EUR(O)VISION NEWS
OPERATED BY EBU

EBU NEWS REPORT 2024

TRUSTED JOURNALISM

in the Age of Generative AI

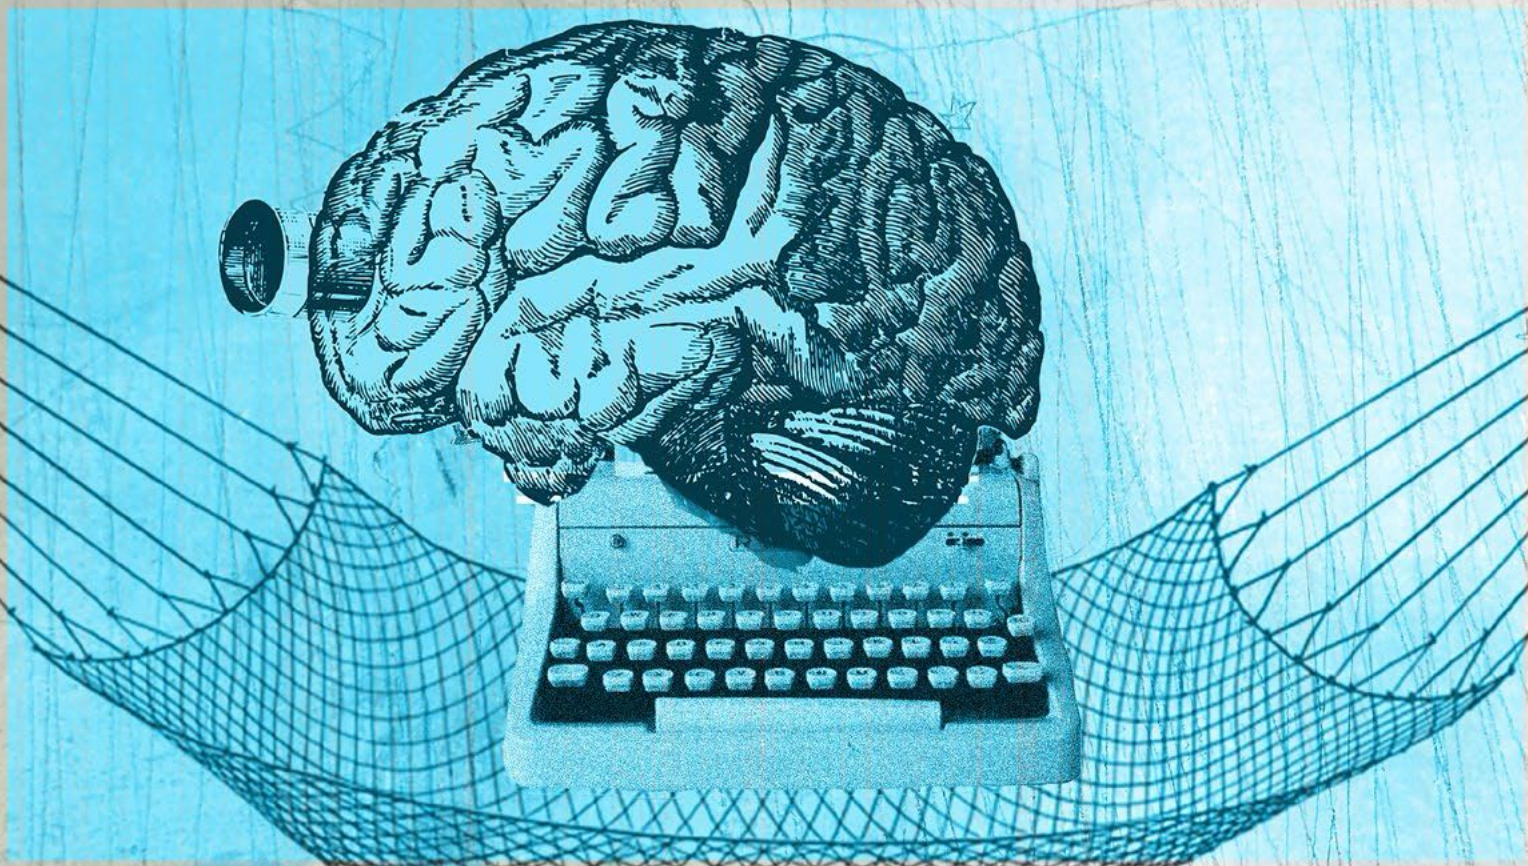

#SALMANQADIR

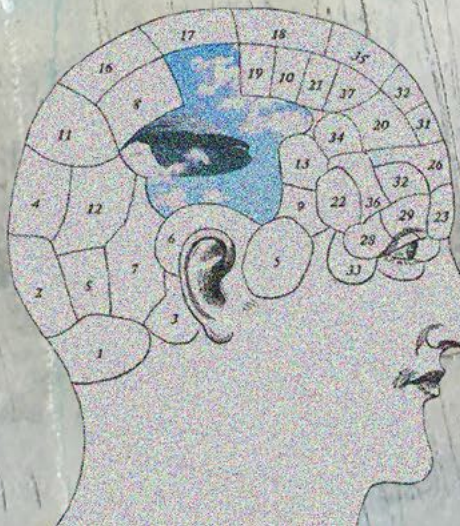

FOREWORD

"It was the best of times, it was the worst of times," as the opening lines of Charles Dickens' novel *A Tale of Two Cities* goes. In the world of AI and news, we can say the same. We are drowning in AI-generated fake news. Yet does this not also increase the value of authentic, human-checked news? And will it open the door to a new golden age for news?

In an increasingly personalized and on-demand information universe, where citizens no longer rely on the same sources of news, where synthetic media and the impact of influencers are on the rise, our mission as public service journalists is clearly more crucial than ever.

Our responsibility is to deliver excellent, trustworthy news. We must be vigilant, 24 hours a day, to deserve this moniker of 'excellence'. It has to be earned; it is never a given.

The arrival of AI has transformed the context of how we produce and deliver news. With its promise of increased efficiency comes the risk of amplifying manipulative news to unprecedented proportions, destabilizing democratic debate and freedom of expression. Among the fast-growing lists of challenges we face today are synthetic media, automatic disinformation, and copyright.

Our mission is to continue to be the guardians of free, independent, fair and verified information in an era where generative AI is redefining the boundaries of reality.

In the information sector, destabilized by more than fifteen years of the negative impacts of out-of-control social networks in the public sphere, there is a mixture of excitement and concern. Excitement because the time saved will enable editorial teams to concentrate on the essential; concern because this new tool can also be used to twist reality, manipulate facts, even invent them.

Yes, generative AI is transforming the way people access information and create content. But this new AI, while highly creative, doesn't investigate, doesn't question sources, rarely gives its own sources and doesn't hold authorities to account.

Journalism needs a solid, widely recognized ethical foundation. We need to further protect the integrity of information in the age of generative AI. Ethics must guide our technological choices, human judgement must remain at the heart of our editorial decisions. We must also help society to distinguish authentic content from synthetic creations and participate actively in the global governance of AI.

It's up to us to cooperate even more closely, to coordinate and to have an impact. Here, size matters. As will the quantity of our high-quality structured data.

It's up to us to stay true to our editorial values, to maintain the highest standards in news judgement, and to uphold our tried and tested journalistic methods in order to stay credible, legitimate and trustworthy in the mission to maintain a common, evidence-based shared reality.

ERIC SCHERER
Director News MediaLab and International Affairs at France Télévisions & Chair of the EBU News Committee

CONTENTS

CONTENTS

CONTENTS

GENERATIVE AI AND JOURNALISM: CHARTING THE TERRAIN

In history's timeline, the public launch of OpenAI's ChatGPT may well be seen as one of those pivotal events: the world before and the world after November 2022. Like the internet, the smartphone, or social media before it, the advent of generative AI has hit news media at its core. A technology capable of creating new, realistic content – text, visuals, and code – at speed and with ease, generative AI replicates many of the tasks journalists do, day in, day out. And while the previous disruptions challenged news organizations by opening up new channels for distribution, generative AI targets the very act of creation.

This technology has been charged with hopes and expectations and has equally elicited concerns. On one hand, tools driven by generative AI are predicted not only to empower people to do better and different work, but to entirely transform the way humans interact with machines. Many in the media hope that AI will enhance journalism by expanding its capacities – for example in data journalism – and broadening access for a variety of audiences. On the other hand, fears are held that AI-driven content and products will entirely erode trust, undermine existing business models, devalue cherished human skills, and deplete natural resources while further tying news organizations' fates to tech monoliths.

The range of unknowns and the speed of development are formidable. Nevertheless, large language models and other tools which are powered by pre-trained foundation models have mesmerized many in media organizations and beyond for a variety of reasons. Their potential to increase efficiency holds as much promise for journalists, who hope to be freed of mundane tasks and spend more of their time on research, as it does for CEOs who want to cut costs and expand reach. This is why in the second year since ChatGPT's launch, news organizations around the world have been experimenting, testing, and even implementing AI-based solutions while also trying to shield their brands from damage. The temptation to give in to the hype is enormous. And worries that a sea of fabricated, unverified content will erode the public's sense for what is true and what is false are widespread.

Public service media have a special responsibility in this context. While they are bound by the mission and mandate to inform, educate, and connect people, to serve and represent all of society, they are also accountable to the public. This calls for a particularly sensitive and risk-conscious yet strategic approach. The task is to explore how generative AI can help to fulfil public service organizations' mission while not endangering the fragile relationships of trust. In the face of technology that will have a fundamental impact, it will be critical for news media as a whole to continually prove the value and legitimacy of journalism to the public. Additionally, public service media must inform people about how generative AI might change not only the information ecosystem but central institutions like education, workplaces, and the way humans interact with each other. Investigating and explaining the new AI world will be the media's responsibility, too.

This report provides perspectives on what generative AI has in store for the news industry. It explores how this technology can support news organizations in strengthening connections with their audiences, how they might profit from efficiency gains and at the same time protect themselves from risks while living up to their ethical responsibilities. It highlights current debates in newsrooms and in the scientific community.

Crafting this work was not an easy ask. The frenzy about generative AI and the looming business opportunities have inspired countless publications – from studies and reports to podcasts and social media posts. To ensure this report makes a unique and impactful contribution, we have opted for a decidedly human lens. We provide a close-up view of those working with generative AI or researching its effects. We talked to dozens of newsroom leaders and managers in semi-structured interviews about their experiences, reflections, and predictions, and asked experts from beyond the media for their insights. Additionally, we drew on the latest research and many years of combined newsroom and academic experience in our team of authors.

This work cannot provide an in-depth assessment of all the issues that surfaced in our research, since the development and deployment of generative AI touches on so many different areas. Debates about copyright, regulation, the pros and cons of technology solutions or the environmental footprint deserve further consideration. Moreover, the development is extremely dynamic. New announcements and model releases on an almost daily basis make it impossible to provide an entirely up-to-date picture. Still, this report will help you reach a deeper understanding of the important, ongoing debates in the industry and the state of the research that naturally advances at a slower pace than the technology. For those already familiar with the subject, the report aims to inspire curiosity about the insights of your peers.

Our approach is also practical. The report contains checklists for strategy development and ethical guidance as well as an abundance of resources and advice for reporters covering AI. Importantly, it discusses more than a dozen use cases of what news organizations are already doing and exploring with generative AI. These practical applications have been an important part of all five previous editions of the EBU News Report.

However, a general understanding of generative AI is a precondition for its use in the newsroom and any organization. As such, Chapter 1 outlines some basic facts, and the hopes generative AI holds for journalism. It also discusses potential risks, such as its impact on trust, an ever-increasing dependency on tech companies, implications for the labour market, and the environmental footprint.

Chapter 2 bundles all the use cases and groups them into three categories: those that enhance audiences' experiences, those that increase efficiency in the newsroom, and those that support business models for journalism. Chapter 3 gives managers a voice. It highlights issues that arise in strategy development and operations, such as how to balance encouragement with risk management, how to set up processes and draft and develop guidelines. Chapter 4 is again about the bigger picture. It discusses the building blocks of ethical AI in news organizations by highlighting the imperative for public service value, touching on the contested issues of copyright and regulation, pressing for collaborations within the industry and the industry and tech.

The conclusion attempts to paint a picture of the future. It shares hypotheses on what will change and what will stay the same in journalism. And it finishes with a set of currently unanswered questions. While we hope you walk away from this report with at least some answers, asking the right questions is an essential skill in an age where prompting might be the new knowing. Understanding the big unknowns in this fast-moving field is a condition for any responsible use of generative AI – in the media and elsewhere.

Oh and, by the way, this report is entirely human made, including the illustrations. We hope it shows.

INTRODUCTION


CHAPTER 1
NEWS!
BREAKING NEWS!
NEWS!

GENERATIVE AI AND JOURNALISM: EXPECTATIONS, HOPES, AND RISKS

Provocatively put, journalism and generative AI contradict each other. At the core of journalism lie accuracy, consistency, and facts.[1] Some even like to call it 'truth' – or at least the best obtainable version of it. Generative AI, however, is about probabilities. Foundation models and large language models (LLMs) such as ChatGPT, Claude or Llama approximate facts. They are not designed to supply the truth, rather something that looks like it. This makes them both highly attractive and highly dangerous. As Bill Thompson, Head of Future Value Research at the BBC, says: "Journalism is based around the early Wittgenstein's theory of truth and facts, the coherence model of truth.[2] None of this is reflected in the way LLMs work. […] ChatGPT is like a drunk person in a forest staggering along the trees shouting out their names."

Nevertheless, generative AI will be a game-changer for journalism, as it will change practices, habits, and products everywhere. Applied in the right way, it promises more efficiency and effectiveness, in speeding up work routines, increasing access to journalism for a broader public, and helping to better serve a diversity of audiences, to give a few examples. There are many cases where this already happening, as Chapter 2 of this report discusses.

However, the media leaders, academics, and other experts we interviewed disagreed on how exactly the expected changes will play out, where they will impact the most strongly, and how fast they will happen. These experts fall roughly into two groups. The first and larger group expects the technology to affect predominantly journalistic practices and workflows. Niddal Salah-Eldin, Vice President People and Culture at German publisher Axel Springer predicts radical changes in news production: "In the long run, journalistic production will become a by-product, more technically supported and automated." (Read the Q&A with Niddal Salah-Eldin, page 119) But this group also thinks that journalism as such, its status in society, and its core objectives won't be touched. These are about public service, objectivity, independence, ethics, a focus on collecting and sharing facts, telling stories, and helping different voices to be heard. As Dmitry Shishkin, former BBC journalist and CEO of Ringier Media International, puts it: "I am optimistic about journalism as long as it is proving its worth to society. Machines cannot go to interesting people and talk about interesting things."

The second group forecasts a profound change in communication habits and consequently the entire information ecosystem. David Caswell, who has worked in the media and tech industry for many years – including at the BBC – and now consults for news organizations on AI, argues: "The most dramatic change is that generative AI separates information content in news artefacts like texts or videos. In other words,

[1] See the definition of quality journalism in the Council of Europe's Recommendation "Promoting a favourable environment for quality journalism in the digital age," adopted by the Council of Ministers on 17 March 2022: "Quality journalism, with its unwavering commitment to the pursuit of truth, fairness and accuracy, to independence, transparency and humanity, and a strong sense of public interest in promoting accountability in all sectors of society, remains as essential as ever to the health of democracies." Disclaimer: Alexandra Borchardt was vice-chair of the corresponding expert committee. https://edoc.coe.int/en/international-law/11046-promoting-a-favourable-environment-for-quality-journalism-in-the-digital-age-recommendation-cmrec20224.html, retrieved on 26 April 2024.

[2] Put simply, truth in the coherence model is determined by how well a statement ('The apple is green') fits into a system of beliefs or knowledge. In other words, a statement is considered true if it is consistent or coherent with other beliefs or pieces of knowledge that are held to be true.

it separates the semantic information from the way it is communicated. Dealing with these artefacts rather than the unit is a fundamentally new thing. You can then consume information in the form you chose. You are not restricted any more to the way the producer wanted you to consume it. It will reconfigure the information ecosystem." (Read the Q&A with David Caswell, page 31)[3]

These predictions are to some degree complementary. But they are overshadowed by significant concerns about journalism's future business models. It is presently unclear how brands and sources will show up in a changing search environment which is potentially driven by conversations in the form of chats and personalized to individual users – creating 'end products' rather than lists of results. In May 2024 OpenAI announced a model which provides better interaction via speech[4], a move that was followed a day later by Google announcing its Project Astra, AI assistants capable of answering spoken questions and handling multimodel inputs such as videos, images, audio and texts.[5]

While much is in flux, generative AI will likely be used to bring about new business and information environments. Some fear that the most important loyalty-drivers for publishers in recent years – websites and apps – will further decrease in relevance. At the Mobile World Congress in February 2024, Deutsche Telekom was the first company to present a study of a smartphone without apps.[6] Telekom-CEO Timothy Höttges was quoted as saying: "I can tell you that in 5–10 years from now, none of us will use apps anymore." This is a dramatic shift as it undermines publishers' strategies to draw back the audience to their platforms.

Given the high degree of uncertainty, media companies need to prepare for different scenarios. It is essential that they take advantage of new opportunities, ideally with an exploratory yet risk-conscious mindset. At the same time, they need to future-proof their business models. This will be quite a challenge, since many media organizations haven't even fully navigated the first waves of disruption that occurred with the internet and social media. As media strategist Lucy Küng points out: "The media's reaction to generative AI now bears the scars from the painful transition to digital. It took the industry a long time to understand what was shifting, what that meant, and then which changes to make. And it has emerged from that transition in a weaker position." (Read the Q&A with Lucy Küng, page 35)

This first chapter will discuss the opportunities, expectations, and hopes that media leaders and AI experts identify in the context of generative AI. It will shed light on what our interviewees singled out as their biggest concerns and their perspective on the risks. These include fears of internal disruptions, and dependence on a powerful set of tech giants who define the rules of the game while regulators try to steer the technology into safer waters. The effects of generative AI on the entire information ecosystem are of considerable concern to many as new opportunities to create and amplify misinformation and disinformation might further erode trust in journalism. In fact, 70% of international media leaders surveyed in December 2023 for a Reuters Institute report thought so.[7] At the same time, newsrooms are worried about audiences being further exhausted by information overload and consciously avoiding news as a result. Nevertheless, many interviewees

[3] See also: David Caswell, "AI and Journalism: What's Next?", Reuters Institute for the Study of Journalism, 19th September 2023. https://reutersinstitute.politics.ox.ac.uk/news/ai-and-journalism-whats-next, retrieved on 29th February 2024

[4] Hern, A. (2024, May 14). OpenAI's new GPT-4o model offers promise of improved smartphone assistants. *The Guardian*. Retrieved from https://www.theguardian.com/technology/article/2024/may/14/openai-gpt-4o-model-offers-promise-of-improved-smartphone-assistants

[5] Vallance, C. (2024, May 14). Google's new AI can find lost specs. BBC News. Retrieved from https://www.bbc.co.uk/news/articles/cq5nnlrwn9do

[6] Reuters, "Deutsche Telekom showcases app-less AI smartphone concept", 26th February 2024. https://www.reuters.com/technology/deutsche-telekom-showcases-app-less-ai-smartphone-concept-2024-02-26/, retrieved on 29th February 2024.

[7] Nic Newman, "Journalism, Media, and Technology Trends and Predictions 2024", Reuters Institute for the Study of Journalism, January 2024. https://reutersinstitute.politics.ox.ac.uk/journalism-media-and-technology-trends-and-predictions-2024, retrieved on 7 April 2024.

feel that the very technologies causing this disruption may also have the potential to remedy at least some of these ills.

Some think – and even hope – that all of this will lead to a long overdue general debate about the role, legitimacy, and value of journalism for the public in a trade that has been practised around the globe with surprisingly similar rules and routines for many decades. Bill Thompson of the BBC says: "It leads to the question: What is the role of journalism in society? The internet, web 2.0, social media, all have challenged journalism in different ways. Generative AI is just the latest wave of this transformation. It is a game-changer, because it asks why do you need journalism? Generative AI doesn't answer that, but the question is a good starting point."

Anne Lagercrantz, deputy Director General of Swedish Television SVT, thinks the advent of generative AI provides an opportunity for the media to reassess its strength: "It will fundamentally change journalism but hopefully not our role in society. We have to work on the credibility of the media industry. We need to create safe places for information."

Understanding generative AI

This report focuses on journalism and the media industry. There are other publications more suited to discuss the meaning, significance, and potential of generative AI in all its facets (see list of reading recommendations, page 177). But to understand the current hype around this technology, it's important to offer some general remarks, and a note of caution about its capabilities. Plenty of otherwise highly educated people use generative AI applications without having a basic understanding of how these systems work or what their limitations are. In the summer of 2023, the case of an experienced New York lawyer made headlines after his attempt to defend a client citing an array of precedents which didn't exist. As he explained to the judge, he had used ChatGPT to write his brief, assuming the system was a powerful search engine.[8] With all the hype that surrounds generative AI, it is easy to overestimate its proficiency and make unwise, even dangerous, decisions.

Melanie Mitchell, a computer scientist and Professor at the Santa Fe Institute, works at the intersection of cognitive science and AI. She defines artificial intelligence as "computers doing something that would require humans to use intelligence". One of the core questions Mitchell is exploring is how humans acquire and use intelligence, and how this is – or is not – mirrored by computers and current AI systems: "LLMs are doing well on certain benchmarks but they are doing badly on others. I am trying to understand how humans do these things that seem to be complicated for machines without huge amounts of training data and how to get machines to do this. What I am interested in is human adaptability and robustness, our ability to apply knowledge successfully to situations we have never seen before." One example she likes to share concerns a Tesla that kept slamming on the brakes without obvious reason. "It turned out it was about to pass a billboard, with a policeman holding up a stop sign. The car – that is full of tech – thought it was a stop sign. Any toddler without much previous knowledge can tell reality and images apart. We are quite robust in our abilities to deal with the world."

To understand the intersection of generative AI and journalism, it is important to distinguish between generative AI and 'regular' AI. Regular, 'traditional' or 'narrow' AI includes a diverse range of applications and techniques with different levels of complexity, autonomy, and abstraction, chipping away at narrowly defined tasks and problems. Examples of narrow forms of AI include applications of machine learning and its subfield, deep learning, as well as various forms of natural language processing which often build on machine learning approaches.

[8] Benjamin Weiser, Nate Schweber, "The ChatGPT Lawyer defends himself", The New York Times, 8 June 2023, https://www.nytimes.com/2023/06/08/nyregion/lawyer-chatgpt-sanctions.html, retrieved on 24 February 2024.

In 'narrow' AI, a computer programme or system learns directly from examples, data, and experience with algorithms trained on large amounts of data, thus improving the system's performance on a narrowly defined task over time.[9]

Many media companies have been using these forms of AI for quite some time, to report on earnings, elections, and minor league sports results, or to optimize paywalls and content moderation, for example. Notable ventures in this field have been made by well-resourced organizations like the Associated Press, Finland's public broadcaster Yle, and the Washington Post, but also by companies like the Canadian *Globe and Mail* or the regional public broadcaster Bayerischer Rundfunk (BR) in Southern Germany. In Sweden, local news organizations have been driving reach and subscriptions with AI-based products for quite some time. The EBU News Report 2019, written by former Yle News Director Atte Jääskälainen and Maike Olij, documented some of the then state-of-the-art usages by public service media and beyond.[10] In the same year, data journalism and AI expert Nicholas Diakopoulos, professor at Northwestern University, discussed artificial intelligence and journalism in his widely acknowledged book Automating The News.[11]

Generative AI started to become publicly known in the wake of the public launch of OpenAI's ChatGPT in November 2022. At its core are so-called 'foundation models' which are trained on massive amounts of text, images, code, mathematical equations, and the like. They can generate new, realistic forms of data resembling this training data. In their present configuration, these systems often take the form of chatbots which can be operated using regular human language dialogue. They produce output in response to a question, which in generative AI language is called a 'prompt'. For text-based foundation models, this output is not a factual representation of truth. Instead, the system uses complex statistics to predict the most likely word to follow each word or 'token'[12] based on massive amounts of training data and producing what appears to be a probable answer to the posed question. Which data the application uses and can handle depends on the foundation model (see Figure 1). The latest are multimodal and can generate video, audio, and text. This makes them particularly attractive for journalists who are specialized in one of these categories.

Luciano Floridi, philosophy professor and Founding Director of the Digital Ethics Centre at Yale University, explains: "A lot of what generative AI improves in digital services is not vertical but horizontal. In the past, search took you deep into silos. Generative AI will be more of an interface between services, like a glue that puts things together: text, images, sound, code. (...) In the past, large language models couldn't link to search engines, that's why their results were always outdated. Now they can. Generative AI is creating a seamless environment." (Read the Q&A with Luciano Floridi, page 156)

Content produced by some generative AI systems can outperform humans with average, or even above average, language or coding skills. This makes it particularly valuable – and potentially risky – for those working in foreign languages or in fields beyond their expertise. With the help of generative AI, everyone who knows how to prompt and use the right tools can, for example, build a website, compose an infographic, or produce an illustration without having to consult a specialist. This is why many organizations expect huge potentials for cost-savings.

[9] Felix Simon, "Artificial Intelligence in the News: How AI Retools, Rationalizes, and Reshapes Journalism and the Public Arena", Tow Center for Digital Journalism, 2024, page 10. https://www.cjr.org/tow_center_reports/artificial-intelligence-in-the-news.php

[10] Atte Jääskälainen, Maike Olij, "The Next Newsroom: Unlocking the Power of AI for Public Service Journalism", European Broadcasting Union, November 2019, https://www.ebu.ch/news/2019/11/ebu-news-report-urges-public-service-media-to-work-together-to-face-challenge-of-platforms, retrieved on 24 February 2024.

[11] Nicholas Diakopoulos, Automating the News – How Algorithms Are Rewriting the Media, Harvard University Press, 2019.

[12] Systems like ChatGPT use 'tokens', linguistic units that might be whole words, or components of words like 'pre' or 'ing' or 'ized'. See: Wolfram, S. (February 14 2023). What Is ChatGPT Doing ... and Why Does It Work? Stephen Wolfram. https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work/

Figure 1: The elements going into foundation models.

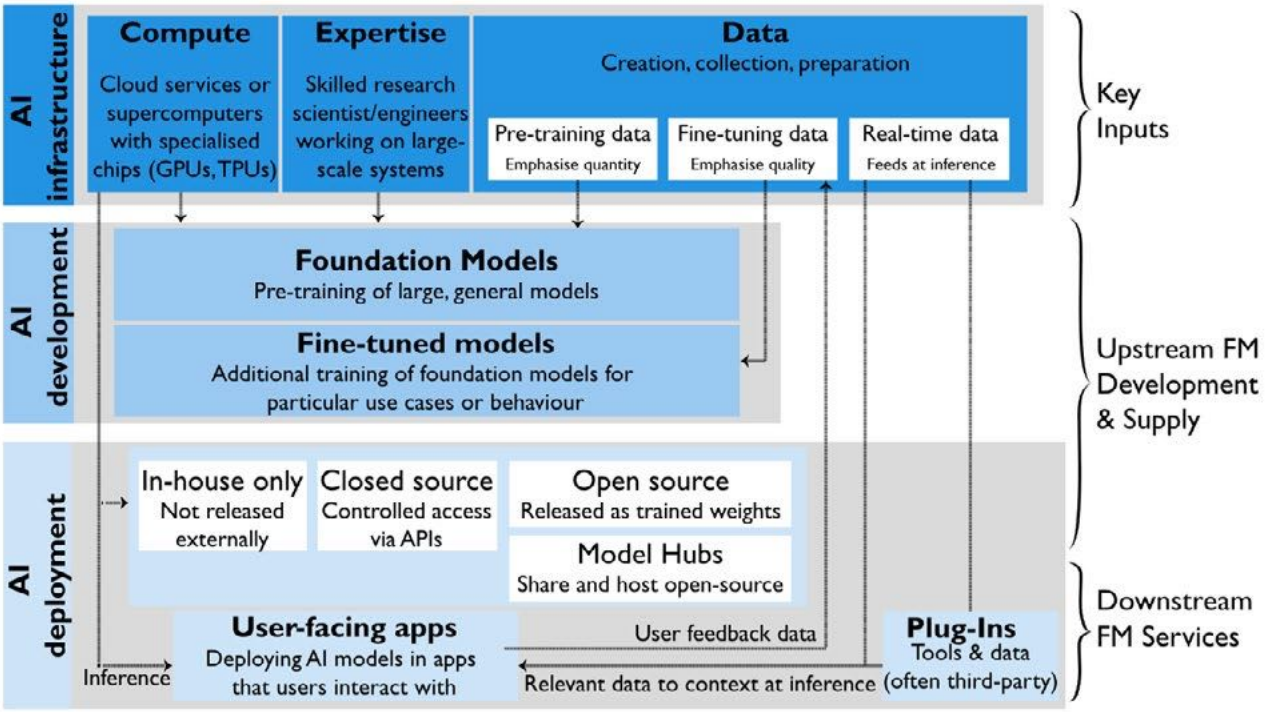

Source: Competition and Markets Authority, AI Foundation Models Review (2023): https://assets.publishing.service.gov.uk/media/65045590dec5be000dc35f77/Short_Report_PDFA.pdf [accessed 14 December 2023]

Johanna Törn-Mangs, Director and Editor-in-Chief of Svenska's Yle in Finland is responsible for the broadcaster's AI guidelines. She describes it as a breakthrough innovation even in her tech-savvy organization: "We at Yle have been using AI for decades, including generative AI. But ChatGPT was a game changer. It gives people direct access to AI capabilities with an easy-to-use user interface, and makes it possible to achieve higher quality, richer content, and more efficiency... Journalists and programmers have started using and developing tools faster than before," she claims. As Mattia Peretti, AI Consultant and former Manager of the London School of Economics' Journalism AI Project puts it: "Generative AI has turned AI from something technical that only a few could deal with to something everyone can use."

Natali Helberger is Professor of Law and Digital Technology at the University of Amsterdam and Co-Director of the AI, Media and Democracy Lab. She also describes this accessibility as the key factor that makes generative AI transformative: "Society might not have realized how good it is. Generative AI allows people to engage with information, with written or spoken text in very different ways. For example, with ChatGPT, you can talk to a document and ask it questions. New generations of chatbots and personal assistants will hopefully make life easier. So far it has been used to optimize for commercial goals, but it can be optimized for people's needs." Numerous opportunities for professionals in the media industry, whose expertise has traditionally been rooted more in text, voice, and visual productions, are already emerging.

A significant benefit for any type of knowledge work could be that by using generative AI, people can play to their strengths and fill the gaps on what they are not as good at or not directly skilled to do. Gina Neff is a Professor at Cambridge University and member of several boards which deal with the potential for and impact of AI on businesses and the workplace. She studied the roll-out of automated tools for technical construction planners, and describes how the lessons from this research are applicable to generative AI, too: "There I saw the sheer creativity people can develop when they bring tools into their work for tasks they are not good at doing. The tool

Figure 2: Many players develop tools, and many more involved in AI

Source: EBU MIS Report: Artificial Intelligence - Public Service Media Leveraging AI.
https://www.ebu.ch/research/membersonly/report/artificial-intelligence---public-service-media-leveraging-ai

didn't find solutions but freed up space for people to do their best work. [...] Highly skilled workers will benefit, people who are not programmers can automate jobs, it will change how people get their work done. But it will take a lot of work to produce results."[13]

The big players who are driving the current activity around AI are mostly household names: OpenAI, which launched ChatGPT, was founded by a small tech circle in Silicon Valley, but it is strongly connected to the ecosystem of existing tech behemoths. At the time this report was written, Microsoft held a 49% stake in OpenAI. It was also shareholder of French start-up Mistral which has been praised for providing a European alternative to OpenAI's ChatGPT.[14] Google has been developing Gemini, its own suite of generative AI models which it has been integrating into its search function. However, not only is the technology developing at a speed that cannot be captured by a report like this, the business ecosystem around it is transforming, too. This combines to create a dynamic and competitive process in which smaller players often become entangled in the networks of power and interests of existing market leaders (see Figure 2).

Frictions between technology companies and publishers have emerged primarily around copyright issues, the big question being: how should publishers be acknowledged and compensated for the content used to train foundation models? Much of this data has been drawn from the work of journalists and others in the creative industries. In response, news organizations including *The New York Times* have gone to court; others have struck deals with AI companies. Meanwhile, an increasing number of organizations have attempted to block their content from being used for certain generative AI applications. A factsheet by the Reuters Institute published in February 2024 revealed that in a sample of 10 countries, an average of 48% of news websites blocked OpenAI's AI crawlers from accessing their content, while 24% blocked Google's.[15] It is unclear how content blocking will impact the quality of LLMs in the long run, or how effective it will be.

These and other conflicts and considerations will be dealt with in Chapter 4, which discusses the ethical implications around the use of generative AI. We suspect that quite a few more issues will emerge in the debates surrounding the dependence of the media industry on large technology companies as the technology advances and public adoption accelerates. Tech companies are already playing favourites by giving certain media companies permission to test certain tools in exchange for access to content.[16]

[13] Dossick, C., Osburn, L., & Neff, G. (2019). Innovation through practice: The messy work of making technology useful for architecture, engineering and construction teams. Engineering, Construction and Architectural Management. https://www.emerald.com/insight/content/doi/10.1108/ECAM-12-2017-0272/full/html

[14] The Economist, "Meet the French startup hoping to take on OpenAI", 26 February 2024. https://www.economist.com/business/2024/02/26/meet-the-french-startup-hoping-to-take-on-openai, retrieved on 29 February 2024.

[15] Richard Fletcher, "How many news websites block AI crawlers?", Reuters Institute for the Study of Journalism, 22 February 2024, https://reutersinstitute.politics.ox.ac.uk/how-many-news-websites-block-ai-crawlers, retrieved on 24 February 2024.

[16] See the discussion of Google's activities in this podcast episode: "Mystery AI Hype Theater 3000. Episode 29: How LLMs are breaking the news," published on 3 April 2024. https://www.buzzsprout.com/2126417/14807430

Whether generative AI will ultimately deliver on the high hopes currently invested in it is far from clear. The outcome depends on advances of the corresponding technologies as well as public uptake. As Melanie Mitchell warns: "Every new technology goes through that cycle of hype and disappointment. Currently, there is an enormous amount of hype [...] Anyone who tries to predict things will get it wrong."

How journalism could profit

Generative AI has created quite a buzz in many media organizations. Even though its full implications are not yet clear, some journalists have high hopes for a future in which generative AI can enhance quality, be it in research, editing, format development, or data journalism. As Ezra Eeman, Strategy and Innovation Director at Dutch public broadcaster NPO says: "With generative AI we can fulfil our public service mission better, it will enhance interactivity, accessibility, creativity. AI helps us to bring more of our content to our audiences." (Read the Q&A with Ezra Eeman, page 127)

Niddal Salah-Eldin of Axel Springer points out: "It's giving us incredible opportunities to streamline our processes – making our newsrooms more efficient – and create new experiences. This enables us to focus more on the heart of journalism, creating compelling stories and digging deep into investigations. Plus, AI enables us to offer a wider range of content, from special interest pieces to more localized stories, which opens additional avenues for advertising revenue."

A survey conducted by the Journalism AI project at the London School of Economics among 105 news and media organizations from 46 countries revealed that 75% of the surveyed newsrooms were already using AI applications (without differentiating between generative AI and 'regular' AI). However, almost half felt that their organizations were not ready or not quite ready for the challenge.[17] Nevertheless, in 2023 one third of such organizations reported that they were working on an AI strategy.

Agnes Stenbom, Head of IN/LAB & Trust Initiatives at Schibsted and founder of the Nordic AI Journalism Network together with Olle Zachrison, says: "I have big hopes. I don't think generative AI will change journalism all at once, but it can play a part in future-proofing journalism. It is a tool with immense potential, but it is only a tool. Currently, the promises are a bit over-hyped. But my hopes are that it will enable us to create content in ways that engage a diverse audience." Verena Krawarik, Head of Innovation at Austria's news agency APA, thinks similarly: "I believe that it is a gift, because we journalists have to handle a massive content production output. Today you have to produce many stories for different channels. Many of us have run out of steam in recent years. This helps us to reach our readers much better. It's also a game changer because it frees up resources for sourcing. It will help with investigative journalism."

Together with her team, Uli Köppen, Head of AI and Automation Lab at Germany's Bayerischer Rundfunk, has been working on AI for many years. For her, generative AI is definitely a game-changer: "People now can lay their hands on AI and can try out something themselves. With those platforms you don't always need a developer anymore. You don't even have to think like a developer anymore. You can use your simple plain language to talk to AI, and you can get outputs like texts and pictures. This possibility has changed the mindset in the newsroom and in the management. This is a positive thing because we don't have to persuade anyone anymore that this is really an important topic that needs resources."

Jane Barrett, Global Editor Media News Strategy at Reuters, expects generative AI to have an impact on three fronts: workload reduction, augmentation of existing work, and transformation. Barrett: "Replacing routine

[17] Charlie Beckett, Meera Yaseen, "Generating Change: A global survey of what news organizations are doing with AI," London School of Economics, September 2023. https://static1.squarespace.com/static/64d60527c01ae7106f2646e9/t/656e400a1c23e22da0681e46/1701724190867/Generating+Change+_+The+Journalism+AI+report+_+English.pdf, retrieved on 29 February 2024.

tasks at scale may take more development work but already AI can speed us up and help us do more with the resources we have. (...) AI skills might help us re-version a story into social posts, a video script, a quick summary for busy readers, a translation. Or AI could augment our work by helping find stories in data dumps or write explainers from our archive." (Read the Q&A with Jane Barrett on page 134)

Some expect journalism to fundamentally change from a push activity, where media organizations serve audiences with news products, to a customer-driven pull experience where audiences decide for themselves the format they want to interact with at any given moment. Anne Lagercrantz of SVT is convinced: "Journalism will become a pull activity. Today we use search engines, in the future we will use answer engines." Bill Thompson of the BBC takes an even more radical stance: "I think the web is over, don't you? A web page with predictable content for everyone who visits it, these days are gone. The future will be more like: When you go to a URL and frame a question, it will generate something for you, based on your history, the current context, the editorial approach of the site you are visiting. It will be made on the fly with different images. Give it ten years' time, and the web might be over."

Agnes Stenbom is more cautious: "I think we can see both. The need for pushing news will remain. But we will also see a need for a more dialogue-based news experience. It means making an opening for users to have an active say in what information they would like to have and in what formats. Making it more a two-way relationship is required for users who are born digital." Niddal Salah-Eldin says: "I think this shift is nothing new. It has always been the purpose of journalism to make stories so relevant that people actively pull them into their lives."

A note of caution is necessary, as the assumption that people want a lot of say in how they experience news has been wrong in the past. Many seem to enjoy the

Figure 3: How people in six countries say they use generative AI, including for news.

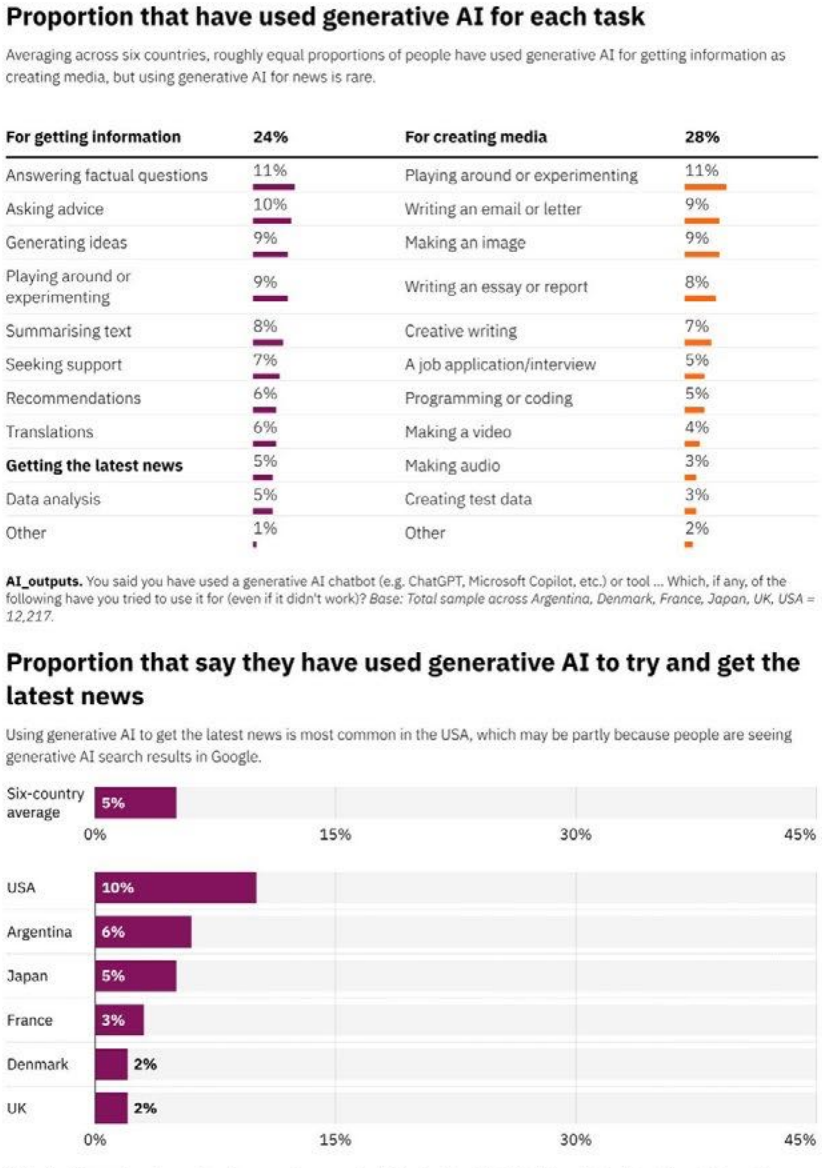

Source: Data from 'What does the public in six countries think of generative AI in news?,' published in May 2024


With generative AI we can fulfil our public service mission better, it will enhance interactivity, accessibility, creativity. AI helps us to bring more of our content to our audiences.
EZRA EEMAN, Strategy and Innovation Director at Dutch public broadcaster NPO
NEWS!
BREAKING NEWS!
NEWS!

role of a 'passive' consumer, with others deciding what might be of interest. They don't want to comment on news as much as expected. They've often felt overwhelmed by complicated (and expensive) storytelling endeavours. And, despite the popularity of podcasts and video, particularly with young audiences, text is still a popular and efficient format of news consumption, as the Reuters Institute's Digital News Report has repeatedly revealed.[18] Agnes Stenbom points to the experiments with gamification as something which journalists might have had more fun developing than audiences were interested in engaging with: "People thought that gamification could lower the threshold to news consumption, but it made the threshold even higher." So far, surveys of audience attitudes towards generative AI tools seem to support these more cautious views. In a representative survey of six countries led by Richard Fletcher of the Reuters Institute at the University of Oxford, only 5% said that they currently use generative AI to search for news, although news-adjacent behaviour such as getting answers to factual questions ranks more highly (see figure 3).

One of the major features of generative AI is that, in its current form, it lowers thresholds to almost any interaction with content. Whether and how people will use chatbots for their information and news needs is, however, impossible to tell. Those who are well educated, tech savvy and started to experiment with these systems early might find it difficult to imagine, but most people don't pro-actively use generative AI. Further, many communities have no stable or affordable internet access, low media and tech literacy and a reluctance to use digital services, particularly among older generations and the socially disadvantaged both in the Global North and South.

From the perspective of resource-starved news organizations, however, generative AI is widely regarded as a gift. Many hope that it will enable such organizations to broaden their range of offerings without hiring additional staff. Expensive and scarce resources such as graphic design and programming, which are out of reach to today's newsrooms, may no longer be needed – or needed less – to produce quality products. Styli Charalambous, CEO of the South African *Daily Maverick* which employs around a 100 people, says: "The game changer for us is that generative AI allows us to play with tools and technologies that would have been previously very, very difficult for us to get our hands on and to be able to do this across so many different parts of the journalism value chain. It does democratize access to technology for smaller and midsized organizations."

David Caswell, who has interviewed media leaders from all over the world, concurs: "I don't see any geographic differences. I see differences by the size of organizations. Small newsrooms are disproportionally empowered by this. The functionality is so accessible. The only barrier for small organizations is imagination and a couple of hundred dollars' fees for tools."

Risks, dangers, and side effects: Between hallucinations and new dependencies

Dwelling on the risks of generative AI doesn't seem to be popular among media managers as they endeavour to adjust to the new reality. Admittedly, our interview sample might be positively skewed towards hoping for new solutions to old problems. But many of our interviewees said that the media is now better equipped to focus on the gains than in previous waves of innovation. This was different when the internet emerged as a new platform for content distribution about 30 years ago. Back then, many in the media had hoped it would somehow disappear or paid it scant attention, to their detriment.

When social media disrupted news again a decade later, resistance wasn't as pronounced, even though there has been ongoing scepticism around certain platforms – much of it justified. For example, with its Chinese ownership, TikTok is a platform many media

[18] Nic Newman, Digital News Report 2023, Reuters Institute for the Study of Journalism, Digitalnewsreport.org

companies would rather do without. But it has recently been viewed as the most likely path to reach young audiences who are otherwise not engaging with news. This is why many news organizations consider a TikTok presence to be a must. Public service media, which are mandated to serve all of society, feel a particular responsibility to produce news for the platform, while commercial players can afford to be more reluctant.[19]

Our interviews made it abundantly clear that the industry's outlook on tech has changed. The current generation of decision makers are in many ways more realistic about the benefits and risks of new technologies than their predecessors. They know this technology will not go away and many adjust their strategies accordingly. As Ezra Eeman of NPO puts it: "We have a moral duty to be optimists and convey a sense of opportunity rather than despair." But it would be naïve to gloss over the risks and outright dangers of generative AI, just because it might spoil the party.

