities. A global survey of journalism and artificial intelligence.* Polis. London School of Economics. https://blogs.lse.ac.uk/polis/2019/11/18/new-powers-new-responsibilities/

companies such PLX AI and AlphaSense for their financial reporting. More recently, news organisations have begun experimenting with LLMs, such as OpenAI's ChatGPT for various tasks. LLMs can generate realistic-looking text and images, inspiring news organisations to explore their usefulness in the context of e.g., automated content production, various text editing tasks (punctuation, missing words, simplification of language, adaptation to house style, etc.) or in assisting news workers around tasks like summarisation, translation, transcription, and coding. News organisations are also experimenting with LLMs to generate new ideas that could follow a news story, optimise headlines, or re-format news content.

AI systems also play an increasingly central role in tailoring content to users' interests, enhancing the distribution process, and improving audience engagement. News organisations, including some of those present, use AI-powered systems for content curation and recommendations, for example improving their article recommendations, automating newsletter content, and operating dynamic paywalls.[11] Additionally, AI-driven audience metric systems help media professionals analyse content performance and, in some cases, guide editorial decisions. Just how vast the existing and potential use cases are can be seen in a recent whitepaper by Anna Hansen and colleagues produced for the European Union's AI4Media project, which aggregated findings across countries.[12] We summarised tasks that speakers said can be partially or fully automated with AI in the table below.

Production and distribution process	Tasks for AI systems
Access and observation	• Information discovery • Audience and trends analytics • Prompting for new ideas following from a news story
Selection and filtering	• Verification, claim matching and similarity analysis for fact-checking • Content/document categorisation • Automated collection of structured data (e.g., financial, sports) • Coding for various tasks • Transcription and translation of audio and video • Search in archives/metadata
Processing and editing	• Brainstorming structure • Content production (writing, editing) • Formatting (for online, social media, print, broadcast) • Summarisation, translation
Publishing and distribution	• Personalisation, content curation and recommendations • Audience analytics • Content moderation

[11] Newman, N. (2023). *Journalism, Media, and Technology Trends and Predictions 2023* (Reuters Institute Report). Reuters Institute for the Study of Journalism. https://reutersinstitute.politics.ox.ac.uk/journalism-media-and-technology-trends-and-predictions-2023

[12] Hansen, A. S., Helberger, N., Blanke, T., & Bočytė, R. (2023). *Initial white paper on the social, economic, and political impact of media AI technologies* (p. 124). AI-4Media - A European Excellence Centre for Media, Society and Democracy. https://www.ai4media.eu/reports/initial-white-paper-on-the-social-economic-and-political-impact-of-media-ai-technologies-2/

Many media organisations have worked on developing their own AI tools, most notably Bloomberg which recently announced the development of its own large language model BloombergGPT.[13] Another one of the organisations represented at the symposium has also trained language models for tasks such as sentiment analysis to understand text, for instance, determining whether someone is expressing support or opposition regarding a topic such as 'job cuts'. Such systems also can help surface useful information more quickly for reporters.

However, as various speakers explained, the development of AI systems and applications can be a demanding endeavour, requiring significant investments of time and resources. One speaker highlighted that the process of developing and implementing AI in-house can go through several cycles of inefficiency, followed only later by eventual efficiency and modest productivity gains. In addition, the tools developed may not always meet initial expectations, or fail to work as intended altogether. For instance, one organisation conducted experiments on summarising content that produced incorrect information which was not included in the original material. They also attempted to automatically tag their archive of news stories with the help of AI, an endeavour that ultimately proved to be unsuccessful.
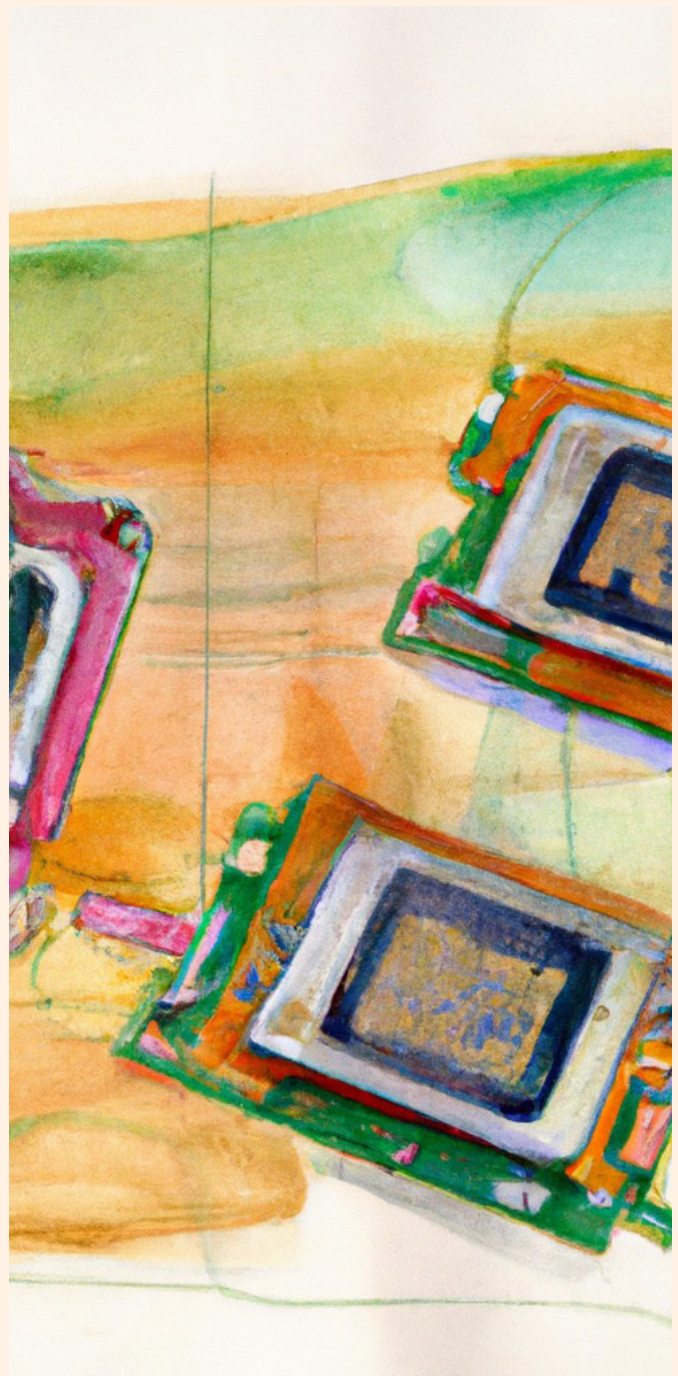
[13] Bloomberg Professional Services. (2023, March 30). Introducing BloombergGPT, Bloomberg's 50-billion parameter large language model, purpose-built from scratch for finance. Bloomberg. Retrieved from https://www.bloomberg.com/company/press/bloomberggpt-50-billion-parameter-llm-tuned-finance/