Madhav Chinnappa, who was one of the founders of the Google News Initiative after a career at the BBC, has concerns. "A difficulty with any technical tool is, that we in the news ecosystem tend to think about it from what I call the good actors' perspective. But having worked in tech for 13 years, what you realize is you've got to think about it from the bad actors' perspective as well. I think about generative AI in a few buckets. One is: what can generative AI and all the incumbent tools do for a newsroom? And this is kind of positive. But then what can it do for the information ecosystem? And I'm worried that generative AI will be better for the bad actors than it is good for the good actors."

From a journalistic perspective, the risks can be categorized into three groups. Firstly, the risks built into the technology itself. These include the propensity of generative AI systems for hallucination, reinforcing stereotypes, violating copyright and data protection – all this depending on the underlying foundation models. The potential impact of generative AI on energy and resource consumption belongs in this category as well.

The second risk group relates to the way these systems are employed. Will users fact-check content and sources before publishing or sharing? Will the sheer possibility of producing new content and products almost instantaneously lead to massive overproduction? As Ritu Kapur, Founder and

Figure 4: Concerns around generative AI among a non-representative sample of international journalists[20]

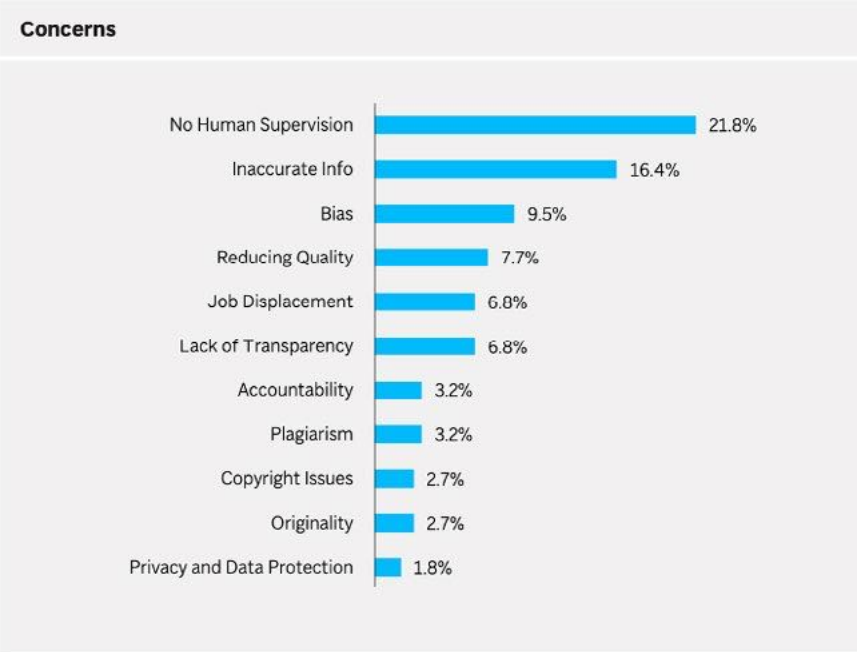

Source: Data from "Generative AI in Journalism: The Evolution of Newswork and Ethics in a Generative Information Ecosystem", published in April 2024

[19] For a closer look at strategy and leadership in a multi-platform media world, see the EBU News Report by 2021 Alexandra Borchardt and Felix Simon, "What's Next? Public service journalism in the age of distraction, opinion, and information abundance", European Broadcasting Union, November 2021. https://www.ebu.ch/publications/strategic/loginonly/report/news-report---whats-next-public-service-journalism-in-the-age-of-distraction-opinion-and-information-abundance, retrieved on 29 February 2024.

[20] Diakopoulos, N., Cools, H., Helberger, N., Li, C., Kung, E., & Rinehart, A. (2024). Generative AI in Journalism: The Evolution of Newswork and Ethics in a Generative Information Ecosystem (p. 48). Retrieved from https://hdl.handle.net/11245.1/9a0679dd-7bf8-4e01-9ba0-ba10409afe80, p. 24

CEO at Indian news brand The Quint, says: "A lot of the conversation around generative AI is about speeding up the news process, about increasing the volume of publishing, increasing traffic, search engine optimizing, and therefore putting on steroids everything that tech algorithms have been doing for news. All of this worries me because speed and volume and tech optimization have been what has killed a lot of good journalism." Further, an array of external factors is involved. Will tech companies worry about their carbon footprint or the increases in energy consumption? To what extent will bad actors use generative AI to create disinformation, deep fakes, fraud, and can other tools successfully counteract these efforts? Which jobs will be replaced with the help of generative AI, and which rules will be applied to ensure it is developed and implemented responsibly?

The third group of risks is structural: How will dependence on a few major technology companies shape the way people access information and communicate and connect with each other? How will an abundance of artificially (co)fabricated content affect overall trust in news and information? And, most importantly: how will individual rights and freedoms fare in this new environment?

Much of this depends on how the public will engage with the available tools – which is still a big unknown. Verena Krawarik of Austria's APA says: "We have to explore more about how all this affects the consumption and the presentation of content. How will we consume media in the future in which formats? How do regular people interact with it? It is always good to look at what happens in schools and universities. ChatGPT is no journalistic research tool, but people use it like this anyway, whether we like it or not. We need to find answers to this."

Generative AI, of course, doesn't exist in a void. It exists in societies already impacted by political polarization, inequalities, and discrimination – to name just a few issues. As the UK's House of Lords' Communications and Digital Committee summarised it in a February 2024 report: "LLMs may amplify any number of existing societal problems, including inequality, environmental harm, declining human agency and routes for redress, digital divides, loss of privacy, economic displacement, and growing concentration of power." The committee concluded that the integration of LLMs in critical infrastructure needs to be questioned. The report did not address media organizations explicitly, but many news publishers would probably define themselves as 'critical infrastructure' – public service media most certainly are. The committee recommended the situation be reassessed further down the road: "LLMs continue to hallucinate, exhibit bias, regurgitate private data, struggle with multi-step tasks, and pose difficulties for interpreting black-box processes. […] Improvements to bias detection, memory, complex task execution, error correction and interpretability are major areas of research and some improvements within three years are highly likely."[21]

Risks to the brand: Separating right from wrong

The most obvious direct risk is that AI systems, in particular LLMs, fabricate content and sources that don't exist in response to requests for accurate information. Andrew Strait, Associate Director of the Ada Lovelace Institute in London and responsible for the institute's work around addressing emerging technology and industry practice, is outspoken on this: "Hallucination is not a bug but a feature. Generative AI is predicting what the most likely answer is. The way it mimics behaviour is disturbing. After using it for a year I sometimes can't tell fact from fiction." Strait's somewhat discouraging assessment is: "Everywhere where truth doesn't matter it is okay." He says this should encourage decision makers to think twice when implementing products that are based on generative AI. "I'm still not very bullish on this. Perhaps this is a trauma from seeing this playing out for

[21] House of Lords, Communications and Digital Committee, "Large language models and generative AI", 2 February 2024, pages 50 and 16. https://publications.parliament.uk/pa/ld5804/ldselect/ldcomm/54/5402.htm, retrieved on 26 February 2024.

the last 15 years I have been working in this industry. The expectations are insane. We need to be very cautious when we replace sensitive human relationships. I'm talking about tax advisors, teachers. You can imagine this moving into mental health and psychotherapy. The pace at which untested and unevaluated systems are being deployed in various areas of life is alarming."

Jane Barrett of Reuters says newsrooms definitely need to learn how to deal with potentially made-up content: "We have to tread very carefully because as journalists we deal in facts and generative AI models are prone to hallucination. I liken today's gen AI models to a Formula One car. However well you drive, you need to train to get behind the wheel of an F1 Ferrari and not crash. And you need a team of excellent technologists – and in AI, data scientists – around you to get to where you want to go safely. It's not a silver bullet or a quick solution to our problems."

Matt Frehner, Head of Visual Journalism at the Toronto-based *Globe and Mail* says that even though the *Globe* has used plenty of AI in the past – the title even developed its own AI system to curate its home page – it has been much more restrained with generative AI: "We are cautious and we are not in a rush to do things. We are not going to use every flashy new thing. If we get caught with one photo that is AI-manipulated, that calls everything into question. One mistake can really damage the brand, and for the *Globe*, it is all about the brand. That's why people subscribe. Look at *Sports Illustrated*. That has become a skeleton brand." In November 2023 the sports magazine was caught publishing fabricated images and stories with fake author bylines that didn't exist.[22] The resulting controversy led to firings, a change of publisher for the title and litigation.[23]

It is particularly tricky when hallucinations occur in systems people trust and work with on a daily basis – like Microsoft Office. These systems are increasingly equipped with functions building on foundation models, but users don't necessarily know this, nor are they informed about the corresponding risks. For example, the most popular office software package in the world includes a co-pilot 'assistant' function which was apparently still prone to fabricating facts in early 2024. According to a report published by Sherwood Media, the software suggested quotes by Russian president Vladimir Putin about the death of Alexei Navalny.[24] However at that time, Putin was yet to comment on Navalny's death at all. Confronted by the reporter, a Microsoft spokesperson was quoted saying, apparently nonchalantly, they were working on the "quality of the responses."

It is common practice in the software industry to release a version and then update it on the fly. But this is dangerous territory when individual rights are concerned. This is one of the reasons why granting victims of algorithmic miscalculations fast access to remedies was one of the key points in a Recommendation by the Council of Europe's Expert Committee on Digital Technologies and Freedom of Expression adopted in 2022.[25]

The trust question: Between deep fakes and general doubts

Many journalists worry about deep fakes, which seem to be proliferating exponentially. Blathnaid Healy, Executive News Editor for Growth, Social and Delivery at the BBC, says: "For a long time, deep fakes haven't been that large an issue for newsrooms. We can see this changing now." Enabling colleagues to verify material has become an essential task in many organizations, particularly as 2024 is an election year in many parts of the world.

[22] PBS, "*Sports Illustrated* found publishing AI generated stories, photos and authors," 29 November 2024. https://www.pbs.org/newshour/economy/sports-illustrated-found-publishing-ai-generated-stories-photos-and-authors, retrieved on 26 February 2024.

[23] https://www.bbc.com/news/world-us-canada-67619015. He is also suing the company: https://www.thewrap.com/ousted-ceo-of-former-sports-illustrated-publisher-sues-for-20-million/. Also the title was removed from the company that published it: https://www.nytimes.com/2024/03/18/business/sports-illustrated-magazine.html.

[24] Rani Molla, "Microsoft's Copilot AI search is making up fake Vladimir Putin quotes from press conferences that never happened," Sherwood, 22 February 2024. https://sherwoodmedia.com/news/microsoft-copilot-ai-search-chatgpt-is-making-up-fake-vladimir-putin-quotes/, retrieved on 26 February 2024.

[25] Council of Europe, "Recommendation CM/Rec(2022)13 of the Committee of Ministers to member States on the impacts of digital technologies on freedom of expression," adopted on 6 April 2022.

A report by the Global Investigative Journalism Network described how audio deep fakes could pose "the chief threat in this election cycle" and prepared a tip sheet on how these can be identified and investigated.[26]

But this is not only about the most disturbing instances of deep fakes simulating politicians' speeches or pornographic images of celebrities as happened to Taylor Swift – a fake picture that appeared on the platform X had more than 47 million views before it was blocked.[27] Deep fakes are also used for more mundane purposes, including commercial interests. Anne Lagercrantz reports, for example, that SVT saw one of their most popular news anchors appearing in a fake newscast promoting internet casinos.

Jean-Marc Rickli, Head of Global and Emerging Risks at the Geneva Centre for Security Policy, is sceptical about journalists' capacities to keep up with the technological developments: "Journalists whose job is to provide verified information will increasingly be faced with a race that is completely asymmetrical with machines that are able to produce content that looks really legit but is actually wrong, and we have already been seeing it during the Gaza war." (Read the Q&A with Jean-Marc Rickli, page 45)

This could have effects far beyond misleading people on key issues. Many interviewees voiced grave concerns about a further decline of trust in information, institutions, and even in one another. Ezra Eeman of NPO says: "My biggest concern is that it will decrease the trust in information systems even more. The feeling that you cannot believe your eyes any more will also reflect on trusted brands." SVT's Anne Lagercrantz observes: "The mere presence of deep fakes will make people doubt everything. After the gruelling Hamas attacks people dismissed images as fake but they were authentic. We might develop a feeling that we cannot trust anything." Cambridge Professor Gina Neff shares a similar view: "I worry about bad information flooding out the good. […] But my meta concern is: I worry about chipping away at a social contract of trust, of how people trust one another. Experiments have shown that trust in humans declines when there are AI agents involved. We are in the midst of a big social experiment at a time when trust in Western societies is declining. The problem is more fundamental than the question 'Do we have enough fact checkers and journalists?' It is about how people's expectations of trust will change in the next five years."

Nic Newman, Senior Research Associate at the Reuters Institute, authored a report which revealed that more than two-thirds of 300 media leaders globally said AI will most likely decrease trust. He thinks that media plays a role in this, since coverage of deep fakes and misinformation influences people's views on trust. Newman: "They might develop the perception that you cannot trust anything. But conversely a flood of unreliable synthetic content may make people want to seek out someone they can trust. Trust in some brands might increase. […] Much will depend on what the platforms will be doing. It is actually in their interest to promote trustworthy content and keep their platforms as clean as possible during this shift." (Read the Q&A with Nic Newman, page 38)

It's worth remembering that the views and worries of media executives, keynote speakers, and thought leaders are not always borne out by reality. While the concerns raised about AI-generated content and deepfakes are valid, it is an oversimplification to suggest that journalists can't keep up with technological advancements. For one, the existence of more false information does not necessarily imply that more people will get to see it.[28] Increases in the quality of false information are concerning, but according to fact-checkers, most false content is still produced more crudely or uses materials

[26] Rowan Philip, "How to Identify and Investigate Audio Deep Fakes, a Major 2024 Election Threat", Global Investigative Journalism Network, 26 February 2024. https://gijn.org/resource/tipsheet-investigating-ai-audio-deepfakes/, retrieved 1 March 2024.

[27] David Mouriquand, "Pornographic deepfakes of Taylor Swift spark calls for new AI legislation", Euronews, 29th January 2024. https://www.euronews.com/culture/2024/01/29/pornographic-deepfakes-of-taylor-swift-spark-calls-for-new-ai-legislation, retrieved on 26 February 2024.

[28] Simon, F. M., Altay, S., & Mercier, H. (2023). Misinformation reloaded? Fears about the impact of generative AI on misinformation are overblown. Harvard Kennedy School Misinformation Review, 4(5). https://doi.org/10.37016/mr-2020-127

out of context (even though the proportion of AI-generated content is slowly increasing).[29] Responsible news organizations have rigorous fact-checking processes and editorial standards in place to verify information from credible sources. These work just as well for AI-generated content, and such news organizations continue to play a vital role in educating the public on media literacy and identifying misinformation. Additionally, the ethical and transparent use of AI technology can actually aid journalists in data analysis, research, and pattern recognition, complementing efforts in addressing the problem, even though these methods have gaps and shortcomings.

It is also important to not underestimate the audience's ability to identify misinformation, particularly when they are used to it. Erik Roose, Chairman of the Board of Estonian Public Broadcast, says that for people in his country who were raised in the Soviet era misinformation is the norm. "We consumed misinformation on a daily basis in the 1970s and 1980s. Our newspapers were 100% misinformation, so you had to read between the lines."

Remaining visible: The old and new dependence on platform monopolies

Generative AI will likely make all industries even more dependent on the large, often oligopolist technology companies, as well as their infrastructure, products, and services. This handful of players shape the most popular AI systems faster than even the most well-resourced news organizations can develop their own. Only the biggest publishers will be able to develop their own foundation models and fundamentally rebuild their content management systems from scratch. And even they will most likely rely on the ubiquitous software packages that have and will continue to shape how people collaborate and communicate. As Felix Simon wrote in a report for Columbia University's Tow Center: "The complexity of AI increases platform companies' control over news organizations, creating lock-in effects that risk keeping news organizations tethered to technology companies. This limits news organization's autonomy and renders them vulnerable to price hikes or shifting priorities of technology companies that may not align with their own. [...] Platform companies increasingly control both the means of production and connection in the news."[30]

Cambridge University sociologist Gina Neff says: "I'm not a techno pessimist. But we do need capacities to govern this. The trend is that large companies are doubling down on their ability to capture the market. They will keep out as many competitors as they can. Data, tools, cash, they are building the infrastructure. 18 months into the generative AI hype, they are taking our eyes off cloud infrastructure, energy, data centres. They are growing to build a new market just to build a new application. We are globally depending on a handful of companies. How many companies will have the cash to build foundation models?"

The power of tech platforms and winner-take-all dynamics have already shaped much of the current news ecosystem, as Rasmus Nielsen and Sarah Anne Ganter outlined in their book The Power of Platforms.[31] Most media companies strongly depend on the algorithms of search and social media platforms which they cannot control, particularly for their connections with young audiences. This has given English language brands an advantage internationally. Local media have been the most fragile in this system because they – by their very nature – have to rely on a geographically limited market. This has increased the responsibility of public service media to serve people in communities that would otherwise remain without access to independently researched and curated information.

The proliferation of generative AI systems will add another layer to digital platforms which could make media brands increasingly invisible to their customers. As the London School of

[29] https://reutersinstitute.politics.ox.ac.uk/news/bbc-expert-debunking-israel-hamas-war-visuals-volume-misinformation-twitter-was-beyond

[30] Felix Simon, "Artificial Intelligence in the News: How AI Retools, Rationalizes, and Reshapes Journalism and the Public Arena", Tow Center for Digital Journalism, 2024, page 4.

[31] Rasmus Nielsen, Sarah Anne Ganter, "The Power of Platforms – Shaping Media and Society", Oxford University Press, 2022.

Economics' media professor Charlie Beckett said in a 2023 interview, "AI that is based on language models will answer all questions without people ever encountering the source of the information. This is a massive problem for business models. What kind of regulation will be needed, what commercial agreements, what about copyright? Frankly, I've never seen industry executives so worried before."[32] Beckett is also the founding director of the Journalism AI project, which is funded by Google, and has contributed to building probably the largest global network of AI-conscious journalists.

Recent developments suggest that such fears may prove correct. Google recently announced that it would roll out AI-generated overviews – summarized answers to user questions made possible through a combination of Google search systems and its in-house suite of 'Gemini' AI systems – to everyone in the United States, with more countries following in the future. It's something many industry observers have expected for months and will likely usher in a new era of how people search for information online, given Google's monopoly position in online search.[33] For now, it is impossible to tell which searches will be impacted, with early data suggesting low numbers – but this may very well change. Some industry experts fear that up to 25% of traffic to publishers could be negatively affected by the de-emphasis on website links[34] with knock-on effects on ad revenue and subscriber numbers. This comes on the back of already declining referral traffic from platforms.

While Google claims that the links included in these AI overviews get more clicks than if the page had appeared as a traditional web listing for that query, this claim has not been independently verified. Apart from the question who will be linked to, including news publishers or also other sources, time will tell whether AI-generated overviews ultimately direct the same numbers to news publishers' websites – or if traffic will materially decline. As Felix Simon writes, the question for publishers will be how they can ensure that people will still find and consume their content in an environment where major platforms such as Google cater to their informational needs in an ever more targeted manner.[35] Or, as technology journalist Reed Albergotti writing for Semafor put it more bluntly: "Publishers [...] need to make tough decisions. How can they build businesses that they alone control?" But it would be fearmongering to solely emphasize the power of the few. Just as social media and platforms like Substack or Steady have helped community media and individual authors build relationships with their audience, generative AI tools will likely make it easier to start news outlets without initial investments in specialized talent. Agnes Stenbom of Schibsted prefers to paint a nuanced picture: "Obviously there are benefits of scale. We can see a clear concentration of power and profit. Then again generative AI can be really empowering to individuals, too. But we have all the unanswered questions about ethics and using these tools. We have to figure this out before we can fully use the toolbox."

Several media managers emphasize the importance of engaging with tech platforms early on in conversations about their needs. They seem to be more confident this time around as AI companies depend on quality content to train their foundation models. Verena Krawarik of APA says: "The media has to enter into a much stronger dialogue with big tech and regulators and place demands, like 'I'm not happy with the way Google is trying to present results here, if sources don't appear, then we have a problem.' We can only solve this together. I am convinced that foundation models without journalism content are absolutely worthless." Tech companies might also need to partner with media organizations to test tools they are developing for unwanted side-effects. And

[32] Alexandra Borchardt, interview with Charlie Beckett, "Google may no longer be the great power we thought it was", Medieninsider, 14 September 2023. https://medieninsider.com/ai-expert-charlie-beckett-interview-english/18558/, retrieved on 26 February 2024.

[33] Reid, L. (2024, May 14). Generative AI in Search: Let Google do the searching for you. Google Blog. Retrieved from https://blog.google/products/search/generative-ai-google-search-may-2024/

[34] Liedtke, M. (2024, May 15). Google unleashes AI in search, raising hopes for better results and fears about less web traffic. AP News. Retrieved from https://apnews.com/article/google-search-ai-overviews-internet-traffic-ebb6bbbde17ed29a5f7b630d9e5e285b

[35] Simon, F. M. (2024, May 16). AI search & some implications for publishers and the public sphere. Medium. Retrieved from https://medium.com/@felixsimon/ai-search-implications-for-publishers-and-the-public-sphere-db931b6c8274

Figure 5: Publishers are not optimistic about funding from big AI companies[36]

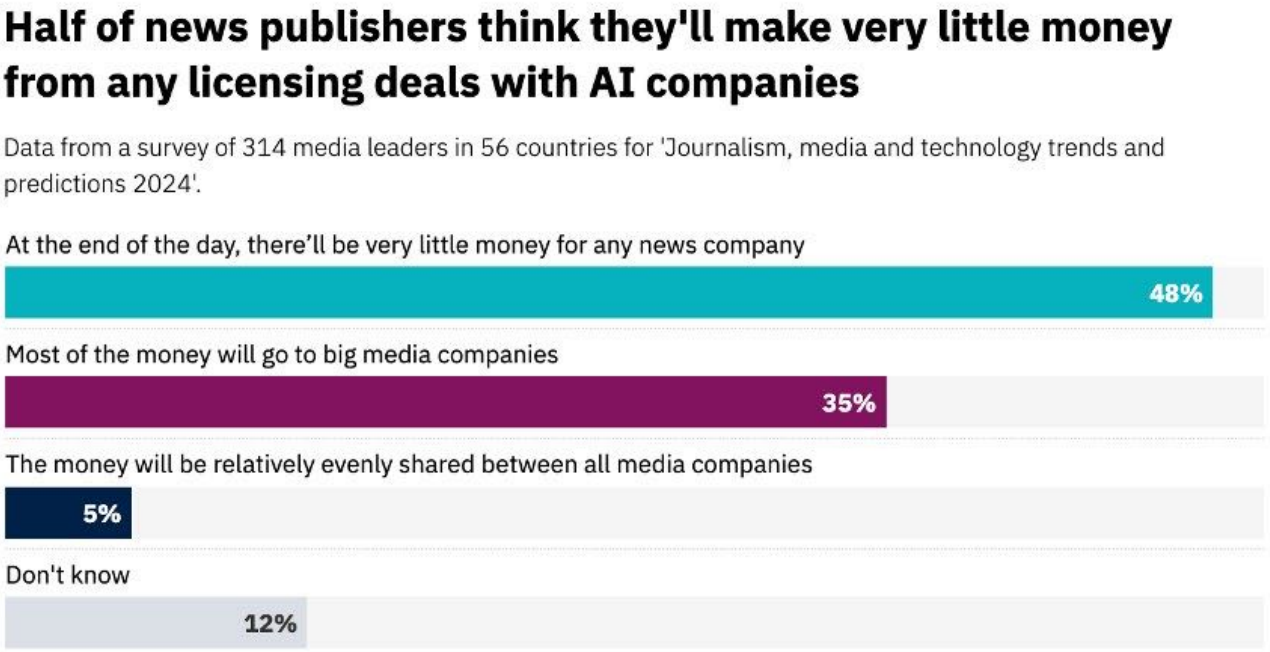

Q16. When it comes to negotiating licensing deals with AI companies around news content IP (intellectual property), which of the following do you think is the most likely scenario? Base: 289

Source: Reuters Institute for the Study of Journalism

they'd probably prefer to avoid run-ins with the competition authorities in Europe and the US.

Dmitry Shishkin, CEO of Ringier International, says: "You need to be in a conversation with the tech companies. Monetary compensation should be much stronger as it has been with social media and search. I'm a little bit more optimistic here. Of course, there will be some small organizations that will never have the power to negotiate. Governments need to do something to protect local and regional news." The sceptical part of this assessment is echoed by findings from the Reuters Institute's media leaders survey. In December 2023, only a minority believed there would be sizeable compensation for the industry. (see Figure 5).

Labour markets and the digital divide

How exactly generative AI will affect the job market, particularly in the already battered media industry, is largely unpredictable. Professor Carl Benedikt Frey of Oxford's Martin School, co-author of the often-quoted 2013 study on automation's effects on labour markets, predicts a significant shift:[37] "I think it's going to be the first AI tool that really has a widespread impact on the labour market." Those new to a job and those with below average talent will most likely profit most, he says, because the algorithms have been trained on vast amounts of data and are therefore geared towards the average. "And so, if you're below average, you're likely to benefit more. On the other hand, that doesn't mean that everybody will benefit, because there will be more competition at the entry level." (Read the Q&A with Carl Benedikt Frey, page 41)

Scott Galloway, business professor at New York University, is convinced that significant job cuts are already being implemented as a result of AI, even though CEOs might shy away from saying so. While profits at major US companies have been increasing, headcount has been reduced, he wrote in a blog post, listing examples.[38] Also, companies will be a lot more

[36] Newman, N. (2024). Journalism, Media, and Technology Trends and Predictions 2024 (Reuters Institute Report). Reuters Institute for the Study of Journalism. https://reutersinstitute.politics.ox.ac.uk/journalism-media-and-technology-trends-and-predictions-2024#header--3

[37] Carl Benedik Frey, Michael A. Osborne, "The Future of Employment: How susceptible are jobs to computerization?", Oxford Martin School, 17th September 2013 https://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf

[38] Scott Galloway, "Corporate Ozempic", 23rd February 2024. https://www.profgalloway.com/corporate-ozempic/, retrieved on 1st March 2024.

hesitant to hire new staff, since many roles won't be needed any longer, according to Galloway. On the upside, this could help many start-ups survive longer, he predicts. Galloway points to a Goldman Sachs study of 2023 calculating that one quarter of today's jobs could eventually be automated. According to the report, this applied to 26% of professionals in the media, arts, and entertainment industry, with one third of all management roles across all industries at risk.[39]

For journalism and the media industry, there are two factors that make predictions about the development of labour markets difficult. Some observers expect further drastic job losses, since smaller outlets and those who perceive less risk to their reputation will most likely incorporate generative AI driven tools in much of their production. Some have already gone into offense mode and confidently include AI reporter avatars in their staff lists. For example, the German daily *Kölner Stadtanzeiger* proudly presented an artificial columnist named Klara Indernach as "our new AI colleague".[40] Unions hold hopes that the technology will free reporters and editors from mundane and repetitive tasks, enabling them to do more on the ground reporting and in-depth research. But it is unclear how publishers and broadcasters will implement this, given the pressures on revenue and cost.

The Brussels-based European Federation of Journalists (EFJ), which represents journalism associations and unions, had no statistics on AI-related job cuts as of early 2024. But its director, Renate Schroeder, said that the launch of ChatGPT had exacerbated worries about job losses. For example, in June 2023 the German publisher Axel Springer announced plans to cut an estimated 200 jobs at its *Bild* tabloid, Germany's biggest newspaper, in a move to go entirely digital.[41] Axel Springer's CEO Mathias Döpfner had intimated this shortly after the launch of ChatGPT, when he was quoted saying: "Artificial intelligence has the potential to make independent journalism better than it ever was – or simply replace it."[42]

Prospects like these are rather dire. On the other hand, automated content production and AI-supported journalism might be the only way some smaller news organizations can produce sufficient output to secure their survival. "A 10% gain in efficiency can be an existential question for them," as LSE Professor Charlie Beckett notes.[43] But it is not only about cost. In recent years, local newsrooms in particular have been struggling to find the talent they need.[44] In a commentary for Digital Journalism, Alexandra Borchardt pointed to the talent crisis: "While local and regional publishers struggle to create unique products their audiences are willing to pay for, they increasingly face the reluctance of those who have to do the creating. Where higher demands are paired with diminishing returns, particularly those tend to quit who are needed the most: talent that is able to do both, digital and journalism. (...) But if local journalism is to survive to stabilize democracies from the bottom up, affordable, accessible, and easy to use AI-solutions for local news publishers are needed – and fast."[45]

Renate Schroeder of the EFJ echoes these sentiments: "Specifically local journalism can take so much advantage of AI, if it is being used in a responsible way. It could safeguard local

[39] Joseph Briggs, Davesh Kodnani, "The Potentially Large Effects of Artificial Intelligence on Economic Growth", Goldman Sachs Economics Research, 26th March 2023. https://www.key4biz.it/wp-content/uploads/2023/03/Global-Economics-Analyst_-The-Potentially-Large-Effects-of-Artificial-Intelligence-on-Economic-Growth-Briggs_Kodnani.pdf

[40] Wan-Ifra, "AI and robot writer 'Klara' key to Kölner Stadtanzeiger Medien's tech future as it switches of its presses", 10th November 2023. https://wan-ifra.org/2023/11/ai-and-robot-writer-klara-key-todumonts-kolner-stadt-anzeiger-mediens-tech-future-as-it-switches-off-its-presses/, retrieved on 27th February 2024.

[41] Anna Coban, "Germany's biggest newspaper is cutting 20% of jobs as it prepares for an AI-powered future", CNN, 21st June 2023. https://edition.cnn.com/2023/06/21/media/bild-germany-newspaper-ai-layoffs/index.html, retrieved on 27th February 2024

[42] Deutsche Welle, "Axel Springer to cut jobs, warns AI could replace journalism", 28th February 2023. https://www.dw.com/en/axel-springer-to-cut-jobs-warns-ai-could-replace-journalism/a-64846886, retrieved 27th February 2024.

[43] Alexandra Borchardt (2023), interview with Charlie Beckett for Medieninsider, see above.

[44] Read about the talent crisis in A. Borchardt, S. Kieslich, J. Lück, T. Schultz, F. Simon: "Are Journalists Today's Coal Miners? The Struggle for Talent and Diversity in Modern Newsrooms – A Study on Journalists in Germany, the United Kingdom, and Sweden", Reuters Institute for the Study of Journalism, Johannes Gutenberg University Mainz, July 2019. https://reutersinstitute.politics.ox.ac.uk/our-research/are-journalists-todays-coal-miners-struggle-talent-and-diversity-modern-newsrooms, retrieved on 29 April 2024.

[45] Alexandra Borchardt (2022) Go, Robots, Go! the Value and Challenges of Artificial Intelligence for Local Journalism, Digital Journalism, 10:10, 1919-1924, DOI: 10.1080/21670811.2022.2149584

journalism, where we see plenty of news deserts already. But it can also be misused. We will see more municipalities using AI tools for content production themselves, eroding independent journalism even more."

However, job issues cannot only be discussed as statistics. Some media organizations, including *The New York Times*, have more employees today than before they started to innovate digitally, including an array of new roles, with others discarded. In the book covering his tenure as editor-in-chief at the Washington Post, Martin Baron describes how he had to ceaselessly defend editors' jobs from owner Jeff Bezos' attempts to eliminate everything that wasn't original reporting.[46] And it can be assumed that only those who are willing to embrace working with AI will have a future in media organizations, given editors and media managers are increasingly saying as much publicly. Reuter's Jane Barrett quotes their CEO when describing how the news agency communicates with staff who fear losing their jobs: "Generative AI won't take your job, but someone who knows how to use it will."

Anne Lagercrantz of SVT is worried about the disappointment and friction this will inevitably create in the newsroom, just as digital transformation has done: "This will most likely devalue work skills that have been highly valued for decades. I grew up in a TV industry, now TV skills are not worth as much any longer, that has been extremely painful. And this will happen again." She also worries about an increasing technology literacy gap, in society as well as the newsroom. "We need to talk more about the digital divides and how we can help our co-workers understand the technology and the changes we are entering."

It would be a mistake to assume that young people, so-called 'digital natives', naturally grasp everything about new technologies and the corresponding opportunities and risks. Like many others, Matt Frehner of *The Globe and Mail* doesn't see this as a simple generation gap. "This is different for everyone. Some are more tech savvy; some are less attuned to this. Lots of senior reporters are very interested in how the stuff works. We have a writer who is 70 and writes a lot about AI, he is very sophisticated."

Costs to the environment, society, and human well-being

The environmental cost of large-scale computing is immense, and it is not talked about nearly enough. One early figure suggested that a search answer generated by an LLM uses an estimated five times the energy of a 'regular' Google search.[47] While this figure might reduce, the number is still staggering. Data centres have been calculated to contribute about 3% to global carbon emissions – more than the airline sector – and the figure is predicted to grow. Further, the server parks that keep AI systems running use tremendous amounts of fresh water, one of the resources forecasted to dwindle in the context of global warming.[48] Microsoft admitted that because of AI and the associated expansion of its data centres its carbon emissions rose by 30 percent over the last four years.[49]

Needless to say, the costs need to be weighed against the possible benefits. Otherwise, to put it provocatively, the most effective strategy against global warming would be to get rid of humanity altogether. In addition, many claim that the battle against climate change could profit considerably from AI-based solutions, with Microsoft unsurprisingly arguing in its favour.[50]

[46] Martin Baron, "Collision of Power: Trump, Bezos, and the Washington Post," Flatiron Books, October 2023.

[47] This figure is quoted frequently, for example in Renee Cho, "AI's growing Carbon Footprint," Columbia Climate School, 9 June 2023. https://news.climate.columbia.edu/2023/06/09/ais-growing-carbon-footprint/, retrieved on 7 April 2024.

[48] Read, for example, Karen Hao, "AI is Taking Water from the Desert – New data centers are springing up every week. Can the Earth sustain them?", The Atlantic, 1 March 2024. https://www.theatlantic.com/technology/archive/2024/03/ai-water-climate-microsoft/677602/, retrieved on [29] April 2024.

[49] Akshat Rati, Dina Bass, "Microsoft's AI Push Imperils Climate Goal as Carbon Emissions Jump 30%", Bloomberg, 15 May 2024, https://www.bloomberg.com/news/articles/2024-05-15/microsoft-s-ai-investment-imperils-climate-goal-as-emissions-jump-30?cmpid=BBD051624_GREENDAILY&utm_medium=email&utm_source=newsletter&utm_term=240516&utm_campaign=greendaily

[50] Brad Smith, Melanie Nakagawa, "Accelerating Sustainability with AI: A Playbook," Blog published on 16 November 2023. https://blogs.microsoft.com/on-the-issues/2023/11/16/accelerating-sustainability-ai-playbook/. See also: Victoria Masterson, "9 ways AI is helping tackle climate change," World Economic Forum, 12 February 2024. https://www.weforum.org/agenda/2024/02/ai-combat-climate-change/, both sources retrieved on 29 April 2024.

However, costs and benefits are rarely debated in the current context of generative AI, given all the hopes and expectations it bears and the commercial opportunities it presents in the race for new solutions.

This report cannot provide an environmental impact statement on generative AI. But it is puzzling that so many media leaders who have emphasized the need for improving climate journalism and sustainability in recent years seem to have mostly forgotten about it when it comes to AI. While in the research for the 2023 EBU News Report '*Climate Journalism That Works: Between Knowledge and Impact*', many stressed the importance of ramping up climate coverage and sustainability programmes, hardly anyone in the media even mentions sustainability when discussing AI-driven solutions.[51] At industry conferences for journalists and media executives, AI and sustainability run on separate tracks.

A 2024 report issued by the Associated Press came to a similar conclusion: none of the media leaders in the survey who were asked about potential ethical concerns mentioned the environmental and ecological footprint of AI or the danger of reinforcing social injustice. Instead, it reads, there was a focus on daily journalistic practice. "Considering that journalism does have an important role in informing the public discourse, there is a need to explore whether this lack of concern for broader ethical issues is the result of a mental disconnect, lack of awareness or the way the questions were framed," the authors write.[52]

In our interviews, academic experts seemed to be much more aware of the back-end of artificial intelligence and the costs and harms involved. The climate impact is highlighted, for example, by Cambridge Professor Gina Neff when asked to consider what is missing from the AI conversation: "First the energy and water costs it takes for large scale computing. Some hope these costs will come down. But we cannot afford the current kind of growth from the climate point of view. There are environmental costs we are missing." Melanie Mitchell of the Santa Fe Institute believes current approaches to generative AI need a much higher level of sophistication. "Only relying on what we have now and throwing more data on it will not do it. We need innovation in training methods. Using these huge amounts of training data is just not sustainable. We have to develop systems that don't use that much computer power. Now it is a brute force approach. We need to get something in a more energy efficient way."

Sarah Spiekermann, Professor of Information Systems at the Business University of Vienna, says: "IT-companies often claim they are CO_2 neutral, because cloud centres are placed next to hydrogen power plants. But the ecological balance is not positive. You have to start with the resources, the minerals. For one ton of rare earth, you produce 75,000 litres of acid water, this is a lot of waste. The ecological balance sheet must also include shipping, manufacturing. The total global ecological balance sheet of IT starts there, not to speak of CO_2 emissions. [...] Before shifting everything to generative AI, a balance sheet needs to be developed: how much investment of digital transformation is scalable?" (Read the Q&A with Sarah Spiekermann, page ..)

Political awareness of the intersection of AI and climate is limited. In February 2024, several US Senators introduced a bill for the *Artificial Intelligence Environmental Impacts Act of 2024*. While AI could be used to help mitigate environmental harms and climate change, it takes a significant toll on resources and the efforts to curb greenhouse gas emissions, they argued. "For example, increasing AI use could contribute to data centre electricity demand doubling by 2026, leading to more carbon emissions. Demand for water to cool data centres is already creating concerns about water supplies, and the chips needed to run AI software are contributing to a growing

[51] Alexandra Borchardt, Katherine Dunn, Felix Simon, "Climate Journalism That Works: Between Knowledge and Impact", European Broadcasting Union, March 2023.

[52] Nicholas Diakopoulos, Hannes Cools, Nicholas Diakopoulos, Hannes Cools, Charlotte Li, Natali Helberger, Ernest Kung, Aimee Rinehart, "Generative AI in Journalism: The Evolution of Newswork and Ethics in a Generative Information System," 9 April 2024. https://www.aim4dem.nl/out-now-generative-ai-in-journalism-the-evolution-of-newswork-and-ethics-in-a-generative-information-ecosystem/ retrieved on 9th May 2024.

mountain of electronic waste," the bill reads.[53] While the chances that this will become law are slim under current circumstances, this could at least spur a public debate.

Tamara Kneese, Project Director at the Data & Society Research Institute, believes that the entire computing sector should be subjected to ongoing environmental impact assessments. But she asserts that this is not even part of the current AI conversation. While tech companies in the past had emphasized their commitment to becoming 'net zero' at certain points in time, almost all top management seems to have forgotten about this in the context of generative AI, Kneese observes. "It feels like the fascination with especially generative AI and powerful AI systems and the massive amount of funding going into that right now is displacing the conversation that was happening around net zero and decarbonization as goals within the computing industry."

Kneese suggests that tech companies should be much more mindful about the location of large data centres and foster community participation. Companies could also put more effort into considering energy use when training large models, for example, "the time of day that they're training a particular model that requires a lot of computing and the availability of renewable energy on the grid." Further, all actors should be more conscious about the deployment of generative AI. "There are a lot of fantasies being projected into this. The weighing of costs and benefits doesn't seem to happen anymore. But a lot of use cases are not going to work. There need to be conversations about what will generate real value and what is just a silly waste of energy."

Kneese has studied not only the environmental impact of AI but also the human toll its rapid development takes. The training of LLMs still needs human supervision. Like content moderation on social media platforms, this is often done by workers in parts of the world where wages are low and labour laws are lax. Kneese's findings are summarised in the report '*Climate Justice & Labor Rights*'.[54]

In their widely quoted paper "*On the dangers of Stochastic Parrots: Can Language Models Be Too Big?*" a group of researchers around Emily Bender argues that disadvantaged communities will bear the brunt of the cost of large language models while benefitting the least.[55] They ask the tech industry to take research in new directions to develop smaller, less harmful models.

In search of a holistic approach

This report aims for a balance between examining the opportunities and highlighting the risks generative AI poses for journalism and the media. There is a tendency in large organizations to carve issues into pieces and delegate the responsibility to separate departments. But with large-scale developments that permeate everything – from business models to the impact on the natural environment and social justice – this doesn't work. A holistic, big-picture approach is needed.

Chapter 3 will shed light on how managers are approaching the opportunities and dealing with their biggest concerns.

Chapter 4 will shift to the bigger picture of ethical questions. However, to understand the challenges we must understand the opportunities – and vice versa. As such, we will dive into highly practical use cases which media organizations are exploring today, in Chapter 2.

[53] "Markey, Heinrich, Eshod, Beyer introduce legislation to investigate, measure environmental impact of Artificial Intelligence," 1 February 2024 https://www.markey.senate.gov/news/press-releases/markey-heinrich-eshoo-beyer-introduce-legislation-to-investigate-measure-environmental-impacts-of-artificial-intelligence, retrieved on 27 February 2024

[54] Tamara Kneese, Climate Justice & Labour Rights (2 August 2023). Available at SSRN: https://ssrn.com/abstract=4533853, retrieved on 27 February 2024.

[55] Emily Bender, Timnit Gebru, Angelina McMilllan-Major, Shmargaret Shmitchell, "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?", conference paper presented in March 2021. https://dl.acm.org/doi/pdf/10.1145/3442188.3445922, retrieved on 15 May 2024.

Q&As

Q&A

> "We will probably still be telling stories for thousands of years to come."

DAVID CASWELL

AI Consultant

How is generative AI a game-changer for journalism?

The most dramatic change is that generative AI separates the information content in news from the artifacts that contain that information, like articles or videos. In other words, it separates the semantic information of news from the way it is communicated. Dealing with these semantic characteristics separately from the artifact is a fundamentally new thing. It enables newsgathering at vast scale. It enables personalized production of content. Consumers can consume information in the form they choose and are no longer restricted to the way the producer wanted them to consume it. All this will radically reconfigure our information ecosystem. Generative AI also opens a set of possibilities that the news industry has been exploring for a while but that haven't worked that well yet, like voice agents, voice bots, automated newsgathering and automated journalism. These innovation efforts now become much more possible with generative AI.

Do you think journalism will develop from being a push activity, where news is directed to the audience by the media, to a pull activity, where people choose customized news to fit their needs?

I have my doubts. A lot of people think it will. But I don't know if there is evidence yet that most audiences will pull news. For routine news consumption the cognitive effort that is required to pull news will be a barrier. The alternative scenario is that we might develop relationships with bots, similar to what we have already seen with AI avatars as boyfriends and girlfriends, maybe as our own personal news presenters, but at this point I have not seen the evidence.

Experience shows that people don't necessarily want a very personalized experience, but they want to choose between and compare alternative news sources and compare.

The universality requirement may be a huge opportunity for public service news. Universality is almost a myth in the digital world, where everyone can usually get exactly what they want. People do have very different consumption behaviours, however, and so being able to publish common stories to different audiences in very different formats and media might enable digital news to become universally accessible, or at least more accessible. There is already a big set of early opportunities to do this using AI, for publishers but also for audiences. If you look at the news

avoidance crisis, it's clear that a lot of newsrooms are producing something that audiences don't consume.

Most publishers have in their guidelines that there needs to be a human in the loop. We assume you agree, as someone who worked for the BBC.

Editing of the output is necessary for large publishers, and I think it is likely to remain so for the foreseeable future. But that is not going to be the case for lots of other players in the information ecosystem. For example, Rappler, the digital news brand in the Philippines founded by Nobel Laureate Maria Ressa, is publishing article summaries directly from generative AI. They label those very, very clearly as AI generated and that they may contain errors. But I think that for most public service news organizsations who choose to use these tools, and for premium publishers like the *New York Times* and *The Guardian*, a human in the loop is going to will continue to be required.

You published a widely recognized report 'AI and Journalism: What's Next?', which describes different scenarios for journalism. Can you outline them briefly, and which scenario do you think is most likely?

I wouldn't call those scenarios, and most of that stuff is actually already happening. The article breaks down the options available to newsrooms from AI into specifics: what are the strategies? How can you approach this in a systematic way? What are the specific projects you could do, what are the risk structures? Most people in news begin by looking at efficiency-enhancing strategies – tagging and SEO are examples. Then there are opportunities around producing products that serve your audience but that you wouldn't do if you didn't have AI. Many of these options are relatively obvious and even relatively easy to implement, but some are probably temporary because the entire ecosystem of news is likely to change.

You have talked to media leaders all over the globe about AI. Do you see significant geographic differences?

I don't see any geographic differences. I see differences by the size of organizations. Small newsrooms are disproportionally empowered by this. The functionality is so accessible, and so the only barrier for small organizations is imagination and a couple of hundred dollars in fees for tools. Small newsrooms will be able to do amazing things very quickly. We have already seen examples of that in the 12 small teams from all over the world who participated in the Open Society Foundation's AI in Journalism Challenge in 2023. Large organizations seem much more encumbered by bureaucracy, legacy assumptions, and risk avoidance.

What will be the biggest challenges in managing AI in news organizations?

Even in that question there is an assumption that there is a necessary management structure in an organization. But these tools essentially give independence to very small teams and individuals, potentially making them spectacularly productive and much less dependent on other parts of the organization. They don't need the graphics people. They don't need the data journalism people. They don't need the video production people. They don't need skilled writers. So, you will likely see very

productive little islands throughout the organization. If you play that out, you end up with an organization that is much more flattened, decentralized, federated. The organization provides the brand, the process, technology, certification, the capital, but the actual production is done by small and relatively independent AI-empowered teams or even individuals.

Will this improve the journalism?

In this kind of scenario, the alignment of these smaller teams is that they serve particular audience segments, even very small audience segments. Take the BBC, for example. Maybe 10 years from now, you could see two or three people that take the output of the BBC and use these tools to create news for, let's say, young people in rural areas in parts of Wales or Scotland. And they know that audience viscerally, the needs and concerns and values and situations of that audience, and so they can act as a sort of translator between the information environment and the audience. In a world in which anyone can produce any kind of media almost instantly, knowing what to produce for a specific audience might become extremely valuable.

Is this bad news for large organizations?

It could be bad news for the whole middle management layer. The introduction of AI into media organizations is likely to diminish power for those who move information around the organization and to enhance the power of those who originate it. There might be quite a battle for control as the originators of content and information become more autonomous and less dependent on managers or on other parts of the organization.

How would you define ethical AI?

I would define it as simply AI that makes society better. Societies will be going through a very rapid and difficult transition because of AI and a news organization's job will be to help society navigate that transition. That is a huge responsibility, and it means more than just reporting about it; it means helping people to be part of it. An example is in the BBC educational programming for schoolchildren, which has opportunities around developing a curriculum about large language models, about learning how to learn with them, about the best practices for self-education in an AI environment. Another example is about helping underserved populations to appreciate what they can do with the tools. One of the most impressive use cases is improving writing skills of people, allowing them to become more articulate in writing, and enabling them to interpret writing in inaccessible writing styles. There are obviously many risks from AI in media, but realizing the opportunities are part of an ethical response too.

Many news organizations have developed ethical guidelines for AI. Do you see any that are exceptionally good?

I am pretty impressed with the speed that the industry has developed guidelines. There is a lot of cross-pollination and what seems to be a fairly solid set of ethical best practices has emerged. In contrast, there is this large part of the media industry that is not driven by ethical frameworks, especially among newsrooms that are very

focused on clicks and programmatic ads. Some of these have done very aggressive things in a very loose way and have done substantial damage to their brands as a result.

What is missing from the conversation?

People are using these tools with the assumption that existing workflows and the competitive environment will stay the same. And that is not going to happen. The changes that are made possible by these tools, these are too fundamental for the existing information ecosystem to remain as it is. What will all this look like five years from now? When apps and browsers and even operating systems enable users to have different consumption experiences? What will be the unit of journalism, if articles are broken apart by generative search and conversational interfaces? What should journalists produce in that environment? These are the big questions that I think are inevitable.

What's your view on these questions?

I did a scenario planning process recently with a large number of people to try to understand what the possible visions for the long-term impact of AI on news might be. We identified a range of fairly dramatic ways that all this could play out, and so I've been thinking a lot about what a long-term response to those scenarios might look like. If I had to guess, I would point to some kind of structured narrative – a replacement for the article or video in news. I think a computational version of a news narrative could become the core of journalism in the future. But this is just an educated guess at this point.

What are the things that won't change?

The fundamental information needs of humans won't change. I think the cognitive biology of human beings is biased towards narrative. We've been telling each other stories for tens of thousands of years, and we'll probably still be telling stories for thousands of years to come, regardless of the tools we use to do that.

Q&A

> The media industry gave away the keys to the kingdom once – that shouldn't happen again"

LUCY KÜNG

Independent Media Strategist and Advisor

Is generative AI a gamechanger for journalism and the media?

Yes, but not just for journalism and the media. It is such a profound shift; the media industry will be caught up in the changes. The media's reaction to generative AI now bears the scars from the painful transition to digital. It took the industry a long time to understand what was shifting, what that meant, and then which changes to make. And it has emerged from that transition in a weaker position. There has been an explosion of innovation, and profound changes in business models, but overall, the sector has emerged structurally challenged in a way it wasn't before the advent of digital technologies.

In your recently published book *Strategic Management in the Media* you wrote "the fruits of two decades of painful disruption are being themselves disrupted." Your confidence in the media's ability to change seems to be limited.

I think the media has exhibited an extraordinary ability to change – it has transformed itself radically. I don't doubt the sector's ability to change – it's more that it is extremely unclear at this point what the intelligent response is. We want to act, to equip ourselves to deal with what's coming, what that looks like is less obvious. If we layer on to that deep concerns about AI's impact, plus an AI fatigue arising from 18 months of worrying then we get to a place of 'threat rigidity'. When we are stressed and concerned it's hard to be innovative, to be objective. I see it at conferences, with consulting assignments. People want predictions, they are anxious and want their anxiety reduced. But this is very difficult because we are at such an early stage in such a big transformation. We are at the dial-up modem stage of the internet. Things will change fast.

Do you think the industry is better prepared this time?

Mentally yes. Definitely. But the scope and profundity of these changes mean responses are hard to design. We can't wait for something to happen, for things to become clear, but this is a new foundational general-purpose technology. It's a platform shift, a new layer that products and businesses are built on. The last shift for the media emerged from the coming together of mobile devices, social media,

and the cloud. Generative AI is an amazing tool. It will get cheaper, easier to use, and applied more and more broadly, and other tech advances will get bundled with it. This is all to come.

How can organizations prepare for these big unknowns?

We know that, as Mustafa Suleyman put it, "generative AI is 'climbing the ladder of cognitive abilities' fast. That makes it a 'discontinuous tech shift' – it can sweep away the value of some existing competencies and knowledge." So, some labour will be displaced. But new tasks will emerge – particularly around bridging and blending synthetic and human input.

At one step removed, however, it is also what the late Harvard professor Clayton Christensen called a 'disruptive technology' – it allows new players to enter the market with new products produced in different ways and different value propositions. To respond to this threat media organizations must become more agile, and open to rethinking what they offer and how they produce it.

Are you worried or excited about this development?

From the perspective of a researcher, it is fascinating. I can't wait until I have a generative AI personal assistant … But I am concerned. For the media business, it's a very uncertain time. Competition is already fierce, and this brings in new actors which will increase competition further. A flood of synthetic content can worsen the decline in trust that quality media are already struggling with. And for commercial media, more products mean less opportunity to raise prices and ARPU (average revenue per user).

So, you think, media organizations should slow down a bit?

No, the issue is not to panic. But it is also critical to do the right things at different levels of the company. At the top it's about setting strategy. Looking at developments in a broad sense, at what matters most in your organization, and identifying which activities lend themselves to generative AI, where is the biggest potential leverage. On top of that, there's setting policy: developing guidelines, codes of practice, and putting guardrails in place – where would the application of generative AI tools create risk? Then there's protecting the value of your content assets. Large language models need two things, computing horsepower, and quality content. Media organizations have that content, so have an asset that LLMs need. They have leverage and need to ensure they use it. It would be great to see more collaboration.

You were talking about what needs to happen at the top. What needs to happen at the middle level?

As Charlie Munger said, "just learn, learn, learn all the time" – one of the positive impacts of generative AI is that it has put organizational learning at the top of the strategic agenda. We need to experiment. By using these tools, we build up knowledge. So, in the middle it is about building the expertise, finding efficiencies

and testing use cases, letting knowledge bubble up from the periphery. Start small. Identify the things that lend themselves to this technology, get people to try the tools and find the places it can move the needle.

Is that why everyone is appointing AI directors now, just to shed responsibility at the top?

I think it's more a case of orchestrating the response, ensuring all bases are covered and mitigating risk. With digital, we had digital directors for a while too, but eventually those roles disappeared: every role had a digital component.

Content creation is at the very heart of media companies. Could this make them react more slowly than other industries, just as they were too slow to react to classified advertising migrating to platforms because it was at the heart of the publishing business model?

The media industry is well-set up for a speedy response. Generative AI has huge potential for sectors that do 'knowledge work' and have well-structured data sets. Both are true for the media. Plus, they have learned from digital it is better to move fast than wait to see what emerges. In this respect I am optimistic.

Do you see any red lines for media organizations?

It needs to be made super clear internally, when generative AI is used, what it is used for, how it is flagged to audiences. It needs to be very intentional. You must be extremely careful what data you load up to ChatGPT. Also important is to know the strategically critical areas where there could be reputational damage. And, of course, be smart in dealings with the tech majors. The media industry gave away the keys to the kingdom once. That shouldn't happen again.

Will public service media be better equipped to deal with the challenges? They have the advantages of size and trust.

Public service media have an amazing strategic advantage because their brands are so well known. They stand for quality and are deeply anchored in the communities they serve. But they are constrained. Public service media are large complex organizations under high levels of scrutiny with relatively constrained strategic options – it's harder for them to roll with this, to try out the new technologies and see what happens. Smaller players can simply give things a go.

Q&A

> "The public does not want AI labels everywhere, only when it is materially important"

NIC NEWMAN

Senior Research Associate at Reuters Institute for the Study of Journalism, Lead Author Digital News Report

Is generative AI a game-changer for journalism?

I'm a bit more sceptical than six months ago. It's a huge transformation in terms of the potential for automation and better personalization. But where it moves the dial and how quickly I'm not sure. In the short term it won't change that much.

What made you more sceptical?

Experience tells me: it's one of those things that at first looks like magic. But when you really think about what journalism does, generative AI will only do part of the job. When you look at what it can do now, it's great at speed and scale, aggregating things, summarizing things, but a lot of the output is mediocre. Journalism aims to be better than that and connect people as well as inform them.

You interact with media leaders a lot. What do you sense as the general mood?

It depends on who you talk to. Most journalists are not that engaged but on the business and innovation sides, people are hugely engaged. Many hope for more efficiency. But they are also worried about what this is going to mean for jobs. On the business side some look at this as an existential threat. Will it further destroy business models? Then there are disinformation concerns.

You tend to be fascinated by new tech. What fascinates you most with generative AI?

It offers the opportunity to make journalism more relevant. How can we use AI tools to battle news avoidance? For example, to create content that will engage young people by turning a text story into a video story. Then again, we'll see a lot more of everything, lots more content, more marketing messages too. In that sense an overload of content could make news avoidance a lot worse because it will be even harder to find journalism that is relevant. My other concern is that the promised efficiency gains may not be that great. And publishers might not invest the savings in journalism but just take the money.

What has surprised you most?

The speed. The change from AI-created photographs that looked a bit weird into photorealistic images. That happened within a year. This is unbelievably exciting but also worrying. Anyone can now use AI to create/fabricate a picture of a child sitting amidst the rubble in Gaza with a sad look in their eyes, and it's hard to distinguish from the real thing.

Some say journalism could evolve from a push activity to a pull activity. What do you think?

People have talked about this for a long time: oh, with AI you could change the ending of a Netflix drama to your liking. In reality, high-quality linear drama has become even more important in a world that becomes fragmented and confusing. So, the outcomes here are not binary and there will still be both push and pull in news.

What about custom-made journalism?

We clearly haven't made customization in news work. Most people don't want the news agenda to be customized, because they 'want to know what they don't know'. But the personalization of formats is a different matter and more promising, but how this will happen is not obvious because the user interfaces are complicated. Any friction makes it harder.

We see an ever-increasing dependency on tech. Will there be anything in this for publishers?

We would hope that publishers have learned something from previous changes and especially the importance of the interface, which is likely to become an even bigger factor with AI. Publishers should have learned how to develop direct relationships with their customers including building great interfaces and platforms of their own. There will also be a significant amount of money from licensing in the new world of AI, but who will get it? Probably the big companies, including news agencies. Probably not smaller or local players.

What research have you done around AI for this year's Digital News Report?

We've conducted qualitative and quantitative research this year which uncovers high levels of public scepticism about AI and the news – much more in Europe compared with the United States. But we also show that the public is much more comfortable with back-end tasks where journalists remain in control. There's much less comfort in general with use of AI around politics or other hard news topics compared with entertainment or sports – and especially where there aren't enough human checks. Interestingly, transparency is important but the public does not want AI labels everywhere, only when it is materially important.

BACK TO CONTENTS

Your audience research very much centres around trust. Will AI destroy even more of it?

In our recent *Trends and Predictions Report*, 70% of news executives said AI will most likely lower trust. There's a lot of media coverage of deep fakes, for example around Joe Biden or Taylor Swift. That makes people more sceptical and more worried. They might develop the perception that you cannot trust anything. But conversely a flood of unreliable synthetic content may make people want to seek out someone they can trust. Trust in some brands might increase.

You're talking about the Covid effect, when trust in media spiked?

There are lots of unknowns here, but it is possible. Much will depend on what the platforms will be doing. It is actually in their interest to promote trustworthy content and keep their platforms as clean as possible during this shift.

What is the particular task for public service media? They have a big reach and enjoy trust. Will they be winners of all this?

Their big challenge is how to get young audiences consuming them with visibility and attribution, given their preference for platforms. This could be done with some kind of regulation, for example to prioritize PSB 'public service broadcasting' content on certain platforms. They could be the big winners – or big losers if their content is harder to find or AI further flattens content.

What is missing from current conversations?

The audience's perspective on this. So far, we've mainly had a debate about technology or about the business perspective, about licensing. We also need a longer-term view. Currently, everyone is experimenting, but we need to figure out what we really want from AI strategically. Public service media should really be at the heart of these debates because they are looking to act in the interest of all audiences, not just the privileged few.

Q&A

> "If you are below average, you are likely to benefit more"

CARL BENEDIKT FREY

Dieter Schwarz Associate Professor of AI and Work, Oxford Internet Institute

Do you think that generative AI is a game-changer in AI?

It certainly has democratized AI for the first time thanks to the interface that OpenAI has provided to make it more accessible. I think it's going to be the first AI tool that really has a widespread impact on the labour market. Yes, we've previously used AI in various contexts and we've been very exposed to it as consumers in various ways. But I think it's the first technology that is going to find widespread use among producers and normal people as well.

Is that mainly because it's much easier to use for most people compared with previous forms of AI?

Yes, but it's also better and it can be used for a greater variety of tasks. So previously, you had issues like transfer learning; you teach an algorithm to play Go on a square board, and then you know if you do the same thing on a rectangular board, you have to retrain it from scratch. Now we have a foundational model which you can use and tweak for a variety of activities.

Are you delighted or worried about this development, or do you take the classic position of every economist saying it depends on the context?

I think I'd do the latter. There are some reasons to be optimistic. A number of studies have found that it's primarily novices and low-skilled workers that benefit from generative AI, whether it's in terms of writing, coding, or repetitive customer service, and it makes sense because if you're not a great writer, well, with the help of technology like ChatGPT, you can become an average writer. So, if you're below average, you're likely to benefit more.

On the other hand, that doesn't mean that everybody will benefit. When you have reduced barriers to entry into a profession, that will mean more competition. In content creation, we're now seeing barriers to entry being reduced. I think it's likely that there will be many people who take a pay cut unless demand for content creation expands very significantly, and I find that hard to see. Think about it this way: if Netflix becomes much cheaper and much better, how much more content would you watch? The answer is probably not that much because your day is still

constrained by the 24 hours that it has, and you're not going to have time to watch a lot more content.

In your view, how is generative AI challenging ethical principles, and what should be the red lines, if any? Is there anything where you'd say regulators have to step in and do something?

Well, first of all, I think we should know more about the material that it has actually been trained on to see whether it can be regarded as fair use in terms of copyright legislation. I also think we obviously need to have safeguards in place to make sure that you can't use technology to perform acts that are illegal and potentially harmful. And we need to be able to stress test these models to make sure that that is the case. We can't entirely rely on the goodwill of companies in that regard.

Do you think the stakeholders in this game are up to the task of developing these systems in a way that is beneficial to society at large? And where do you see room for improvement?

Well, these companies are providing a product, and for them to be able to sell that product, it needs to be of some use to consumers. But because it's a general-purpose technology, it can be used by those consumers in a very wide variety of ways. It can be used to generate artificial images of people doing things that they've never done in their lives, which could be harmful, and we need to safeguard against that. It's not clear to me where the liability should lie going forward and what the right trade-off is. There are clearly certain use cases that we may decide are generally harmful, like facial recognition technology, except perhaps in tasks like counterterrorism. And we can regulate those without stifling innovation.

But at the same time, I do have concerns over over-regulating. Large corporations are much better placed to deal with compliance costs and regulation than smaller businesses. And so, it's not surprising that well-intended regulation like the GDPR has benefited large technology companies more than anybody else. So, we also need to factor in the effects that regulation might have on things like market concentration, and for the most part, those who have a seat at the table are most likely to shape the regulation. Startups who haven't been founded yet or are just in the early days of setting up are very unlikely to be part of that conversation, unfortunately.

What do you make of some of the discussions we've seen crop up around existential risks and the narrative that artificial general intelligence is just around the corner?

I don't make much of it, to be honest. There are compelling incentives to engage in that debate, especially for companies that are concerned about competition from open source, for example. So, if you're a bit cynical it seems to me that there are incentives to push that narrative. I think the conversation should focus more on the potential for human harms using AI than robots taking over the world.

Indeed, there are existential risks for people in this very moment, which brings us to your famous 2013 study with Michael Osborne where you argued that about 50% of jobs in the US are at risk from AI and automation.[1] How would you assess that study in the light of generative AI?

We did the reassessment of the bottlenecks to automation last year. Basically, in the original paper, we argued that there are three key bottlenecks to automation that persisted, and those relate to complex social interactions, creativity, and perception and manipulation tasks. We argued, for example, that the automated cleaner was one of the last things that we are likely to see simply because homes are unstructured; there are so many different objects in them, and we don't all share the same setups. That would be a very hard task to automate going forward, compared to, let's say, warehouse or factory automation where you can structure the environment purposefully for robots to be able to navigate it.

I think when it comes to complex social interactions, we've clearly seen a lot of progress in basic text communication, but also virtual audio assistance, etc. I think the potential scope of automation has clearly increased in that particular space. On the other hand, though, as AI becomes better at virtual communication, I think the returns to in-person communication are going to go up simply because if you're doing sales, for example, and you have two companies both using AI, those algorithms are likely to perform quite similarly, and so to distinguish yourself in that competitive race, I think in-person meetings are going to become more important.

When it comes to creativity, I think there's been some progress too, but the reason that technology can write a letter in the style of William Shakespeare is that Shakespeare has existed, and that means that you have clear benchmarks and AI can be used to rehash existing ideas and combine different concepts in ways that sometimes might be novel, but for the most part, progress has been made with clear benchmarks, like in video games where the objective is always to maximise the game score. It's been relatively straightforward, but when you do truly creative work, what do you maximise for? So, I still think that bottleneck broadly holds.

And then the last bottleneck, perception and manipulation tasks were constrained by advances in robotics, which have been less dramatic than those in AI. And in addition to that, we are also much less willing to accept risks, right? There's been a lot of progress in driverless cars and autonomous vehicles, but if an AI produces text that is incorrect, it's not the end of the world. We can edit it. If you have a car that causes an accident, unfortunately, you can't reverse that. So the rollout is much slower, but broadly speaking, I think those bottlenecks have held up quite well.

Let's pick up on one thing you said there, the bottleneck of human creativity. The readers of this report will be quite happy to hear that there's still some premium on human creativity because that's something that journalists,

[1] Frey, C. B., & Osborne, M. (2016). The future of employment: How susceptible are jobs to computerisation? Technological Forecasting and Social Change. https://doi.org/10.1016/j.techfore.2016.08.019

humans, are quite good at: finding the things that are not necessarily contained in the training data. Would you agree with this?

Absolutely. And I think a key concern with AI in general is that we don't know how it performs in novel situations, and news, by definition, should be relatively new. And take Go. We all know that AlphaGo won against Lee Sedol a couple of years ago, but humans actually made a quite astounding comeback recently with the help of quite normal computer equipment, and it turns out that AlphaGo doesn't understand certain concepts of Go. It's basically doing statistical approximation, and humans have been able to exploit that quite effectively. So even when we have AI that performs very well against the training data and then yields some good results, it's not always clear that they will continue to do that when novel circumstances emerge.

What do you generally think about AI and the news industry, both as an expert and as a consumer? What do you think is going to happen in that area?

The news industry has obviously faced a number of headwinds in recent years from other platforms emerging, and I think generative AI is basically a continuation of that trend. It reduces barriers to entry into content creation. It's also unclear, I think, to what degree people will use technologies like ChatGPT or new forms of search going forward, and if this ends up taking a large share of news organizations' search traffic, with knock-on effects for things like digital advertising. Finally, AI is also a tool that allows you to increase productivity and potentially allows you to produce more stories with fewer people. So it's also potentially a tool for cost-cutting within news organizations in response to those competitive pressures.

What do you think is missing from current conversations?

Something that might be missing is what AI does to the organization of work. The people who are actually working and developing these technologies have better information, so you actually might need to delegate more, and that might lead to greater decentralization. I also think AI is likely to increase complexity. If you have 20 artificial research assistants, you're going to do more projects, but you're also going to have a lot of additional complexity that you need to manage. So I think people are focusing more on which type of tasks and activities AI can be used for and not so much on the implications for their organization of work. I think people are also using AI to automate things that they are already doing and trying to be more productive in those activities, rather than thinking about what it is that they now can do that they previously couldn't do with the help of AI.

Q&A

> These technologies will allow perfect manipulations, and you might no longer be able to identify that you are being manipulated"

JEAN-MARC RICKLI

Head of Global and Emerging Risk, Geneva Centre for Security Policy

Do you consider generative AI a game-changer for journalism?

Yes, on several counts. It has great potential for productivity and efficiency gains. For example, I was asked to write an op-ed for a newspaper. I wrote the text. But then I used ChatGPT to format it according to the op-ed style of that newspaper. So, it saved me a lot of time. The ideas, the arguments and the structure of the paper are mine, but all the formatting in accordance with the newspaper's style was done by the algorithm.

There are also lots of potential risks to be aware of that go beyond just journalism and target the essence of democratic societies. For the news industry, Generative AI (GenAI) is definitely a game-changer because it basically addresses and challenges a key factor of the job of a journalist, which is the identification of legitimate sources and fact reporting. Beyond this, the fact that GenAI is now directly competing with reporting as a key function of journalism, that it has a global reach that few journalists have, and that false information is spreading more rapidly than true information, is a risk to democratic societies as it further deepens the erosion of trust that people have in democratic governments.

But we like to think GenAI cannot do what reporters do.

In news production, humans will increasingly be outpaced by algorithms in terms of efficiency and the speed at which they work. Not only do algorithms work way faster than humans do, but they can process all the information available on the internet 24 hours a day, 7 days a week. A recent experiment published in the magazine *Wired*[1] showed that with GenAI, it was possible for anybody to create the equivalent of a news agency that would automatically generate articles on any topic from fake journalist accounts, and all this for $400. Disinformation tools at an industrial scale, with global reach, are now accessible to anybody. GenAI algorithms are also known for hallucinating, which is when the information that is provided, though at times sounding accurate and legitimate, is completely wrong. In that op-ed piece I was referring to,[2] I argue that ChatGPT is the first-ever WMD.

[1] https://www.wired.com/story/400-dollars-to-build-an-ai-disinformation-machine/

[2] https://www.weforum.org/agenda/2023/08/does-un-needs-watchdog-fight-deepfakes-ai-threats/

Short for weapons of mass destruction ...

Yes, but in this case, I'm using WMD to stand for Weapons of Mass Disinformation because GenAI is able to generate content that looks very legitimate, but in fact is often wrong. This creates a major concern for how people relate to information as the amount of fake information increases exponentially. The consequences affect especially democracies, whose working principle relies on accurate information conveyed by trusted sources such as journalists so that people can form legitimate opinions, for example to elect their representatives or vote on specific issues. Journalists whose job is to provide verified information will increasingly find themselves in a race with machines that is completely asymmetrical. We have already witnessed the use of GenAI through deepfakes during the Israel-Gaza war. Some have even been unknowingly published by legitimate media sources.[3]

How do you see the potential of GenAI to verify information?

It has become almost impossible to come up with failure-proof detection and verification methods. And if deepfake output is becoming reliable (meaning, no hallucinations, which it is not the case today), it could then directly threaten journalists' jobs in their very essence.

You are an expert in geopolitics and security. What are the implications of GenAI you see in your field?

The immediate concern is unquestionably disinformation and the undermining of the social contract between citizens and the legitimate authorities in democratic countries. The second major issue is the democratization of knowledge that should not be democratized because this information has huge security implications, such as the identification of new pathogens or chemical compounds that could be weaponized, or the democratization of complex malware creation methods. As a security analyst who closely follows what terrorist organizations have done over the last two decades, the democratization of knowledge in bio or chemical weapons, or the proliferation of swarming algorithms, is becoming a major security headache.

Could you elaborate on the concept of 'subversion control' that you have been working on recently?

With GenAI and social media, it is becoming way easier to lock people into information bubbles. As a result, we're getting closer and closer to not only controlling the flow of information but also increasingly having an impact on the target's response. We're entering the era of cognitive warfare that is controlling how and what people think in order to control how they act. Augmented and virtual reality technology will bring this to a new level.[4] But it is the developments in neurotechnology that could be a real game-changer. Brain computer interfaces (BCIs) are devices that interfere directly with your brain, either as wearables or implanted directly on the brain. The combination of AI and neurotechnology has opened the door to the possibility of mind reading. Algorithms are now being developed to

[3] https://www.euronews.com/my-europe/2023/10/24/israel-hamas-war-this-viral-image-of-a-baby-trapped-under-rubble-turned-out-to-be-fake

[4] https://www.gcsp.ch/publications/our-digital-future-security-implications-metaverses

interpret brain signals such as electrical activity and then reconstitute, for instance, images of what the person sees or thinks. The next step, once we've mastered mind reading, will be mind writing: the ability to directly interfere with brain activity. These technologies are still in their infancy, but mind-reading technologies have made some impressive strides forward lately.

What can be done to protect humans and societies from this?

The international community should carefully invest into the concept of subversion control, designing a system of laws and regulations to prevent the undermining of legitimate authorities through different means and strategies. At the individual level, one such development would mean to prevent malevolent agents from interfering with our brain integrity, and preventing influencing the way we think and act. Chile was the first country to issue constitutional legislation enshrining neurorights via the protection of brain activity and data, in 2021. But this is the exception rather than the norm, as there has been little action from governments in terms of legislation or governance in this area so far. These technologies offer fantastic prospects to treat trauma, psychiatric disorders or paralysis. However, my concern is that if we're able to develop devices that are able to read people's thoughts and maybe one day rewrite what they think, then there'll be ways to weaponize this technology.

This sounds a bit like science fiction.

If you think about GenAI and social media and you add a dose of immersive technologies, these technologies increasingly get us towards perfect manipulations, because they target peoples' cognition by providing them with information that resonates with their existing belief systems and locks them in a specific information ecosystem inside which they can easily be manipulated without even noticing it. The other issue is that these technologies are developed by privately owned companies that have financial and research capabilities that few governments have.

Do we need more regulation?

Monitoring what happens on social media is very important as this technology permeates everybody's life. For instance, the US government is on the brink of banning TikTok because it considers it a threat to its national security. What TikTok is doing (like any other social media company) is locking its users into a specific information ecosystem that resonate with the users' interests in order to generate as much data as possible that can then be monetized. TikTok and more generally social media represent an early example of how malevolent agents could wage cognitive warfare. The more immersive the technology becomes, the more effective it will be.

Many news organizations think they need to be on TikTok because it seems to be the only way to reach a large segment of young news consumers.

TikTok is a social media company that is indeed very popular among the young generations. News organizations have to be on every platform that reaches different generations. This also implies using different communication strategies. Beyond this, what the TikTok case highlights, however, is the clash with democratic principles such as freedom of speech. The banning of TikTok highlights the legal hurdles that

democratic governments face in dealing with the weaponization of the very principles that underpin democracies. In this case, resorting to national security is a way to overcome this, but it's a slippery slope. Democratic governments must conduct a cost-benefit analysis between the democratic principle of freedom of speech (or using other means of communication) and the way this principle is being weaponized against democratic values in their essence. We have to be much more realistic about the type of challenges that we're facing in terms of how enemies of democracies can weaponize digital tools within free and accessible information ecosystems.

What should governments do?

One way to deal with that is for governments to invest much more in supporting national media and legitimate sources of journalism. At the end of the day, one of the functions of journalism is to be a gatekeeper of legitimate information and therefore a key enabler of the democratic process. Well-informed citizen with legitimate and true information at their disposal are the backbone of democracies.

What would be your advice to media industry leaders?

Don't compete on speed, but on accuracy and the legitimacy of the information you provide. The job of media companies is to filter information and create legitimate and verified knowledge that supports people's decision-making. Also, there is no point in just stating facts because anyone can find them online. Rather, triangulating sources and providing context and explanations is crucial for any legitimate media.

CHAPTER 2

‘JOBS TO BE DONE’ AND HOW AI CAN HELP: USE CASES

While we ended Chapter 1 with a note of caution, Chapter 2 should inspire some fun. It’s all about use cases: ways of using AI which news organizations have already been testing and implementing. Consciously thinking about how AI could add value to serving audiences is crucial for any kind of media organization. Generative AI provides endless possibilities, so it can quickly become overwhelming. And in a world of information abundance, simply adding to the stew won’t make it more satisfying.

For example, AI-powered tools might battle news avoidance. They could make content more accessible by presenting it in different formats to a variety of audiences and customize the news experience by reducing the sheer volume of news. Both outcomes will affect news consumption habits. Nic Newman, lead author of the *Digital News Report*, certainly hopes the technology will have a positive impact. AI tools could be used, “for example, to create content that will engage young people by turning a text story into a video story. Then again, we will see a lot more of everything, lots more content, more marketing messages, too. In that sense an overload of content could make news avoidance a lot worse, because it will be even harder to find journalism that is relevant.” (See Q&A with Nic Newman, page 38) A systematic approach is needed. Good advice for dealing with disruptive technologies is to focus on the ‘jobs to be done’. This harks back to the work of the late Harvard Professor Clayton Christensen, who developed the theory of disruptive innovation and applied it to the news industry in an essay for Nieman Reports.[1]

The article, published in 2012, anticipated the ‘audiences first’ approach, which many news organizations follow today. Christensen and his co-authors recommended crafting journalism to address audience needs: “The basic idea is that people don’t go around looking for products to buy. Instead, they take life as it comes and when they encounter a problem, they look for a solution – and at that point, they’ll hire a product or service.”

Applying this approach to the use of generative AI in news means starting with an examination of the strategic problems AI might solve. This keeps publishers focused on delivering for their audiences while being open to the new possibilities that the technology provides. It also helps them to stay resource conscious in every sense of the term. This is why each of the use cases we present in this report starts with a reflection on which problem is being solved.

The experimental research project AI in Journalism Challenge 2023, which was led by David Caswell and funded by Open Society Foundations, used a similar framework.[2] Applicants were asked to state what kind of journalistic impact they were hoping to achieve through the programme and how they would measure this. The competition gave rise to an interesting array of use cases, ranging from the Brazilian Agência Pública which experimented with tracking the impact of its journalism, to a team that worked on a Chinese language monitoring tool which was designed to ‘read between the lines’ of content and in doing so circumvent censorship filters.

[1] Clayton Christensen, David Skok, James Allworth, “Breaking News: mastering the art of disruptive innovation in journalism,” Nieman Reports, October 2012, retrieved on 13 February 2024. https://niemanreports.org/articles/breaking-news/

[2] See final report with 12 examples of AI experiments conducted by small news organizations around the world: David Caswell, “AI in Journalism Challenge 2023,” Open Society Foundations, April 2024. https://www.opensocietyfoundations.org/publications/open-society-s-applied-ai-in-journalism-challenge?utm_campaign=osfacct&utm_medium=referral&utm_source=threads, retrieved on 18 April 2024.

Most of the projects were in the early stages of development but had one thing in common: their starting point was a problem the team wanted to solve.

Traditionally, news organizations have not been trained to put potential customer needs first when thinking about innovation. When the internet disrupted journalism in the 1990s, most newsrooms responded by simply transferring their print logic to their websites, as did the advertising departments. Their starting point was what they did and liked to do, not what their users might need to make a buying decision or – even better – build a habit. Many failed miserably. Even big tech has not been immune to this producer-centric thinking. Google Glass and the Metaverse are expensive testimonials to developments that might have sounded fascinating but have so far failed to meet customers' needs.

This should serve as a warning to news organizations who are preparing for the future. Before they start investing in the implementation of generative AI tools just because it can be done, they should carefully identify user and audience needs and work from there. Dmitry Shishkin, CEO of Ringier Media International and user-needs evangelist, says: "Journalism in the age of generative AI still needs to be done in a way that is helpful and meaningful to people. Everything else is just a gimmick. If you don't have a connection to your audiences, you don't have loyalty, if you don't engage people, it will be of no use."

Sometimes audiences' needs are very different from their surface appearance. One example is personalization. (Chapter 4 discusses its ethical aspects.) While it is tempting to assume that people want personalized news, it is well worth testing user behaviour to check if this assumption holds true. A news timeline is different to a Spotify playlist where users want to listen to the same songs over and over again. Most journalism items will serve their purpose only once. That's why surprising people with something fresh and unknown will remain one of the major jobs to be done by journalism – and it's not necessarily easily achieved with technology alone.

Figure 6: Uses of AI systems in news organizations.[1]

Production and distribution process	Use of AI systems
Access and observation	• Information discovery • Audience and trends analytics; story detection • Prompting for new ideas following from a news story
Selection and filtering	• Verification, claim matching, and similarity analysis (e.g., for fact-checking) • Content and/or document categorization; analysis of datasets • Automated collection and analysis of structured data (e.g., financial, banking, and sports data) • Coding assistance for various tasks • Transcription and translation of audio and video • Search in archives and/or metadata
Processing and editing	• Brainstorming and ideation • Content production (writing of draft text or articles; editing of news content) • (Re-)formatting of content for online, social media, print, broadcast (e.g., summarization, simplification, stylistic changes; text-to-video, speech-to-text, text-to-speech translation) • Copy editing, adaptation to house style • Tagging of content, headline, and SEO suggestions
Publishing and distribution	• Personalization and recommendation • Dynamic paywalls, audience analytics • Content moderation

Table 2: Common applications of AI in news organizations

[1] Simon, F. M. (2024). Artificial Intelligence in the News. How AI Retools, Rationalizes, and Reshapes Journalism and the Public Arena (p. 46). Tow Center for Digital Journalism, Columbia University. https://www.cjr.org/tow_center_reports/artificial-intelligence-in-the-news.php, p. 13.

To identify these jobs in more detail, we need to look at the biggest challenges for media today. These can be divided into three groups: The jobs journalism has to do – and do better – for society, the jobs that need to be done to improve the performance of news organizations and the jobs necessary to sustain journalism as a business. Naturally, all three are connected. Newsroom performance and business performance are not ends in themselves. Their ultimate goal should be to serve society with better journalism. But while the first group of challenges is audience-facing with the focus on news consumers, for the second and third group it is newsroom staff and management who need to feel the advantages when using a particular tool. Figure 6 provides an overview of activities that can be supported with AI tools.

Styli Charalambous, CEO of *Daily Maverick*, describes it like this: "We think about innovation in two ways. One is exploitative and one is explorative. Exploitative means to get better and more efficient at your current processes. And explorative is to go to new places, to create new products, to do new things, to find new ways. Generally, exploitative is internal and explorative is external. And news organizations have traditionally focused on exploitative innovation. So, it's about improving processes and efficiencies and workflows, which has all been internal stuff. But we have to be conscious about making sure that we are doing both." Defined this way, most use cases we describe here are exploitative, as tools replace or enhance actions that have been performed by humans before. But quite a few experiments and pilots deliver services that would have been impossible or far too expensive to be generated by editors or other staff.

PART 1: AI support for journalism's jobs in society

The basic role of journalism for people and society hasn't changed much over the years. Some of the most important include providing facts as a basis for decision-making and public debate, holding power to account, explaining, educating, giving voice to and reaching the disadvantaged, connecting and empowering people, connecting the general public with elites and vice versa, and helping people navigate their daily lives. The mission, strategy, and legitimacy of public service media are built on these tasks. In the modern information and communication environment, journalism faces several challenges to getting these jobs done. The most fundamental are (in no particular order) news avoidance and lack of interest, declining trust, misinformation and disinformation, lack of access and comprehension, underserving of significant parts of the population, and competition with alternatives to news. Some media companies have already updated their strategies to tackle these challenges and to redefine their value proposition in an AI driven context, most have not.

Generative AI is likely to provide an abundance of tools and opportunities to support news organizations in battling these challenges. For example, targeting and serving different audiences with specific products and content will likely contribute to them feeling acknowledged in their needs and interests. This could build trust and increase news consumption.

Additionally, the increasing availability of data generates opportunities to serve local audiences better with hyper-local content, including specific weather reports, real estate listings, traffic reports and the like.[3] Research shows that local news brands tend to be among the most trusted media in many countries – only topped by public service media, which often have a regional focus, too.[4] But catering to the needs and appetites of people in specific communities can be prohibitively expensive. It requires reporters to put in many hours of effort with no opportunity to scale their output. This has become even harder over the years, with the collapse of the advertising business model in many regions. Generative AI may make it

[3] See Alexandra Borchardt, "Go, Robots, Go! The Value and Challenges of Artificial Intelligence for Local Journalism," Digital Journalism (Volume 10, 2022, Issue 10), https://doi.org/10.1080/21670811.2022.2149584

[4] The Reuters Institute's Digital News Report, the world's largest ongoing digital survey on news consumption, shows this again and again in its trust rankings.

possible to produce a critical mass of content that makes brands visible. For example, to serve people in different communities better and around the clock, the BBC and German RBB have used synthetic voices for automated traffic and weather updates. (Read Case #1, Automated Reporting, page 59) In the summer of 2023, NewsCorp in Australia reported having produced 3,000 local news stories a week using generative AI for 75 hyperlocal mastheads about topics like weather, fuel prices, traffic conditions, and death notices.[5]

Philippine news organization Rappler, founded by Nobel Laureate Maria Ressa, won the AI in Journalism 2023 (AIJC) project. The Rappler team had worked on AI-generated means of transforming traditional online news content into formats that are more likely to engage young audiences, including short videos and comics. The report states: "Of all of the AIJC projects, Rappler's TL:DR probably came the closest to showing how some of the most threatened news organizations globally – mid-sized regional or metro publishers –might productively and pragmatically use AI to extend the reach and impact of their work."[6]

AI can also help to lower barriers to news use and comprehension – for example through automated translation, text-to-speech, or speech-to-text tools, or visualizing written content. This is particularly important for people with different native languages or audiences with impaired hearing or vision. But it can also be used to enhance experiences, for example at large events when the sound is compromised. When the Belgian public service broadcaster VRT, in cooperation with the EBU, used live subtitling at a major charity event, audience feedback was overwhelmingly positive. In a follow-up survey, about 85 percent of the 1,200 respondents were satisfied with the result, even though the transcription was not always perfect. (Read Case #2, Live subtitling, on page 61)

Sveriges Radio learned that people with hearing impairment don't expect perfection when relying on the transcription of audio news. Christian Gillinger, who is responsible for accessibility at SR, says: "Through our dialogue with organizations representing people with hearing loss we know that text supported audio is favoured even if it contains errors. People want and deserve to have equal access to our content." (Read Case #3, Automated transcription, page 63) Not only does automated transcription of audio news make audio content more accessible, it also makes it searchable and easier to retrieve quotes.

Johanna Törn-Mangs of Finnish Yle talks about her favourite examples of use cases supported by generative AI: "I really like how we presented the news in Ukrainian when the war started. We only had one person who spoke the language. The news was made by AI but checked by this person. We have also implemented an automated weather forecast on local TV. We can't have a news anchor for this, so we just used pictures before. Now we have them explained by a synthetic voice. This is important for audiences with special needs, for people who cannot see but can hear it." (Read Case #4, Pop-up news in Ukrainian, page 66)

Automated translation also helps content to travel further and expose stories from regions where the dominant languages have low reach to a broader audience. *A European Perspective*, a project developed and led by the EBU, serves as an example. It increases the reach of relevant journalism done in the participating countries through automatic translating and showcasing.[7] It is also an important research tool where public service media journalists can search and understand translated news from all over Europe, with 3,000 new stories added daily.

Importantly, AI can help with customizing the news experience for different tastes, habits, and individual learning profiles. While some people

[5] Amanda Maede, "News Corp using AI to produce 3,000 Australian local news stories a week," *The Guardian*, 31 July 2023. https://www.theguardian.com/media/2023/aug/01/news-corp-ai-chat-gpt-stories, retrieved on 30 April 2024.

[6] The case is described in David Caswell, "AI in Journalism Challenge 2023", pp. 27-28.

[7] Read a case study on 'A European Perspective' in the 2021 EBU News Report: A. Borchardt, F. Simon, "What's Next? Public Service Journalism in the Age of Distraction, Opinion, and Information Abundance". https://www.ebu.ch/publications/strategic/loginonly/report/news-report---whats-next-public-service-journalism-in-the-age-of-distraction-opinion-and-information-abundance, retrieved on 30 April 2024. Latest developments can be found here: https://www.ebu.ch/eurovision-news/european-perspective.

“Journalism in the age of generative AI still needs to be done in a way that is helpful and meaningful to people. Everything else is just a gimmick.”

DMITRY SHISHKIN, CEO of Ringier Media International and user-needs evangelist

prefer to digest information visually, others enjoy reading or listening. The Norwegian news outlet *Aftenposten* and Swedish SvD, both owned by Schibsted, have been among the front runners using cloned voices enabling users to listen to news articles (Read Case #5, Cloned voice to read text articles, page 68). The Brazilian newsbrand Agência Pública experimented with cloned voices reading investigative stories to extend their reach (Read Case #6, Broadening The Reach Of Investigations Using AI Cloned Audio, page 70). AI can enhance the possibilities for creating visual content, ranging from video to cartoon production. In February 2024, OpenAI launched Sora, an AI application that turns text into video for those who'd rather watch than read.[8] We expect news organizations to adapt tools like this to their own needs over time.

Additionally, generative AI will increase opportunities for interactivity. Chatbots will be able to explain news, answer questions, or serve as personal news assistants. Deutsche Welle's Editor-in-Chief, Manuela Kasper-Claridge, says: "As and when chatbots become the main way that people find their information, their relationship with news will change. It's likely they will be able to ask questions of news events and stories much more easily, and more context will be at everyone's fingertips." (Read Q&A with Manuela Kasper-Claridge, page 130) The BBC has tested exactly that with a chat format serving audiences with automated explainers. (Read Case #7, Tell me more, page 72)

Finally, generative AI will help with data journalism, an underexplored but important field, not only in investigative journalism. Zach Seward, *The New York Times*' Editorial Director of AI initiatives, listed an array of traditional AI uses in investigative journalism projects in the US In his first public appearance at SXSW media and tech gathering in Austin in 2024. They ranged from the detection of tax havens (Quartz), abandoned oil wells in Texas (Grist, *Texas Observer*), and spy planes (BuzzFeed News) to reporting on the war in Gaza by the analysis of satellite images (*The New York Times*) and on which books are banned in state prisons (The Marshall Project).[9]

Generative AI driven tools are potentially useful for recognizing patterns across different types of documents. As Seward summarized: "Faced with the chaotic, messy reality of everyday life, large language models (LLMs) are useful tools for summarizing text, fetching information, understanding data, and creating structure." But human guidance and oversight were indispensable. "In all of these cases, it's humans first and humans last, with a little bit of powerful, generative AI in the middle to make the difference." However, Andrew Strait of the Ada Lovelace Institute says that the assumption that generative AI will save time depends on the task, and that many of the promises around the technology have to be taken with a grain of salt. With respect to the summarization capabilities of LLMs, Strait pointed out that many of the early promises around the reliability of these systems in this respect do not hold up in later studies.[10]

Clearly, AI-informed or generated products for public consumption require careful reflection, human input and genuine transparency. This might increase workloads instead of reducing them, particularly as research projects will be conceived that wouldn't have been possible before. But publishers need to consider the potential risk to the brand before implementation. *The Guardian* is among the publishers who are particularly cautious with rolling out tools in the newsroom, focusing instead on developing and testing. According to Chris Moran, their head of Editorial Innovation, his team has built and tested tools others have already deployed. "For example, an '*ask the Guardian*' tool." But they felt it was not reliable enough for release. "We are developing it as an internal research tool instead."[11]

Well-known brands might not want to risk appearing on a list like Zach Seward's. He gave

[8] Read here about Sora: https://openai.com/index/sora/

[9] Zach Seward, "AI news that's fit to print," 11 March, 2024, https://www.zachseward.com/ai-news-thats-fit-to-print-sxsw-2024/, retrieved on 29 April 2024.