SYMPOSIUM SUMMARY

AI and Journalistic Work

'What can we give up doing?'

A recurring theme echoed by the panellists was the potential of AI to enhance efficiency and productivity in newsroom operations. 'We're always trying to do more with less', said one of the speakers during the symposium, who also emphasised that AI's capabilities have inspired their organisation to look for ways to save time and costs while also producing more and better content that can reach both new and existing audiences. However, while acknowledging the potential of AI, the speakers also recognised its limitations: 'It's like having infinite interns', said one of them, rephrasing a famous quote by Benedict Evans.[14] Put differently: in practice, the work it produces can sometimes be surprisingly good but sometimes also require a lot of additional editing. One connecting theme was that claims regarding efficiency and productivity gains through AI should be approached with caution, as the extent of these gains varies depending on the specific task and context. While efficiency gains may be more straightforward in e.g., coding or transcription and translation, they are not as easily achieved in more complex tasks such

as writing or reporting (which in fact are a collection of different tasks). In fact, the use of AI can sometimes lead to additional work for journalists and news workers.

At times, speakers' views on the impact of AI were strongly at odds with each other. While one of the speakers remarked that AI 'is a superpower', another warned that 'it is not a magic bullet; it is not magic'. However, these views in many ways reflect what the adoption of AI has felt like for these newsrooms and publishers: unleashing these tools often comes with both benefits and significant challenges. For many new organisations the integration of AI into workflows has brought forth a multitude of tensions and questions. During the symposium, participants highlighted many of these challenges, some of which are inherent to the use of AI itself, while others are indirectly related to it – like concerns over violations of privacy[15] or copyright[16] laws.

One significant issue lies in the structural limitations and biases of many AI systems. One of the participants pointed out that many of them are biased towards the Anglosphere and the United States context in particular. This is due in part to the abundance of freely available content in English which is used to train these models. Moreover, two participants warned that AI systems, and especially LLMs can exhibit unexpected behaviours like misplacing quotes, hallucinating or confabulating, making them more difficult to work with. In addition to technical challenges, the integration of AI tools has also raised concerns on the possible

---

[14] The actual quote by Benedict Evans is 'AI gives you infinite interns'. Evans updated it in December 2022 saying 'Old: AI gives you infinite interns. New: AI gives you infinite bullshitters.' See: https://twitter.com/benedictevans/status/1598425986973855745.

[14] As mentioned during the symposium, there have been cases of AI tools leaking personal data. See, for example: Burgess, M. (2023, April 4). ChatGPT Has a Big Privacy Problem. *Wired UK.* Retrieved from https://www.wired.co.uk/article/italy-ban-chatgpt-privacy-gdpr

[15] Appel, G., Neelbauer, J., & Schweidel, D. A. (2023, April 7). Generative AI Has an Intellectual Property Problem. *Harvard Business Review.* Retrieved from https://hbr.org/2023/04/generative-ai-has-an-intellectual-property-problem

disruption of the dynamics within newsrooms: the social environment, an important element of news production, could be negatively affected if people are more invested in working with AIs than in speaking to each other or working collaboratively, as one speaker worried. However, one participant offered a different perspective. According to them, the biggest risk for news organisations lies not in the problems described above, but in not fully making use of everything that AI has to offer: 'The single biggest risk is complacency.'