[10] See for example Kim, Y., Chang, Y., Karpinska, M., Garimella, A., Manjunatha, V., Lo, K., Goyal, T., & Iyyer, M. (2024). FABLES: Evaluating faithfulness and content selection in book-length summarization. arXiv preprint arXiv:2404.01261. https://arxiv.org/abs/2404.01261

[11] See a demonstration on the panel "AI showcase", recorded at the International Journalism Festival 2024 in Perugia. https://www.journalismfestival.com/programme/2024/ai-showcase.

some negative examples of publishers who implemented generative AI solutions with business interests prioritized over audience needs, commenting: "There are some common qualities across these bad examples: the copy is unchecked; the approach, as lazy as possible; the motivation, entirely selfish; the presentation, dishonest and/or opaque."[12]

Many of the early experiments show that generative AI can create content that looks and sounds like news, is faster and more comprehensive. Bots are already helping journalists to come up with interview questions, different hypotheses, counter arguments, and data that informs stories. Humans and machines have even been working together on screenplays – which seems like an obvious use, as facts don't matter as much in fiction. One of the major reasons Hollywood writers went on strike in 2023 was the fear of being replaced by AI.[13] But the tools will not be able to replace journalism's core activities anytime soon: eliciting human voices, eyewitness accounts, original thoughts, unique stories that haven't been told before.

The big question is: how will audiences react to AI-generated content, and will they even notice? Research indicates that audiences don't mind automated writing, though they might find articles written by humans more comprehensible, as Sina Thäsler-Kordonouri, who researches AI's effects on audiences at Ludwig-Maximilians Universität in Munich, contends.[14] However, as generative AI currently is shaking up the industry, "whether and to what extent generative AI will improve the quality of news largely will depend on how newsrooms implement its use," she says. Audience perceptions might also depend on the type of content. A cookbook written and designed by an LLM might still do its job, as long as the recipes work, while a longform report or an essay might not.

Most likely the way the media reports on AI will also shape audience perceptions. Nic Newman,

Figure 7: **Public Attitudes to use of AI in News Production.**

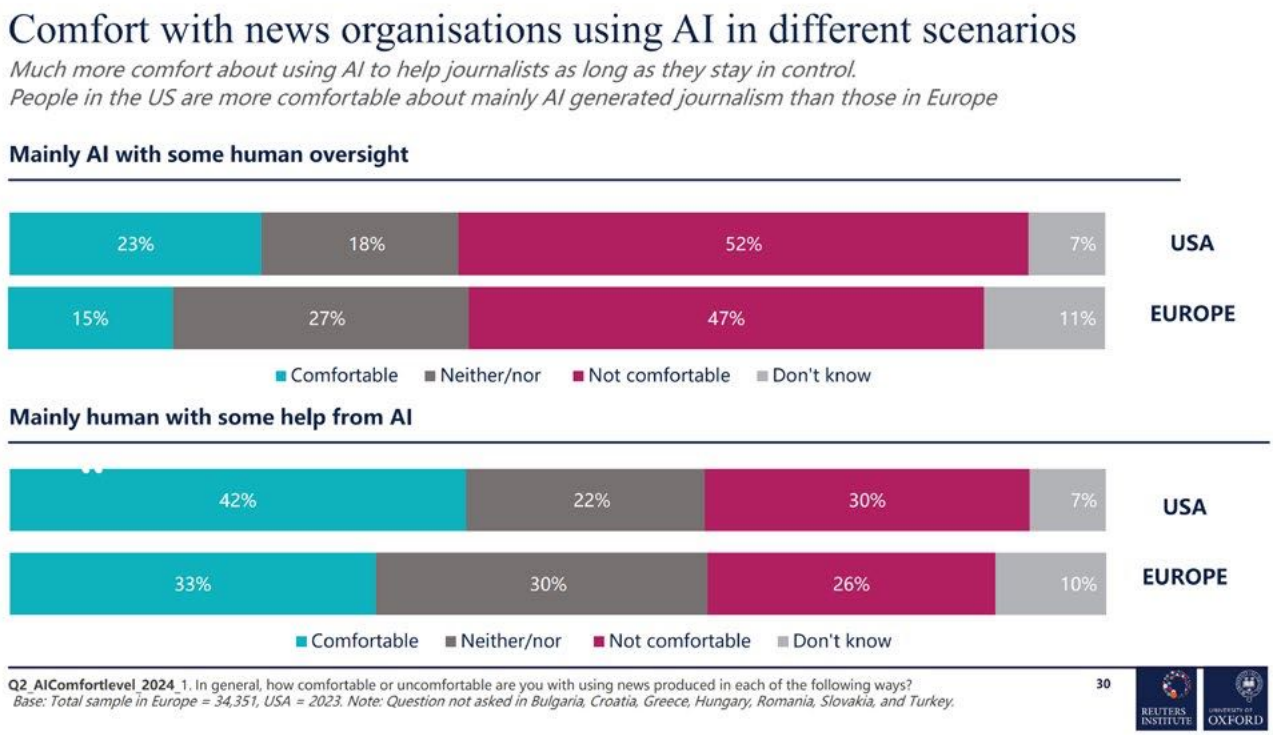

Source: Reuters Digital News Report 2024

[12] See above Seward (2024)

[13] Jake Coyle,"AI is one of the reasons that Hollywood writers are on strike: 'Too many people are using it against us and using it to create mediocrity." Fortune, 5 May 2023. https://fortune.com/2023/05/05/writers-strike-hollywood-ai-scripts/, retrieved on 2 March 2024

[14] See also Stalph, F., Thurman, N. & Thaesler-Kordonouri (2023), "Audience reception of news articles made with various levels of automation – and none: Comparing cognitive & emotional impacts," paper presented at The Joint Computation + Journalism European Data & Computational Journalism Conference 2023, 22-24 June 2023, Zurich, Switzerland. https://openaccess.city.ac.uk/id/eprint/30699/1/, retrieved on 18 April 2024.

Figure 8: Public Attitudes to use of AI for News Categories.

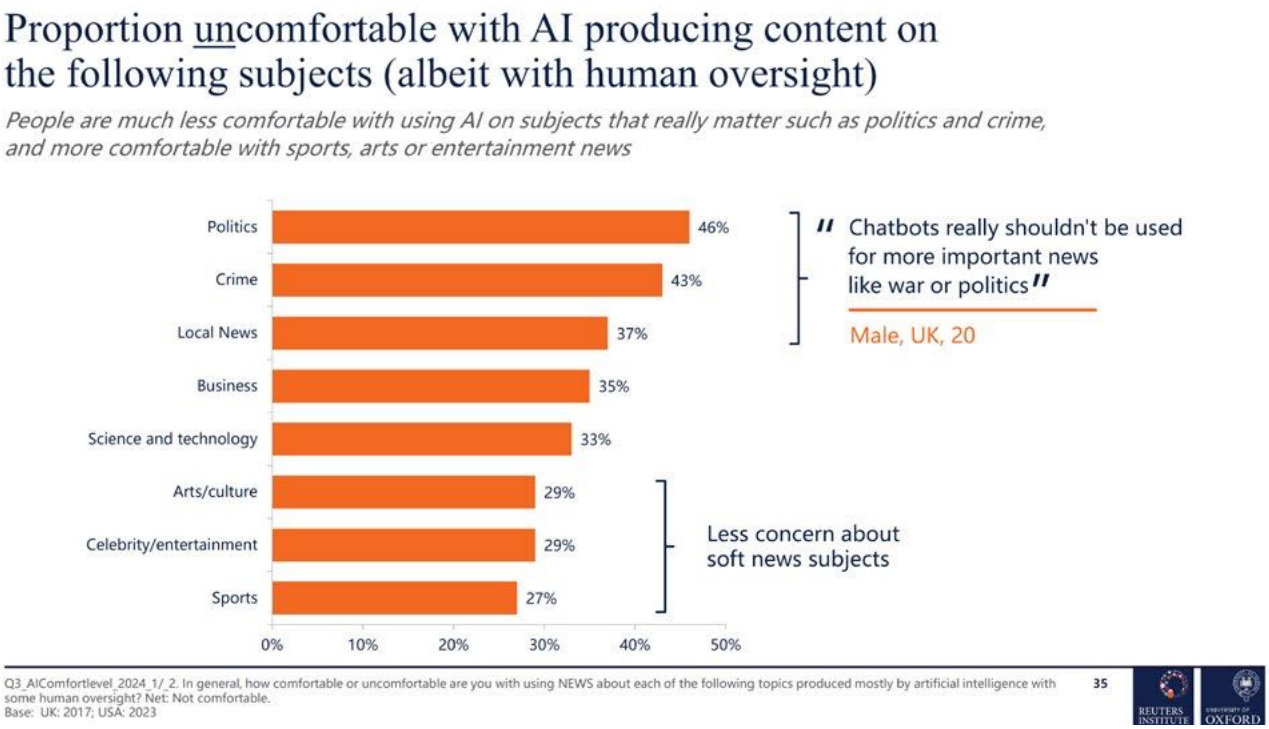

Source: Reuters Digital News Report 2024

lead author of the Reuters Institute's Digital News Report, shares some of the findings of the 2024 edition: "We've conducted qualitative and quantitative research this year which uncovers high levels of public scepticism about AI and the news – much more in Europe compared with the United States. But we also show that the public is much more comfortable with back-end-tasks where journalists remain in control. There is much less comfort in general with use of AI around politics or other hard news topics compared with entertainment or sports." (See Figures 7 and 8, read the Q&A with Nic Newman page 38).

CASES

CASE

WEATHER AND TRAFFIC UPDATES WITH SYNTHETIC VOICES

BBC (UK) AND RBB (GERMANY)

What was the problem to be solved?

The demand for localized weather updates in audio form is large. The BBC Weather app was claimed to be 'the fastest growing BBC app of all time' after its launch in the summer of 2013 and has a massive reach in the UK. However, the demand for fast, accurate and local weather forecasts makes it impossible to provide updates several times an hour for multitude of locations. The answer for several public service media has been to introduce automatic weather services presented with synthetic voices, in addition to forecasts presented by humans. These services are now provided by the BBC and by Rundfunk Berlin-Brandenburg (RBB), the German public service broadcaster.

What was done?

The BBC offers a local audio forecast with updates four times a day, every day. The synthetic voice was trained at the BBC on 3,000 spoken sentences to create a natural-sounding voice clone which can read out a weather forecast, and other articles if needed. Signed-in users of the Play BBC News smart speaker service for Alexa devices also hear the synthetic voice weather forecast for their postcode, after the news from their local radio station.

Since 2021, RBB has given app users AI-supported weather updates and takes it a step further with traffic information every 15 minutes. Both companies are using state-of-the-art text-to-speech models. In RBB's case, the voice of a well-known presenter has been cloned to make it more familiar and trustworthy, although it is clearly annotated in the player that a synthetic voice is used. This training was done in partnership with Microsoft Azure.

Both companies include a disclaimer that weather reports are read by a synthetic voice in their products.

What was learned?

At RBB, one learning is that the service needs to be cost efficient. Otherwise it is not attractive, feasible or scalable to implement in times of pressed budgets.

"This project also helped making AI-based services for our audience – and our colleagues – more tangible. Over the past three years our editorial department has grown more positive towards the use of AI services such as text-to-speech," says Anne Ain, Technical Project Coordinator at RBB.

BBC adds that the automated service, perhaps the most well-known AI service from the public service giant, proves that text data provided by meteorologists can be converted into a rich audio weather forecast.

What were the results?
Both companies have noted increased engagement in the app, as service info like weather and traffic is in such high demand. It has also been a cost-efficient way of offering added value as they have been relatively cheap to develop with new AI models that work around-the-clock.

Particularly suited for:
Weather, traffic and other services. Neither company uses it in news nor has any plans to.

Resource needs:
Substantial training data in audio and text. Text-to-speech information flow that allows updated text information to be converted to audio files. Project management, resources for product integration/UX and, if needed, partnerships with presenters and external tech providers.

Future plans:
Neither BBC nor RBB have any plans to evolve the weather offer. BBC states that the synthetic voice could be used for reading up other content, if there are future needs for that.

Illustrations:

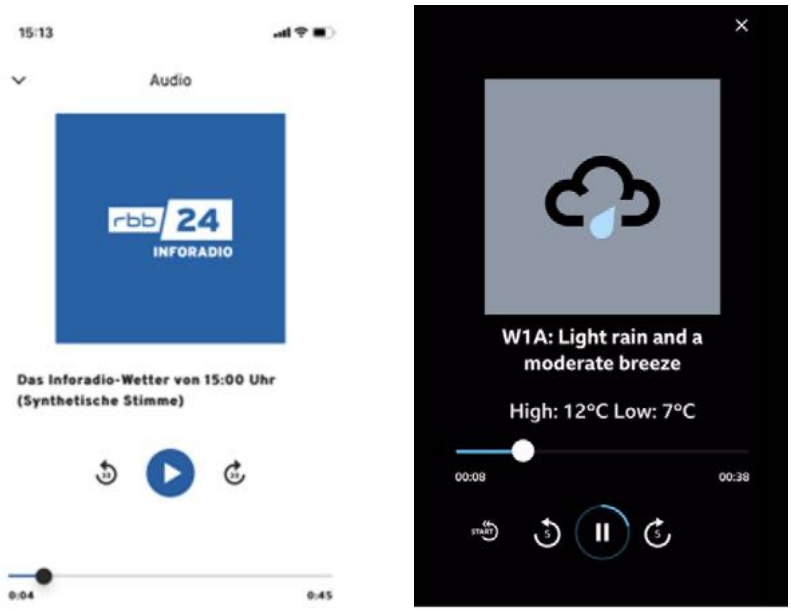

Link: https://www.bbc.com/weather/articles/c7219x55vygo

Contact person BBC: tom.roles@bbc.co.uk

Contact person RBB: anne.ain@rbb.de

CASE

AI LIVE SUBTITLING OF A MAJOR PUBLIC EVENT

VRT (BELGIUM) & EBU

What was the problem to be solved?

Every year VRT (Flemish Public Service Media) hosts *De Warmste Week* (The Warmest Week) – a 7-day live radio broadcast from 't Zand Square in Bruges, raising money and awareness for a range of different charities and causes. As well as their radio transmission the programmes are live streamed on the digital platform VRT Max. The ambition of De Warmste Week is to be 100% accessible and so they needed a way to provide live subtitles for the shows. Could automatic subtitling using AI provide a solution?

What was done?

VRT collaborated with EuroVox, the EBU transcription and translation AI tool, to provide a live captioning service for the events. This was the first time EuroVox was used in a live production and the captions were available on VRT Max online and through connected TVs. The technical setup involved establishing a clean audio signal from Bruges to Brussels, which was then routed to the EuroVox platform maintained by the EBU. At VRT, the live transcripts received from EuroVox were automatically integrated into the video signal on VRT MAX.

What was learned?

The feedback was overwhelmingly positive. In a follow-up survey, approximately 85% of the 1,200 respondents expressed satisfaction. Most people appreciated having the service available even if the transcription was not always perfect. For EuroVox, the insights gained from this pilot project are invaluable for enhancing the transcription service so it can be used on more live content in the future.

Particularly suited for:

This innovative approach is particularly well-suited to live broadcasting and streaming, whether in video or enhanced audio format.

Resource needs:

Collaboration between EBU developers and VRT's local teams was essential. The consumer service at VRT involved six developers and took slightly over a month to develop. Given the complexity of *De Warmste Week*'s production, involving various teams from Radio and Television, clear stakeholder roles were critical in managing the subtitling experiment. The tight window between finalizing production in Bruges and going live necessitated precise coordination and a clean signal of the output transmitted to VRT headquarters in Brussels.

Next steps and developments

For EuroVox to use the data gained to improve the accuracy of its transcriptions and to facilitate its use on other live events. For VRT to examine how the technology could be more fully integrated into its production model and to examine how it might further expand accessibility.

Illustrations:

Contact person VRT: klaas.baert@vrt.be

Links: EuroVox: https://tech.ebu.ch/eurovox

VRT's De Warmste Week: https://www.vrt.be/vrtmax/themas/de-warmste-week/

CASE

MAKING NEWS AUDIO MORE ACCESSIBLE THROUGH AUTOMATED TRANSCRIPTIONS

SR (SWEDEN), ABC (AUSTRALIA)

What was the problem to be solved?
The Australian Broadcasting Corporation (ABC) and Swedish Radio (SR) provide transcripts for podcasts and online audio, to improve both accessibility to audio content and discoverability.

What was done?
ABC has developed its own in-house AI transcription platform, which is integrated into its content management system. This platform uses machine learning to continually improve transcription quality over time. The process aims to overcome some of the shortcomings of international transcription service in accurately representing strongly accented Australian content and overcoming the use of nicknames and indigenous names within it.

ABC has been piloting this platform since mid-2023 and continues to improve both its accuracy and functionality. According to the ABC, the diversity of the content has created a wide variety of challenges for the development team. ABC manually corrects all its transcripts.

For Swedish Radio, their project was particularly aimed at improving the accessibility of their content for audiences with hearing impairments.

Swedish Radio research has concluded that 1.5 million Swedes, out of a population of 10.5 million, have some kind of hearing difficulty. It is an important part of the SR public service mission to increase accessibility for these audiences to the journalism contained in the audio formats.

SR uses its own model SR STT for the transcription. This is based on an open-source model further trained by The Swedish Royal Library. The latest quality improvements have been made using the Faster Whisper model from OpenAI.

A major difference between the approaches is that in contrast to ABC, SR provides automatic transcription within their app without human supervision. The texts come with a disclaimer that they may contain mistakes.

According to Christian Gillinger, who is responsible for accessibility at SR, this approach has been carefully considered:

"Through our dialogue with organizations representing people with hearing loss we know that text supported audio is favoured even if it contains errors. People want and deserve to have equal access to our content."

What was learned?
Transcription accuracy has significantly improved in the last 12 months according to both companies but there is still a large amount of manual work required to sub-edit and format this output into a publishable state.

"As we expand this pilot to include daily radio programmes and more kinds of content, we continue to discover new workflow challenges we had not anticipated," says Tanya Nolan, Managing Editor, Audio News and Current Affairs at Australian Broadcasting Corporation (ABC).

"Even in this early state, we do see evidence that the accessibility support is valued by our audience. When transcripts are missed, or production has ceased for periods, listeners will contact us to complain. We have not been able to find any evidence of search optimization (SEO) benefits or greater discoverability to date. But our transcripts are currently only available on our website, and on the ABC Listen app."

At SR they discovered some secondary benefits for the audience in searching their content as the technology enables users to navigate through longer chunks of audio.

Particularly suited for:
ABC finds that the transcripts have been most useful for short, pre-recorded reports and for shorter one-on-one interviews. Content featuring many voices, crosstalk, more production elements like music, sound effects and archival tape can be a lot more work.

SR notes that generative AI models like Whisper improves the quality of the transcripts particularly for live broadcasts. The traditional models work best for pre-recorded podcasts and programmes.

Resource needs:
Creating this platform was largely the work of the ABC's AI development team. They estimate that the time taken to produce a transcript is roughly one and a half times the length of the content.

SR has a whole development team dedicated to text where automatic transcription is one of the main tasks.

Illustrations:

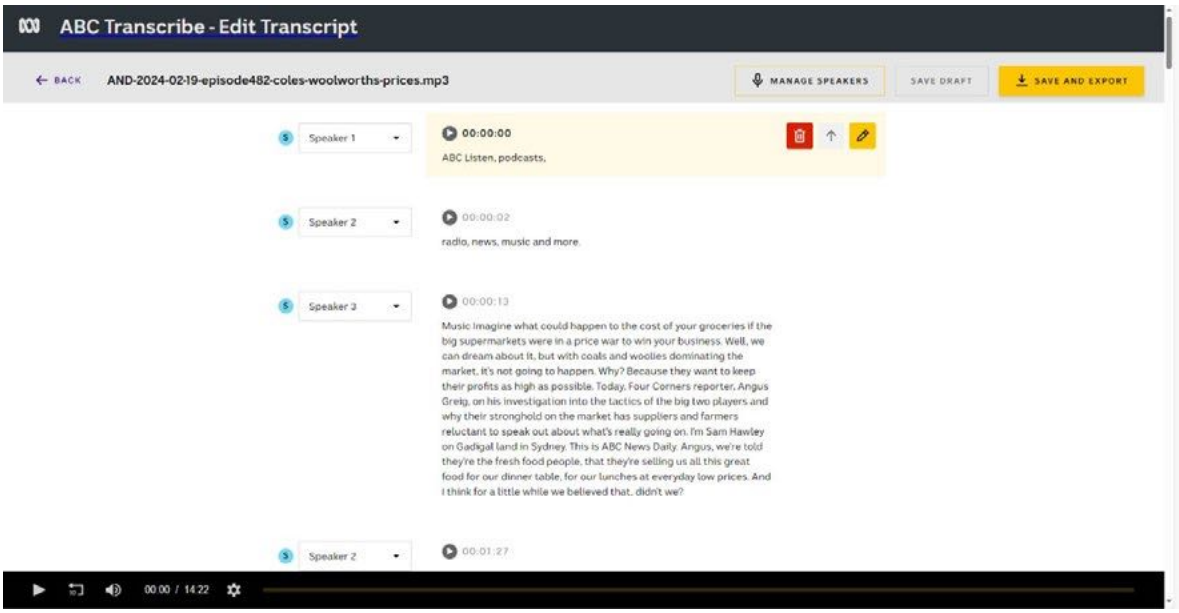

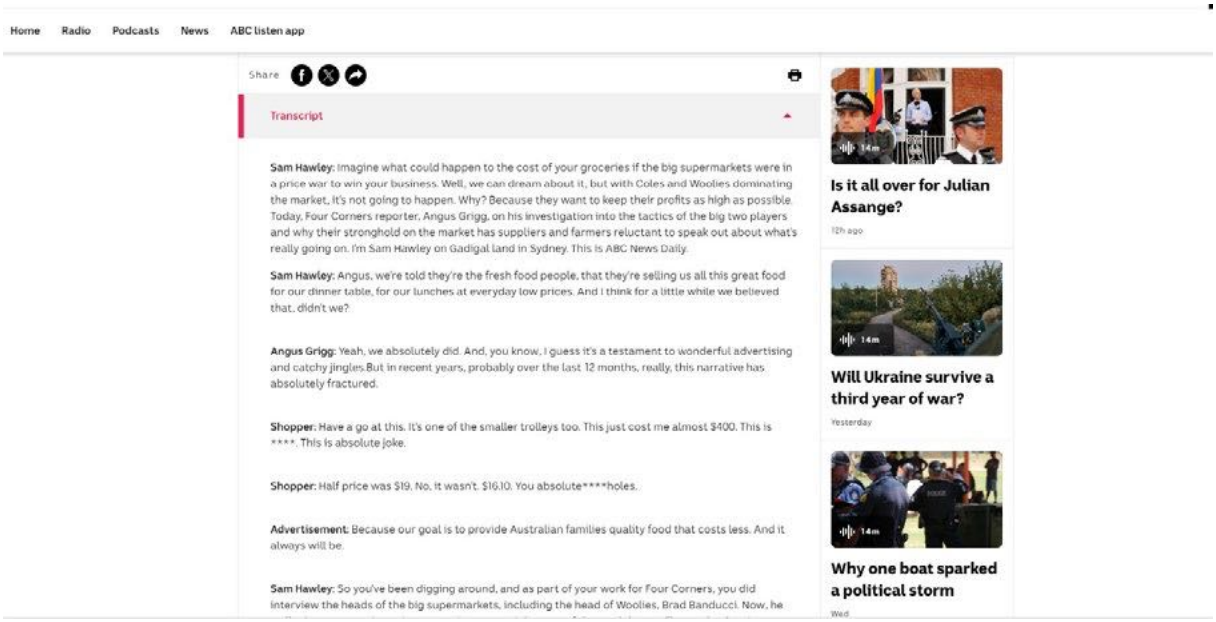

The ABC Transcribe tool where you upload the audio.

Published transcripts are found below the player of the ABC audio episode pages online. When users click the drop down arrow they can access the transcript.

Additional example: https://www.abc.net.au/listen/programs/abc-news-daily/nikki-haley-vs-trump-s-republican-party/103529548

Contact person ABC: nolan.tanya@abc.au.net
Contact person SR: christian.gillinger@sr.se

CASE

POP-UP NEWS SERVICE IN UKRAINIAN USING AI TRANSLATION

Yle (FINLAND)

What was the problem to be solved?
Russia's attack on Ukraine has forced millions of Ukrainians to flee their homes. According to estimates, around 62,500 refugees are living in Finland as of 2023. Finland's national public media company Yle wanted to facilitate the integration of Ukrainians into Finnish society by providing a news service in their native language.

The Ukrainian news page was added to other services for language minorities including news in Swedish, Sámi, Roma, and Karelian, as well as in English, Latin, Russian, and plain Finnish.

What was done?
The news is based on the Finnish and sometimes English content from Yle News, an approach that enabled the service to be launched in just a couple of weeks after management's decision. The news is translated into Ukrainian using machine translation services provided by EBU's Eurovox tool. Translations are always checked by a Ukrainian-speaking journalist before publication.

The most important Ukrainian news is published in text format every weekday on https://yle.fi/novyny and distributed on Telegram https://t.me/ylenovyny

The news selection focuses on the main national news as well as stories of special interest to Ukrainian refugees. The aim has also been to provide not only news, but useful information.

What was learned?
For Yle the initiative is a clear indication that AI-based translation can provide content in an increasing number of languages by reducing and speeding up translation work.

"We have tested machine translation services with good results. But the computer doing the translating needs the help and supervision of a human," says Aki Kekäläinen, Head of Democracy and Digitalization at Yle News Lab.

Particularly suited for:
News services for underserved languages. Speedy pop-up projects when resource efficiency is key.

Resource needs:
At least one journalist proficient in the relevant native language. Support in news selection by the larger newsroom. Tech competence to refine the transcription and use more sophisticated language models.

Illustrations:

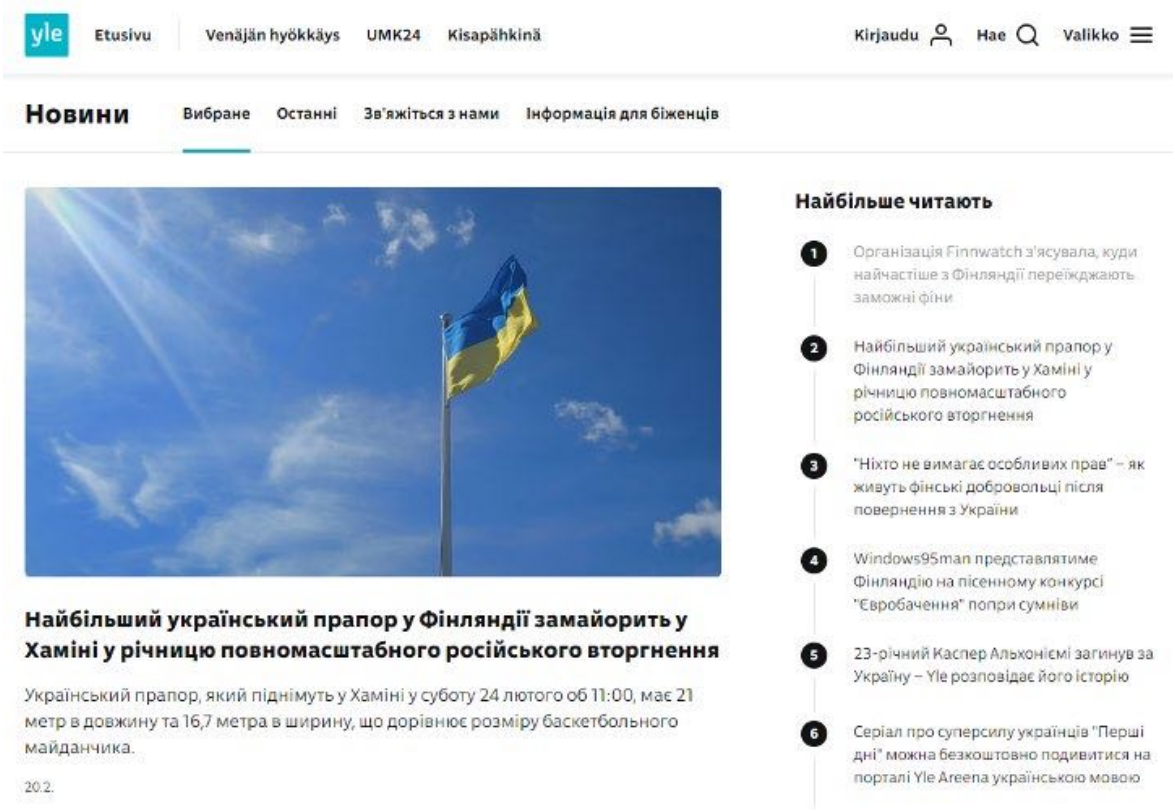

Link: https://yle.fi/aihe/a/20-10002751

Contact person Yle: tove.myllari@yle.fi

CASE

CLONED VOICE TO READ TEXT ARTICLES

AFTENPOSTEN (NORWAY) & SVD (SWEDEN)

What was the problem to be solved?
Newspapers naturally focus on text. However, in a multimedia landscape, the audience expects to access content in a variety of ways. Both *Aftenposten* and *SvD*, digital-first newspapers owned by media group Schibsted, wanted to increase accessibility and reach by enabling users to listen to the text-based articles. Value for uses can be added by delivering content in audio format. Audio formats can be consumed at any time of the day. Content in audio format also increases accessibility.

"Some users simply prefer listening rather than reading, some users also have reading challenges. We also know that a younger audience prefers listening to content rather than reading. An AI-voice can deliver on these drivers in a cost-efficient way," says Ebba Linde, Senior Product Manager at *SvD*.

What was done?
Aftenposten's pioneering voice clone is based on the voice of well-known podcast presenter Anne Lindholm. She went into the studio and spent 34 hours recording 6,812 sentences extracted from relevant news articles. Together with an external technology partner and through collaboration with a linguist they saw they are improving the quality of the voice on a weekly basis. *Aftenposten* is also working on a male voice using slightly different technology from that used to clone the first voice.

"It all started as an accessibility project for kids and youngsters in Norway. *Aftenposten*'s Junior edition is used in Norwegian schools, but a requirement from the school authorities was that it should be just as easy to listen to as to read all the news content," Lena Beate Hamborg Pedersen, Senior Product Manager *Aftenposten*, explains.

From the spring of 2023, the voice application has been rolled out on almost all *Aftenposten* articles. By pressing a play button on the website or in their app the consumer can listen to the full article. The success of the experiment has led to other newspapers in the Schibsted following suit.

Swedish *SvD* has made great progress in cloning a voice having learnt a lot from the *Aftenposten* experience. *Aftenposten* identified some important factors: focus on the best user experience when using the text to speech feature, carefully select what formats to start with and how the interface looks. Swedish online giant *Aftonbladet* is also planning to

introduce its own cloned voice application. In all cases the aim is to be transparent about the fact that the voice is synthetic.

What was learned?

Aftenposten noted a steady increase in listeners but not a high proportion of use compared to its reading numbers. By January 2024, around 10% of subscribers were listening at least once a month. They concluded that a single audio article lasting 3-4 minutes wasn't enough, as research showed people wanted to listen longer. *Aftenposten* addressed this by automatically playing a new, related piece after the first one and offering a curated playlist of interesting audio articles.

Some advice from Schibsted: "The industry evolves at a fast pace, new providers emerge quickly. Have an open mind on what partners to work with as technology is in rapid development."

Particularly suited for:

Converting text articles into natural sounding audio. The potential uses of this application for publishers is expanding as the voice clone quality gets better.

Resource needs: A journalist ready to clone their voice. Partnership between journalists and the in-house development team and, if needed, external expertise for oversight and quality control. A selected technology partner that can perform the voice clone.

Future plans:

In 2024, *Aftenposten* will make it easier to use and discover the text-to-speech opportunity. They will improve the information given to new and existing users about all their audio offerings. They will also focus on making an audio-friendly version of the original text and further improve the quality of the voice.

Illustration:

Contact persons:

ebba.linde@schibsted.com
lenabeatehamborg.pedersen@schibsted.com

CASE

BROADENING THE REACH OF INVESTIGATIONS USING AI CLONED AUDIO

AGÊNCIA PÚBLICA (BRASIL)

What was the problem to be solved?
At Agência Pública, the Brazilian investigative journalism agency, the use of a synthetic voice to narrate certain of its articles had several aims, including the chance to get hands-on experience with AI-based tools. But the main aim was to widen the impact of the long-form investigative reporting that Agência Pública produces by making it more accessible and digestible. Not everyone has the means, ability, will, or time to read long or complex material, even if it concerns them directly.

Agência Pública had already experimented with the technology when preparing an investigation into a corrupt politician who was abusing his power to push people off their land. They sent an audio version of the article, created with machine reading technology, to one of the affected farmers they had interviewed, to help him respond to the investigation. This proved to be a hit – it was listened to by the farmer and others in his local community and shared with others affected by the story.

This success alerted them to the opportunity to reach a broader audience by using a new format and creating a new channel for distribution.

What was done?
Agência Pública is now producing audio versions of several stories. The articles are 'read' by a synthetic version of the voice of one of their reporters using technology from ElevenLabs.

What was learned?
They found that the ElevenLabs application, while being one of the market leaders in the field, had more difficulty with Portuguese that it does with English. This meant Agência Pública had to incorporate 'tricks' in parsing the original article to make it more easily digestible by ElevenLabs. Still, the process of creating audio versions of articles cannot be fully automated, as the output needs to be checked by a human, and sometimes sections need to be reproduced to smooth out glitches.

According to Babak Fakhamzadeh, Developer at Agência Pública, the synthetic voices create a good, functional standardization. "It doesn't mean it's a huge time saver. However, costs, compared to producing human-read versions of articles, are still lower," he says.

Particularly suited for:
Agência Pública focus on creating audio versions of stories that are more 'evergreen' so they can be shared over a longer period of time, and that have a strong 'human' touch so people will be compelled to share the material.

Resource needs:
After the initial creation of the synthetic voice, someone needs to produce the audio versions of articles.

Future plans:
According to Agência Pública, it is still too soon to assess how effective the production of these audio versions is. They have not yet decided whether to expand the project.

Illustrations:

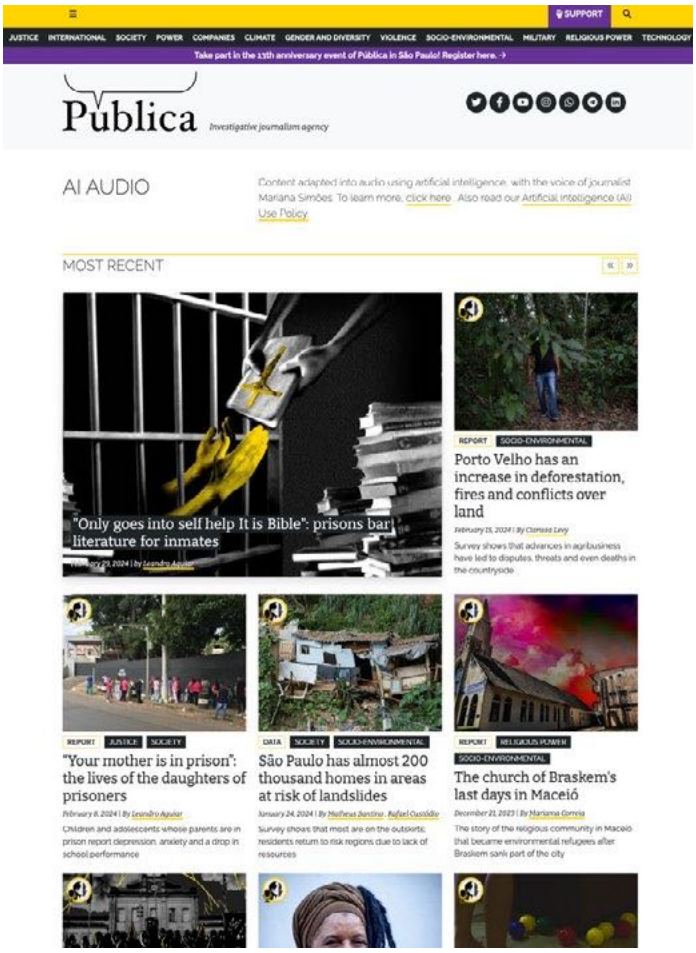

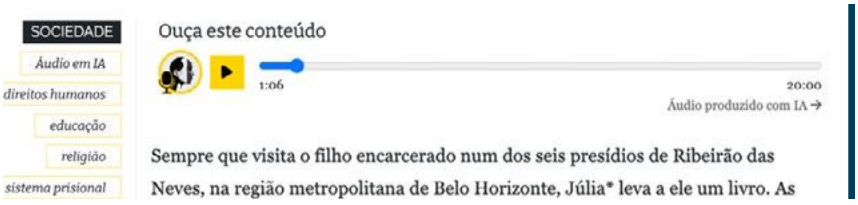

Contact person: babak@apublica.org

CASE

'TELL ME MORE', USING AI TO EXPLAIN COMPLEX TERMS

BBC (UK)

What was the problem to be solved?
BBC News articles need to be accessible to all audiences, whether they are new to or familiar with the topic. When readers come across unfamiliar terms which are not clarified, they may feel left out, which can lead to a loss of engagement.

What was done?
A prototype called 'Tell Me More' was created. It started from an ideation session aimed at providing explainers and context for long running news items. The prototype uses automation and machine learning, particularly GPT-3, to identify complex terms within an article, and generate draft explanations. It sources the information for the explanations from relevant published BBC content. These drafts can then be refined by journalists and included in an explainer block giving optional context to the articles.

What were the results?
The prototype was successful in identifying complex terms within articles, and drawing on relevant published BBC content to generate draft explanations to specific character lengths .

What was learned?
Editorial feedback suggested that when used at scale, features like this would enable more consistency across the BBC's coverage, but that it is critical to provide visibility of the sources used to generate the explanations. All the source material was indeed already published and verified BBC content and the final output had editorial supervision. Audience feedback indicated that the presence of the short form explanations on the page helped to make news stories feel more accessible and useful.

Particularly suited for:
The prototype is particularly suited for the coverage of complex and current topics that are relevant to the target audience, such as the ongoing cost of living crisis, or information relating to democratic processes during elections.

Resource needs:
For the next steps, required resources would include full analysis on the quality of the explanations, audience validation at scale, and development of formal editorial oversight processes.

Contact person:
miranda.marcus@bbc.co.uk
https://bbcnewslabs.co.uk/projects/tell-me-more/

The prototype identifies complex terms from articles and generates explanations for them.

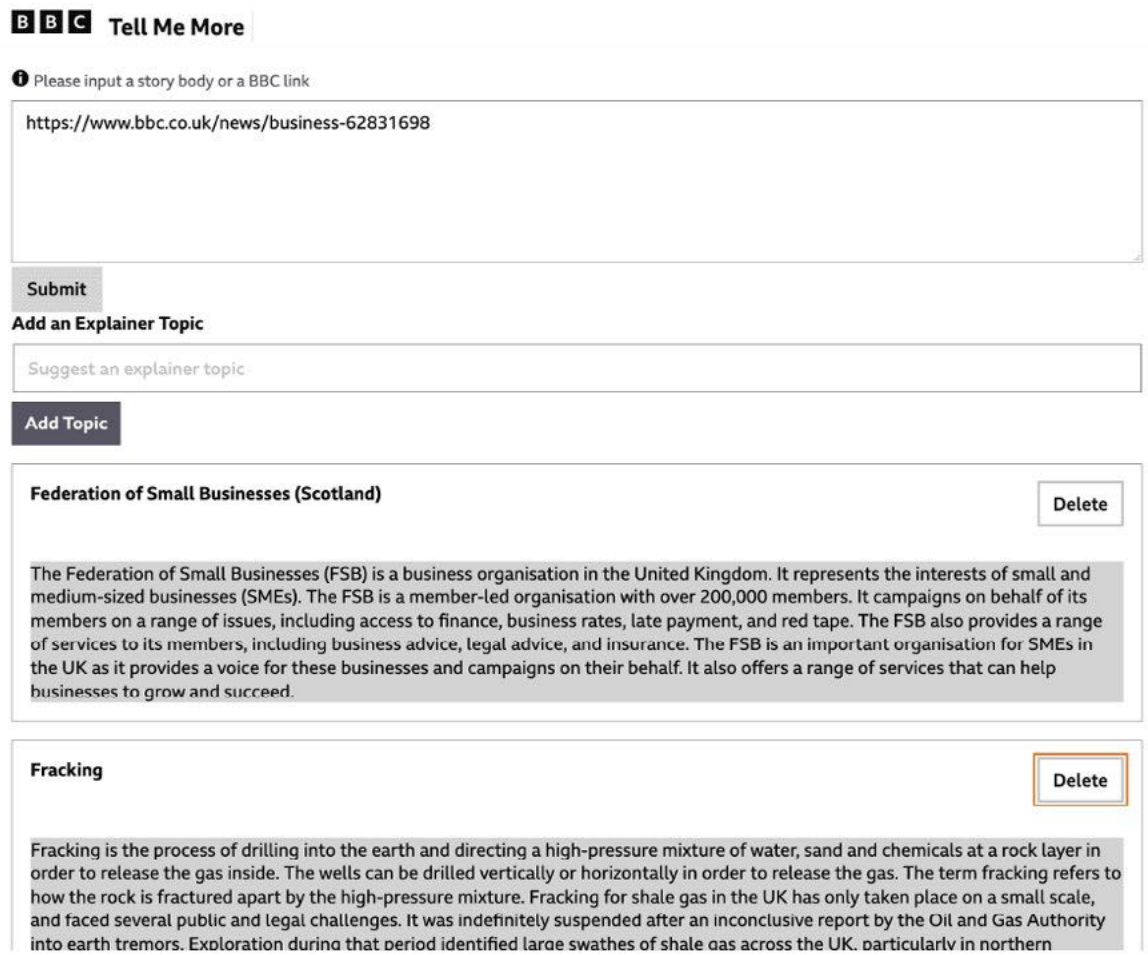

PART 2:
AI support for jobs in the newsroom

Increasing efficiency is among the major hopes of publishers and newsroom managers alike in a resource-strapped media environment. And there are many 'jobs to be done' in news organizations of which hardly anyone is fond. The hope is that generative AI can help with copy editing, headline writing, search engine optimization, transcription, churning out social media posts and moderating comments. And as a consequence, free up journalists to invest more time in original reporting – provided that cost-conscious managers are supportive. Many journalists can't wait to have visual tools create illustrations for topics that don't lend themselves naturally to original photography. The struggle of business desks to avoid images of men in suits populating web pages and print products is legendary.

Many news organizations already use generative AI around a variety of activities that affect language, be it transcription, translation, or voice-overs. Anne Lagercrantz of SVT calls them 'everyday AI', applying the term from the consultants at Gartner to the news context. She refers to tools that don't transform what news organizations do but make existing activities easier – as opposed to game-changing AI.[1] She says: "There is so much potential, for example with subtitles. We are also producing investigative news with synthetic voices now instead of using an actor, with great results, though it is not saving time yet." But using synthetic voices can also help battle the negative side effects of AI. As France Télévisions has experienced, ill-meaning actors can use generative AI to identify sources in investigative research whose voices or faces have been technically anonymised to protect their identity. Using synthetic voices, in contrast, protects interviewees from being found out. (Read #8, Unmasking sources: how AI is undermining anonymization and how to combat it, page 78)

Some news organizations are attempting to minimize risk when their staff experiment with generative AI without knowing the origins of the content, its reliability and truthfulness, copyright restrictions or data protection issues. This is why the EBU has developed *Neo*, a closed base for low-risk research. (Read Case #9, How generative AI can increase database value: the Neo project, page 80) Similarly, Danish publisher JP Politikens HUS has created Magna, a system containing a database with content of their publications that makes it safer for their employees to work with generative AI.[2]

Having LLMs edit, rewrite, and shorten stories still requires some competencies. For example, a team at German Ippen Media saw off-the-shelf generative AI struggle with the concept of quotes and tried several options to teach it how to respect their integrity.[3] However, the AI in Journalism Challenge 2023 showed that even teams starting out with limited technology knowledge can achieve considerable progress in a reasonable amount of time by simply practicing and iterating. As David Caswell concluded in the report: "This transformation from novice to relative expert seemed to be the default experience across almost all the teams, and it suggests that the opportunity to change how news providers think about and apply AI may be easier to realize than might be assumed. It certainly appears to be far easier, for example, than thinking about and applying the previous generation of AI, based on highly technical concepts, data science talent, and engineering resources."[4]

Some news organizations are already working with AI assistants which are trained in their internal style guides. Finnish Helsingin Sanomat, for example, is testing a personal AI assistant that helps its reporters enhance their writing. It is named Hennibot – reminding staff of a particularly diligent editor everyone in the newsroom respects.[5] Verena Krawarik, Head

[1] Gartner provides a framework for analysing AI-related activities here: https://www.gartner.com/en/newsroom/press-releases/2023-10-16-gartner-says-cios-must-prioritize-their-ai-ambition-and-ai-ready-scenarios-for-next-12-24-months, retrieved on 30 April 2024.

[2] Watch this video recorded at the Nordic AI in Media Summit where Tore Rich presents Magna: https://www.youtube.com/watch?v=hh4db4SlWIk.