Amidst the concerns and questions surrounding the integration of AI in news organisations, where do publishers draw the line? Many participants emphasised that maintaining reliability and trustworthiness is a guiding principle in their use of AI – something that has made them cautious when it comes to the technology: 'A newsroom is naturally a very conservative place because we've got to get things right. We've got to be very, very careful.' In practice, this has meant that various organisations have set self-imposed limits: for instance, some of them do not use a lot of what is currently considered AI, while others have decided to employ other forms of AI but not to use LLMs, especially to mass-produce content. Others have in the past paused the use of AI, in some cases for longer periods of time, to make sure they can be repeatedly and rigorously tested until they are deemed to be good enough at the tasks they are meant to do.

Nonetheless, even within these constraints, some speakers remain optimistic about the future uses of AI in the news. One speaker

pointed out that AI could be a good thing for their organisation and the news at large as long as they keep learning from what others are doing and apply their existing ethics and standards. Another participant was optimistic about AI allowing their organisation to create more value by using the technology to adjust their content (usually directed at people who are more educated, male, and older than average) to newer audiences, which is something that would be impossible to do manually at scale. However, they also recognised that they 'are not there yet'.

This optimism also extended to the potential benefits of AI for local news organisations and publishers in the Global South who are often disadvantaged due to e.g., a lack of resources in case of the former or difficult political and economic circumstances in case of the latter. Various speakers spoke about AI technologies as one avenue to narrow the gap, particularly through the use of off-the-shelf tools. With respect to regional and local outlets, one participant argued that 'these tools [would] disproportionally empower small, tiny newsrooms', a view shared by some of the other speakers who hoped for AI to allow smaller organisations to produce more and more valuable content.

Will AI Take Away the Jobs of Journalists?

> *'ChatGPT gives you the median opinion of the Internet, which is usually very boring.'*

There was also, of course, the question that is on many peoples' minds these days: Will AI

take my job away? The recent announcement of the layoff of 200 workers (about 20% of the team) of BILD, Germany's biggest tabloid newspaper, was in parts justified with the increasing functionalities of AI. For many, it confirmed fears that the use of AI would lead to job losses for many people currently working in the news.[17] However, the symposium participants had a more nuanced view on the labour impact of AI.

While the speakers recognised that AI systems can be very good and even outperform humans at certain tasks (e.g., fixing commas and general copy-editing, making transcriptions, analysing financial statements at scale and even helping with coding), they all highlighted the need for people in the newsroom – not least because the day-to-day work of journalists is more complex and cannot simply be substituted with automated processes and AI. Jobs in journalism consist of a combination of tasks, only some of which can be automated. Many speakers, for example, were convinced that AI cannot match a journalist's writing skill (two of them called AI's writing 'boring' and bland). Likewise, the ability to report or identify new facts and opinions are skills that will likely remain hard to automate. Other participants highlighted that humans are unbeatable at understanding the meaning and context of information

– something that together with accuracy, however, remains the core of good journalism.

Large language models in particular have no true capacity for either. In the words of several leading AI researchers, including Timnit Gebru, they are stochastic parrots[18] (although this analogy is also somewhat problematic, as e.g., philosopher Luciano Floridi points out: 'parrots have an intelligence of their own that would be the envy of any AI but, above all, because LLMs synthesise texts in new ways, restructuring the contents on which they have been trained, not providing simple repetitions or juxtapositions.')[19]. Yet, the overall point stands: LLMs do not 'think' as a human would but instead mimic our ability to communicate, producing output that is based on the predictions of 'tokens' – educated guesses about what will likely come next in a sentence.[20] But for many journalistic tasks, the human ability to think and reason –drawing intuitive inferences about other inferences[21]– remains of the essence. Questions such as 'Why does this story matter?', 'Why is this important?', and 'What does this mean in a given context?' can only be asked and properly answered by a human (even though there is a role for technology to aid them in this task), according to one of the speakers, rendering the same, and therefore humans, more valuable in the future.

[17] Henley, J. (2023, June 20). German tabloid Bild cuts 200 jobs and says some roles will be replaced by AI. *The Guardian.* Retrieved from https://www.theguardian.com/world/2023/jun/20/german-tabloid-bild-to-replace-range-of-editorial-jobs-with-ai

[18] Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🦜. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 610-623). Association for Computing Machinery. Retrieved from https://doi.org/10.1145/3442188.3445922

[19] Floridi, L. (2023, February 14). AI as Agency Without Intelligence: On ChatGPT, Large Language Models, and Other Generative Models. *Philosophy and Technology.* Retrieved from SSRN: https://ssrn.com/abstract=4358789

[20] Shanahan, M. (2023). Talking About Large Language Models. arXiv preprint arXiv:2212.03551, https://arxiv.org/abs/2212.03551

[21] Sperber and Mercier describe reasoning as a 'mechanism for intuitive inferences about one kind of representations, namely, reasons'; in other words our ability to represent 'not only things and events in our environment but also our very representations of these things and events'. Sperber, D. & Mercier, H. (2018). T*he Enigma of Reason.* Penguin.

AI and Big Tech Dependency

'It's hard to see how they can be more dominant than they already are.'

During the discussion on news organisations' dependency on the technology sector for AI, participants engaged in a lively exchange of ideas. All participants recognised the influence that big tech companies already exert over the news and the news industry. As one of the speakers explained, companies such as Google and Facebook largely control the gateways to news content and audiences (according to them, over 50% of the traffic their sites get comes from search engines). Platform's decisions around what content to show and suggest to their users on popular social media and video platforms as well as search engines gives them shaping power over a significant portion of news consumers' experience. 'Users have voted with their feet already', gravitating towards the convenience and user-friendly experiences offered by these platform companies, solidifying their position in the news ecosystem. Moreover, they own tools and infrastructure used in the process of news production and distribution, are dominant in the digital advertising market upon which many news organisations still rely, fund research, industry conferences, and support schemes and lobby regulatory bodies on news-related issues. In the words of one of the participants, 'everything leads to big tech'.

It is not hard to then see why, as one of the speakers put it, that 'the technology industry is driving this [AI] full on'. A key question here is if the influence of these companies will increase with the growing adoption of AI technologies in news organisations.[22] Two of the speakers indicated that this might just be the case, with AI further cementing the power of major technology companies, given their head start around the development and deployment of AI, as well as their control over key infrastructure.[23] One likened large platform companies to 'landlords' who exercise control over crucial resources for AI, such as data, computing infrastructure, and substantial financial and human capital, thus both shaping and controlling large swathes of the AI ecosystem. In contrast, smaller start-ups like OpenAI were described as 'tenants' – somewhat independent but still heavily reliant on these larger firms for their AI development.

Two panellists showed themselves particularly concerned about the risk of vendor lock-in, where news organisations become so strongly tied to a particular AI system or the infrastructure to run the same, that switching to other competitors become very costly and difficult. However, one of them argued that while the risk of vendor lock-in is real, it is not necessarily problematic in any context. Assessing where this will matter and where it will not, is a 'management challenge' according to them, with news outlets having to ask themselves which tools and use cases are so fundamental to what they do that they cannot afford becoming locked in – and in which cases such concerns are unwarranted. One participant gave the example of AI-powered

[22] Simon, F. M. (2022). Uneasy Bedfellows: AI in the News, Platform Companies and the Issue of Journalistic Autonomy. *Digital Journalism, 10*(10), 1823–1854. https://doi.org/10.1080/21670811.2022.2063150

[23] See also: Ahmed, N., Wahed, M., & Thompson, N. C. (2023). The growing influence of industry in AI research. *Science, 379*(6635), 884–886. https://doi.org/10.1126/science.ade2420

transcription services: 'There are now hosts of companies doing transcription services, competing with each other on price, often probably using the same underlying technology. And that's a pretty good position for news organisations because it's much cheaper than it used to be, widely available, commoditised. I don't mind applying a technology that is owned by one company as long as the price is right, and the competition works.'

Nevertheless, many participants acknowledged that the development and implementation of AI across various domains could further consolidate the power of platform companies – and extend their control over the information ecosystem at large. One speaker argued that while they thought that there might be an opportunity now for new companies to break through and end the dominance of technology giants, most of them were likely in a strong enough position to survive at least for now. Perhaps unsurprisingly, there were diverging opinions on the appropriate course of action when confronting platform power and their ability to shape the news environment. Suggestions included strengthening competition and data protection laws, although it was acknowledged that regulators often struggle to keep pace.[24] Additionally, panellists stressed the importance of news organisations confronting these challenges directly and proactively – by being more creative and experimental, investing more strongly in research and development as well as product development and becoming 'more distinctive' by producing formats 'that people really want.'

As one participant put: 'We have to be more distinctive.'

The Future of AI and the Information Environment

What are the long-term implications of AI for the news industry, information ecosystems, and the public discourse?[25] During the symposium, a few speakers made predictions on the future of AI, most of which were optimistic.

AI Technologies Will Improve, but not Kill Us

There was a strong consensus that AI technologies but more specifically LLMs will get better over the next couple of months and years. As an example, one speaker mentioned the Alpaca Project at Stanford University, an open source LLM, that was developed at a fraction of the price of ChatGPT.[26] The same speaker also saw great potential in open-source developments and models: while they might not be able to match the best proprietary models anytime soon, they would likely lead to a greater democratisation of these tools. They argued, however, that the downside of this is that attempts to prevent bad actors from getting hold of AI technologies will likely also fail.

When speaking about the possible problems AI could generate, one speaker explained that experts in these fields usually made contrasting predictions to the ones made by super forecasters, and that the risk of an extinction event was overstated. 'It's going to get better,

---

[24] See also Seipp, T. J., Helberger, N., de Vreese, C., & Ausloos, J. (2023). Dealing with Opinion Power in the Platform World: Why We Really Have to Rethink Media Concentration Law. *Digital Journalism, 0*(0), 1–26. https://doi.org/10.1080/21670811.2022.2161924

[25] While this final section of the symposium was a conversation with Tom Standage, Deputy Editor of *The Economist*, this summary also includes the views of other speakers on the long-term implications of AI.