[3] Anika Zuschke, Citation 5.0: Exploring why AI struggles to grasp the essence of quotes - and learning from it", AI@Ippenmedia, 14 February 2024. https://medium.com/ai-ippen-media/citation-5-0-exploring-why-ai-struggles-to-grasp-the-essence-of-quotes-and-learning-from-it-ddc7d590d179, retrieved on 2 March 2024.

[4] David Caswell, "AI in Journalism Challenge 2023", pp. 31-32.

[5] Conversations with Esa Mäkinen in Oxford, January 2024, and Copenhagen, April 2024.

of APA's Medialab says: "The biggest topic for us is the text assistant. It doesn't write stories; we designed it as an assistant. It helps with title suggestions, lead suggestions. We deliver it in such a way that suggestions come with it. The tool also suggests summaries, abstracts, and proofreading. Our media clients also use the character-accurate shortening to prepare online editions for print. That took many hours of work at prompt engineering. The first feature was the production of social media posts, which is also used by our marketing departments." Radio France uses generative AI mostly for digital editing, including headline writing and automated transcription. (Read Case #10, AI transcription and metadata for digital news editing, page 82)

Dmitry Shishkin of Ringier thinks that generative AI could help tremendously with adding perspectives to news stories: "The big effects will be on workflow aspects. You will be able to create a co-pilot at your side that will help you to become a better editor. You won't replace an editor, that will never happen. But in the morning meeting, you will already be prompted by your co-pilot to have better ideas. Imagine also, you are writing a story, and your AI-assistant will help you transform that story to another user need." Danish *Politiken* and Swedish *Expressen* are already using a whole array of generative AI powered newsroom tools that aim to make life easier for their editors. Jakob Wagner, Head of Digital and Editorial Development at *Expressen* says: "The initiative shows that AI can significantly boost journalistic creativity and efficiency, provided that the tools are user-friendly." (Read Case #11, Newsroom tools powered by generative AI, page 85)

And then there are the tasks nobody enjoys doing. Uli Köppen of German BR describes an AI tool that filters user comments and displays the most useful ones. Editors are particularly happy. "That was a very fast prototype that has been useful and has helped the newsroom for some time already." Product Manager Jörg Pfeiffer adds: "These tools show the potential for AI-supported journalism that fosters dialogue with the community of users and positively contributes to the editorial output. They allow value to be extracted from a large volume of material by integrating innovative solutions into existing workflows." (Read Case #12, "Dein Argument": AI-powered community management, page 88)

Making it easier for journalists to do their jobs also means reducing the number of interfaces they have to engage with. Manuela Kasper-Claridge explains how AI already informs Deutsche Welle's content management system: "As an international broadcaster that publishes in 32 languages, the rapid development in AI-supported translation and voice recognition is very exciting. It has the potential to save us a lot of time translating and revoicing our journalism from one language into another. Still, these translations and voiceovers would need to be checked by an editor. We developed an AI-powered content adaptation platform, plain X, which helps with this. It is integrated into our editorial systems, bundles various tools in one interface and offers lots of options for transcription and subtitling of videos, as well as other AI-based services."

This can even work with little-spoken languages – like the native language in Greenland. In the North American autonomous territory of Denmark, law mandates that everything be translated between the native language and Danish. At the Nordic AI in Media Summit 2024, Masaana Egede and Lars Damgaard Nielsen described AI that delivers just that – at a fraction of the cost and time needed before.

Styli Charalambous of *Daily Maverick* is confident that generative AI will do a great deal to help his multilingual staff with news production in their non-native English. "We have 12 different official languages in South Africa. It would be great to provide our staff with something like an on-demand personal assistant editor that they could query and ask questions of the copy that they generated. Like: am I missing anything, are there any holes here? Does it comply with the press code? Would there be additional concerns? I think that's ultimately where we want to get to: that every person in the organization has a personal assistant that is tailored to their job."

How generative AI will affect headcount in news organizations and which jobs are up for elimination is impossible to predict. In theory, entire TV news programmes will be able to be run using avatars instead of news anchors. One example is Channel 1, a prototype AI news channel from veteran producer and director Scott Zabielski and tech entrepreneur Adam Mosam that was launched as a proof-of-concept in late 2023.[6] The same holds true for radio. The start-up Futuri AudioAI produces AI-generated content for small radio stations to fill content gaps. But it is not only about cost-management, CEO Daniel Anstandig claims. AI tools also help to better connect with people: "We can see what's trending in any local market and predict what audiences are going to be talking about over the next 4–6 hours." (Read Case #13, AI generated radio - from voices to content, page 90)

In public service media, approaches for 'replacing' human voices with avatars vary. Spanish RTVE, for example, uses an AI-generated avatar host for its podcast Hiperia.[7] In an April 2024 EBU strategy document, Swiss RTS's radio channel Couleur 3 claimed that it had "made history with a world first: a full day of radio entirely generated by AI."[8] In contrast, Sweden's Sveriges Radio has a 'no voice clone' policy. In SR's new guidelines on generative AI it is stated that the audience should be able to trust that human voices heard on SR's platforms should always come from real people. However, SR adds that that principle could be changed with further AI development and changes in audience attitudes.

In these cases, news organizations need to carefully think about the needs of the particular programme. Audiences might not mind an automated, avatar or voice-clone driven product if it is for simple news updates or night-time programmes that don't require much more than a few headlines, traffic, and weather updates. But for other programmes, the personality of the anchor and the human touch is core to why people tune in. The job then is rather to provide a feeling of connection and trust, independent of the content of the news. This means that public service media in particular have to rethink the traditional approach of having presenters – to put it bluntly – behave like avatars. Particularly in major news programmes, the idea prevails that presenters should act in an authoritative and distanced manner to convey impartiality. It might become more important to showcase personalities that are very much alive, engage in conversations, or ask interview partners critical questions, to set them apart from the avatars that are likely to turn up everywhere.

Diminishing resources make it increasingly necessary for newsrooms to focus on what works rather than what journalists think works – a luxury that has defined the industry for a long time. As Felix Simon writes: "What I believe we are witnessing is – to a degree – a further rationalization of news work through AI, as work processes that traditionally relied on human intuition are increasingly becoming suffused with or replaced by a technology imbued with ideas of rationality, efficiency, and speed – some of which it does indeed deliver."[9]

[6] David Cronshaw, "Channel 1: Are GenAI avatars the future of video news?", 17 December 2023. https://www.linkedin.com/pulse/channel-1-future-tv-news-david-cronshaw-lwb0f/, retrieved on 2 March, 2024.

[7] See EBU, for members only: https://www.ebu.ch/video-talks/membersonly/2023/05/hiperia-the-iaai-presenter-radio-3-of-rtve

[8] Examples for the use of AI by public service media can be found in this members only document "Public Service Media AI Strategies" published in April 2024: https://www.ebu.ch/guides/membersonly/report/public-service-media-ai-strategies

[9] Felix Simon (2024), p. 37.

CASES

CASE

UNMASKING SOURCES: HOW AI IS UNDERMINING ANONYMIZATION AND HOW TO COMBAT IT

FRANCE TÉLÉVISIONS

What was the problem to be solved?
Editors at France Télévisions were aware that their sound anonymization methods were no longer secure, so they implemented countermeasures. Then they discovered a new challenge to their visual anonymity processes. Up to now, the blurring techniques used to protect the identity of sources was considered effective and secure. But with the development of new AI technologies, they discovered through testing and consultation with experts that the security could be breached. According to one of the experts: "AI technology will be able to reconstruct realistic faces. The danger is tangible; the first capable algorithms are already publicly available. Training AI is straightforward, and advancements could be swift. Constant vigilance is necessary. is presents an evident challenge for France Télévisions," according to Titus Zaharia, Head of ARTEMIS (Advanced Research and Techniques for Multidimensional Imaging Systems) at Telecom Sud Paris.

What was done?
Immediate creation of a best practices charter for all in-house and external productions to ensure anonymity from filming to post-production with advanced masking techniques.

Development of a compulsory e-learning module on best practices for approximately 5,000 internal staff members, representing over half of the France Télévisions' workforce.

Prompt restriction of access to archives and replays internally and throughout INA (the national archive for broadcasting in France), which has ceased all sales of content that requires anonymization.

Initiation of research into image technology solutions that will reinforce the best practices mandated by the newly established charter.

What were the results?
Today, France Télévisions is confident that the implementation of these best practices has improved the protection of the anonymity of these sources.

"Editorial managers will therefore have to assess risks as they always have, with the support of the legal department, but now with the advancements of artificial intelligence in mind. Gaussian blurring may suffice if the assessed risk is low." (Excerpt from the France Télévisions Anonymization Charter)

What was learned?
In addition to the technology challenges of protecting anonymity which are still being addressed, France Télévisions research has concluded that sources can be identified by factors other than facial features or other distinctive elements. They must now consider environmental details, specific gestures, and patterns in speech, among other subtleties.

Particularly suited for:
The issue is particularly fundamental in the need to the anonymity of people whose lives are at risk if they are exposed, for example, in times of war and in situations where time for filming and post-production is reduced due to urgency.

Resource needs:
News MediaLab for overarching project management.

- Digital Innovation Department responsible for testing through its R&D team.
- Consulting with external academic experts for authoritative insight.
- Information Department Head charged with drafting the policy and endorsing the Online Learning module.
- Enactment of protocols both internally and in collaboration with external entities.
- Liaison with external parties, including INA.
- France Télévisions University tasked with the development and oversight of the Online Learning module.

Next steps:
Coordination with France TV Technologies and collaboration with Radio France to streamline and scale up the solution.

Illustrations:

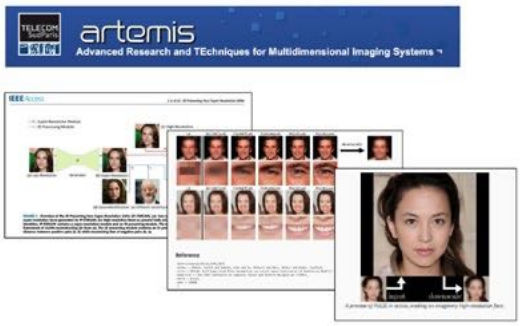

Contact persons:
Christophe de Vallambras, News MediaLab France Télévisions
christophe.devallambras@francetv.fr

Pascal Doucet-Bon, Deputy Director News France Télévisions
Pascal.doucet-bon@francetv.fr

CASE

HOW GENERATIVE AI CAN INCREASE DATABASE VALUE: THE NEO PROJECT

NEO PROJECT, EBU

What was the problem to be solved?
The EBU wanted to get more value from its database of 2,500 news stories and enable quicker and more effective searches by journalists. The goal was to create a more natural way for journalists to engage in a dialogue with the database to answers specific questions.

What was done?
EBU Technology and News developed a prototype using the News Pilot/ European Perspective database of news stories provided by EBU members. They then applied the EBU's recommendation algorithm, PEACH, with a vector database to facilitate the rapid semantic retrieval of elements, to identify the content most relevant to answer the search questions and present it to the user. Then, using the power of large language models (OpenAI's GPT), the system formulates an answer for the user. This occurred within the closed and secure EBU system using the verified news content of its members thereby reducing error and inaccurate responses.

What were the results?
There is a working prototype which responds very quickly to the users' requests. The quality differs significantly depending on the nature of the query (the prompt). The user can make up their own mind about their research by exploring the retrieved stories.

What was learned?
The project was deemed useful in exploring the potential use of generative AI within the context of public service media while maintaining their core values and trust requirements. Project leader Sebastien Noir explains: "As the world evolves, we observe increasing demands from third-party organizations and major technology companies for reliable news content. This content may be presented in a modified form, including transformations, summarizations, or adaptations, as part of the discussions they aim to facilitate for their users. The EBU and its members will have to define if they see this as an opportunity to connect with these audiences – or a threat since attribution, prominence and trust might be undermined. Building this prototype was helpful to understand new services that could be developed, and how companies will be using these technologies to transform the user experience."

Particularly suited for:
Operating within closed databases, as this can reduce of hallucinations in responses and limit the exposure of proprietary content in terms of both input and output.

Resource needs:
A combination of several technological components: a database of indexed content, the EBU's proprietary recommendation system, PEACH, an LLM API access, and editorial teams to test the solution.

Next steps:
In future developments, implementing an agnostic solution that can access multiple LLMs as required could be useful. The RAG approach is expected to facilitate this kind of change with relative ease and limits the data exposed to OpenAI. Later, another model may be employed, particularly an open-source model, to ensure full mastery of the process and data.

Illustration:

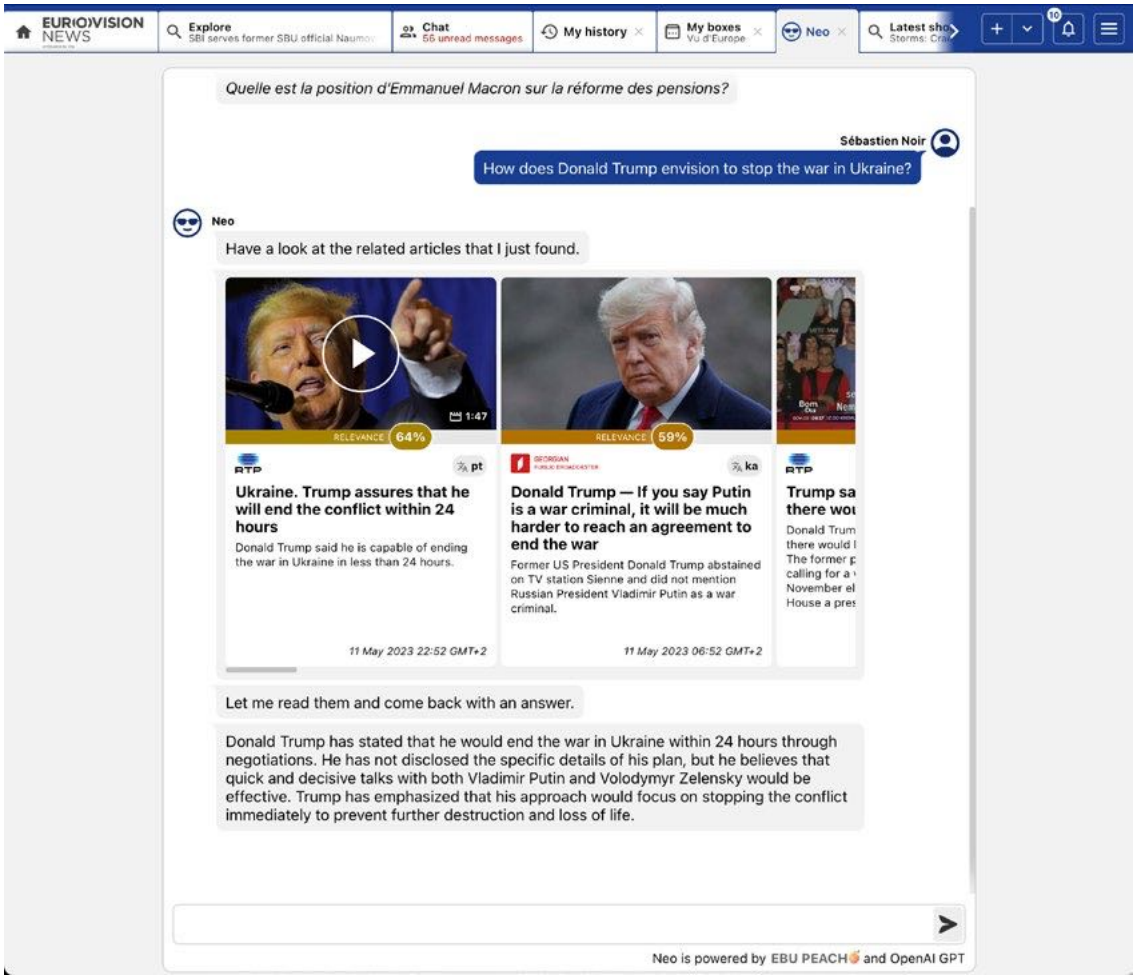

Contact person:
Sébastien Noir, EBU noir@ebu.ch

CASE

AI TRANSCRIPTION AND METADATA FOR DIGITAL NEWS EDITING

RADIO FRANCE

What was the problem to be solved?
Public broadcaster Radio France (RF) aims to streamline their news editing process by empowering digital editors with smart features including audio content transcription. The transcription needs to be easily accessible and ensure a reasonably high-quality level. Swedish Radio (SR) have adopted a similar system for their digital editing.

Additionally, RF seeks to automate tagging, thematic categorization, and identification of individuals featured in each article. Currently, manually searching and attributing the right metadata in RF's extensive database is a very time-consuming task.

What was done?
To achieve this, RF has opted to deploy the open-source version of Faster Whisper on its hosting infrastructure. This choice minimizes costs for the tool which are basically just inferred by data usage expenses. This deployment strategy ensures the required monitoring and confidentiality, as the content remains inaccessible to OpenAI.

SR uses its own model, SR STT, which is based on an open model, retrained for Swedish by the Royal Library in Stockholm and enhanced with Faster-Whisper-Large v2.

RF is progressively rolling out the functionality and has begun with the audio content of France Inter, Franceinfo (100% news), France Culture (cultural channel), France Musique (classical music) and Mouv' (the urban channel for age group 15-24). Transcripts from six local radio stations of the France Bleu network (44 local radio stations overall) is underway. Swedish Radio is transcribing all news clips (360 per day) and all podcasts that do not have external copyright holders.

Key features include:

- **Transcription and Speaker Detection:** AI model Faster Whisper handles both transcription and speaker detection.
- **Automated Proposals:** After analysing the transcript, the system generates automatic suggestions for tags, themes, and relevant individuals.
- **Article Enhancement:** RF is experimenting with a customised GPTs built on OpenAI technology to propose article titles, summaries, and longer pieces.

In the SR tool the text is highlighted line by line when playing the audio.

What was learned?

The transcription feature has become indispensable for RF's digital publishers. Additionally, there has been positive feedback for the functionality of the automatic proposal component.

"The last feature is currently limited to a select group of digital publishers, but already at early stages it shows promising results. I also want to stress that our audio transcription is done exclusively for published audio on demand, not for audio rushes or unpublished material," says Matthieu Beauval, Director of Innovation Acceleration.

However, the current, basic prompts yield unsatisfactory results. The next step involves refining prompts to achieve better outcomes.

Particularly suited for:

The RF service is particularly well-suited for two purposes:

- **Efficient publication of long audio content:** It enables faster and more efficient publishing of lengthy audio content.

- **First text captions from unedited content:** Specifically, it is used to extract text descriptions from audio content that has not undergone any digital editing. Examples include news bulletins, shows produced by smaller teams, or those with limited digital skills.

  Resource needs: Open-source tool Faster Whisper for transcription and speaker detection. Cooperation between editorial staff and developers for model implementation and interface development.

Illustrations:

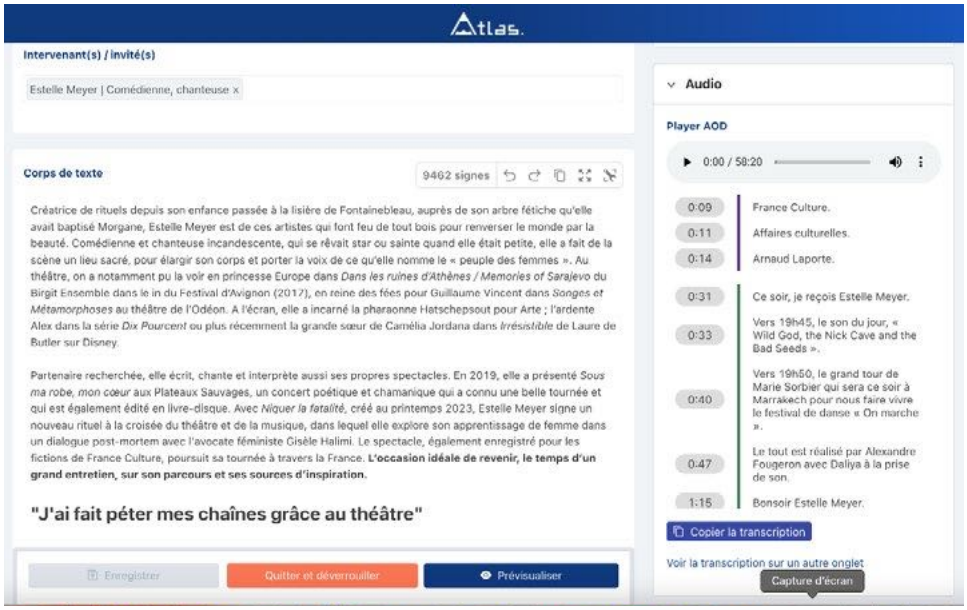

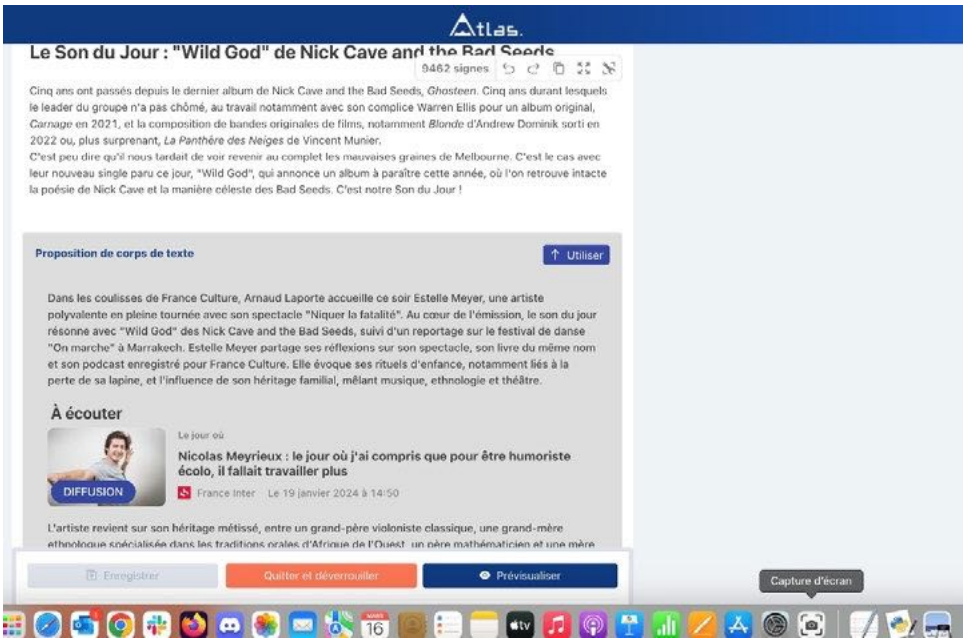

Swedish Radio's tool hitta.sr.se

Contact person RF: matthieu.beauval@radiofrance.com
Contact person SR: ann-sofie.ottosson@sr.se

CASE

NEWSROOM TOOLS POWERED BY GENERATIVE AI

JP POLITIKEN (DENMARK) & EXPRESSEN (SWEDEN)

What is the problem to be solved?
MAGNA is an online editorial tool in the newsroom that leverages generative AI to assist journalists at the Danish media group JP/Politikens Hus. The primary goal of MAGNA is to streamline work processes and unveil novel opportunities in journalism, thereby enriching the news experience. Simultaneously, MAGNA strives to familiarize employees with the use of AI and provides a safe platform for exploring AI capabilities aligned with editorial principles. Furthermore, MAGNA enables Danish media group JP/Politiken to keep data in-house and oversee and evolve the use of generative AI.

Currently a lot of media companies are launching in-house built interfaces where staff members can try and start using AI-supported journalism. Finnish Yle is an example, with its YleGPT. Swedish commercial newspaper *Expressen* is also trying to improve efficiency with a new AI-powered tool. According to Jacob Wagner, Head of Digital and Editorial Development:

"The best headlines, social media texts, and push notifications often emerge from discussions between journalists. By providing the newsroom with a creative AI assistant, they could swiftly generate such high-quality content elements."

What was done?
During late 2023, a test version of MAGNA, with editorial tools for the journalists at *Ekstra-Bladet*, was introduced. The tools enable the use of generative AI in various ways:

A) Articles: In the Articles category, MAGNA is designed to assist journalists by acting as a 'sparring partner'. To date, eight features are available, each tailored to meet specific needs that journalists have identified for their everyday tasks:

- Bullet point summary
- Headline suggestions
- Fact Box
- Shorten article
- Proof reading

- Improve article
- Rewrite to final article
- Find similar articles

B) Live coverage: MAGNA supports live news event coverage, which is an important time of engagement between the journalists and the audience, by generating key point summaries, headlines and analyzing readers' input.

C) Archive: From the use of semantic search tools, various Retrieval Augmented Generation (RAG) tasks can be performed, such as generating a fact box, a summary of the case, or a draft article. In the free text field, a journalist or editor can decide which task to perform. The task is performed solely based on the content of the articles that appear in the search.

MAGNA's features are aligned to the editorial values of JP/Politikens Hus:

- Human in the Loop for overview
- Built-in auto-evaluation for every AI-generated text. To ensure the relevance and factual accuracy of the auto-generated content, they employ automatic evaluation metrics that assess each sentence in the generated text against the most similar reference sentence.
- Adherence to strict editorial guidelines to maintain journalistic integrity and quality.

On 1 March 2024, the new MAGNA was extended to all employees at JP/Politikens Hus. The next step is MAGNA Custom where the MAGNA platform includes AI tools geared towards the specific brands in the media company.

Expressen's tool, built with tech from OpenAI, provides journalists with AI-generated content suggestions, such as headlines, push notifications, social media texts and article summaries, directly within the editorial workflow while also tailoring these to meet specific audience tones and preferences.

What was learned?
Both *Expressen* and JP/Politikens Hus say they have experienced substantial results, even if they are in early stages.

Jakob Wagner: "The initiative shows that AI can significantly boost journalistic creativity and efficiency, provided that the tools are user-friendly. Moreover, incorporating Expressen's editorial standards into the AI tool ensured the generated content remained high-quality and aligned with the newspaper's voice."

According to the Product Manager at JP/Politikens Hus they have not had great difficulties with inaccuracies: "In the initial beta versions of MAGNA, we experimented with displaying the results of the auto-evaluations to MAGNA users. However, it turned out that the generations were so effective in practice that the auto-evaluation did not find problems with factual accuracy," says Tore Rich.

He attributes this to a combination of improved language models and robust prompt engineering. Factual accuracy issues are minimized when not relying on the internal knowledge of GPT-4, but using GPT-4 to process external text.

Particularly suited for:
As a media corporation, JP/Politiken find the RAG system to be extremely valuable in granting employees easy access to our extensive article archives. The RAG's semantic search feature not only reveals previously hidden knowledge but also saves time for journalists.

Both media organizations think the internal AI 'co-workers' are particularly suited for aiding the creation processes, improving efficiency without compromising on quality or editorial standards.

Resource needs:
ML engineers, ML specialist, front-end developer, product manager. Expressen underscores that developers and machine learning engineers are essential for developing the AI tool and integrating it into the CMS. It must also include a manual or automated feedback loop to continuously fine-tune prompts and models.

Future plans:
For the moment, JP/Politikens Hus is focusing on the textual aspect, but the intention is to gradually develop the workflow and integrate various forms of media modalities, such as audio and images. *Expressen* will be improving and extend the use of the tool throughout the larger Bonnier News group.

Illustrations:

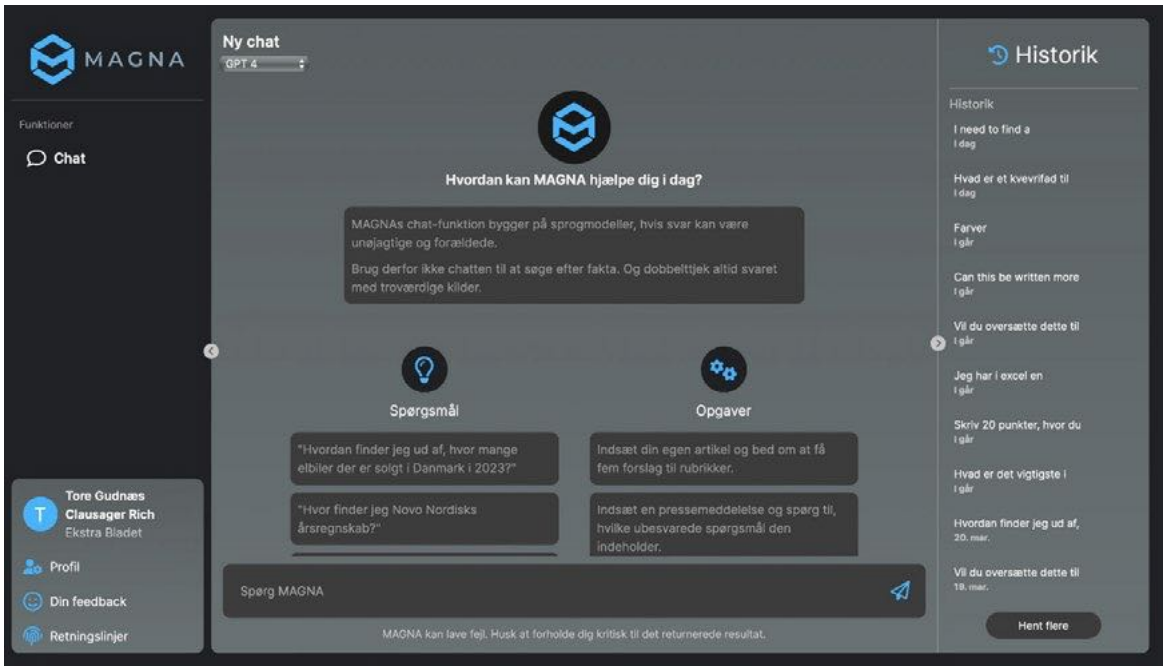

Contact persons:
jakob.wagner@expressen.se
tore.rich@jppol.dk

CASE

DEIN ARGUMENT (YOUR ARGUMENT) – AI POWERED COMMUNITY MANAGEMENT

BR (GERMANY)

What was the problem to be solved?
German public broadcaster Bayerischer Rundfunk (BR) wanted to engage more effectively with its online users and sought to enhance its community management, improving interactions and dialogues with the audience. However, with over 4,000 comments daily, it is tricky to identify the ones specifically directed at the editorial team with questions or leads and suggestions for follow-up stories. It is also a challenge for the community managers to respond promptly to user inquiries.

What was done?
The BR24 News Lab introduced a new editorial format called *Dein Argument* ('Your Argument') which selects comments made to the editorial team which could be used as user questions and comments and presented together with existing BR24 articles on its site.

To facilitate this, the AI and Automation Lab developed two tools:

1. Comment Recognition Tool: This tool scans BR24's comment section for user comments that directly address questions or nudges the editorial team. When the AI identifies such inquiries, it automatically forwards them to the *Dein Argument* editors via Microsoft Teams.

2. Answer Research Tool: The second AI tool assists the community team in swiftly researching answers to specific user questions in BR24's vast database. It subsequently suggests relevant text modules for potential responses.

What was learned?
Both AI systems are currently in an early beta stage and serve as experimental assistance tools for the team. All results require editorial supervision and undergo rigorous scrutiny. But even at this early stage, the results are encouraging according to Product Manager Jörg Pfeiffer:

"These tools show the potential for AI-supported journalism that fosters dialogue with the community of users and positively contributes to the editorial output. They allow value to be extracted from a large volume of material by integrating innovative solutions into existing workflows."

Particularly suited for:
This system is particularly valuable in using tech when collaborating with an editorial team to respond promptly to community questions in a qualitative manner, thereby enhancing user loyalty.

Future plans:
In the next phase, BR aims to transition the system from its early beta stage to regular operation. Later, further functions will potentially be included.

Illustrations:

Cover image with the *Dein Argument* icon:

Link: https://www.br.de/nachrichten/autoren/br24-dein-argument,7b847f1b-7928-42a8-a0a2-44356e8d5545


Häufig kein Notenschutz und Nachteilsausgleich

Mitdiskutieren lohnt sich: Die folgende Passage hat die Redaktion im Rahmen des BR24-Formats "Dein Argument" ergänzt. Hintergrund ist ein Kommentar des Users "KarlSchn", welche Bundesländer Regelungen zum Nachteilsausgleich in Schulen haben.

Screenshot of the *Dein Argument* notice that a user input was added to the article

(It's worth joining the discussion: The following passage was added by the editorial team as part of the BR24 format *Dein Argument*. The background is a comment by the user KarlSchn about which federal states have regulations on compensation for disadvantages in schools.)

https://www.br.de/nachrichten/wissen/dyskalkulie-was-bei-rechenschwaeche-im-gehirn-ablaeuft,U6Lnt2z

Contact person: joerg.pfeiffer@br.de

CASE

AI GENERATED RADIO – FROM VOICES TO CONTENT

FUTURI (USA)

What is the problem to be solved?
In an age where the appetite for content is never-ending but where budgets and teams are shrinking, the solution for an increasing number of media companies and creatives - from CEOs at broadcasting companies to screenwriters - is to use AI in their workflows.

"AI is not a replacement for human creativity. It's a tool to augment it - and manage the cost of content," says Daniel Anstandig, CEO of Futuri.

Futuri developed RadioGPT (now Futuri AudioAI) for this purpose - to help overworked creative teams with limited budgets produce more content. The product uses AI-generated voices to develop and read scripts with hourly updates on weather and other news. Some radio stations are using it to fill in segments for their on-air personalities using their own cloned voices.

Futuri AudioAI addresses two key challenges:

1. Creating personalized content at scale.

2. Providing live and localized content 24/7.

What was done?
Futuri AudioAI incorporates natural language processing and machine learning tools as part of its AI driven solutions.

Here's how it works:
Monitoring Local Trends: Futuri's TopicPulse uses Machine Learning to monitor real-time trends on Instagram, X (Twitter), Facebook and over 250,000 news sources.

According to Daniel Anstandig, their aim is provide news that interests the audience.

"We can see what's trending in any local market and predict what audiences are going to be talking about over the next 4–6 hours," Daniel Anstandig claims.

AI-powered content generation:
Based on what is currently trending, LLMs are used to create content that is fit for broadcast. When the service initially launched, GPT3.5 and GPT4 were used, now Futuri tunes to multiple LLMs. In its purest form, the service can be fully automated directly to the audience, which is challenging for publishers that stress editorial supervision of all published content.

Voice cloning:
The Futuri AudioAI tool can also replicate the voices of on-air hosts in a way that sounds humanistic and natural on the air.

What was learned?
Linguistic Diversity: Every language possesses its unique characteristics, including rhythm and intonation. These elements are important in conveying emotions, emphasis, and intent in speech, making them fundamental in AI voices.

Emotional Intelligence:
Teaching AI to distinguish between questions and statements is just the beginning. The ability to detect and express sarcasm, as well as to indicate the speaker's mood or feelings based on content, is a significant leap forward.

Dynamic Speech Generation:
Overcoming the challenge of monotonous speech, typical in older text-to-speech technologies is a key focus to produce speech that is varied, expressive and reflective.

What is the product particularly suited for?
Futuri AudioAI is particularly suited for radio stations, podcast producers, and digital audio platforms looking to scale their content production and enhance personalization. The tool has been used to:

- Produce up-to-the-minute citizen service elements, like weather and news, automatically.
- Instantly generate commercials and spec spots.
- Host full shows as an AI DJ.
- Generate podcast episodes without human involvement.

Resource needs:
To use the Futuri AudioAI service, customers need to have an assigned budget and a team that can enter a collaborative process with Futuri to make it meet desired needs. To integrate the system into their offering, Futuri provides an API that ensures compatibility with pre-existing systems.

Contact person: sharialexander@futurimedia.com

PART 3: AI support for the business side: conversion and retention

This report focuses on the editorial and production side of media organizations, but journalism needs business models to survive. While Chapter 1 discussed some of the fears held by media executives about a future that is increasingly based on AI search systems, there are hopes, too. Generative AI will make it easier and less costly to target users with personalized content or marketing offerings, to send out large volumes of mails intended to help with retention, or to increase reach. Madhav Chinnappa, ex-Google, says: "I am quite interested in the business model side of it, how you can use AI and machine learning to understand audiences, their appreciation for you, their propensity to pay and why they pay at certain times. It can also help with experimentation around the kind of visual prompts or calls to action needed that help to influence people's behaviours in certain ways, for example to make them commit to a membership or subscription."

For example, the UK's largest regional publisher, Reach, has been using a bot called Guten to – among other tasks – rewrite copy for its regional editions, according to a report by the industry publication Press Gazette.[1] The idea behind it is that simply copying content negatively affects Google search rankings; texts that are slightly altered circumvent this barrier. Reach titles mostly operate on an advertising business model. The example of News Corp in Australia mentioned above shows that AI can also help to generate subscriptions. At the 2023 World News Congress, the company's Executive Chairman Michael Miller described how 55% of all new subscriptions were driven by hyperlocal news, and every local masthead was staffed by just one journalist overseeing content created by AI. Miller: "If that single journalist can generate seven new subscriptions a week, then their salary is covered."[2]

South Africa's *Daily Maverick*, which runs on a membership model, has profited enormously from a home-made custom GPT that produces drafts for marketing emails targeted at different audiences. According to CEO Styli Charalambous, conversion rates have significantly improved while ineffective communications have been reduced. "It has also saved us hiring a copywriter," he says. (Read Case #14, Growing Reader Revenue, page 95)

It is important, however, not only to try to anticipate large changes but to keep an eye on lessons from current experiments, even if they fail. Agnes Stenbom of Schibsted: "Generative AI in news can be completely transformative, but it could also lead to incremental innovations in ways we haven't expected. One example is news summaries. We had the hypothesis that young people would read summaries. But what happened is that all audiences read summaries. And when they do, they are likely to finish the full article. This is an incremental small change that had big impacts." There is evidence in the industry that people are more likely to subscribe if the time they spend reading increases. In this case, generative AI could be an indirect driver of subscriptions. (Read Case #15, News Article Summaries, page 97)

The same could hold true for models that make sure news articles have a certain tone of voice or content that is likely to attract more loyal users and subscribers. At the Constructive Institute in Aarhus, Peter Damgaard and Gustav Aarup Lauridsen have been working on an algorithm that analyses whether news stories are constructive – in the sense that they provide perspectives and solutions. The foundation model has been built on news articles that were rated by the Institute's fellows. The working hypothesis is that constructive news draws more subscriptions.[3]

[1] Bron Maher, "Reach using AI to speed up 'ripping' and use same article on multiple sites," Press Gazette, https://pressgazette.co.uk/publishers/nationals/reach-ai-guten/, retrieved on 4 March 2024.

[2] Dean Roper, "Michael Miller on how NewsCorp Australia has transformed its journalism and business," Wan-Ifra, 12 July 2023. https://wan-ifra.org/2023/07/michael-miller-on-how-newscorp-australia-has-taken-a-stand-and-transformed-its-journalism-and-business/, retrieved on 30 April 2024.

[3] Conversation with Peter Damgaard and Data Analyst Gustav Aarup Lauridsen in February 2024. Alexandra Borchardt is on the Board of Constructive Institute.

For larger media organizations, the development of in-house tools might be a business opportunity itself. As APA's Verena Krawarik explains: "APA is not just a news agency. It is a company that has a large newsroom but an equally strong IT subsidiary. We serve different clients in media, PR managers and companies. We always build the tools in such a way that we can sell them to media. Secondary and tertiary exploitation is always involved. That's why we take a very explorative approach to everything, also in cooperation with companies."

But generative AI will not only benefit those who are household names in the industry. The media industry will see plenty of start-ups that seek to use generative AI as a basis for an improved news experience. One example is Particle, an AI-powered news reader founded and developed by former Twitter engineers.[4] It is impossible to predict which of these ventures will prevail. For example, the originally well-received news app Artifact, developed by the co-founders of Instagram, had to shut down after less than a year because it failed to attract a critical mass of users. Its technology might survive, though: in April 2024, Yahoo bought what was left of it.[5]

Generative AI creates abundant opportunities for news gathering, production, product development, and distribution. But how should news organizations get started? How should they develop their strategies and organize progress, who do they need to include, and how does generative AI affect workflows and responsibilities? This is what Chapter 3 is all about.

[4] Sarah Perez, "Former Twitter engineers are building Particle, an AI-powered newsreader, backed by $4.4M", Techcrunch, 29 February 2024. https://techcrunch.com/2024/02/29/former-twitter-engineers-are-building-particle-an-ai-powered-news-reader/?guccounter=1, retrieved on 4 March 2024.

[5] Sarah Perez, "What happened to Artifact", Techcrunch, 18 January 2024. https://techcrunch.com/2024/01/18/why-artifact-from-instagrams-founders-failed-shut-down/, retrieved on 4 March 2024, and David Pierce, "Yahoo is buying Artifact, the AI news app from the Instagram co-founders", The Verge, 2 April 2024. https://www.theverge.com/2024/4/2/24118436/yahoo-news-artifact-acquisition, retrieved on 18 April 2024.

CASES

CASE

GROWING READER REVENUE WITH A CUSTOM GPT

DAILY MAVERICK (SOUTH AFRICA)

What is the problem to be solved?

The independent South African digital news platform *Daily Maverick* uses a business model based on membership revenue. The main path to conversion is email marketing. This requires significantly more copywriting than on-screen texts. It has also become clear that different types of audiences respond to different kinds of marketing messages. This requires a high volume of sophisticated copywriting.

What was done?

The team composed a list of the best performing email marketing messages it had sent out to audiences soliciting members. From this, it developed a copywriting style guide and a manual for those engaged in membership marketing and built a custom GPT around that. The person responsible for writing the mailers now draws on this to get a first draft of copy for a membership appeal. This can then be customized for different audience segments identified by the email newsletter system.

What was learned?

The membership marketing team has massively increased its effectiveness because it can turn around a lot more customised copy in different formats. *Daily Maverick* CEO Styli Charalambous explains: "Instead of sending one generic appeal, we can now send three different versions of that same underlying message but customised to three different audiences. There might be a business audience, there might be one that is particularly fond of our climate change work. And there might be a third audience that we've identified as our most highly engaged newsletter subscribers. So, it's a three times improvement for us to be able to turn that around in the same amount of time."

Daily Maverick's own analysis indicates that this has improved conversion rates by an estimated 15 to 20 percent. In addition, there are cost savings, according to Charalambous: "It has also saved us hiring a copywriting assistant. Now a GPT for $20 per month does the job."

Particularly suited for:

This can be used by any organization dependent on direct marketing or membership campaigns for reader revenue. But the experiences could be applied to the editorial side when the task is to provide different versions of stories in order to target distinct audience segments from the same base content.

Resource needs:
A custom GPT needs to be trained with foundational data, in this case examples from well performing marketing messages. Thus, measuring the success of copy and drawing conclusions from it is essential. Once all the inputs have been collated and organized, training the model is as simple as writing some contextual information and playing around with prompts to test the quality of the outputs. If all the inputs have been prepared, e.g. best-performing campaigns, copywriting frameworks, target audiences, then uploading and training the model can be done in 30 minutes. Training the copywriter to use the GPT can be done in another 30–60 minute session to get comfortable with using the tool.

Future plans:
The *Daily Maverick* aims to build LLM-based customized 'personal assistants' for all its editorial staff to fact check their copy, make suggestions for the improvement of content and to help them better navigate English, as for many of them it is their second language.

Illustrations:

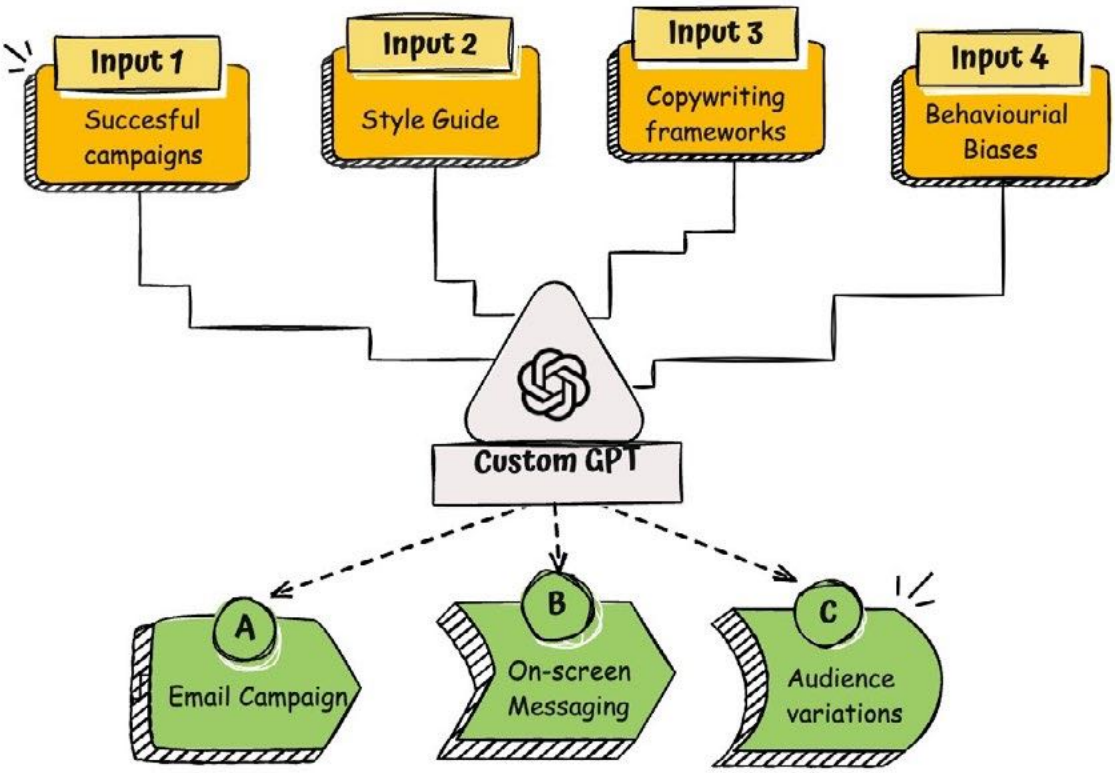

Contact person: Styli Charalambous,
CEO *Daily Maverick*, Styli@dailymaverick.co.za

CASE

NEWS ARTICLE SUMMARIES WITH GENERATIVE AI

AFTONBLADET (SWEDEN), VG (NORWAY)

What was the problem to be solved?
Tabloids *Aftonbladet* and VG (Verdens Gang), are part of the Schibsted media group and major online publishers in Sweden and Norway. Since 2023, they both offer their online audiences quick summaries of longer articles using generative AI.