[26] Taori, R., Gulrajani, I., Zhang, T., Dubois, Y., Li, X., Guestrin, C., Liang, P., & Hashimoto, T. B. (2023, March 13). Alpaca: A Strong, Replicable Instruction-Following Model. Retrieved from https://crfm.stanford.edu/2023/03/13/alpaca.html

and we all need to get out heads around it […] Yet, the people we should not be listening to on whether AI is going to kill us are the people building it.' According to the speaker, experts, regardless of the field, consistently overestimate the likelihood of significant events, whether positive or negative, occurring within their respective domains. This tendency, they said, is explained by the fact that experts' deep immersion in a particular subject often leads them to perceive their areas of work as the most important thing, skewing their predictions.

A Return to Trusted Brands…and an End to the Current Business Model

'How do we make our journalism more human?'
'What's up for grabs is the front door of the internet.'

According to some participants, the anticipated widespread use and abuse of AI in content creation could be an opportunity for news organisations to reinforce their position as trusted brands which provide an important service to the public. Many participants reasoned that a growing amount of information – thanks to the reduced cost of producing the same at scale – would make it more difficult for people to find reliable information. In such a scenario, publishers could play a key role as trusted sources that help audiences orient themselves. The responsible use of AI, for example in exploring new audience-focused formats and an emphasis on the 'human element' in journalism were seen as important conditions to keep and gain peoples' trust in this context.

A major concern was how the adoption of AI, especially by platforms, could change the way the public gets to see and access information, including news. Relatedly, participants worried about what such developments would mean for the reach and the viability of publishers. As one speaker explained, for over 20 years, Google has been the main face of the web and a key route to content for users, as well as a main provider of traffic for news organisations. Yet, increasingly, platform companies are pulling back from news content, as some of them see news as 'toxic' and damaging to their products and brands. These worries are compounded by recent developments which suggest that platform companies such as Google and Bing are considering search experiences where users are presented with a short summary – so called 'zero-click results' – instead of individual search results. This, some speakers worried, could drastically change news organisations' visibility and business models. 'Click-driven journalism is going to fail', as one of the participants put it. Another argued that 'it's kind of panic stations', with many news organisations very concerned about a possible demise of a core part of their current business model.

Will the Information Environment Change?

The question of whether our information environment would undergo significant changes was a source of great uncertainty during the symposium. Given the early stage of AI integration in news organisations and beyond, many found it challenging to predict the nature of these changes – not least because they depend on a flurry of decisions by technology companies, news organisations,

regulatory bodies, and how the public will come to accept and use AI systems. One common view was that the rise of AI would likely lead to a further re-balancing of power within the information ecosystem, especially between news organisations, platform companies, and newer entrants, although it remained unclear who would benefit most in this regard.

Many participants believed that between news organisations, those that embrace AI early and strategically would potentially emerge as winners – strengthening their position as trusted brands and their financial viability, if they use the technology to increase efficiencies, build new products and find improved ways of attracting new and retaining existing audiences. However, some speakers also reasoned that AI would end up strengthening the grip of the technology sector over communication infrastructures and the public sphere, given their experience in developing and implementing the technology, their vast financial resources, and their role as providers of key infrastructure. Various speakers underlined that to navigate these future challenges successfully, news organisations had to invest in and develop more robust R&D capabilities. This should entail fostering a culture of experimentation and embracing the possibility of failure. Likewise, continuously cultivating the necessary expertise and skills to harness the potential of AI effectively was seen as one way forward.

WAYS FORWARD: RECOMMENDATIONS AND BLIND SPOTS

In this summary, we present key suggestions and observations from the symposium regarding the necessary actions and potential strategies for maximising the advantages of AI for news organisations and the broader public sphere. These are not meant to be comprehensive[27] but reflect key issues that emerged during the research phase leading up to the symposium as well as during the discussions on the day of the event. Additionally, we shed light on some of the blind spots that became clear during the symposium, exposing gaps in current discussions on the subject.

Recommendations

Scrutiny of old and new technology companies: Large technology companies serve as both infrastructure providers and investors in the field of artificial intelligence (AI). However, the current AI landscape is marked by significant disparities. The future trajectory of new start-ups, such as OpenAI, remains uncertain. It is unclear whether they will eventually surpass platform companies, maintain a state of interdependence or co-dependence, or be acquired by them.[28] Both start-ups and platform companies operate with their own

---

[27] For a broad overview of policy recommendations for the use of AI in the media sector see: Krack, N., Dutkiewicz, L., Yildirim, E. O., Papadopoulos, S., Tsalakanidou, F., Kompatsiaris, I., Schjøtt Hansen, A., Sassatelli, L., Cortés, A., Garcia-Gasulla, D., Aichroth, P., Steenpass, A., Cuccovillo, L., Gray, B., Baris Schlicht, B., Bocyte, R., Krack, N., Dutkiewicz, L., Yildirim, E., … Georgiev, C. (2022). *Pilot Policy Recommendations for the use of AI in the Media Sector* (p. 104). AI4Media - A European Excellence Centre for Media, Society and Democracy. https://www.ai4media.eu/wp-content/uploads/2022/10/AI4Media_D2.4_PolicyRecommendations_final.pdf