"For a long time, we have wanted to offer different options for the audience to consume our content. But due to a lack of resources, we have not been able to do so. Now AI summaries open that possibility," says Martin Schori, Deputy Editor at *Aftonbladet*.

In promoting their project VG points out that young users in particular "often emphasize in user interviews that they want shorter, more summarized content."

What was done?
Aftonbladet presents its AI-generated summaries at the beginning of articles just after the lead paragraph. The feature is integrated into the CMS. The summary is proposed by ChatGPT and then approved by an editorial staff member.

The editorial policy is that the AI summaries should be used on all published news articles over 2,500 characters. However, according to Martin Schori this not always the case, due to pressure on the sub-editors.

At VG, they present a summary after the introduction to each article so that the reader can choose a long or short form. The summary is generated by ChatGPT and checked and processed by VG journalists before publication.

What was learned?
Both *Aftonbladet* and VG are very content with the feature and say that they can see increased digital engagement as a result.

"The audience loves them! Between 30–40 percent choose to expand the summaries, there is an even higher percentage among young people. Those who expand the summary also continue to read the full article more than those who do not, which surprised us," says Martin Schori of *Aftonbladet*, who also learned how to introduce AI in the newsroom.

"This type of tool needs to be sold to the newsroom staff, using data."

Both companies are alert to the importance of monitoring the content as sometimes the summaries can go wrong.

In September of 2023, VG published a summary that seemed to be saying that Manchester City's football star Erling Haaland had been shot in Norway, which was completely false.

"Photographer David Yarrow has taken unique pictures of soccer player Erling Braut Haaland, who was shot on billionaire Arne Fredly's property on Bygdöy," the AI summary of the article read.

The mistake probably occurred as the AI confused being shot by a gun and being shot by a camera. VG's Sports Editor Frode Buans was apologetic in responding to the error:

"I have to apologize to Erling Braut Haaland. Although Haaland almost every week wrecks the opponents' defence in England, it was in no way intended to kill him."

VG now puts even more emphasis on manually checking and editing the summaries prior to publication.

Particularly suited for:
Can be applied to all articles and a useful addition as long as the accuracy is protected. Found to be a particularly attractive feature for younger audiences.

Resource needs:
At *Aftonbladet*, it took six weeks to develop the tool, and takes about 30 seconds to insert the summaries into an article.

Illustrations:


AFTONBLADET TISDAG 5 MARS 2024
Dagens namn: Tora, Tove

Nyheter / Kriget mellan Israel och Hamas

FN: Minst tio barn i Gaza har dött av svält

Nivette Dawod

Publicerad 2024-03-04 10.47

+ Följ Mejla Dela Spara

Minst tio barn i Gaza har svultit ihjäl, uppger FN.
Siffran kan vara högre.
– De hade inte behövt dö, säger Petter Gustafsson, pressekreterare på FN:s barnfond Unicef.

SNABBVERSION

- Minst 10 barn har svultit ihjäl i norra Gaza, rapporterar FN. Antalet kan vara högre.
- Mat finns bara någon mil bort över gränsen till Israel. Nödhjälp levereras endast till södra Gaza, medan norra Gaza lider av svält.
- UNICEF, Världshälsoorganisationen och FN:s kontor för samordning av humanitär hjälp har lyckats leverera nödhjälp, men betonar att det är en droppe i havet jämfört med vad som behövs.

ⓘ Sammanfattningen är gjord med stöd av AI-verktyg från OpenAI och kvalitetssäkrad av Aftonbladet. Läs vår AI-policy här.

VG summary:

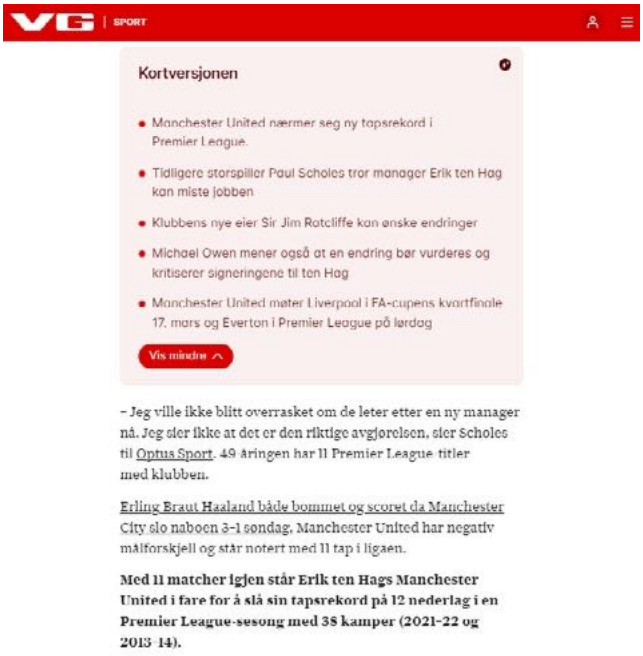

The summary where a significant mistake about Erling Haaland was presented and quickly corrected.

Contact person *Aftonbladet*: martin.schori@aftonbladet.se


Lesson 1 2 3
L'ART DU PROMPT
人工智能
OPEN
FREE
ACCESSIBLE
FOR
ALL
human in the loop

CHAPTER 3

MANAGING THE AI-CONSCIOUS NEWS ORGANIZATION

Generative AI is increasingly seeping into various industries and parts of life. News organizations will not be alone in having to manage this transformation, even though their commitment to accuracy and facts gives them a special responsibility. There's no doubt about it, the pressure is on. As Verena Krawarik of APA says: "The game changer is that AI has now made it into the journalism sphere. Before, the technology departments had taken care of it, for example to improve advertising performance. Now it's on the table of all CEOs and editors-in-chief, and everyone has the feeling that if I don't do something about it, I'll miss the next leap." The process may be easier for those who have been developing and pushing the topic in the background for years. At the same time, many people who didn't necessarily care about it before suddenly have an opinion – including top leadership.

Uli Köppen of German Bayerischer Rundfunk (BR) remarks that since the launch of ChatGPT, it has become eminently clear that investment is needed. Köppen: "We started as a bottom-up movement. Now, a few years after building the AI Lab, we have a very clear view of what we started out of intuition. You need those innovation cores that are doing the work, that are experimenting, and that are rooted in the newsroom culture to understand how you can use this technology. We have full prototypes and we have drawn strategy knowledge out of those prototypes. Now we have to combine it with a top-down movement, because we have to think about how we can use AI and automation in a prioritized and structured way within the whole company, not only in our journalistic silo. So, this will be the next step: How can we combine our bottom-up movement with top-down strategy?"

This chapter will discuss the methods leading managers in the news industry are using to deal with the challenges of generative AI. How do they build strategy, what kind of mindset do they encourage, who participates in decision-making, how does AI impact workflows, and what are the main considerations when developing AI guidelines? This will hopefully help practitioners to clarify their own thinking. Obviously, most questions that arise in the day-to-day management of this ever-evolving technology don't have clear-cut answers. They will very much depend on an organization's mission, remit, and company culture, its previous experience with innovation, the skills and adaptability of its staff and the resources and technology it has available. This is why exchanging experiences between different enterprises – large and small, public service media or commercial – is essential in weighing the options.

Rule number one: strategy first

The key question is and will always be: is there a strategy – and how does AI fit into it? The archaic world of one-size-fits-all media is long gone. News media is already fighting for relevance and attention in a disrupted digital world where audiences have boundless choice. With an ever-increasing number of platforms available to news consumers and an abundance of ways to serve their needs, news organizations need to put a lot of thought into which audiences to serve where, and when. Strategy development must precede the implementation of any technology and tool, generative AI assisted or otherwise. This is all too often forgotten in the frenzy of new tech, with eager sales teams constantly pitching exciting tools. As Uli Köppen describes, cutting through the noise can be challenging:

"It's really hard to differentiate between people who just want to sell something and people who are saying something we should listen to."

On this note, in its *PSM AI Strategy* published in April 2024, the EBU cautions its members: "Don't be spooked into a 'panic response' by the power of Big Tech," it warns, and advises: "Fear of missing out can trigger hasty trials of new AI tools, but this risks having many disconnected initiatives open at once. It's more resource-friendly to consider how AI can solve priority problems before acting."[1] Chris Moran, Head of Editorial Innovation at *The Guardian*, also advises against rushing into implementing AI-driven tools without being confident they work. "The first news organization to deploy generative AI meaningfully isn't going to 'win' journalism**,**" he says.

Bill Thompson of the BBC believes generative AI could provide new opportunities to increase public service value: "The BBC should embrace generative AI, because it will help public service with its mission: inform, entertain, educate. It is not about making money out of it but making a better world out of it." He is less optimistic for commercial media. "If organizations are responsible to shareholders, this is a threat."

The bigger commercial players don't seem to agree. Niddal Salah-Eldin of Axel Springer says the publisher made AI a top priority at the executive board level immediately following the launch of ChatGPT: "We ensure that our strategic initiatives are aligned with our long-term goals. This means we continuously iterate and update our approach. Year one of Gen AI was focused around creating momentum across the organization. This was centred around a three-pillar approach comprising education and exchange, optimization, and exploration in a hybrid setup which combines central offers and initiatives from the HQ and empowers decentralised initiatives steered by the business units." (Read the Q&A with Niddal Salah-Eldin, page 119)

In contrast to Axel Springer, *Daily Maverick* employs only around 100 people. Styli Charalambous describes their strategy development path: "We are in the business of creating the journalism that society needs, not the journalism that society wants. So, we have to understand what those needs are. And we have to have the frameworks to help us decide and prioritize how we allocate resources. Once we've created the journalism that society needs, our secondary job is how do we package this in a way that it is more attractive to more people? If you start with what is the journalism that people want and you have no consideration for what they actually need or what your own vision and mission are, then you can't build strategy around it. You become driven by some arbitrary metric."

In a public service media context, the core of this mission is simple: to reach everyone with content that is valuable, if not essential, for their lives. Ezra Eeman of NPO suggests constantly asking: "How can we deliver value with this, what is really necessary, what is our ambition and vision with this? Our aim should be to deliver value in a more granular way to those people we are missing out now. We still have a long way to go from a broadcast model to a model that is involved in people's lives." (Read the Q&A with Ezra Eeman, page 127)

Blathnaid Healy, Executive News Editor for Growth, Social and Delivery at the BBC, feels that the advantage with the current wave of disruption is that media organizations are better prepared: "We are dealing with more mature digital businesses now. 15 years ago, we had no product teams in the organizations. Back then we were not able to tackle the opportunities the way we can tackle them now. We talk more as an industry now with more confidence." The BBC's approach is to not rush into anything but start from where users are. Healy: "We need to understand our audiences. When would they turn to ChatGPT for news coverage, for example? We are very intentional about how we approach each audience depending on which platforms they

[1] EBU, "PSM AI Strategies – Lessons learned so far," EBU Strategy Services, April 2024. Document for members only: https://www.ebu.ch/guides/membersonly/report/public-service-media-ai-strategies

are on. We did that in the past with this lens, for example on TikTok. The gigantic caveat with generative AI is that it hasn't settled. It is changing behaviour, but we don't yet know how." In the spring of 2024, the BBC had less than a dozen pilots running on generative AI tools and in an update laid out its guiding principles for which would ultimately be continued. "As a reminder, our principles commit to harnessing the new technology to support our public mission, to prioritising talent and creativity, and being open and transparent with our audiences whenever and wherever we deploy Gen AI," it reads.[2]

Verena Krawarik of APA agrees that being prepared has helped. And support from research grants in an earlier stage of innovation development has proven to be crucial now: "We founded the APA medialab in 2017. That has helped us to deal with such topics. In 2020 we got a research grant from the Austrian Ministry of Technology to explore AI and its conditions for success. That has led to an incredible wealth of knowhow about, what is possible, what works. Many of the challenges that were carried out back then can be applied now. In 2021 we ran a project: 'Visual.Trust.AI' was about building an algorithm for face recognition for our picture database. On the back of this project, we created the first AI guidelines. When ChatGPT came along, it was clear this is something big, we could build on good things and produce an updated version.[3] It helps to have units that are fundamentally concerned with the future."

Nevertheless, Cambridge Professor Gina Neff, who serves on several advisory boards supporting different industries with AI, says that strategy development will remain quite difficult for the time being: "It is going to take a huge amount of work to get useful use cases out of these tools. Right now, there is a lot of play and experimentation, lots of low hanging fruit. But companies are struggling to make a sensible strategy out of AI." As with all strategy development, it is essential to develop (or revisit) the vision first, define what success could look like and work from there – revisiting the leadership guide in the 2021 EBU News Report might help.[4] The questions listed below can provide a guiding framework on what to discuss in the context of AI.

A guide to check your AI and journalism strategy:

- **What do we want to achieve with our journalism?**
- **Who are the audiences we want to serve, and which jobs do we want to do for them?**
- **What would success look like and how can we measure it?**
- **Which jobs will AI (potentially) be better at doing than our staff?**
- **Which jobs can our staff not do that AI can?**
- **Which jobs can only our staff do?**
- **Which risks are we willing to take with AI and what are our red lines?**
- **Have we defined rules for transparency about the use of AI – internally and audience-facing?**
- **How do we monitor the use of AI in our organization?**
- **Who is responsible if something goes wrong?**
- **How do we make sure our newsroom gets sufficient training, guidance, support, tools?**
- **Do we have advice for journalists who report on AI?**
- **Do we communicate all of the above sufficiently and clearly?**

[2] Rhodri Talfan Davies, "An update on the BBC's plans for Generative AI (Gen AI) and how we plan to use AI tools responsibly," 28 February 2024. https://www.bbc.com/mediacentre/articles/2024/update-generative-ai-and-ai-tools-bbc#:~:text=A%20responsible%20approach%20to%20using%20AI%20at%20the%20BBC&text=As%20a%20reminder%2C%20our%20principles,wherever%20we%20deploy%20Gen%20AI., retrieved on 4 March 2024.

[3] APA Guidelines for AI (in German): https://apa.at/wp-content/uploads/2023/07/Leitlinie-zum-Umgang-mit-kuenstlicher-Intelligent-2023-2.pdf

[4] A. Borchardt, F. Simon, "What's Next? Public Service Journalism in the Age of Distraction, Opinion, and Information Abundance", EBU News Report 2021-2022, https://www.ebu.ch/guides/open/leadership-and-innovation-for-public-service-media-journalism, retrieved on 18 May 2024

The mindset: testing and learning

Practically all media leaders we interviewed emphasized that creating a mindset of openness towards AI in their organization was essential. Staff needed to embrace the joy of experimentation and be willing to learn from it, also from misguided efforts and failures. Dmitry Shishkin of Ringier says: "I'm a great believer in attitude rather than in specialism. You can teach people skills, but you cannot teach them the right attitude. Tell them: Just be curious about this, start with something."

Agnes Stenbom says that generative AI was a good match with the prevailing newsroom culture at Schibsted: "On a general level, Schibsted newsrooms are very optimistic with a lot of local experimentation and innovation going on. We have worked a lot with internal training on AI opportunities and risks, which I think has contributed to our newsrooms being both risk aware and adventurous."

ARD Chairman Kai Gniffke describes the variety of responses newsrooms are likely facing: "Some people's eyes light up and they say, great, things are moving forward here. But there are also those who say it's very dangerous. My job then is to say that if we don't deal with it, we'll leave it to others who won't feel committed to democracy or a sense of togetherness and community." The biggest challenge was "getting the people who want to work with it to do so. They shouldn't have to wait until they get instructions from the top." (Read the Q&A with Kai Gniffke on page 123)

But the managers we interviewed varied in the degree of caution they advised their teams to exercise. Matt Frehner's employer *The Globe and Mail* embraced 'regular' AI applications much earlier than many other news organizations worldwide. The launch of ChatGPT still took the newsroom by surprise. "We all expected this stuff to progress slowly within a three-to-five-year horizon. Then ChatGPT changed everything. The first response was to come up with a policy framework for the newsroom, setting rules, guard rails, think about tools. There are so many possibilities but also potential pitfalls. We wanted a balance between being open to using the tools but also being cautious about the risk from a brand point of view."

SVT's Anne Lagercrantz says the Swedish broadcaster has been innovative but also cautious: "We tried to make people aware of the risks and to be careful, to be aware of data leakage, for example. Maybe we made people too careful. Now we are trying to make even more people to lean in and we have created safe environments for experiments. We should have done that earlier on."

Some feel that the early enthusiasm could soon be followed by disappointment, with the hype-cycle leading to what Lagercrantz terms a 'cycle of despair' when it comes to the hard work of implementation. BR's Uli Köppen agrees: "The enthusiasm right now is not the same as really integrating this technology into workflows in the newsroom. That is a whole different game, because of course we do have those wonderful tools, but we are still in the middle of figuring out how we can really enrich workflows with them. You need to step back a little bit and see how you can use it at the same time as people enjoy this whole revolution."

Agnes Stenbom of Schibsted has observed a similar tension between initial excitement and actual delivery: "People are drawn to the sparkly things, the exciting and hyped-up cases. But the business needs often sit further down the pipeline, for example related to data quality and control. Another challenge is that we have a lot of great ideas, but struggle to scale them into our content management system. Our reality is a big backlog; the bottleneck of production. And with that comes the risk of people losing interest."

Dmitry Shishkin says it is important to start implementation quickly: "My big thing is: try once and scale it everywhere else very quickly. This was my experience at the BBC. Give teams a remit to try something. If five of ten products bring good results you scale them. You can have endless alignment meetings without actually doing something. Media organizations need to become like the best startups in the form

delivery of things." Shishkin also recommends that companies should not waste too much time on developing custom-made tools or systems. "Utilize what third parties are really good at, do not rebuild. In the debate about buy versus build, I am on the buy side. Why bother recreating something when others have done it better already? If you use a third-party product, you are going to be much more effective. You will have more hours in the day to do things that will really stand out."

These might be tricky recommendations for public service media which, by their very nature, cannot be as agile as commercial enterprises. They are held accountable by public bodies which might require more checks and controls. Also, issues like data protection play a larger role. Third-party products might just not fulfil the standards they have to uphold. Then in-house smaller organizations in particular need to be cost-conscious and cannot invest millions in new tools and technology just to find out later it was misspent. Erik Roose, Chairman of the Board of Estonia's public broadcaster says: "If you are in a small market, you just don't have all the resources to go into all the niches. In a large country, one percent of an audience can still be a million people. In Estonia it is 10,000, you will never break even." This is where according to him management skill comes in: "You need to be really good at defining what you should do by yourself. What is important and what just a bubble? And what is something that after six months Apple or someone else will provide you with something similar that costs you 100 per month. Evaluation ability and analytics is really key here. And if you fail there, you can easily spend millions, but don't get anything."

Organizing AI usage across the organization

In many news organizations, AI was primarily dealt with by special teams operating in niches, more on the business than on the editorial side. Talks about AI in the newsroom were well received at conferences on innovation. But the topic took a back seat to other debates, for example which platforms to publish on and which products to develop. With the launch of ChatGPT, this has changed. Most media organizations now have interdisciplinary task forces, some have appointed AI directors to anchor all efforts with one senior role. Industry collaborations have also started to emerge. The EBU, for example, has made AI a strategic priority, bringing members together to showcase best-practice examples, share advice, and complement it with the School of AI launched by its inhouse EBU Academy.

Additionally, more media organizations are providing both basic and in-depth training for their staff. A recent example is Swedish Radio, which launched an ambitious online course on the basics of AI in journalism in May 2024, and made it compulsory for all 2,000 employees. This includes self-assessment, ethical dilemmas, an overview of AI guidelines, inspiring testimonies from colleagues and a checklist for fake detection.

Bill Thompson, who defines his role as "preparing the BBC for a future it doesn't understand," describes the perspective shift from AI being something just for nerds to essential in his organization: "Many years ago a number of interested people formed an organization inside of the BBC: the 'Machine learning ethical design working group'. This is a boring name, there was no political advantage to being part of this group. Consequently, only those interested joined. Over time it has increased to a couple of hundred people, whose job it is to think about how the BBC's values intersect with these technologies. It was a grassroots initiative. That meant when generative AI started, we had muscle memory."

Getting interdisciplinary teams together to tackle AI is a must for all organizations, as there are so many aspects to consider, ranging from editorial to business and legal considerations. As *The Globe and Mail*'s Matt Frehner describes: "Parallel to our work in the newsroom we have had a corporate wide committee. This includes legal, finance, UX teams, data teams, our call centre. We need to understand that this is going to touch every

department. We needed a strategy across the company. We also tried to put some kind of framework around investment, what should we purchase, what should we build ourselves? The next phase will be evaluation."

Austria's APA has a similar approach, having built a task force of nine people across the organization. One goal was to spread knowledge fast. Verena Krawarik: "We held a basic training course for all employees relatively early on, five months after the launch of ChatGPT**,** where we explained what it is and how it works. There's no point in telling people what they can't do, we tell them what they can do. We provided a kind of playground: Everyone who wanted to test an idea was given access. We also invented new formats like hackshops – a hybrid between a hackathon and a workshop with our colleagues. Ten people were put together to work on certain issues. It was about trying things out, getting feedback. There are 500 people in the APA Group, they are all being put through these basic programmes."

SVT's Anne Lagercrantz emphasizes how important it is to reach everyone in the organization – and to invest in creating safe environments: "We have an AI council. It includes an ML-engineer, an editor-in-chief, the head of our kids programming, persons from the ethical department, production, business intelligence, a lawyer, a communications specialist, our chief information security officer and a controller. It has developed our guiding principles and AI policy. We have also created a central LLM-team. Becoming AI literate is the responsibility of everyone, but we needed to create this safe environment. We have now three, four people working full time on this. The timing isn't good, because we have to save money, but we hope to save money when things get more efficient."

Naturally, coordinating everything around AI becomes more complicated as organizational complexity increases. German ARD is particularly distributed, with semi-independent public broadcasters in different regions. Kai Gniffke describes the struggle to keep everyone on board while allowing for independence: "We also set up a competence network. Because, of course, employees in different roles, various departments and broadcasters are experimenting with AI. They should know about each other so that they are not just busy reinventing the wheel. We need one power centre, not nine. It's all about sharing experiences. There are different speeds within organizations, of course, the knowhow of the front runners should spread quickly. What we have also set up at ARD level is a second competence centre for reporting. This is where journalists who report on AI and connected trends and developments meet to keep each other up to date." (Read: Dos and Don'ts of Covering AI, page 174)

Niddal Salah-Eldin says training management and the newsroom is crucial, but it shouldn't be reduced to transmitting knowledge and skills. Cultural change towards an openness to technology and lifelong learning is crucial: "In this spirit, we are sending executives and experts to existing and emerging AI hubs worldwide for 3 to 6 month fellowships to enable them to immerse themselves into the local AI scene to explore opportunities for Axel Springer. Locations include San Francisco, Seoul, Tokyo, and Singapore. Having boots on the ground and being close to the builders will help our businesses."

Most news organizations won't be able to afford these kinds of expensive programmes. But the report on the AI in Journalism Challenge 2023 (AIJC) suggests that the investment needed most is not funding but giving teams time to learn – and offering coaching. The newsrooms that participated in the programme were small to mid-sized, without significant resources and spread out across the globe, far away from Western tech hubs. Still, they learned fast. The report concludes: "Most of these barriers did not appear to be significant factors in limiting what the AIJC teams were able to achieve. The geographic dispersal of the teams proved essentially irrelevant, as all teams were fully engaged in the global digital community and at ease with accessing and using digital tools and using digital resources for learning

and troubleshooting. The smaller size of the newsrooms appeared to be an advantage, rather than a disadvantage, because the teams were able to make quite significant implementation decisions in a largely self-contained way or with minimal bureaucratic overhead."[5]

Does AI require new talent?

It is unclear what kind of new roles will emerge in the wake of generative AI, since the technology affects everyone. In 2023 and 2024, many organizations appointed AI directors to ensure a power hub. But the dynamic of the change means CEOs and editors-in-chief cannot look away nor shed responsibility.

Uli Köppen of German BR has a special role in mind: "I'm always insisting that every newsroom needs a tech translator. This is a person who is rooted in the newsroom, and understands its strategy, but has the time and the capacity to study what's going on in the outside world, to translate those technology trends into the newsroom strategy and start projects. This can only be one person, but this person is very, very important. This person can combine his or her skill sets with the skill sets of other people in the newsroom. Like this you can really do interdisciplinary work. This person doesn't necessarily have to be a journalist, this can also be a programmer, a product person, you name it. Important is only that this person understands the USP of the newsroom and the strategic way forward."

Mattia Peretti looks at staffing and upskilling as a process: "New roles will be evolving and reskilling will be needed. If you decide to build inhouse, you might need new people, of course. There has been a frenzy about prompt engineers, but I believe we should frame the role as prompt editors instead. It is not a technical skill but a human skill. We need to keep editorial values in mind when writing prompts. There will be cases when the prompt will be in the background and you just click a button, for example to write a summary."

During the AI in Journalism Challenge 2023, it became clear that prompting will be a key skill but developing the skills to do so will be complex and time consuming. According to the project report, the biggest hurdle for many participants was developing a frame for how their particular problem could be solved with the support of AI. During the project, teams developed more self-confidence and were able to improve results, but they also acknowledged the limits of their skills: "(...) many of them came to question whether the insufficiency of a prompt was due to the model or due to their communication with the model. It was in some ways like watching a newly promoted manager gradually becoming familiar with managing a new junior employee, and exploring whether unsatisfactory work was due to the employee or due to the way they had been instructed."

Many editors hope that technological advances will enable them to streamline their work responsibilities. Ideally, generative AI will make it easier to design simple infographics, code, translate, and automatically moderate online comments. If the technology proves to be as reliable as hoped, newsroom complexity might decrease rather than increase for the first time in many years. But there may be a retention challenge too, as AI-skilled staff is in high demand elsewhere. As David Caswell notes in the AIJC report: "We have already seen several participating teams lose key members of their teams, often to technology companies rather than other news organizations."

Getting it done: Workflows and responsibility

Experiments are about envisioning a future; task forces are about navigating the reality. But once the experiment proves worthwhile and decisions have been made, efforts have to shift towards integration into daily workflows. Bill Thompson, the BBC's Head of Future Value Research, talks about the challenges of implementing generative AI in the newsroom: "There are three things: for the BBC in particular it is getting senior

[5] David Caswell, "AI and Journalism Challenge 2023." The following quotes are taken from pages 32, 38, 44.

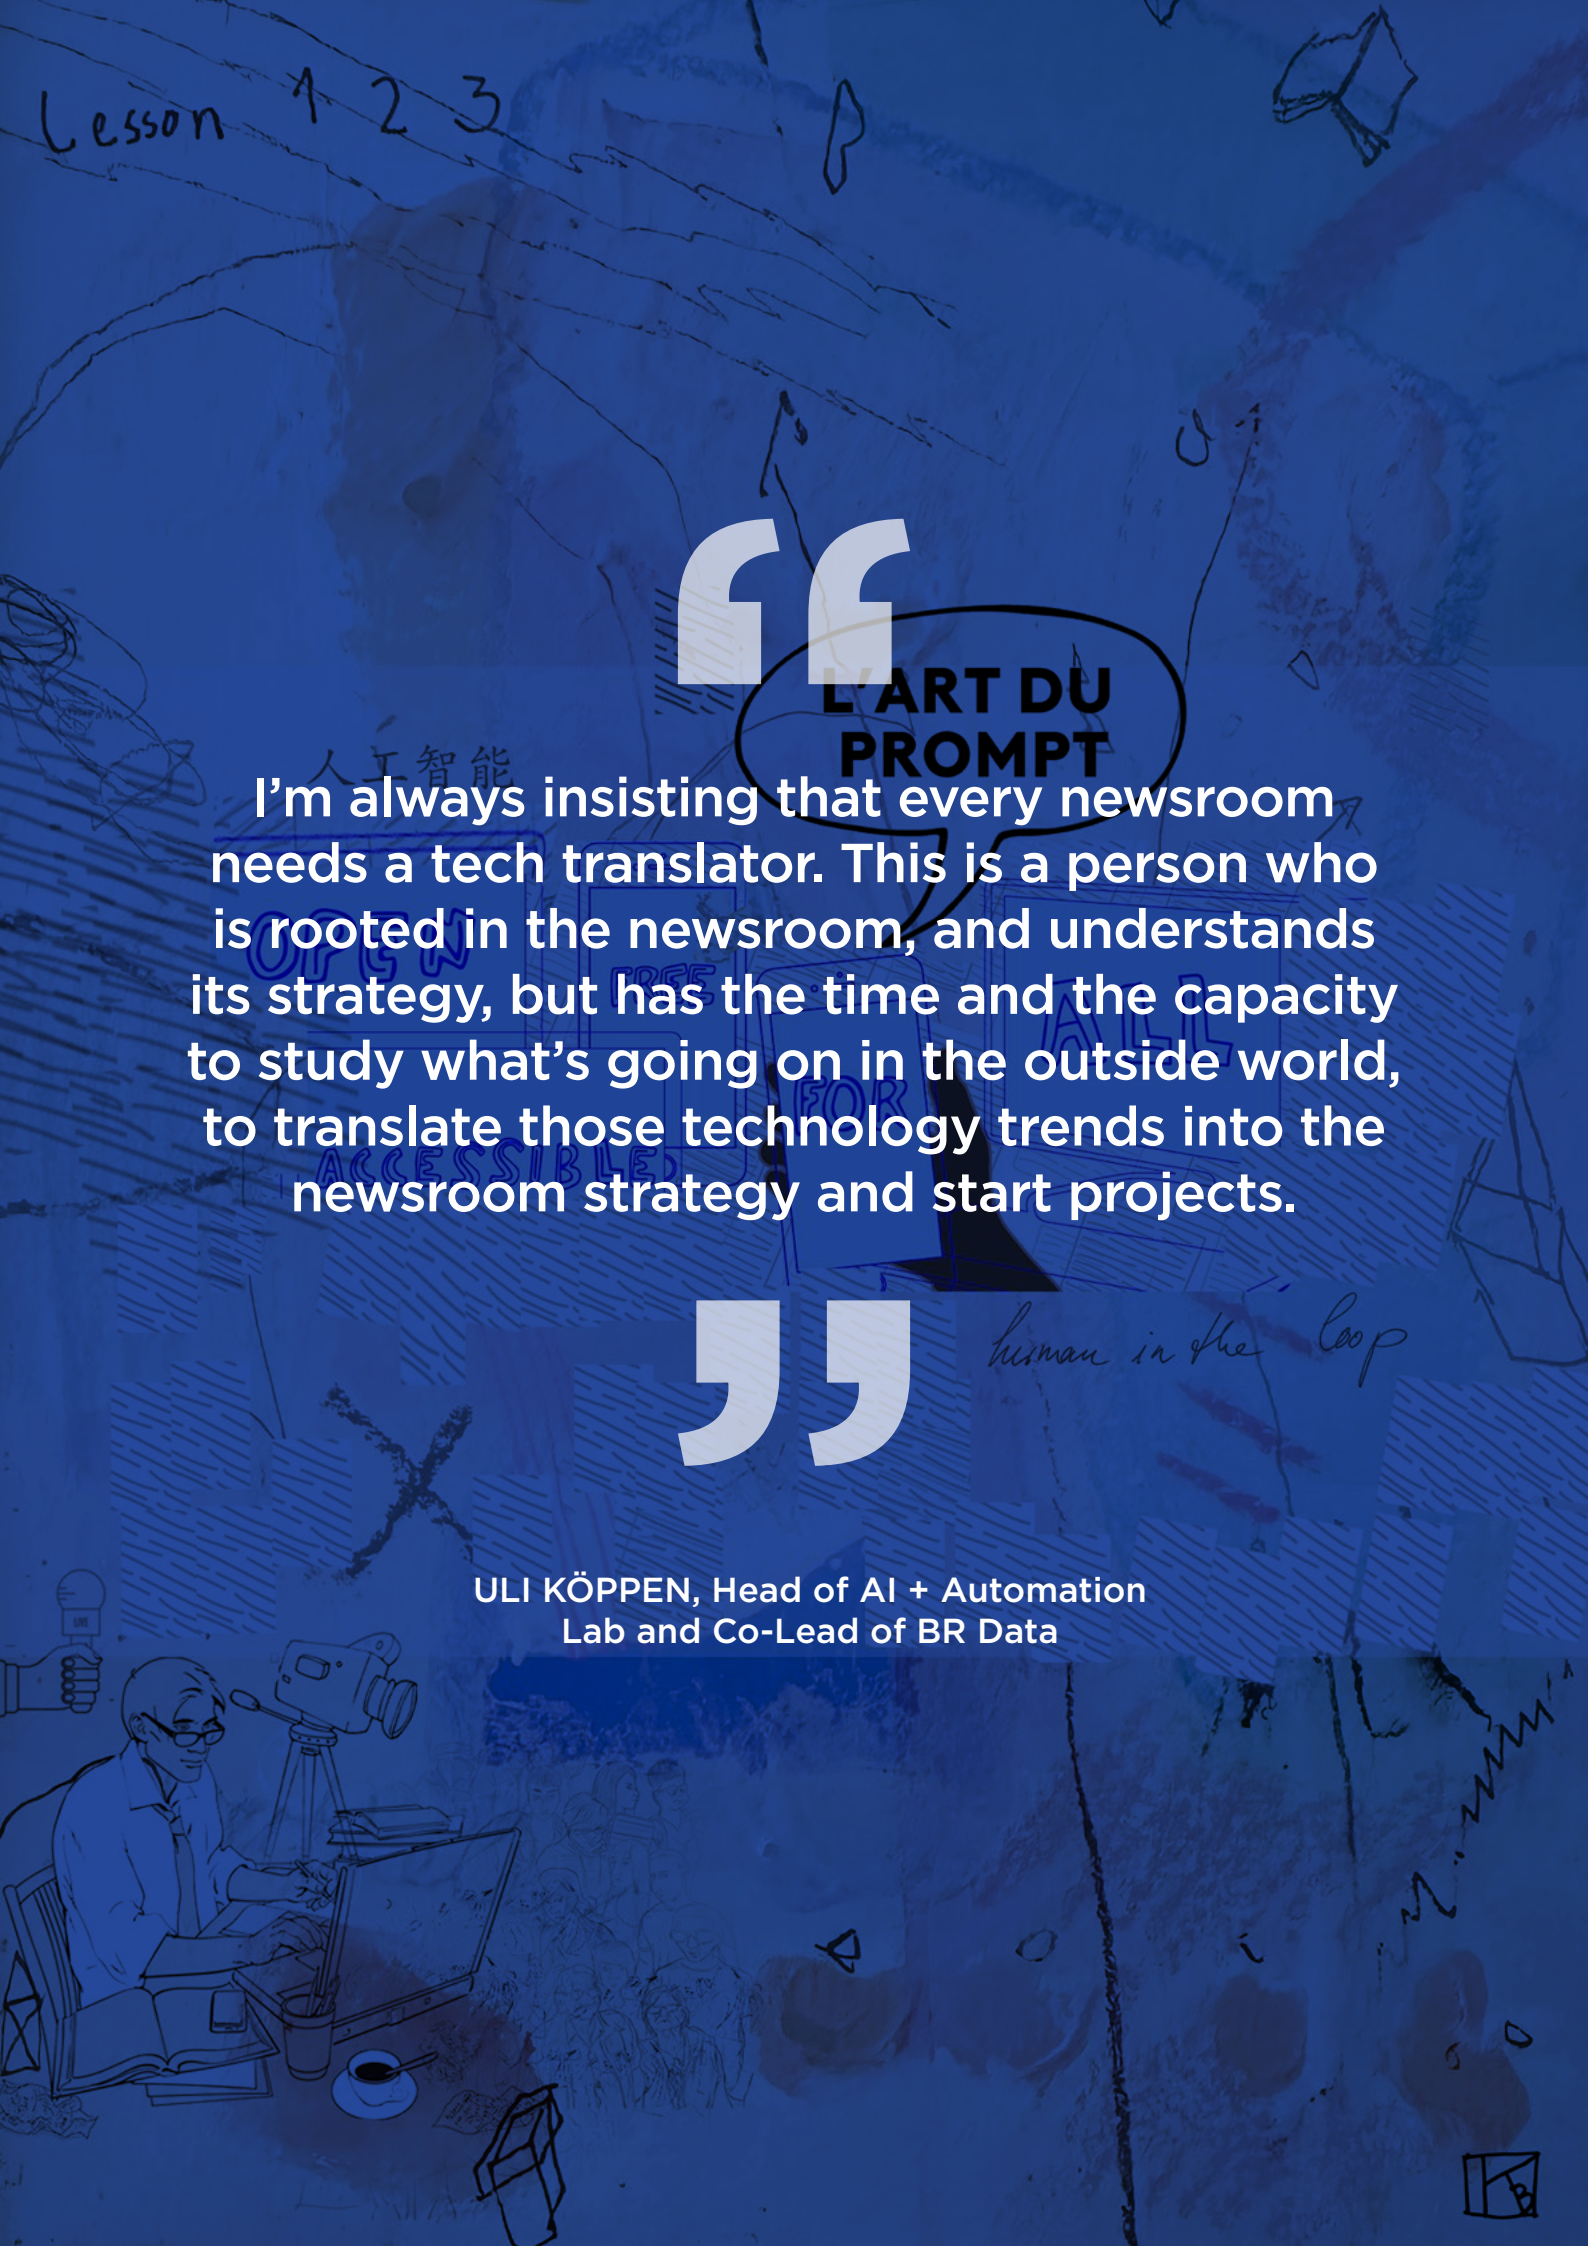

“I’m always insisting that every newsroom needs a tech translator. This is a person who is rooted in the newsroom, and understands its strategy, but has the time and the capacity to study what’s going on in the outside world, to translate those technology trends into the newsroom strategy and start projects.”

ULI KÖPPEN, Head of AI + Automation Lab and Co-Lead of BR Data

stakeholder support in a useful way. Getting the right person to get excited in the right way is pretty hard. In an organization people will not do something if they don't see a political advantage in this. And they shouldn't do just the flashy stuff. They need to be invested in representing the technology to other stakeholders. Second, different parts of the business move at different rates, and that creates tensions, the technically competent bits tend to move quite quickly, the policy and the platform people less so. Also, you need to make sure that everybody knows about it but relaxes a bit. When people ask, 'why can't I do this?' you need to answer, 'because your work will be useless in six months' time.' Third, the thing with the BBC in particular is that it is very closely observed. The consequences of everything are high. That's why experimenting is very hard. You need to put the tricky stuff in R&D, not on the BBC."

Most media companies are not exactly front runners when it comes to research and development. Styli Charalambous of the *Daily Maverick* says: "We are not traditionally an industry that spends on R&D. I read somewhere that furniture companies invest more in R&D than media companies.[6] But it is not just about money, you can burn a lot of money very quickly. It is about the execution. Actually, some AI and large language model (LLM) tools are very cost-effective. The investment is more on the time side." One of the investments that require time is work on the data infrastructure. "You need to get your data in order, if not you cannot play this game," according to AI consultant Louise Story, formerly at the *Wall Street Journal.*[7]

Even when experimentation is encouraged, a minimum level of oversight is needed. To ensure leadership knows what is going on, *The Globe and Mail* developed a project framework which everyone could follow, led by the vice president of data. Matt Frehner: "They need to explain: What's the reasoning behind this and identify what is the problem you are trying to solve? Is it automating a workflow, doing HR evaluations, is there a tool or process that could benefit that theme? What is the cost of the investment, what is the opportunity, what is the risk? There might be an investment of 1 million dollars for saving ten minutes of an editor's time. The other thing is security and legal process. What data are we putting in, what data are we getting out?"

But even standardized processes might not prevent staff from certain actions, simply because they can or because they do it anyway at home. Frehner: "The challenge is keeping track and looking at how people are using the tools in the newsrooms. I cannot approve it for every single assignment: what's the tool you are using, is it appropriate? There are so many tools out there, I'm sure people in the company don't even ask but use them."

Johanna Törn-Mangs describes this as the biggest challenge in managing AI: "Yle has always been a tech-savvy company. But now when the usage has increased so much it is hard to know of all the usage and control all the decisions. Decisions are made in all parts of the organization, a small decision can actually be quite big, for example buying a new tool. A decision of a programmer can influence the journalist without us even knowing. You need to map this out, and you have to tell people if you make these decisions. You cannot just let everybody do whatever they want. People want to do the things that are cool and not always the things that are strategically relevant. Cool things sometimes don't matter in the big picture. You still don't want to kill the creativity."

Deutsche Welle's Manuela Kasper-Claridge says that communication is the most challenging part: "We say that we want to work quickly and flexibly, but we have over 3,000 employees, as is the nature in large organizations, we sometimes aren't able to

[6] Rasmus Kleis Nielsen, Director of the Reuters Institute, made this argument for NiemanLab: https://www.niemanlab.org/2021/12/invest-in-tools-and-talent-and-newsrooms-can-finish-the-job/

[7] Notes taken at a webinar hosted by FT Strategies and the Google News Initiative on 10 October 2023. You can also watch this panel with Louise Story and Dmitry Shishkin at the International Journalism Festival 2024: https://www.journalismfestival.com/programme/2024/ai-in-newsrooms-how-to-do-it-how-not-to-do-it-and-how-to-remember-your-audiences-needs-as-you-go-along

BACK TO CONTENTS

start projects or react to developments with the agility that we would like. Communication can be difficult. Obviously not everyone is working on AI projects, and some people sometimes hear about what is going on through the grapevine. This can be unsettling for colleagues who have heard about how AI might be coming for everyone's jobs – which I do not think is the case." (Read the Q&A with Manuela Kasper-Claridge, page 130)

Communication is also crucial to lessen the impact of another inevitable digital divide: those who are used to working with AI-based tools on a regular basis and perhaps even trained on these tools at university will meet those who are reluctant or even afraid to do so. It is therefore important for newsrooms to create a technical environment where these tools are embedded in content management systems for a smooth user experience. The EBU Neo News Pilot project, shared as a use case in Chapter 2, aims at doing just that. (Read Use Case #9 EBU Neo, page 80)

Madhav Chinnappa agrees that embedding AI in the systems everyone uses is crucial. "Someone said recently that AI is like sex in high school: Everybody is talking about it, but very few people are actually doing it. There are just pockets of action. I understand why people aren't leaning into it more. The technology is at a very nascent stage and ChatGPT is not connected to a workflow. So, if you're in a newsroom, you can open ChatGPT, then work over there and then bring the result back in. (...) The technology is not where the journalist is. It is only going to take off once it is embedded in the systems journalists work with on a daily basis."

Verification, fact-checking, and avoiding mistakes

One major challenge is avoiding severe mistakes. A content management system (CMS) that provides guardrails can help.[8] But technology will only ever be part of the solution, particularly in an industry that is inherently about creativity. Humans will have to perform checks. Most of the media leaders we interviewed work for organizations that follow the 'human in the loop' principle. That is, nothing is published without an editor checking for accuracy. As long as LLMs are prone to hallucinating, this will not only be necessary to prevent brand damage. It's also a question of legal or at least ethical responsibility, particularly for public service media.

In the automotive industry, for example, the proliferation of self-driving cars seems to be more hampered by liability disputes than by technological hurdles (although these persist, too). Since publishers will be held responsible for their output, for the time being, humans will need to conduct the final check. This might change. In the airline industry, the widespread use of autopilots and collision control systems has made flying on average a lot safer compared to the times when most of the decisions were made by pilots and air traffic controllers.

Accuracy has always been a core aspiration of responsible journalism. As journalists from the age of print know, it doesn't take more than a few typos for readers to lose trust in a media brand. Attitudes about inadequate spelling might relax over time, with people becoming used to near misses in subtitles and words that are made unrecognizable by autocorrect modes. However, when it comes to factual errors, they might not be so forgiving. As the BBC's Bill Thomson says: "You don't look at LLMs in isolation, you always look at the system they are embedded in. As a journalist you need to embed it in a system that can be made accountable. Human journalists have flaws as well. But there has always been a whole process that allows for the fact that human beings make mistakes." An article on *Newsweek*'s use of AI describes the publication's rules for transparency about mistakes. But as *Nieman Lab* noted, the boundaries were already blurring when

[8] Guardrails in the context of AI is not only a metaphor but a specific safety structure to guard against the display of restricted content. But it is also used in a more philosophical way, as for example in the book by Urs Gasser and Viktor Mayer-Schönberger, "Guardrails – Guiding Human Decisions in the Age of AI", Princeton University Press, 2024.

it came to attributing corrections to AI or human error in the US magazine.[9]

Newsrooms will also have to consider the increased pressures journalists might find themselves under when being forced to fact-check and oversee production at increased speed. Renate Schroeder of the European Federation of Journalists says: "There needs to be concern about mental health. AI will help journalists to do the boring stuff. But with content production running for 24 hours, there will be more pressure to fact-check and get it right. It is impossible to compete with the speed of AI."