[28] Waters, R. (2023, June 29). AI boom forces tech companies to make M&A choice. *Financial Times.* Retrieved from https://www.ft.com/content/ef7ad4c8-be0c-4ad8-b829-a105e6965f5f

distinct agendas, which may not align with those of news publishers. In some instances, conflicts have already arisen, such as disputes over the use of content for training purposes. Other conflicts might only emerge over time, for example around the traffic directed towards news sites. Their interests may also not necessarily align with those of the general public and with what is best for the functioning of the public sphere. For instance, they may not prioritise the development of an information ecosystem that benefits all individuals, enables plurality and freedom of expression, while simultaneously curbing the spread of harmful speech.[29] This uneven and complex landscape underscores the importance of a careful examination of the motives and actions of the players in the technology space – both old and new – and the potential consequences of the same, to assess the risks and potential harms around the implementation of AI systems in communication infrastructures.

Investment in Research and Development: To ensure the future viability of news organisations and harness the benefits of AI, it is imperative for publishers to bolster their investments in technology and talent. Unfortunately, the news industry has fallen behind other sectors in these crucial areas.[30] By strategically allocating resources to R&D initiatives and comprehensive training programs, news organisations can enhance their ability to adapt to the evolving media landscape and maintain a competitive edge in the market. Moreover, it is essential to challenge the misconception that only large and successful companies have the means to afford such investments. In reality, a number of smaller news organisations operating in challenging markets have successfully demonstrated that it is indeed feasible to allocate resources to product development and foster professional growth.[31]

Strengthening industry collaboration: The challenges publishers face around AI show the importance of increased collaboration among news organisations.. Notably, several prominent publishers – including The New York Times, News Corp, Vox Media, Condé Nast, and Axel Springer – have already initiated discussions regarding the establishment of a coalition that looks at AI together.[32] However, to ensure the success of such collaborative efforts, publishers must identify common areas of concern, even in the face of competitive dynamics and conflicting priorities. It will also be crucial to involve smaller and international publishers in such efforts. Failure to include these stakeholders could otherwise reinforce existing winner-takes-most dynamics[33], consequently neglecting other crucial interests within the industry.

Collaboration with researchers: Collaboration between news organisations and researchers

[29] Rau, J., & Simon, F. M. (2022). *Digital Turbulence: Building a democratic society in times of digital turmoil.* IPPI | Heinrich-Böll-Stiftung. https://il.boell.org/en/2022/02/16/digital-turbulence-building-democratic-society-times-digital-turmoil

[30] Nielsen, R. K. (2021, December). Invest in Tools and Talent, and Newsrooms Can Finish the Job. Nieman Lab. Retrieved from https://www.niemanlab.org/2021/12/invest-in-tools-and-talent-and-newsrooms-can-finish-the-job/

[31] Ibid.

[32] While specific details and commitments were not yet finalised at the time of writing, the effort seems to revolve around concerns how publishers' content is used for training purposes (without compensating publishers) and how AI could potentially eliminate the need for users to access original sources. See: Bruell, A. (2023, June 28). Big news publishers look to team up to address impact of AI. *The Wall Street Journal.* Retrieved from https://www.wsj.com/articles/big-news-publishers-look-to-team-up-to-address-impact-of-ai

[33] Newman, N., Fletcher, R., Eddy, K., Robertson, C. T., & Nielsen, R. K. (2023). *Reuters Institute Digital News Report 2023* (Digital News Report). Reuters Institute for the Study of Journalism. https://reutersinstitute.politics.ox.ac.uk/sites/default/files/2023-06/Digital_News_Report_2023.pdf

at universities and research institutes in the development of AI for journalism holds benefits for both sides. Currently, such collaborations predominantly occur within public broadcasters, which often perceive it as integral to their mission. However, even within these organisations, barriers to collaboration exist. In contrast, commercial publishers often display considerable reluctance in granting researchers access, engaging in collaborative research endeavours, or sharing data. This situation presents a certain irony considering the frequent calls from the news industry for platform companies to share their data, while simultaneously exhibiting a similar reluctance to do so themselves.

Responsible use of AI: Various news organisations are already considering the ethical implementation of AI, as demonstrated by the rise in published guidelines on the topic.[34] For instance, the BBC's Machine Learning Engine Principles[35] have served as an influential resource for both public service and commercial news organisations worldwide. These principles emphasise the importance of human oversight and regular review of AI applications. Similarly, Porlezza et al. have proposed recommendations specifically for public service media but with broader applicability.[36] Their guidelines advocate for upholding plurality, respecting human rights, and applying responsible AI practices throughout the organisation. Transparency

policies, such as explaining the impact of AI on user experience, and clear accountability for AI-generated output are also emphasised. The authors further recommend the implementation of internal auditing processes. Yet, the effectiveness of the industry's current self-regulatory regimes and the potential need for more regulation from political bodies remain important considerations. While it is essential to approach regulatory efforts with caution, especially given their potential impact on news organisations' independence, there may be value in encouraging self-regulation. Additionally, addressing privacy concerns and other non-journalistic challenges associated with AI should also be part of the discussion.