Andrew Strait of the Ada Lovelace Institute says the question of liability will frame the intensity of the pressure. Staff might prefer to blame technology if something goes wrong, but with 'human in the loop' policies this will not always work. Strait: "The biggest red line is accountability. Individuals will feel the weight of that." He recommends using a test when deciding whether a generative AI-based tool should be used: "You always need to ask: how are you planning to use generative AI, for what purpose? If this system makes a mistake, is the outcome any worse than the situation was before? (...) There are worse-case scenarios when getting it wrong is not bad. But there are other scenarios when you know: If you get that wrong, it is killing your credibility."

Erik Roose of Estonian ERR says that his news organization made a promise to the Estonian public that every output will be checked by humans who bear the final responsibility. "Not every reporter is happy about this, because it still means that they are responsible and must invest a lot of time and effort into controlling all this data. But we told the public we will never accept excuses like 'this wasn't me, this was some AI incident'. It is like this everywhere, even in the military. A human needs to be in control."

Generative AI will also bring new challenges to battling misinformation purely because of the scale and speed at which fake images, audio and video can be produced. As was stated in a report by the UK's House of Lords: "The most immediate security concerns from LLMs come from making existing malicious activities easier, rather than qualitatively new risks."[10] Then again, there is hope that AI will also improve automated misinformation detection. Media organizations, particularly news agencies, will further have to step up their verification efforts, and LLMs can potentially help doing this. Ritu Kapur of *The Quint* says: "During election campaigning in India, this year, there has been a trend of dead leaders being 'AI revived' via deepfake videos, where they are 'saying' things they never really did. We are very mindful of that. And our fact-check unit is using tools like True Media or Hive to do the fact-checking." Jean-Marc Rickli of the Geneva Centre for Security Policy has doubts about the effectiveness of technological factchecking solutions though: "It has become almost impossible to come up with failure proof detection and verification methods."

The debate about misinformation has fuelled plenty of funding in this field. A whole ecosystem of start-ups has evolved around it. In an ideal world, verification would be as easy as a spellcheck. Gaute Kokkvoll of the Stavanger-based startup Factiverse, which creates automated fact-checking solutions, knows from experience: "If journalists have to log into yet another front end, it is a hurdle. You need to integrate it into the CMS. It should be more about finding credible information, not about going through fact checking." But Verena Krawasik of Austrian APA warns: "We media won't be able to tackle the disinformation problem on our own. We need regulators and the platforms for this."

However, waiting for technical solutions will not serve as an excuse to remain idle. Upskilling employees or helping them build the tools to battle misinformation is essential. For example, in its School of AI, the EBU Academy offers a course titled 'Build a Fake News Detector'.[11] Axel

[9] Andrew Deck, Newsweek is making AI a fixture in its newsroom", Nieman Lab, 17th April 2024. https://www.niemanlab.org/2024/04/inside-newsweek-ai-experiment/, retrieved on 19th April.

[10] House of Lords (2024), page 42.

[11] See https://academy.ebu.ch/schoolofai, retrieved on 30 April 2024.

Springer's Niddal Salah-Eldin also emphasizes the need to train staff on spotting fabricated content or lies: "Our employees need to be able to recognize misinformation. At the Axel Springer Academy of Journalism and Technology we offer training sessions on OSINT [Open-Source Intelligence] and fact checking where employees learn techniques for verifying AI-generated content. It includes identifying suspicious patterns and inconsistencies, checking internal and external logic, and exposing manipulation and misinformation."

Many of our interviewees worry about the effort and energy it takes to verify content, energy that could instead be invested in producing relevant journalism. Johanna Törn-Mangs of Yle says: "Debunking misinformation takes time and a lot of resources; it is a huge challenge. Of course, we are being super, super careful about our own material. If there is even a slight risk something is wrong, we don't do it. We don't need to be the fastest. It is very important to always be very transparent about mistakes."

Agnes Stenbom of Schibsted expects that traditional news organizations will hesitate to implement automated processes, because they have a reputation to lose and are bound by journalistic principles and standards: "The news industry is not interested in fully automated workflows. We talk about 'human in the loop'; we will have editors in the loop. This makes me less worried about misinformation. But I am worried about new competitors in the information market who won't have editors in the loop."

Ethical guidelines: between empowerment and red flags

Until recently, only a few newsrooms had drafted specific guidelines on the use of AI and under what circumstances and constraints it can and should be used. The launch of ChatGPT changed that. It motivated many media organizations to come up with rules, both general and specific. In today's journalism enterprises, AI guidelines are as abundant as mission statements. A paper co-authored by Felix Simon and published as a pre-print in September 2023 compared 52 of these documents, finding more similarities than discrepancies.[12] This is not surprising. On the one hand, all media leaders we spoke to wanted to encourage a spirit of openness for experiments and therefore keep the number of red lines to a minimum. On the other, generative AI is a fast-moving technology; being too specific impedes progress as today's rules can be outdated tomorrow.

Further, guidelines are not user manuals. They can serve several purposes. First, their mere existence signals that an organization has recognized the significance of a topic. They can also help to convey expectations and guardrails to staff. Additionally, they are an important feature of self-regulation and help to clarify accountability, even though only a few organizations seem to police their implementation. Northwestern University Professor Nick Diakopoulos, who analysed a set of 21 guidelines together with Hannes Cools of the University of Amsterdam, concluded: "At least for the externally facing policies, I don't see them as enforceable policies," he was quoted saying. "It's more like principal statements: 'Here are our goals as an organization.'"[13] This doesn't mean that guidelines are useless. As Becker, Simon, and Crum conclude: "And while the notion that AI guidelines in and of themselves will somehow magically resolve the intricacies of AI implementation and its attendant challenges is questionable, they can potentially make an important contribution in ensuring the responsible, ethical, and effective use of the technology in the news."

In November 2023, following the initiative of Reporters Without Borders (RSF), 17 media organizations from around the world passed the Paris Charter on AI and Journalism, the

[12] Kim Björn Becker, Felix Simon, Christopher Crum, "Policies in Parallel? A Comparative Study of Journalistic AI Policies in 52 Global News Organizations", SocArxiv, 6th September 2023. https://doi.org/10.31235/osf.io/c4af9, retrieved on 15th March 2024. (Quote page 24)

[13] Quoted in a summary of guideline analyses by Clark Merrefield, "Researchers compare AI policies and guidelines at 52 news organizations around the world", The Journalist's Resource, 12th December 2023. https://journalistsresource.org/home/generative-ai-policies-newsrooms/, retrieved on 15th March 2024.

first set of industry-wide guidelines for the use of AI. The EBU participated in its creation. It centres around accountability, human agency, and transparency and demands the inclusion of journalism support groups in the governance of AI systems. RSF highlighted the document with its ten principles as "the first global ethical benchmark for AI and journalism."[14] (Read the Paris Charter, page 179) Many EBU members have crafted guidelines for their organizations and made them available as a resource to others.[15]

In order to create useful guidelines, AI consultant Mattia Peretti recommends a participatory process: "It is tricky to strike the right balance and to not go too much on the side of concerns. But you have to treat people in the newsroom as adults not as children. If the latter is the case, maybe there was a problem in the newsroom before. Guidelines need to focus on the use cases and on general principles. You need to state: These are our values; this is our general approach. You need to rally the organization around this. Creating guidelines needs to be a participatory process. Pair it with some trainings. You want the newsroom to be informed."

German public service media BR, which is part of ARD, has some experience with ethical guidelines. The broadcaster was among the first news organizations worldwide to develop a set of rules for the use of AI which have been widely consulted internationally. These guidelines have contributed to the development of guidelines in the wider ARD group, which only started to develop AI guidelines recently. Uli Köppen and her team at BR have meanwhile developed internal guidelines which are more specific: "This is about how you use platform-based AI. What can you upload? Where do you have to be careful? What kind of data can you easily give to those platforms? And what kind of data please do not. A very important point I've added to our guidelines is how we report on AI. We don't report on AI as if it was human-like. AI doesn't think. AI doesn't do. AI is not a person. AI is not threatening AI is a tool." Given that debates led by journalists significantly influence the perception of decision-makers and the public, educating the newsroom about how to report on AI (and how not to) belong in all general guidelines used in an editorial context.

With its considerable AI experience, Yle also works with two sets of rules, as Johanna Törn-Mangs explains: "When we created these rules, we had different groups. A small management group with representation of all units, and a large group. Our AI Forum consists of 60 people, this includes not only journalism but also HR and other supporting functions.. One difficult thing was the level we put the rules on. They needed to be not too detailed but also not too general. We now have general rules and for some areas detailed instructions." She says that preserving public service value has been a key intention in developing the rules: "We do not just want to maximize our use but do what is the responsible way for a public media company. It needs to be aligned with our mission and values. We do not use everything that is cheap and easy. When we create boundaries for employees regarding which tools to use, it frees up a lot of energy, because they are not confused or unsure about what to use anymore. It helps to focus more on people than on technology. Trust is the most important of our values."

Jane Barrett of Reuters describes the news agency's attempt to make it simple: "We have four basic guidelines: first, it's a great opportunity for our journalists and journalism. Second, Reuters is always responsible for our output, whether or not generative AI was used in its production. Third, we will be transparent about where we've used generative AI. Finally, we will be increasingly sceptical given the rise of synthetic media. We said we will tweak the guidelines as new insights emerge. For instance, we are now fleshing out what

[14] "RSF and 16 partners unveil Paris Charter on AI and Journalism", 10th November 2023. https://rsf.org/en/rsf-and-16-partners-unveil-paris-charter-ai-and-journalism, retrieved on 5th May 2024.

[15] A collection of public service media companies' guidelines can be found here (some are accessible to EBU members only): https://www.ebu.ch/groups/ai-ethics#relatedPresentations-8a6ff688-cef9-4dd5-88d7-729213fceec2

human oversight of AI means in practice." (Read Q&A with Jane Barrett, page 134)

Many news organizations have decided to define a few red lines, particularly around how to deal with news photography and exclusive content. Verena Krawarik of APA says: "We don't use tools that are freely available. We always choose the paid versions and look at the terms of service because we have to make sure that the data will not be used as training material. In the past, tools were used without much thought. Also, we need to think about security. We analyse: What new artifacts are created in an AI-oriented workflow? Many artifacts like interview transcripts are now being created along the chain. It's about collecting them somewhere and not letting them float around freely in space. (...) We don't create images in our professional services and we don't manipulate them just because it's sometimes difficult to find a picture. Someone has to tell the truth, that's us."

Matt Frehner of *The Globe and Mail* highlights the importance of keeping a human in the loop: "If you publish something, you are responsible for it, regardless of the tools that you use. Using an AI tool to publish something that is not verified by a human being is completely unacceptable. We don't put any prepublication material into any language tool because we don't know how it is being used. We wouldn't use a tool to cut 200 words from an unpublished story investigation. On the visual side, we wouldn't use any AI photo generation for news photography. Anything that interferes with the truth we don't do." Mattia Peretti confirms: "Using generative AI tools for investigative journalism is risky. We need to be extremely careful what we type into a chatbot. My advice is always: If you wouldn't put it on a LinkedIn post, don't type it on ChatGPT."

Ritu Kapur of *The Quint* describes how guidelines need to match newsroom culture: "We have a very young and tech-savvy newsroom. It is the generation that was possibly born with devices in their hands. We went into developing the AI guidelines very quickly, because we have a very experimental-minded, very curious journalistic team. For some of the people their first job was at *The Quint*. They haven't gone through the old-style process of journalism. So, there was the risk of them jumping to using generative AI, and a lot of generative AI tools are free. A lot of them are very poor copies of original tools. So, they're very, very prone to error." But leadership encourages the staff to participate in the process and suggest new tools to be included in the range of what is acceptable.

Erik Roose of Estonian Public Broadcast, warns that leaders should be prepared for the new generation of students entering the labour market and joining newsrooms. "For them AI-tools will feel as natural as social media for the young generation that is in newsrooms now. There will be totally new ethical questions. Will they even be aware of having used AI to create something?"

The process of developing guidelines can be quite challenging. If the goal is consensus across the organization, the rules tend to reflect the lowest common denominator, with the downside of not having any teeth. In contrast, the road to more detailed and restrictive rules can be quite conflict prone. Agnes Stenbom of Schibsted describes how guidelines can vary within the news group depending on the publication: "There are different stances by different titles when it comes to the use of generative AI. *Svenska Dagbladet,* for example, agreed to never use photorealistic images." If staff gets confused, some ask a bot. Stenbom: "I really like how *Aftonbladet* has created their own GPT about their AI policy. This is an internal chatbot for their staff to ask any questions about their policy."

How much transparency does the audience need?

One of the major features that come up in most guidelines are transparency requirements. If something is written by AI, it is better to say so in the byline, or so the argument goes. Kai Gniffke of ARD says: "I don't think it's a good idea to limit change processes with red lines. But transparency is important. We must make clear where we have resorted to AI. For example, we

will not use images that have been modified by AI in news programmes such as the *Tagesschau*, Germany's most prominent news programme. And when we do make changes, we will indicate this."

On the other hand, if searches using LLMs become standard practice, it doesn't make sense to declare this each and every time. As Charlie Beckett says: "I think editors are a bit too worried about that. Today it doesn't read under texts: 'Some of the information came from news agencies' or 'The intern helped with the research.' Newsrooms should confidently use transparency notices to show consumers that they want to give them added value."[16] Erik Roose says it is important to divide tools into those that are rather technical and those doing creative work. "We don't tell people when we use something like a spell checker. If we inform people, it must be of some new quality."

Styli Charalambous of the *Daily Maverick* expects transparency labels to have a short shelf-life: "Once we will have AI baked into every single tool, I don't think we'll be making those disclosures anymore, because it will just be expected. We are not disclosing anymore that we wrote this using an electronic word processor after all." The public might be much more relaxed about transparency than some in the industry think. Nic Newman refers to findings from the 2024 Digital News Report that suggests this: "Interestingly, transparency is important but the public does not want AI labels everywhere, only when it is materially important. "Recent research by the Reuters Institute supports this view (see Figure XX), with almost 50% of those surveyed saying that news organizations should, for example, disclose if a text was written with the help of AI but only 32% saying that copy-editing with AI needs to be labelled.

Verena Krawarik says that the value of transparency is very much connected to

Figure 9: Audience attitudes in six countries towards AI-disclosure.

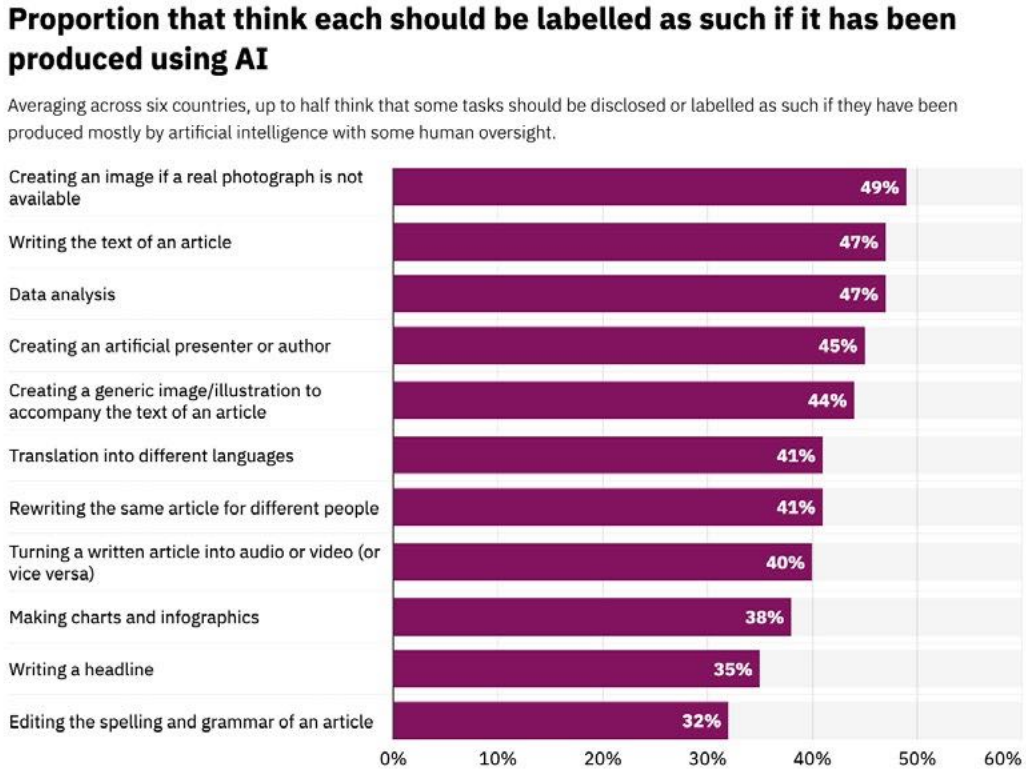

AI_news_labelling. Which, if any, of the following should be disclosed or labelled as such if it has been produced mostly by artificial intelligence with some human oversight? *Base: Total sample across Argentina, Denmark, France, Japan, UK, USA = 12,217.*

Source: Data from 'What does the public in six countries think of generative AI in news?' published in May 2024

Source: Reuters Institute for the Study of Journalism

[16] Alexandra Borchardt (2023), interview with Charlie Beckett.

[17] Fletcher, R., & Nielsen, R. K. (2024, May 28). What does the public in six countries think of generative AI in news? Reuters Institute for the Study of Journalism. Retrieved from https://reutersinstitute.politics.ox.ac.uk/what-does-public-six-countries-think-generative-ai-news#header--10

audience perceptions: "I am very much in favour of transparency, but it is of course bad news if there is a perceived loss of quality when we indicate that co-working with machines has taken place. People out there no longer know exactly what journalism is, how a journalistic contribution differs. That's why I believe that generative AI can lead to an important debate about who we are, what we want to be, what our craft looks like. It is very important that we succeed in this. One of the counter-movements is that we identify our content as journalism-made."

Transparency needs will also depend on the level of tech literacy and culture of the audience in a certain environment. Johanna Törn-Mangs reports how an AI-assisted playwright wrote a play for Finnish Yle's streaming services. "All the actors were synthetic voices, all the dialogue was written by AI, of course, with the help of humans. We were open about the fact that it was made by ChatGPT, people knew. Some thought it was weird, some thought it was surprisingly good. Yle has done different AI-based solutions for the past 20 years. People are used to that. Technology is regarded as a good thing in Finland."

Madhav Chinnappa thinks that public service media have a special responsibility to be transparent: "What I would be advising public service broadcasters to do is what I call total transparency. They should be transparent in how they use generative AI and when in doubt, go further than they need to. But also, I think audiences are more sophisticated than we give them credit for sometimes. They understand when you differentiate things. So, if you had, for example, a lab section on your website, where you're saying: 'hey, this is where we're just testing stuff' they might understand. This may allow you to have a test environment where you could do more and learn more as opposed to doing this behind the scenes, getting it completely perfect and then launching it. I would advise public service media to establish that kind of transparency in their product development and experimentation."

There is no such thing as a blueprint for guidelines, nor is there a recipe for success in the AI-conscious newsroom. The debates within editorial departments and between the editorial and business sides can be heated. Particularly when it comes to which approach to take, as well as what is ethically advisable, legally safe, and desirable from a newsroom point of view. A list of some hotly contested topics can be found below.

Many organizations therefore opt for allowing guidelines to be open for amendment. There will doubtless be plenty of new evidence from practical experiences, research, and technological advances in the near future. Though it might not feel like it for those concerned about the viability of their professional roles, it is still early days in the age of generative AI.

It is high time to take an ethical stance on many of the critical issues that will shape the future of generative AI – for news, for journalism, and for the creative industries as a whole. Media organizations have to equip themselves with basic knowledge to be able to engage in an informed dialogue with the tech industry and develop their own vision of a sustainable, ethics-proof information ecosystem. Chapter 4 will explore some of the ethical building blocks.

Six hot AI strategy dilemmas for news publishers

- **Workflow efficiency VS audience-facing products.**
  AI can enable both. What is the right balance for your company's resource allocation?
- **Synthetic media VS human-only content.**
  AI can increasingly create realistic voice-clones and artificial video/ images. Should these be published?
- **Transparency VS AI as a normal feature.**
  Some publishers and social media identify AI-generated content. Does it increase trust?
- **Scraping blocks VS updated journalism in AI-answers.**
  A 'no' to scraping can make content invisible in AI-search. Does this pose a long-term problem?
- **Partnerships VS no deals with AI-giants.**
  More and more media companies receive funds to allow content training. What will be the consequences?
- **Build own AI-tools VS off the shelf.**
  AI-interfaces ensure more customised tailoring, external tools have state of art tech. Which is the right path for you?

Q&As

Q&A

> "Journalistic production will become a by-product"

NIDDAL SALAH-ELDIN

Member of the Executive Board, Axel Springer

In which ways is generative AI a game-changer for journalism?

It's giving us new opportunities to streamline our processes, making our newsrooms and operations more efficient, and to create new experiences. This enables teams to focus more on the heart of journalism: creating compelling stories and digging deep into investigations. Plus, AI enables us to offer a wider range of content, from special interest pieces to more localized stories, which opens additional avenues for advertising revenue. And let's not forget about data journalism. Recognizing patterns in large datasets is invaluable for uncovering important trends and insights that might have otherwise gone unnoticed.

What are Axel Springer's hopes and expectations for generative AI?

We're convinced that AI offers great opportunities, and we want to lead the way in embracing them. AI will revolutionize journalism and all parts of our business. While we're mindful of the challenges it brings, we're really excited about the possibilities it presents for our core business, which is journalistic creation. Our focus is on researching exclusive news, highlighting personal experience in features, as well as providing original commentary. Thanks to generative AI, in the long run, journalistic production will become a by-product, more technically supported and automated. Understanding and adapting to this shift is crucial to sustain our business.

What kind of mindset and behaviour do you encourage in the company?

Ever since Axel Springer founded his company in a barn, a pioneering spirit has been part of its DNA. We will not wait until all the questions AI has raised have been answered. We want to experiment, encourage and empower our staff instead of observing disruptions from the sidelines.

How are you going about this?

Upskilling continues to stay on top of our global talent and culture agenda to meet the changing conditions and requirements within the industry. This means that we're not only looking at this from a technological angle but also a cultural one. We're focusing on identifying and closing skills gaps. We're convinced that creating enthusiasm for AI throughout the company and empowering employees accordingly will provide a crucial competitive advantage. Lifelong learning and openness to

technology and tools will make the difference. In this spirit, we've sent a handful of executives and experts on an AI expedition to existing and emerging AI hubs worldwide to enable them to immerse themselves into local AI scenes and explore opportunities for Axel Springer. The destinations include San Francisco, Seoul, Tokyo and Singapore. Having boots on the ground and being close to the builders will help our businesses as we share the knowledge gained in these fellowships broadly within the organization so that everyone can benefit from their experience.

The New York Times and others sued OpenAI; you struck a deal. Could you explain the reasons behind that decision? What are your hopes connected with this deal?

Our landmark partnership with OpenAI marks a paradigm shift in journalism which we're proud of. For the first time ever, we're seeing a revenue stream from an AI company to a media company for the use of recent content, establishing the principle of remuneration. Our partnership with OpenAI has opened a path that we hope many other publishers will follow along.

Can you talk about what OpenAI will give you in return?

The deal has strategic value for us. In addition to the revenue stream that we've established with this deal, our partnership with OpenAI will further increase the visibility of the exceptional reporting that our journalists do and introduce their work to new audiences.

Who at Axel Springer makes decisions concerning AI? How have you institutionalized these?

AI is a top priority at the Executive Board level. We ensure that our strategic initiatives are aligned with our long-term goals. This means we continuously iterate and update our approach. Year one of generative AI was focused around creating momentum across the organization. This was centred around a three-pillar approach comprising education and exchange, optimization and exploration in a hybrid setup which combines central offers and initiatives from the HQ and empowers decentral initiatives steered by the business units.

Could you be a bit more specific?

In the education and exchange pillar, we broadly equip employees with the skills needed to be ready for tomorrow's disruptions and connect our experts with each other. Various exchange events throughout the year give our experts the opportunity to learn from each other. Our annual Media & Tech Con brings together over 1,000 colleagues in Berlin for a day of sharing best practices, learning and inspiring each other. In the optimization pillar, we're improving products, processes and rethinking business models, right at the core of our brands and units. And lastly, in the exploration pillar, we're using the Group's global power to advance in generative AI. In spring of 2023, we founded our global generative AI team as a ramp-up and to create momentum. In the autumn, we started the AI ambassador network, where we connect dedicated ambassadors from each of our biggest brands in regular exchanges and best-practice sessions.

What's your favourite generative AI product or use case – in your company or beyond?

There's a wide range of different use cases for generative AI at Axel Springer. We've dedicated an internal website to our use cases that employees can screen to get inspired and find best practices that they can adapt to their own products and processes. So there are many great and creative use cases. One example is how Business Insider Germany increased their efficiency through using AI. The team uses the technology to almost completely automate the production and distribution of their stories. This has enabled them to launch several audio and video formats tailored specifically to younger audiences that will be key for further growth and for ensuring that the majority of their resources go into the core of journalism. Another example is the Content Analyzer, a tool that was developed in-house and is used by editors. It can generate suggestions for headlines, SEO lines, conversion outlines and social posts, or run a text through a "Wolf-Schneider filter" [named after a late German journalism school director well-known for his profound language criticism].

Will these efficiency gains cost jobs? Your CEO Mathias Döpfner was probably the first German media manager to be outspoken about job cuts as a result of AI.

For sure, some jobs will cease to exist in the future, while new jobs and profiles will emerge. This has always been a natural consequence of technological progress. I mentioned before that journalistic production will become a by-product. There are many things AI can do more efficiently than humans. However, what journalists will always do better is research exclusive news, write surprising commentaries, and conduct inspiring interviews. That is why we're concentrating on the core of journalism.

Do you think journalism will develop from being a push activity where news are pushed at people to a pull activity when people will demand customized news that fit their needs?

Absolutely, I think this shift is nothing new. It has always been the purpose of journalism to make stories so relevant that people actively pull them into their lives. The distribution channels have evolved over time and we've adapted to that. And we'll continue to do so. With the help of AI and customization as well as data to better understand our diverse audiences we'll continue working on new products and channels to serve our users in the best way possible.

Many people are worried about misinformation. Are those fears justified or overblown, and have you encountered examples that worry you?

AI is a product of human formulations and algorithms. It's fallible. AI hallucinations are part of reality. That's why it's so important to develop the necessary skills to ensure a responsible use of AI. Our employees need to be able to recognize misinformation. We promote that, for example, in our upskilling seminars at the Axel Springer Academy of Journalism and Technology. There, we offer training sessions on OSINT (open source intelligence) and fact-checking where employees learn techniques

for verifying AI-generated content. It includes identifying suspicious patterns and inconsistencies, checking internal and external logic, and exposing manipulation and misinformation.

Some of the dynamics are beyond the influence of the media industry. In which ways do you think AI should be regulated?

We're aware of the challenges and currently taking a very close look at aspects such as data protection, regulation and fair remuneration for the use of our content as training data. For us, this presents an opportunity to avoid repeating the mistakes of platform regulation and create a fair and healthy ecosystem from very early on. Journalism is a part of the value chain and this needs to be reflected. To achieve this, we need a triad of competition law, copyright law and data protection law. Of course, we are striving for a fair balance of interests between platforms and publishers, and our partnership with OpenAI is a great example of that. A deal like this requires the willingness of all parties involved to achieve good results. There are several initiatives in this regard, and much is evolving.

Q&A

> "We must guard this trust like the apple of our eye"

KAI GNIFFKE

Director General of German SWR and Chairman of ARD

Would you call generative AI a gamechanger for journalism?

It is definitely an incredible accelerator of change. It is now up to us to use this new technology responsibly, especially as public service media.

Are you rather delighted or worried about what it can do?

I always try to stay curious and ask: what are the opportunities? How can AI help our journalists with their research, for example? But of course, it is in our DNA as journalists to be sceptical and consider the risks. I recently started a meeting with a video of myself greeting everyone. I had recorded it and then had an AI read it in six different languages. I admit, I enjoyed watching myself speaking fluent Mandarin and my lips didn't miss a beat. At the same time, I imagined it could as easily be Joe Biden 'announcing' that he was about to attack Russia. That is not a future scenario of what these technologies might be able to do, that is what we are dealing with today.

You are Director General of Southwestern German SWR and chairman of ARD, one of the biggest public service networks of the world. What kind of mood do you sense in your organization when it comes to generative AI?

It varies and is not even a generational issue. Some people's eyes light up and they say: great, things are moving forward here. But there are also sceptics who say it's very dangerous. My job then is to remind them that if we don't tackle AI and learn how to deal with it, we'll leave it to others who won't feel committed to democracy or a sense of togetherness and community as we are.

What do you hope to achieve with AI?

Overall: improvement of quality. First, it could help us detect misinformation. One of our main tasks is to distinguish between fake and facts, truth, and falsehood. This is our job; this is the service to our audience. AI could be the basis for perfect fakes but also for debunking lies and misinformation. Second, data journalism. AI will help us deal with very large amounts of data. I'm convinced that this will deepen and therefore improve our research. Third, regional reporting. With the help of AI, we can cater much better to local needs, for example, tell people in the Eifel region or the Black Forest what a certain development means for the future of their area. Fourth, efficiency gains. AI relieves us of

routine tasks, for example, evaluating content from our archives. AI can transcribe speech to text. It can also recognize and index visual content. Today we still need staff to do that. AI can shorten and summarise texts. That often clarifies them and saves time when studying things. Last, not least: accessibility. We can use AI to generate audio description of visual content.

Are you worried about the role of humans in all of this?

Of course, we are discussing the role of humans in the future. For example, we are debating whether we can use AI to automate weather and traffic reports in night-time radio programmes. We believe this is responsible if it is done transparently. However, it would also be possible to have these reports read with the voices of the most popular presenters. But then we're entering a grey area because that could confuse people. They might think: "What, is she also awake at night when I always hear her on the morning programme?" If you listen to your favourite voice 24/7, it might lose its value.

This sounds like job cuts, too. Will job profiles change?

Journalists have had to deal with technology for many years. Today, they need to know much more precisely than before: what audience are we working for and on what platform? AI adds a new quality. But that won't change the basic virtues. The professional handling of information, conscientious research – that will stay.

What specific actions have you already taken in your organization?

At SWR we have developed guidelines for dealing with AI. We are now in the process of doing this throughout ARD with its nine independent media institutions. We need common standards for all of them. We also have set up a network of competence to bring employees in different roles and in various departments together. All member broadcasters of ARD who are experimenting with AI should know about each other. Not everyone needs to reinvent the wheel. It's all about sharing knowledge, lessons, and experience. There are different speeds within organizations, of course, the know-how of the front runners should spread quickly. What we have also set up at ARD level is a second competence centre for reporting: This is where journalists who report on AI and AI-related trends and developments meet to keep each other up to date.

What is the biggest challenge in managing AI in your organization?

Actually, getting the people who want to work with it to do so. They shouldn't have to wait for instructions by top management.

Are there red lines in the use of AI that your colleagues are not supposed to cross? Some news organizations have clearly defined those.

I'm not in favour of limiting change processes with red lines. But transparency is crucial. We must clearly inform the audience about every instance where we have resorted to AI. For example, we will not use images that have been modified by AI in news programs such as the *Tagesschau*, Germany's most prominent news programme. And when we do make changes, we will indicate this.

Do you think AI will help journalism develop from being a push activity – news is pushed at people – to a pull activity: people will demand customized news that fits their needs? Some hope that this kind of on-demand journalism will help them to reach different audiences, particularly younger ones.

AI offers huge opportunities for personalization as long as we offer our users a variety of perspectives. Then again personalization might lead to isolation. This makes large events where society comes together even more important. I'm thinking about big sports events, major shows and so on. But I'm actually quite optimistic. Just imagine how the media world has changed since the launch of the iPhone in 2007. In a tiny amount of time, we have massively changed our communication behaviour. We take photos everywhere and post them, use social media and apps for everything. And so far, we've managed quite well. The venerable *Tagesschau* is the most successful German media brand on TikTok and Instagram. What we do have to worry about, however, is the incredible acceleration of technological development.

Will the dependence on big tech increase with AI, or can big media organizations even gain something? After all, they are sitting on a huge treasure trove of content.

We have been confronted with this development for many years. Most of us work with Microsoft products. Of course, this makes us dependent on one company. Nevertheless, it is easier if everyone is on Microsoft Teams. The same applies to social media. The platforms with the highest reach belong to Meta. But what would the alternative be? Saying goodbye to the audiences whom we can only reach via these platforms? It will ultimately come down to what regulation will look like. Only the EU can prevent us from total dependency. However, as a major producer of valuable content, it must also be in our interest to make it available to AI, at least for training purposes.

When it comes to copyright, do you lean more towards the German publisher Axel Springer and others, who have struck deals with OpenAI, or the *New York Times*, which has filed a lawsuit against the company?

We are somewhere in between. We certainly won't be going to court.

In most countries, public service media enjoy the highest levels of trust with audiences. In the context of AI, there are two schools of thought: One says that AI will destroy trust in the media altogether, because no one will be sure what is true and false anymore. The other argues that this is a great opportunity for quality media, particularly those brands who enjoy a high level of trust.

I am still under the impression of what the Club of Rome said last year: the biggest threat to our societies was the increasing inability of people to distinguish reality from fabrication, facts from lies. This could destroy societies and communities. I would count ARD among the institutions that people trust in this country. So that they don't question the truthfulness of every video. That they say: "These are big brands, they've never lied to me. If something is important to me, I'll go to them."

That was the pandemic effect: in the first year of Covid, trust levels in traditional media skyrocketed.

That is true. We must guard this trust like the apple of our eye. We must be and remain a reliable companion and verifier for people.

Q&A

> "We have a moral duty to be optimists"

EZRA EEMAN

Strategy and Innovation Director, NPO

In which ways do you think generative AI is a gamechanger for journalism?

Down the line it will certainly impact every aspect of the journalism value chain. It will take a while to move from shiny new things to changes of workflow, but it will redefine who is media, how it is created, who is a media maker, how value is created.

What won't change?

AI is not good at reporting the human-centred side of journalism. AI cannot handle hard and live news very well. It is good at structuring language, that means it is good at text analysis, it can do summaries, service content, it works good for search optimization. We will see workflows that skip the creative process, because AI models can generate output directly from raw data. It impacts personalization, how news is presented. Why present it in an article form when you can have a conversation? Then again people like current experiences and will go back to them. People like newspapers, there is even a revival of magazines.

Are you delighted or worried about generative AI for your company and in general?

I would say I am a pragmatist. We have a moral duty to be optimists and convey a sense of opportunity rather than despair. With generative AI we can fulfil our public service mission better, it will enhance interactivity, accessibility, creativity. AI helps us to bring more of our content to our audiences. My biggest concern is that it will decrease the trust in information systems even more. The feeling that you cannot believe your eyes any more will also reflect on trusted brands.

You are talking about deep fakes and misinformation?

The danger is that the narrative of misinformation itself will impact trusted environments. Distrust then becomes the default mode for any news.

How can news organizations fight this potential loss of trust?

The inherent danger with AI is that it creates more distance. It takes out the human element, the boots on the ground, the assurance of reporters who say, 'I understand your reality, I am here to listen.' That is the role public service media could and should

play. We also have a responsibility to bring everyone along. Technologies create gaps. There will always be a part of the population that doesn't understand it. Explain how you use it.

What kind of mindset and behaviour do you encourage in the newsroom and company?

We usually see five to ten percent of early adopters and geeks, then a broader group that is not very negative but reluctant. At this point we are trying to foster as much understanding as possible. We are doing a lot of training in the company. People get hands on experience. From there we need to have a more strategic dialogue. How can we take away pain points?

Can you tell us a little bit about what you are already using AI for and what you are exploring?

We are using AI on backend processes. This is about productivity, efficiency, we reap the low hanging fruit where can machines do a better job: transcription, archives, meta data, subtitling. The second category is where we can add intelligence. We are exploring how we can unlock some of our video archive of one of our consumer programmes with a generative AI interface that allows you to ask question to that specific archive. The idea is for users to have a conversation with our archive rather than entering search queries. Additionally, we are reaching out to audience groups that we haven't served very well by translating news in simpler language. We use AI tools to improve our workflow by breaking down complex words in easier options. To help hearing impaired kids, we produced podcasts with generated video so they could follow the narrative better. We are also exploring synthetic radio voices but have yet to define where and how they could be of use.

What's your favorite generative AI product or use case – in your company or beyond?

The easy language offer is a very nice example. My favourite is maybe: We had a recent podcast on the murder of JFK. William Altman, a Dutch journalist who recently passed away, had direct leads to witnesses, he kept excessive diaries about this. At the anniversary of the assassination, we recreated his voice to reconstruct the investigation, this was fascinating. We consulted the family first, of course.

What is the biggest challenge in managing AI in your organization?

Encouraging people to experiment but not put it out for production that can be a challenge. We allow for failure and don't expect perfect output. Still, we had a warning shot: in one of our TV news bulletins an image was taken from an image bank, the picture was a generative AI picture. Some of our viewers saw it.

Do you have AI guidelines – and what's special about them?

NPO is an umbrella organization with 13 broadcasters, they are all independent. The newsroom guidelines are set by independent organizations, but we have umbrella principles that define how we want to work with generative AI as public service

media. There are three broad categories. First, we care about our audience by being transparent and don't neglect the human aspect; second, we are committed to quality, like reliability and accuracy and third, to ethical values. This means we make sure that we use it for good and minimize bias and harm. We are worried about the climate impact of these technologies, for example, and have an institute monitoring sustainability.

Let's talk about the business side of this. Do you think companies should do deals with OpenAI or others as the news agency AP and German publisher Axel Springer have done?

This makes sense for these big companies, but there are just a few companies out there who will be able to negotiate these deals.

Do you see a space for the EBU to negotiate on behalf of public service media?

I was Head of Digital at the EBU for five years. You are getting a seat at the table, but it is difficult to have a common agenda. The bigger parties like the BBC or France Télévisions have their own agenda. In some countries the debate about technology is more advanced like in the Nordics. Others are more reluctant.

Some of the dynamics are beyond the influence of the media industry. In which ways do you think AI should be regulated?

It helps for the EU to set certain guardrails. Companies look for the EU to set safety regulations, transparency requirements. It will never be quick enough, but it is good that it is there. The agenda I am more interested in is: how can we shape a stronger European media innovation landscape? There could be European language models, collaboration on data sets, more media innovation funding. With the European elections there is an opportunity to shape that innovation agenda.

There is a huge AI hype going on in the media industry. What is missing from current conversations?

A sense of reality. We are still surfing the hype wave. We talk a little about the nitty gritty details that are needed to go from strategy to implementation. More importantly, we have to ask: How can we deliver value with this, what is really necessary, what is our ambition and vision with this? Our aim should be to deliver value in a more granular way to those people we are missing out now. We still have a long way to go from a broadcast model to a model that is involved in people's lives.

Q&A

> "Human beings will always be in control of our journalism"

MANUELA KASPER-CLARIDGE

Editor-in-Chief, Deutsche Welle

In which ways is generative AI a game-changer for journalism?

As with previous disruptive technologies – the internet, social media, and smartphones – we are expecting generative AI to change people's media use and that will bring new opportunities and challenges. We're aiming to have AI support our work, to automize a number of regular tasks and leave our journalists more time to focus on storytelling and creative work. In some cases, generative AI can help with that creative work too. For example, you can use it as a tool to help you craft interesting, SEO-friendly teasers or as an extra sparring partner to develop story ideas targeted towards specific audiences. But we're also already seeing a rise in the amount and the quality of misinformation.

Are you delighted or worried about generative AI for your company and in general?

I would say the word delighted is definitely going too far. I think we have a broadly positive attitude towards generative AI, while at the same time we are considering the limits we need to impose on its use and analysing the risks it poses to journalism. Those limits include publishing anything generated with AI without it being checked by a journalist or publishing photorealistic images. However, I think you have to find as many opportunities as you can. You can't look at the topic too fearfully.

What kind of mindset and behaviour do you encourage in the newsroom?

We are striving to find the right balance between freeing ourselves to test as much as possible and making sure we use the tools responsibly. Colleagues are always free to make suggestions. However, we have committed ourselves to always having a human being in control of every piece of journalism we produce. We are also paying very close attention to data privacy. Data protection is a very big topic in Germany, and we have to make sure everyone who works with AI tools has undertaken the relevant online training. Our legal department has published guidance on using AI chatbots.

Is this more of a top-down or a bottom-up endeavour?

I want our journalists to try things out. I want people to discover things and tell us what they think works. Our teams that have tested chatbots and AI tools so far have collected a significant amount of information demonstrating what works and where

they see the opportunity to have AI support their work. That feedback is so valuable. We have multidisciplinary teams working on and giving feedback on projects. You need people with different backgrounds and different experiences working on projects and their prioritization together. You also need to share expertise. We have a DW-wide AI Circle that meets once every two weeks and brings colleagues from multiple departments into project groups. My Editor-in-Chief's Council is also following the subject closely.

What's your favorite generative AI product/use case – in your company or beyond?

As an international broadcaster that publishes in 32 languages, the rapid development in AI-supported translation and voicing is very exciting. It has the potential to save us a lot of time translating and revoicing our journalism from one language into another. Still, these translations and voiceovers would need to be checked by an editor. We developed an AI-powered content adaptation platform, plain X, which helps with this. It is integrated into our editorial systems, bundles various tools in one interface and offers lots of options for transcription and subtitling of videos, as well as other AI-based services. The potential to use AI to make more of our content completely barrier free is also exciting. More subtitling is the obvious way to go, but having quality AI sign language could be very useful in the future.

Many worry about bias, particularly in AI-generated images. Do you see a danger in further scaling stereotypes or an opportunity to fight bias with AI?

We need more data on that. In our test cases we see a huge bias in AI generated illustrations. For example, if the topic is domestic violence in the Arabic world, women are always pictured with a hijab.

What is the biggest challenge in managing AI in your organization?

It's a massive topic with a lot going on all at once. People get information about developments from different sources. It's very difficult trying to keep everyone on the same level of understanding, with similar amounts of knowledge. Communication between people and with the wider organization is vital to let people know what we are working on.

Have you made mistakes with AI strategy?

We say that we want to work quickly and flexibly, but we have over 3,000 employees and as is the nature in large organizations, we sometimes aren't able to start projects or react to developments with the agility that we would like. Communication can be difficult. Obviously not everyone is working on AI projects, and some people sometimes hear about what is going on through the grapevine. This can be unsettling for colleagues who have heard about how AI might be coming for everyone's jobs – which I do not think is the case.

What about talent? Some expect that journalism will be researching the facts, storytelling will be done by AI with personalized for different audiences. Do we need different types of journalists in an AI-supported media world?

What we still need is journalists on the ground who talk to real people and deliver stories about humans. These are the kind of stories that AI cannot deliver and are classic skills of journalism we cannot lose. Journalists will have to learn about AI on top. They will need to write prompts, identify sources, understand AI, identify what is real and what has been generated or faked. It is very important that we train our journalists in this. Young colleagues will grow into this world naturally. I'm a mother of three children, they all know how to write prompts.

Do you have AI guidelines – and what's special about them?

We have strategic guidelines that were issued by our Business Management and then my Council and I released editorial AI guidelines. They outline our position on generative AI and explain our rules. For example, we state clearly in the introduction that human beings will always been in control of our journalism, we outline exactly what kind of information may and may not be used in prompts, and we link people to the necessary training. We also outline what will guide our future approach – transparency, control, and data security. As with all of our editorial guidelines, it is a 'living document' that can be updated at any time.

Do you think journalism will develop from being a push activity where news is directed to the audience by the media to a pull activity where people choose customized news and formats to fit their needs.

As and when chatbots become the main way that people find their information, their relationship with news will change. It's likely they will be able to ask questions about news events and stories much more easily, and more context will be at everyone's fingertips.

Many people are worried about misinformation. Are those fears justified or overblown?

I think those fears are very real. It's clear that the quality of fake news and deep fakes will only get better, and they will become easier to produce. It will take effort to counter false narratives as they spread. It will likely take a combination of good journalistic training and helpful technology. We will also need to reassure audiences about what is real and how they can trust our information.

Do you think generative AI will impact audiences' trust in journalism?

I think that in the age of chatbots, being able to show we have reporters on the ground, correspondents around the world, talking to people and telling human stories, will be extremely important for maintaining audiences' trust in quality journalism.

Deutsche Welle is operating globally. Do you see differences in the acceptance and uptake of AI around the world?

The internet is not as fast or as affordable in every region we cover. The biggest divide is age and wealth. If you are younger, you are more open to new technologies, if you are wealthier, you have better access. In Africa, for example, people living in cities have good access to the internet, but it is mostly expensive.

Some of the dynamics are beyond the influence of the media industry. In which ways do you think AI should be regulated?

Transparency is very important, as is human oversight. Ideally, this is what would be a sort of standard in the future of AI, especially for news and media. We want transparency about where AI has been used to produce content and for the chatbots to be able to reliably link to information sources. The EU's AI Act envisions some of this for high-risk AI systems, but we need as many people as possible to be obliged to uphold transparency.

What is missing from conversations in the current hype?