Current Blind Spots

Local and regional perspectives: Local and regional publishers play a crucial role in the news and journalism ecosystem. They provide valuable context, insights, and community-focused content and it is imperative not to overlook them and their needs when it comes to AI. Neglecting these publishers could result in a loss of diverse perspectives, local expertise, and a deeper understanding of community-specific issues – all of which are essential for the health of the public sphere and democracy. AI development in the news industry should actively consider and involve local and regional publishers to maintain a balanced and inclusive media ecosystem.

---

[34] See Cools, H., & Diakopolous, N. (2023, July 10). Towards Guidelines for Guidelines on the Use of Generative AI in Newsrooms. Generative AI in the Newsroom. Retrieved from https://generative-ai-newsroom.com/towards-guidelines-for-guidelines-on-the-use-of-generative-ai-in-newsrooms-55b0c2c1d960 and Becker, K. B. (2023). Neues Spiel, neue Regeln. *Journalistik. Zeitschrift für Journalismusforschung, 6*(2), 142-164. https://doi.org/10.1453/2569-152X-22023-11505-de

[35] Macgregor, M. (2021, May). Responsible AI at the BBC: Our Machine Learning Engine Principles. BBC Research & Development. Retrieved from https://www.bbc.co.uk/rd/publications/responsible-ai-at-the-bbc-our-machine-learning-engine-principles

[36] Porlezza, C., Pranteddu, L., & Mazzoni, P. (2022). The Governance of Artificial Intelligence in Public Service Media: A Comparative Analysis. Report for the Federal Office of Communications OFCOM. Lugano. Retrieved from https://tinyurl.com/4yd2a9sb

Non-Western newsrooms: Discussions around AI in the news often focus on publishers in the Global North, at the expense of news organisations in other parts of the world, including in the Global South.[37] This matters because these publishers often operate under very different, often difficult socio-political circumstances,[38] do not always have the same legacy channels and revenues to sustain their operations and reach audiences, and face a slew of challenges that are not present for publishers in the Global North in the same ways.[39] Future discussions, collaborations and research efforts in and around AI for news should encompass these diverse cultural, political, and geographic viewpoints to better understand what AI can and cannot do for the news.

Audiences and the public: The importance of considering the role of audiences in the development and utilisation of AI within news and communication infrastructures cannot be overstated. Currently, audiences and the public are frequently overlooked, merely mentioned as important without active inclusion. Yet, failure to incorporate audiences and the public in AI development and regulation could result in fragmented news consumption experiences, diminished trust and engagement, and the formation of an information ecosystem that prioritises the interests of a select few rather than the broader public.

---

[37] Gondwe, G. (2023). CHATGPT and the Global South: how are journalists in sub-Saharan Africa engaging with generative AI?. *Online Media and Global Communication*. https://doi.org/10.1515/omgc-2023-0023

[38] Shabbir, N., Posetti, J., and, & Simon, F. M. (2022). How Three Mission-Driven News Organizations in the Global South Combat Disinformation Through Investigation, Innovation, Advocacy, and Education. In *Disinformation in the Global South* (pp. 193–209). John Wiley & Sons, Ltd. https://doi.org/10.1002/9781119714491.ch13

[39] Nielsen, R. K., & Cherubini, F. (2022, October 27). Born in the fire: What we can learn from how digital publishers in the Global South approach platforms. Retrieved from https://reutersinstitute.politics.ox.ac.uk/born-fire-what-we-can-learn-how-digital-publishers-global-south-approach-platforms

CONCLUSION

AI has arrived in journalism, and it is here to stay. In a recent report which polled news executives, 28 percent of the respondents answered that their organisation was using AI regularly, while 39 percent admitted that they were experimenting with it.[40] Recent examples of publishers scrambling for an AI strategy[41] or making announcements about integrating the technology more fully into their products emphasise the current industry dynamic: Publishers have voted with their feet and the irresistibility of the technology, and the current hype will do the rest. The wider adoption of the technology is all but certain.

Yet, what is easily forgotten amid the current AI frenzy is that news organisations will still pursue the same goals and have the same needs as they did before. What will change through AI is mostly the way these are pursued. Many of the tasks stays the same, but the arsenal of tools changes. This re-tooling of journalism will be nuanced and gradual, not least because there will be resistance and hesitancy from within (and for good reason).

The impact of the technology will vary, too. The much-touted efficiency gains (which are difficult to quantify, depending on the context) will be easier to achieve where 100% accuracy is not the most important goal – e.g., in targeting readers with more content. It will be much harder in areas where this is required, for example in news writing. While AI tools exhibit great capabilities, the multifaceted nature of the journalistic work cannot be easily substituted for automated approaches. Human thought, empathy, and understanding will remain irreplaceable.

Not every news organisation will benefit from AI in the same way. Building customised AI systems is easier for large, well-resourced publishers who have the money and expertise to experiment with the technology and make it work for them. Many local newsrooms or smaller, less-well-resourced publishers including some in the Global South will not be in such an enviable position. How AI will or will not affect them and how they will make AI work for them belongs to the open questions in need of further exploration.

One of the biggest open-ended questions revolves around the future of news organisations' business models on the one hand, and the first- and second-order effects of AI's increasing use across news organisations and central communication infrastructures that together make up our information ecosystem. As with any new technology, these effects will likely be neither as positive nor negative as some of the current discourse suggests. What they will be exactly, however, will only become clear in hindsight.

[40] Newman, N. (2023). Journalism, media, and technology trends and predictions 2023. [Report]. Reuters Institute for the Study of Journalism. Retrieved from https://reutersinstitute.politics.ox.ac.uk/journalism-media-and-technology-trends-and-predictions-2023

[41] Smith, J. (2023, July 2). The future of work is here. *Financial Times*. Retrieved from https://www.ft.com/content/4fae2380-d7a7-410c-9eed-91fd1411f977

FURTHER READING

Artificial Intelligence

Broussard, M. (2018). *Artificial Unintelligence. How Computers Misunderstand the World*. (1st ed.). MIT Press.

Jarrahi, M. H., Lutz, C., & Newlands, G. (2022). Artificial intelligence, human intelligence and hybrid intelligence based on mutual augmentation. *Big Data & Society*, 9(2). https://doi.org/10.1177/20539517221142824

Mitchell, M. (2019). *Artificial Intelligence: A Guide for Thinking Humans. Pelican*.

Wolfram, S. (2023, February 14). *What Is ChatGPT Doing … and Why Does It Work?* Stephen Wolfram. Retrieved from https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work/

AI and the News

Beckett, C. (2019). *New powers, new responsibilities. A global survey of journalism and artificial intelligence.* Polis. London School of Economics. https://blogs.lse.ac.uk/polis/2019/11/18/new-powers-new-responsibilities/

Diakopoulos, N. (2019). *Automating the News: How Algorithms Are Rewriting the Media.* Harvard University Press.

Gondwe, G. (2023). CHATGPT and the Global South: how are journalists in sub-Saharan Africa engaging with generative AI?. *Online Media and Global Communication.* https://doi.org/10.1515/omgc-2023-0023

Hansen, A. S., Helberger, N., Blanke, T., & Bočytė, R. (2023). *Initial white paper on the social, economic, and political impact of media AI technologies* (p. 124). AI4Media - A European Excellence Centre for Media, Society and Democracy. https://www.ai4media.eu/reports/initial-white-paper-on-the-social-economic-and-political-impact-of-media-ai-technologies-2/

Nishal, S., & Diakopoulos, N. (2023). Envisioning the Applications and Implications of Generative AI for News Media. In CHI Workshop on Generative AI & HCI.

AI, the Public Sphere, and Technology Companies

Jungherr, A., & Schroeder, R. (2023). Artificial Intelligence and the Public Arena. *Communication Theory.* Advance online publication. https://doi.org/10.1093/ct/qtad006

Nielsen, R. K., & Ganter, S. A. (2022). *The Power of Platforms: Shaping Media and Society.* Oxford University Press.

Seipp, T. J., Helberger, N., de Vreese, C., & Ausloos, J. (2023). Dealing with Opinion Power in the Platform World: Why We Really Have to Rethink Media Concentration Law. *Digital Journalism*, 0(0), 1–26. https://doi.org/10.1080/21670811.2022.2161924

Simon, F. M. (2022). Uneasy Bedfellows: AI in the News, Platform Companies and the Issue of Journalistic Autonomy. *Digital Journalism*, 10(10), 1832–1854. https://doi.org/10.1080/21670811.2022.2063150

BIOGRAPHIES

Felix M. Simon is a communication researcher and doctoral student at the Oxford Internet Institute (OII), a Knight News Innovation Fellow at Columbia University's Tow Center for Digital Journalism, and an affiliate at the Center for Information, Technology, and Public Life (CITAP) at the University of North Carolina at Chapel Hill. He also works as a research assistant at the Reuters Institute for the Study of Journalism (RISJ) and regularly writes and comments on technology, media, and politics for various international outlets. His past and current research focus on AI in the news, political communication in the digital age as well as the changing nature of journalism and the media in the 21st century. He also takes an active interest in the future of mis- and disinformation. Felix holds degrees from Goethe-University Frankfurt and the University of Oxford. He is currently a fellow at the Salzburg Global Seminar and an Associate Fellow of the UK Higher Education Academy. Before returning to the OII for his doctoral studies, Felix worked as a journalist, editor and researcher in London. Past work experience also includes the BBC and Olympic Broadcasting Services (OBS) in London and Innsbruck.

Luisa Fernanda Isaza-Ibarra is a Colombian free speech attorney, researcher, and activist and a Master of Public Policy Student at the University of Oxford. She studied law at the Pontifical Xaverian University in Bogota, where she also taught the class Law and Journalism. Luisa has spent the last seven years defending and promoting freedom of expression in Colombia from multiple angles: from journalist protection and the fight against impunity in the crimes against the press, to online free speech and the right to access public information. Luisa has worked at Colombian NGOs and has been a consultant for international organisations focused on free speech. She is interested in creating and advancing public policies that support online and offline free speech in Latin America and is looking to raise awareness among citizens on the importance of journalism for democracies.

ACKNOWLEDGEMENTS