I think the constructive approach to generative AI is missing too much. The companies developing the large language models are obviously focusing on the positives and the opportunities it offers. At the same time, there are many people who are focusing entirely on the negatives, from the amount of misinformation that may be created to the possibility that AI systems could turn against humanity. We need to have more balanced conversations about it.

Q&A

> "We have to educate ourselves about AI, and then report the hell out of it!"

JANE BARRETT

Global Editor, Media News Strategy, Reuters

In which ways is generative AI a gamechanger for journalism?

It's a massive gamechanger for all industries. It is an entirely new way of human beings interacting with data and information, and we are right at the beginning of the journey. Anyone who predicts today what it is going to look like in five years is probably a fool.

How could journalism benefit in the short term?

I see things in three buckets: reduce, augment, transform. First, how can we use generative AI to reduce our journalists' workloads? What repetitive jobs can AI help us do? We have started experimenting on these quickly. We have prompted GPT-4 to help us do a first edit on a story, extract facts from statements, brainstorm headlines, translate stories better. Replacing routine tasks at scale may take more development work but already, AI can speed us up and help us do more with the resources we have. The second opportunity is augmentation. For instance, we can take the reporting we have now and make it available to more people in the way that they want it. AI skills might help us re-version a story into social posts, a video script, a quick summary for busy readers, a translation. Or AI could augment our work by helping find stories in data dumps or write explainers from our archive.

Number three, transformation, sounds like digital transformation all over again.

It really is. As an industry, we can learn lessons from the past waves of digital transformation and be ready to move more nimbly this time round. Internally, how might we re-think the value of each part of our workflow because of what AI can do – for us, our clients, and audiences? Externally, how will the whole information ecosystem change? Will audiences' expectations and behaviour change again because of how AI shapes the rest of their lives? What does that mean for our business models?

At the same time, we have to tread very carefully because as journalists we deal in facts and generative AI models are prone to hallucination. I liken today's generative AI models to a Formula One car. However well you drive, you need to train to get behind the wheel of an F1 Ferrari and not crash. And you need a team of excellent technologists, and in AI, data scientists around you to get to where you want to go safely. It's not a silver bullet or a quick solution to our problems.

What is Reuters using generative AI for already?

I mentioned some of our experiments earlier. We are now building out some of those, testing and integrating them into our editorial tools. We are also training our staff on prompts and have built a prompt builder to help them do that. That has yielded some good successes, for instance in doing a first copy edit or summarising stories into background paragraphs.

Also, we have a tool in Reuters Connect which provides video transcription, translation, shot-listing and facial recognition. It carries a clear disclaimer that the work has been done by AI and makes our content easier for clients to use.

Are you delighted or worried about generative AI for your company and in general?

I'm generally excited but I'd be lying if I didn't admit some concern. This is another huge disruption for the news industry after the explosion of the internet and mobile, search and social. How does it affect our business this time? I also worry about society, given the lack of trust in journalism and even facts. Just as AI can improve efficiency at Reuters, it can also allow bad actors to create convincing misinformation at scale, either misleading people or just confusing everybody as to what is true.

What is the media's role in this?

We have to educate ourselves about AI, and then report the hell out of it! That is the one tool we have that nobody else does – the power of reporting. Generative AI is going to be one of the seminal changes of our lives and we need to turn all our investigative and analytical power on to it to tell the story, hold these new AI powers to account and inform people about how the tools work.

What kind of mindset and behaviour do you encourage in the newsroom and your company?

My big word is play. We have a private version of ChatGPT so it's a safe playground. Come in and have a go. Do some training, see what is possible. Keep your mind open, share what you found. We have a great cohort of early adopters and others who are keen to get going. Of course, there is always fear about what change will mean for our jobs but again, we just don't know yet. The important thing is to get involved. As our CEO says: generative AI won't take your job, but someone who knows how to use it will.

What is the biggest challenge in managing AI in your organization?

The biggest challenge for me right now is prioritization. What do we take from the experimental phase into production. Our newsroom has come up with so many great ideas. But it takes a lot of work to take something from a basic prompt, test it, integrate it into the workflow. Even more if you are fine-tuning a model or building more complex systems.

Have you made mistakes with AI strategy?

In any change you have to communicate, communicate, communicate. Particularly with something so new and powerful, we can't speak to and listen to our teams too much. Getting that right is critical.

Do you have AI guidelines – and what's special about them?

We have four basic guidelines. First, it's a great opportunity for our journalists and journalism. Second, Reuters is always responsible for our output, whether or not generative AI was used in its production. Third, we will be transparent about where we've used generative AI. Finally, we will be increasingly sceptical given the rise of synthetic media. We said we will tweak the guidelines as new insights emerge. For instance, we are now fleshing out what human oversight of AI means in practice.

Do you think journalism will develop from being a push activity with news directed to the audience to a pull activity where people demand customized news that fits their needs?

I suspect so, for two reasons, First, it has always been a pull activity. Nobody reads newspapers from cover to cover. You choose what to read. Second, it has already changed with search. Search answers, now, are a long list of links. The generative search experience feels like a natural next step. How we watch TV and use audio, the whole world has become much more of a pull world.

Some of the dynamics are beyond the influence of the media industry. In which ways do you think AI should be regulated?

I find it useful to have cross-industry conversations. Generative AI is impacting every business: medicine, law, logistics, finance. Journalism is not exceptional. There are already regulations around data and privacy, copyright and the like. So it will be interesting to see how those develop in the world of AI to start with as well as some of the newer conversations about responsible tech.

There is a huge AI hype going on in the media industry. What is missing from current conversations?

I suspect we're not looking hard enough at the transform bucket. The natural tendency is to want to solve today's problems and it is hard to imagine tomorrow. We need to get out of our bubbles and see the possible. How do high school students or first year students at universities use these tools and interact with information? How is that going to change things? We need to get a good balance between solving today's problems and preparing for tomorrow's world.


CHAPTER 4
THIS ISN'T REAL.
WHAT IS REAL?
HOW DO YOU DEFINE 'REAL'?
HERE

THE BUILDING BLOCKS OF RESPONSIBLE AI

Ethical considerations are inherent to the very concept of journalism. To that end, it's debatable how meaningful terms like 'responsible AI', 'ethical AI', or 'trustworthy AI' are in the context of the media, where ethics should be the default option anyway. Even commercial media players are bound by the values that distinguish journalism from other content, define its legitimacy and justify its constitutional protection. Further, it is open to discussion whether media companies can exert influence over the ethical validity of their AI-driven products when the underlying technology created by major tech companies is not ethical. Products and services developed by these companies shape expectations, habits, and consumer behaviour – often contrary to what other actors might want. It remains to be seen whether the news industry is influential enough to impact the broader debate on this topic.

Critics posit that ethical considerations are generally a mere afterthought for big tech. As Mattia Peretti says: "Responsible AI is a buzzword the tech companies love to use. But everything we do with AI needs to be done responsibly. If the going gets tough at tech companies, the first team to be laid off is the responsibility team." These concerns have been expressed by policy makers and international regulatory bodies. Looking at the European Commission's Ethics Guidelines for Trustworthy AI, Eugenia Stamboliev and Tim Christiaens concluded that these reduce "AI ethics to the supporting role of a fire extinguisher subservient to AI industry's projects."[1]

Nevertheless, media organizations – and in particular public service media (PSM) – have a responsibility, if not a mandate, to use AI responsibly in the public interest. This can also mean not using it at all, or at least restricting or ruling out its use for certain applications, workflows, or processes. As the world's largest association of PSM, the EBU established an ethics group, to evolve the debate. As Erik Roose, Chairman of the Board of the Estonian ERR says: "There is no ethics in AI. Ethics is on this side of the laptop."

Law professor Natali Helberger has participated frequently on the Council of Europe's expert committees that provide recommendations and guidelines for the responsible use of digital technologies and AI. She says an ethics debate in the media

Figure 10: Strategies for ethical use of generative AI among a non-representative sample of international journalists

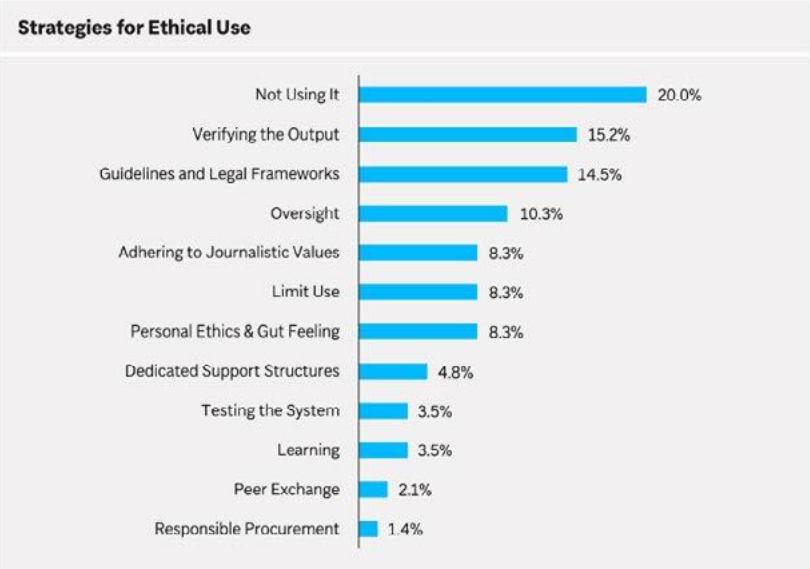

Source: Generative AI in Journalism: The Evolution of Newswork and Ethics in a Generative Information Ecosystem", April 2024.

[1] Eugenie Stamboliev, Tim Christiaens, "How empty is trustworthy AI? A discourse analysis of the Ethics Guidelines of Trustworthy AI," Critical Policy Studies, February 2024, https://doi.org/10.1080/19460171.2024.2315431

industry focusing on digital practices was long overdue. Helberger: "The launch of ChatGPT was finally the incentive for the media to update their ethical guidelines in response to digitization and AI. It has triggered a very vivid discussion on what is responsible and what isn't." Meanwhile, journalists come up with their own ways of dealing with AI in an ethical way (see figure 10).

This chapter discusses the possible building blocks of responsible AI specifically for media organizations. In contrast to the previous chapter, which centred around the development of ethical guidelines in the newsroom, this chapter addresses the bigger issues, some of which are beyond the news industry's direct influence. Nevertheless, media leaders can be outspoken on topics like copyright, data protection, and screening for bias and engage in industry collaborations and lobbying. The themes this chapter covers emerged from our interviews and are particularly relevant for those in leadership roles. The discussion is not intended to serve as a blueprint but instead as a checklist or inspiration for debate and reflection (see Ethics Checklist, page 151).

Community-focus: Public service values at the centre of AI innovation

As discussed in the previous chapters, many in the industry have high hopes that AI can make journalism more productive, inclusive, and help it reach broader audiences. The overall goal is to engage audiences and establish trust, with the dual strategies of battling news avoidance and misinformation. All of this is happening in a changing information environment where news discovery is already becoming a challenge. As the BBC's Blathnaid Healy says: "How will the public find high quality public news and information in that ecosystem? We need to ensure that people have access to the journalism that we are producing. Our greatest concern would be that it remains discoverable in this ecosystem."

This is easier said than done because, as discussed earlier, the big unknown is audience behaviour. According to a representative survey of audience attitudes in six countries by the Reuters Institute, there is still a sizeable minority of about 20-30% who have not even heard of any of the most popular AI tools.[2] New search technologies are just emerging and currently mostly used by early adopters, particularly students and office workers open to new technologies. So, it will take a while to figure out what kind of news experience will make the cut for which types of audiences. Additionally, there is still little evidence for what works in combating news avoidance, or, to frame it positively, increasing engagement. Some view so-called 'slow media' as a recipe against content overwhelm. These are news outlets that consciously bypass breaking news in favour of deeply researched stories with more context and explanation. News organizations like Zetland in Denmark or Tortoise in the UK have explored this concept with success. But it is safe to say that this approach mainly works for the better educated and well-off groups of the public who often have a stronger interest in news.

Others propose that constructive journalism will work best to overcome news fatigue. The Constructive Institute in Denmark was founded in 2015 by Ulrik Haagerup, a former news editor of the Danish public broadcaster DR.[3] Its mission is to advocate for journalism that provides perspectives and solutions, moving beyond the often confrontative, sometimes polarizing style of traditional (political) reporting. The assumption is that audiences are tired of journalism that focuses on conflict and prefer approaches that provide explanation and possibilities for agency. Quite a few PSM organizations have been championing constructive approaches of late, particularly in the Nordics and Eastern Europe. AI could be used to highlight or help reframe news with a constructive angle.

[2] Fletcher, R., & Nielsen, R. K. (2024, May 28). What does the public in six countries think of generative AI in news? Reuters Institute for the Study of Journalism. Retrieved from https://reutersinstitute.politics.ox.ac.uk/what-does-public-six-countries-think-generative-ai-news#header--10

[3] Alexandra Borchardt is a board member at the Constructive Foundation and engaged in project work with the Institute.

Swedish Radio is among those public broadcasters that are using a 'public service algorithm' to curate digital news content.[4] In a series of workshops, staff members came together to define the public service values that are especially important for Swedish Radio in relation to its PSM mission, for example, to deliver on the ground reporting, voices of impacted citizens, reflect minority interests and bring original analysis. This approach has also inspired the exploration of a public service algorithm within the EBU news project 'A European Perspective'. In the current phase, a large language model (LLM) is being used to try to identify news stories with high public value.

As Mattia Peretti says: "PSM by definition needs to be responsible, it needs to respond to diversity, the representation of society. It has a lot to teach to the industry and the rest of the industry has to think about responsibility by design."

Personalization versus collective experiences

One of the big debates about public service content centres around personalization. Specifically: how much catering to individual interests and needs is necessary to attract and engage people while also keeping them informed on what they need to know as citizens, voters, and responsible human beings? Striking a balance here will be a major task for any outlet providing quality journalism, not just PSM. Matt Frehner of *The Globe and Mail* says: "When we do research, people say they want more personalization, people are attuned to that with Netflix and Amazon. But this is also the challenge: We are not Netflix. We are serving content in the hope that it gets consumed. We need to break the algorithm. The public service aspect needs to be core."

Johanna Törn-Mangs of Yle agrees: "People will demand these personalized services. But we need to work on our own algorithm according to our values and mission, we don't believe in personalized news only. We need to think about people not only as customers but also as citizens, we don't want to go to fully personalized news. Collective media experiences are important." Then again, the extent to which media companies will have a say in this is not entirely clear. Törn-Mangs: "We don't know; will there be an AI agent that chooses what kind of news the people will consume. The biggest change will not be within the current media companies. All kinds of companies will be able to produce so much content, the big challenge will be: how will people even find us?"

In case audiences should opt for pull models in news consumption, much of their experience might depend on their own prompting skills. What they will be served by chatbots will depend on their past behaviour but also on their abilities to game the system. In the world of algorithmic news distribution by search and social media, publishers have had little control over which news their customers see. In recent years, the platforms – particularly those owned by Meta – have reduced the share of news content in their offerings. With the use of generative AI to disassemble and remix content to cater to personal preferences, a further atomization of news can be expected.

As the EBU commented in a March 2024 consultation for the European Commission: "In a near future, a popular virtual assistant could provide to end users personalized daily news updates aggregating content from various sources, including PSM and other news publishers. Accordingly, over time, the user could become more loyal to the virtual assistant and less aware of the individual media organizations providing the content. Ultimately, this could jeopardize PSM's relationship with their audience and, more generally, threaten media pluralism."[5] At this point it is unclear how generative AI will reshape the already fragile links between

[4] Olle Zachrison, co-author of this report, has been driving the development of SR's public service algorithm. See the case study in A. Borchardt, F. Simon, "What's Next? Public Service Journalism in the Age of Distraction, Opinion, and Information Abundance", EBU, November 2021.

[5] EBU, "EBU Response to the European Commission Consultation On Generative AI", 11 March 2024. https://www.ebu.ch/files/live/sites/ebu/files/News/Position_Papers/open/2024/EBU_Response_GenerativeAI.pdf, retrieved on 5 May 2024.

news producers and consumers. Much will depend on the producers' abilities to attract users to their own platforms and products.

Responsible use of purchasing power: fairness, sustainability, and safety

Media companies will inevitably remain dependent on third parties for AI tools and technology. In some cases, this will be new technology subject to procurement processes. In many cases, it will be embedded in updates to systems already in use. And increasingly, these third parties will be directly or indirectly linked with the largest technological monoliths who have the resources to control the infrastructure needed to make generative AI work. Some of our interviewees mentioned the importance of procurement decisions when it comes to using AI responsibly. As Johanna Törn-Mangs says: "It is crucial to talk about responsible AI and make choices about what AI to use. This is a challenge for smaller companies. There are lots of difficult ethical challenges involved. The biggest is around trust. Can we trust the big companies like OpenAI and can we trust their solutions are up to our standards?"

Natali Helberger says that media companies need to consider a substantial list of issues when making purchasing decisions, including the environmental footprint of solutions and workers' rights. As outlined in Chapter 1, foundation models are often trained by cheap labour in the Global South, including by underaged workers. "There is a very dark side to this. This all should be important when we procure these models." Helberger says that media organizations need to ask themselves whether they want to remain media companies or would rather invest in becoming tech companies themselves. "If they are media companies, they need to use more external providers of technology. But they can have an impact through their buying power, negotiating and shaping power. The EBU is an actor that has an impact, for example by raising this issue with its members and pushing for responsible procurement decisions." She recommends looking at the guidelines developed by experts for the Council of Europe listing which considerations media companies should take into account when implementing AI systems.[6] "This includes for example the responsibility to share their experience with less affluent local players. For smaller media companies research costs are really, really high", Helberger says.

When it comes to choosing tools, the open-source models are often on the table. In some such models, the source code is open to public scrutiny and can be used, shared, examined, and altered by everyone with access. In contrast, closed models are proprietary and therefore impenetrable to outsiders. This is a choice with almost philosophical qualities. While some consider open access technology to be safer, since diverse sets of users might be more likely to spot shortcomings and help with improvements, other feel that closed models are less prone to being misused by those who want to cause harm.

As the EBU commented in the consultation for the European Commission mentioned above: "Open-source generative AI models offer several competitive advantages over proprietary systems since they (i) benefit from a broad community of developers leading to faster innovation, (ii) have low cost/entry barriers, (iii) address bias and ethical concerns effectively and (iv) may be easily customized." Proprietary, closed AI models were easier for regulators to oversee, however. The rapid dissemination of open-source models could make it "more difficult to track usage and enforce compliance with ethical guidelines and legal standards, potentially leading to misuse or harmful applications of AI technology."[7]

[6] Council of Europe, "Guidelines on the responsible implementation of artificial intelligence (AI) systems in journalism", 12 December 2023. https://www.coe.int/en/web/freedom-expression/-/guidelines-on-the-responsible-implementation-of-artificial-intelligence-ai-systems-in-journalism , retrieved on 13 May 2024.

[7] See above, "EBU Response to the European Commission Consultation On Generative AI," page 14.

The 2024 report published by the UK's House of Lords also discusses the pros and cons of the debate: "Open access models tend to be cheaper and more accessible."[8] The report quotes the managing director of Mozilla.ai, who argued that open models provide a 'virtuous circle' by enabling more people to experiment with the technology. Irene Solaiman, Global Policy Director of Hugging Face, was quoted saying open access is also preferable for its transparency and opportunities for community-led improvements. "Open models have however lagged behind the most advanced closed models on full-spectrum benchmarks and have fewer options to recall and fix harmful products," the report summarises. This means that from an ethical perspective, both options have their pros and cons, and a singular recommendation is difficult. Nevertheless, being informed about the technology behind the models, their limitations, capabilities, and impact is imperative.

However, a choice is necessary. And even where options exist, path dependencies and convenience can make it harder to justify using alternatives. Many will remember the first months of the Covid-19 pandemic when companies scrambled to find video conferencing software that delivered on data protection. A few years later, an overwhelming majority of organizations, public and private, has landed on the all-encompassing Microsoft Teams. Zoom is still widespread and favoured by some as a user-friendly alternative, focused on meetings and conferences. But many other tools that were deemed to excel in data protection have already been forgotten. It might be ambitious to suggest that media companies can make a substantial difference solely through their own purchasing choices, not only because many developments are beyond their influence but also because the industry is relatively small. This is where regulation comes in.

Lobbying for regulation: between innovation and damage control

Media organizations cannot make laws, but they can inform and lobby lawmakers. It is critical that media leaders engage in regulatory debates and keep themselves up to date on what legislation is in the works and when their input could make a difference. While large publishers and public broadcasters will be able to draw on in-house expertise for this, smaller news organizations will most likely have to rely on industry associations to have influence. Regulation can affect the development of technology indirectly, for example through anti-trust action, or directly, as with the EU AI Act. This legislation was passed with an overwhelming majority on 13 March 2024. It was a global first on setting standards governing AI and described as "trailblazing".[9]

Some feel that anti-trust policies, such as the Digital Markets Act, are the most important tool to ensure innovation can take place in a competitive environment. As Marietje Schaake, International Policy Director at Stanford University wrote in an editorial for the *Financial Times*: "From the promise of medical breakthroughs to the perils of election interference, the hopes of helpful climate research to the challenge of cracking fundamental physics, AI is too important to be monopolized. (…) Preventing AI monopolies is part of a healthy innovation climate, and it is increasingly critical for a better public understanding of the technology."[10]

Companies tend to have mixed feelings about regulation. Many are sceptical because it increases bureaucracy and can undermine business models. Others see it having the benefit of granting their products and services a competitive advantage. In AI, a lot more regulation can be expected, as it is early days, and some risks will only become evident with broader adoption. Andrew Strait of the Ada

[8] UK House of Lords, page 18.

[9] Karen Gilchrist, Ruxandra Iordache, "World's first major act to regulate AI passed by European lawmakers," CNBC, 13 March 2024. https://www.cnbc.com/2024/03/13/european-lawmakers-endorse-worlds-first-major-act-to-regulate-ai.html, retrieved on 13 March 2024.

[10] Marietje Schaake, "AI is too important to be monopolized," Financial Times, 13 February 2024. https://www.ft.com/content/1fda45a2-43e0-4c10-b5fb-b6097e3f5c56, retrieved on 13 February 2024.

Lovelace Institute cautions against relying too much on generative AI, exactly because regulation might change: "That is the number one reason why so many companies are so hesitant. They spend years implementing generative AI-based services, and suddenly it becomes illegal under European law. Never rely on these tools alone."

Several experts – often those from countries with a lively ecosystem of AI-start-ups – warn that over-regulation will stifle competition through excessive bureaucracy which can again favour the larger players that have the resources to comply with new requirements. As the UK's House of Lords report states: "Solving the 'Goldilocks' problem of getting the balance right between innovation and risk, with limited foresight of market developments, will be one of the defining challenges for the current generation of policymakers." Striking this balance will be critical though, because "long-term global leadership on AI safety requires a thriving commercial and academic sector to attract, develop and retain technical experts." Regulatory capture – the exercise of excessive influence by those being regulated on regulators – could be a problem, too. "This might occur through lobbying or because officials lack technical know-how and come to rely on a narrow pool of private sector expertise to inform policy and standards. Similar problems may emerge from groupthink. This might lead to regulatory frameworks which favour a select group of commercial rather than public interests, for example by creating barriers to new competitors entering the market."[11]

The quality of regulation very much depends on the competence and knowledge of regulators and on who participates in it. But there are other forces at play, as Yale Professor Luciano Floridi describes: "There is enough knowledge and expertise out there to regulate. But when you look at the efforts – and I have been involved in many – it is getting difficult. There are all these considerations, the awareness of powerful companies for example, or that one country could be offended. We are pursuing many goals at the same time, some of which are not compatible." (Read Q&A with Luciano Floridi, page 156) Melanie Mitchell of the Santa Fe Institute warns that expectations in regulation shouldn't be too high: "I think too many restrictions are dangerous in themselves. This will prevent people from innovating to make systems safer."

The EU AI Act categorizes activities and products according to risk categories, ruling out some entirely and prohibiting certain uses.[12] Floridi suggests that regulation makes most sense further down the value chain. "We should stop thinking about AI as a product and start thinking about it as a service. Today's regulations tend to look like those for commercial product safety. But facial recognition is not a thing. It is a way of using AI to do a certain job. Regulate what it means to use it, for example in schools, at airports, at a stadium, in a nuclear plant, in a prison. An analogy could be, we shouldn't regulate the microwave but the cooking."

Natali Helberger of the University of Amsterdam agrees that the impact on users should be considered – even though the EU AI Act doesn't classify them as high risk.[13] If this were seen through, media organizations would have significant responsibility. "In the Council of Europe Guidelines, we suggest that media organizations planning to use AI in the journalistic process should do some form of impact assessment to understand possible implications for their users, or their democratic role," she says. Helberger worries that the regulatory debate is so full of tech slang that non-experts will have a hard time intervening. "This myth of technology being so complicated is very convenient for its providers. If we exclude society from regulation, we over-rely on providers. We need to involve diverse stakeholders, for example

[11] UK House of Lords, pages 10, 21, 32.

[12] European Commission, AI Act. https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai, retrieved on 15 March 2024.

[13] Quotes here are taken from our interview with Natali Helberger which took place before EU AI Act was passed. Watch Helberger's presentation on the AI Act and the media industry at the Nordic AI in Media Summit here: https://www.google.com/search?client=safari&rls=en&q=nams+24+natali+helberger+youtube&ie=UTF-8&oe=UTF-8#fpstate=ive&vld=cid:4b085bc7,vid:petbXlxN7c8,st:0

from academia and the media. We also have to be more agile when regulating. This is a moving target."[14] She says that the adoption of the AI Act was only the beginning. The potentially more relevant work starts when details need to be carved out and codes of conduct established. "This is where the real regulation will take place, and this is also where we need to bring in expertise on fundamental rights and the democratic role of the media."

The debate about accountability is and will most likely remain fierce. This primarily affects what happens 'downstream', when third parties use models for different use cases. This was also discussed in front of the UK's House of Lords by representatives of Meta, Microsoft, and BT (British telecommunications company). A Meta manager stated that responsibility needed to be at "every level of the chain". A Microsoft manager added that developers would "not be in a position to mitigate the risks of the many different downstream use cases of which they will have little visibility." The report concluded, however, that downstream actors might lack the information to be confident about their responsibilities, because they didn't know what data foundation models were trained on, how they were tested, and what their limitations were.[15]

Critics fear that the effects of regulation will always be limited when systems are not designed at the outset to be safe. Berkeley Computer Science Professor Stuart Russell, one of the leading researchers in AI, was quoted saying that the biggest safety challenges couldn't be addressed by add-on fixes: "The security methods that exist are ineffective and they come from an approach that is basically trying to make AI systems safe as opposed to trying to make safe AI systems. It just does not work to do it after the fact."[16]

Bill Thompson of the BBC says it is worrisome that the founders of the dominant tech companies mostly share a libertarian political philosophy. "This is shaping the way these technologies are designed. A damaging ideology is underpinning them." A variety of investors and founders in Silicon Valley have repeatedly shared thoughts about a libertarian belief system that centres around strict meritocracy and an uncompromising free market philosophy with minimal government intervention. This has been widely criticized.[17]

Picking the right copyright battles

At first it had looked like a smooth ride for big tech. In the summer of 2023, the Associated Press (AP) as the first news organization had negotiated a deal with OpenAI, making part of its archives available for training purposes in exchange for technology. German publisher Axel Springer followed in December with a far-reaching licensing contract, giving OpenAI the right to use paywalled content of its brands *Politico*, *Business Insider*, *Bild* and Welt for news summaries while displaying the source.[19] But when *The New York Times* sued OpenAI for copyright infringement shortly after Christmas in 2023, many news organizations felt they had to take sides.[20] *The Times* argued that ChatGPT was trained on copyrighted material – the journalism it had paid for to be produced. A few months later, the excitement seemed to have died down. More and more publishers – including French *Le Monde* and the *Financial Times* had signed contracts with OpenAI, presumably in the hope of getting at least some share of the pie.[21]

However, in April 2024 a group of US regional publishers owned by Alden Global Capital filed a lawsuit against OpenAI and Microsoft.[18]

[14] We conducted our interview with Natali Helberger before the EU AI Act was passed. For her assessment of the results, watch her keynote at the Nordic AI in Media Summit on 10 April 2024: https://www.youtube.com/watch?v=DFJv3Vqu52M

[15] UK House of Lords 2024, page 61.

[16] UK House of Lords 2024, page 47.

[17] See for example Hallam Stevens, "A new Silicon Valley manifesto reveals the bleak dangerous philosophy driving the tech industry," The Conversation, 6 November 2023. https://theconversation.com/a-new-silicon-valley-manifesto-reveals-the-bleak-dangerous-philosophy-driving-the-tech-industry-216894, retrieved on 9 May 2024.

[18] Katie Robertson, "8 daily newspapers sue OpenAI and Microsoft over A.I.," New York Times, 30 April 2024. https://www.nytimes.com/2024/04/30/business/media/newspapers-sued-microsoft-openai.html, retrieved on 1 May 2024.

According to the brief, they were not only worried about chatbots undermining paywalls and contributing to copyright infringement but also about misinformation being wrongly attributed to their brand and their own research being diluted with other content. "This issue is not just a business problem for a handful of newspapers or the newspaper industry at large," *The New York Times* quoted from the lawsuit. "It is a critical issue for civic life in America." But the news brands might have had a more mundane reason for going to court: a pre-emptive act of self-defence in a country where lawsuits can easily bankrupt companies. A legal precedent would clarify who is liable in the case of misinformation.

Axel Springer, which also owns US publications, doesn't seem concerned. VP Niddal Salah-Eldin likens their agreement with OpenAI to a breakthrough: "Our landmark partnership with OpenAI marks a paradigm shift in journalism which we are proud of. For the first time, we're seeing a revenue stream from an AI company to a media company for the use of recent content. This establishes the principle of renumeration. This partnership has opened a path that we hope many other publishers will follow along." But others are sceptical about the benefits. Most fear that this will once again be a win solely for the big players.[19]

In the context of LLMs, the copyright debate is predominantly linked to the fair use of content for training LLMs or search queries. For the latter, news organizations have some leverage, as the future of generative AI will very much depend on increased reliability of these models. This means they need to be trained with fact-based, up-to-date quality content to ensure that query results improve. This works through a process called retrieval-augmented generation (RAG): A chatbot, when presented with a query, double-checks with a database to establish if the potential response is supported by the latest available facts. Mike Cook of Kings' College London wrote in *The Conversation*: "This is a bit like taking an exam with a textbook open in front of you."[20]

Jim Albrecht, a former Google News director, described why media should have some bargaining power: "The answer, I think, lies in the fact that LLMs tend to hallucinate (...) and that they are so expensive to train that the models are updated on the order of months, rather than days or minutes. (...) Generative AI products tend to rely on a process known as 'grounding', in which the statements made by the AI are checked against relevant source documents to ensure that the AI is not making things up. This process is especially critical if a user is asking about a recent event in which the relevant facts did not exist at the time of the LLM's training. In such cases, the AI can only answer accurately if it retrieves those facts from recent grounding documents. These documents are the essence of the work newspapers do - sourcing and reporting new facts - and the fruits of that labour should reasonably belong to those who perform it."[21]

Madhav Chinnappa thinks this will be key to future collaborations between media and tech: "I think that grounding data is where there should be a licensing model. But this requires some open mindedness and some collaboration and some creativity between news companies and technology companies. Unfortunately, my experience has been over the last 13 years that there isn't enough trust between the two."

Andrew Strait thinks that the copyright debate would profit from being less black and white. US tech companies tend to argue that future innovation would be hampered if there were no exemptions for data mining, harming the American economy and helping China prosper. "Then there is the camp that says: If data mining will occur, the creative sector

[19] Nic Newman, "Journalism, Media, and Technology Trends and Predictions 2024".

[20] Mike Cook: "OpenAI's content deal with the FT is an attempt to avoid more legal challenges and – an AI 'data apocalypse'", The Conversation, May 2024, https://theconversation.com/openais-content-deal-with-the-ft-is-an-attempt-to-avoid-more-legal-challenges-and-an-ai-data-apocalypse-229215, retrieved on 18 May 2024.

[21] Jim Albrecht, "The real wolf menacing the news business? AI," The Washington Post, 6 February 2024. https://www.washingtonpost.com/opinions/2024/02/06/ai-news-business-links-google-chatgpt/, retrieved on 15 March 2024.

will not exist any longer. It is presented as a binary choice of supporting innovation or the creative industry." He suggests the media industry approach this topic pragmatically. "They should ask: 'What is the value of my data, what can I ask for that?' The downside is this will favour large incumbents that can pay."

A much bigger issue will most likely be what happens to those who don't have the size and clout to negotiate deals, ranging from local news organizations to independent journalists, authors, and artists who rely on revenue from copyrighted material. LLMs will be able to use plenty of their content without its producers being aware or able to prove it. Other than plagiarism that can be easily detected, generative AI can rewrite, remix, and recraft any text or artwork without leaving too many traces of the original. In other words, this debate is about the very future of the creative industries.

In the UK's House of Lords' report, image provider Getty Images was quoted as arguing that "ask for forgiveness later" opt-out mechanisms were "contrary to fundamental principles of copyright law, which requires permission to be secured in advance." But in the case of generative AI, copyright holders are often unable to exercise their rights because they cannot access the training data to verify if their works have been used without permission.[22] Renate Schroeder, Director of the European Federation of Journalists, calls on the tech companies to reveal their data's origins: "LLMs need to be transparent about which data will be used. We need remuneration for text and data mining. We are talking money in times when journalism is extremely precarious and fragile."

In many non-English-speaking countries, media organizations feel obliged to contribute to the training of foundation models because the LLM's quality would otherwise be compromised. Schibsted's Agnes Stenbom says that the language models tend to perform poorly in Swedish and Norwegian. "We clearly want our data to contribute to Nordic language models." This is why Schibsted partners with universities to develop 'Nordic' language models. Styli Charalambous of the *Daily Maverick* also emphasizes the importance of sharing quality information for the public good: "I think what a lot of organizations are focusing on now is 'Can we get compensation out of that?' But what people aren't speaking about is what is the long-term impact if we block off the training of these huge systems from journalism? Because where these systems end up is heavily influenced by the quality of the training data."

The positions of PSM organizations vary significantly in the copyright and training data debate. Some were quite vague in our interviews, indicating that debates are ongoing. The EBU strategy advice of April 2024 has no clear recommendation: "There's no single answer, but most EBU Members have implemented a recommendation to opt out and/or prevent unauthorized third parties from using their content and data."[23] Other members were still reflecting. Not so SVT, as their Vice Director General Anne Lagercrantz says: "We are among the few who are not blocking AI crawlers. We need time to learn more about AI, about how can public service news be found, how can we improve the conversation? But there are copyright issues. Of course, we are afraid that our content will be manipulated. We might rethink, but for now this is our position." Nevertheless, helping to fight disinformation with quality content overrides these concerns at SVT.

Navigating bias: Fix the models, the processes, and the recruiting

LLMs scale biases by default. Bias is inherent to the human condition and therefore also embedded in the internet. LLMs are trained on existing documents and calculate the most likely outcome. While quite a few newsrooms have used 'regular' AI to scan their output for diversity – for example alerting journalists

[22] UK House of Lords 2024, pages 72, 73.

[23] See EBU Strategy Services, "PSM AI Strategies", April 2024 (members only).

when only male sources are pictured or quoted – generative AI works on probabilities. Left unattended, this could set back recent efforts to increase the range of perspectives and voices being heard. Jeff Jarvis, Media Professor Emeritus at Craig Newmark Graduate School of Journalism at City University of New York, says: "Generative AI is a mirror to society. It is a warped mirror that reflects just those who had the power to publish in the past." (Read the Q&A with Jeff Jarvis, page 153)

Bias is particularly prominent in images. Deutsche Welle's Manuela Kasper-Claridge says: "In our test cases we see a huge bias in AI generated illustrations. For example, if the topic is domestic violence in the Arabic world, women are always pictured with a hijab." Being a visual editor himself, Matt Frehner of *The Globe and Mail* asks for more conversations about biases and stereotypes within the industry: "If you tell a tool to picture a Muslim, you always get brown men in headscarves. If you ask for a lawyer, you get white men in suits. We need to question what are the inherent power structures that exist within the tools? A lot of reporting needs to be done on biases. Our rules need constant updates." Cambridge professor Gina Neff confirms: "We have many examples about how machine learning models reproduced some terrible bias. It is a lot of work to clean up data sets that are abhorrent."

A controversy around Google's Gemini LLM in February 2024 illustrates the challenge here. According to Google, Gemini was trained with the intention to bring more diversity into imagery. But when it was prompted for images, it tended to produce a variety of ahistorical images like black female popes and Asian Nazis with swastikas.[24] This led to outcries on different fronts. Users leaning to the political right complained that Gemini didn't depict white people, interpreting this as big tech catering to the 'woke community'. People of colour and their allies were offended. Google apologised and paused Gemini's creation of human images as a result.

APA's Verena Krawarik says that the news agency is quite aware of the need to remove bias from LLMs, but there is no easy fix: "Therefore we invest a lot in techniques like Semantic Search and Retrieval augmented Generation that ground in our data. And we are partner of the FAIRMedia research project.[25] This is about training journalistically and ethically correct data sets. We can't avoid thinking about how we would set up an AI ourselves. We have to think hard about how we want to feed it. Who defines what is ethically correct? The journalists, the lawyers?"

But ramping up efforts for more diversity is not only important with regard to bias in content. The variety of perspectives in a team will very much affect the decisions for certain use cases and all aspects of product development. As the BBC's Bill Thompson says: "The first question is always: are you sure you want to use this? The next is, look around your team: is it sufficiently diverse to reflect on how this could affect people?" PSM need to be particularly conscious about which groups to prioritize when deciding to implement products based on generative AI.

This leads to the crucial issue of recruiting. Many larger organizations have already used AI-based staffing solutions for a while, for example when screening applications and searching for talent[30]. This has the potential to increase the quality of recruitment but also can lead to discrimination on the basis of gender, race, skin colour, or personality traits.[31] Many candidates applying for positions online will never learn that their job applications were rejected by an AI filter. Obviously, selection by humans carries others risks. Some in the HR field say that AI has even helped to reveal biases when hiring and has been effective against it by focusing more on skills than on

[24] Sarah Shamim, "Why Google's AI tool was slammed for showing images of people of colour," Aljazeera, 9 March 2024. https://www.aljazeera.com/news/2024/3/9/why-google-gemini-wont-show-you-white-people#:~:text=America%27s%20founding%20fathers%20depicted%20as,flurry%20of%20intrigue%20and%20confusion. retrieved on 16 March 2024.

[25] The project is funded by the Federal Ministry for Climate Action, Environment, Energy, Mobility, Innovation and Technology represented by the Austrian Research Promotion Agency (FFG). https://www.ots.at/presseaussendung/OTS_20231220_OTS0035/fairmedia-ein-innovatives-projekt-fuer-faire-und-vertrauensvolle-anwendung-von-kuenstlicher-intelligenz-in-den-medien-bild, retrieved on 17th May 2024.

personal characteristics.[32] But this is by no means a given. Anne Lagercrantz of SVT is worried about a world where processes get even more standardized by technology: "We need to talk more about how we will recruit for diversity in the future. If we want to be relevant, we need to be diverse." But it is also true that where humans have done a lousy job for decades, AI cannot be expected to be the only remedy.

Educating the public: Covering AI

Contributing to digital literacy is among the wider responsibilities of news organizations. This is not only a matter for newsrooms: companies should ideally let their customers know where they use AI and be transparent about its limitations and potential. Journalists have the special obligation to educate the public about the opportunities and pitfalls of generative AI. Reuters' Jane Barrett says: "We have to educate ourselves about AI and then report the hell out of it!" As the 2023 EBU News Report highlighted in reference to climate journalism,[26] this shouldn't be delegated to a few special correspondents but play a role in every beat – from business and political reporting to health, agriculture, culture, and education coverage. (Read in Resources: List of Recommended Readings on AI, page 177)

While much informed and excellent coverage is produced, reporting on AI suffers from many of the same pitfalls as other issues which necessitate a more holistic approach. The speed of technological development, the lack of data, evidence, and the nebulous character of much of what's happening make reporting on AI a challenge. Additionally, AI coverage is often left to business correspondents, who tend to rely on sources in the business realm, instead of reaching out to a wide-ranging field of scientific experts. Or, the topic is handled by cultural critics who are often drawn to AI's most vocal opponents. Quoting former tech industry executives who predict the end of humanity makes for catchier headlines than relying on a panel of nuanced voices who admit that they cannot really predict where this is going. (Read in Resources: The do's and don'ts of AI reporting, page 174)

We asked all our interviewees what they think is missing from current debates. What are the issues that no one bothers or dares to touch, whether in public discourse or in reporting on generative AI – apart from the environmental and human rights concerns mentioned in previous chapters? This is, of course, simply a snapshot of impressions, which depend heavily on the moment at which the interviews were conducted. Nevertheless, these reflections are thought-provoking and may help journalists refine their approach to a topic that will profoundly shape our institutions, organizations, and personal lives in the future.

Melanie Mitchell of the Santa Fe Institute thinks that the current discussion on AI in the media lacks a variety of perspectives: "Whenever there are stories about how intelligent machines are, journalists always interview AI people. They don't speak to psychologists, developmental psychologists who study how children learn, they don't interview biologists who work on animal intelligence. These people have strong opinions that are different from those of the AI people. I don't know if computer scientists have any idea about what intelligence is." Creating great reporting on AI is inherently difficult because there are many different opinions and high levels of uncertainty. But reporting influences legislators and societies. Jeff Jarvis criticises how reporters draw disproportionately on "young, hubristic white men" as sources when covering AI: "There are many amazing women and scholars of colour who are far more sceptical than the AI boys."

The BBC's Bill Thompson is also not impressed: "We are not having a sensible cultural conversation that is grounded in social sciences and philosophy. We are allowing this to be a technology conversation. The discussion about the broader impact

[26] See A. Borchardt, K. Dunn, F. Simon, "Climate Journalism That Works: Between Knowledge and Impact," European Broadcasting Union, 1 March 2023.

is missing. We are not having that deep conversation about what this digital one-zero culture actually means for us. What does this do about how we construct the human intellect?" Anne Lagercrantz of SVT misses a broader debate on how increasing automation of all kinds of processes – from the workplace to dating – will affect people's values, needs, and general outlook on life. "We need to talk more about how humans interact with change, for example, how they deal with risks in the age of AI. Will there be an expectation that AI could create a risk-free life?"

Luciano Floridi, the philosopher among our interviewees, asks for a more grounded debate on practical consequences: "We need to talk about what is real and stop talking about science fiction. Some solid realistic science-based discussion is missing at the highest levels. One big question is how we cope with risks. There is a classic way of doing this and that is called insurance. You transform a risk into a cost. How much does it cost? If this stuff is really risky, and in a non-sci-fi sense it really is, surely someone out there should take this up as a business. We will probably soon see an insurance against bots' mistakes."(Read the Q&A with Luciano Floridi, page ...)

In a recent blog post, *The New York Times*' Zach Seward highlighted that even the language widely used to cover AI is flawed – starting with the word 'intelligence' itself. "AI is the most anthropomorphized technology in history, starting with the name – intelligence – and plenty of other words thrown around the field: learning, neural, vision, attention, bias, hallucination. These references only make sense to us because they are hallmarks of being human."[27] This could lead to incorrect assumptions about what the technology is and isn't able to do and stop people from using agency when engaging these tools.

Particularly PSM should double down on their mission to dig deeper and investigate such consequential technology and educate the public about what is shaping up to be a huge transition. As David Caswell says: "It is more than reporting about it, it means helping people being part of it. An example is in the BBC educational programming for school children, developing a curriculum about large language models, to learn how to learn with them, what are the best practices for education in an AI environment. It is about helping the underserved population to appreciate what they can do with the tools." (Read the Q&A with David Caswell, page 31)

Finally, coverage of how AI will impact journalism and the media itself is strangely absent of much of AI-related reporting. Journalists are reluctant to write about journalism for many reasons, first and foremost because they cannot claim to be impartial. But this might prove to be a mistake. People need to know how generative AI can support, interfere with, and harm the important institutions that are tasked with safeguarding democracy.

Agnes Stenbom of Schibsted calls for more industry-wide discussions about structural issues, for example the potential impact of AI on the creative industries. "There are big potential implications we are not talking enough about right now. What will be the broader implications for the media ecosystem, the creator economy, the role that media plays in fuelling both knowledge production and creative work?" These are debates that stretch far beyond the media industry with many professions having a stake in this.

Pressing for collaborations – within the industry and between industry and tech

Only a few media organizations are big enough to tackle the challenges of generative AI independently. Shaping guidelines, training models, developing newsroom technology, negotiating compensation, lobbying for regulation, or simply exchanging experiences around use cases – all of this requires collaborations, ranging from structured

[27] Zach Seward, "AI is not like you and me," blogpost published on 2 May 2024. https://www.zachseward.com/email/73743cee-03c5-44a4-8331-cd064b8adfa5/?ref=zach-seward-newsletter, retrieved on 3 May 2024.

dialogues to fully blown partnerships. Not everywhere is the willingness to cooperate as pronounced as in Sweden, where 13 major publishers in collaboration with the Nordic AI Journalism network and The Swedish Publishers' Organization, representing hundreds of media companies, agreed on guidelines for AI transparency in March 2024.[28]

The EBU advises its members to "find scale in collaboration". Its April 2024 strategy document states that "most EBU Members cooperate with other national media organizations to exchange knowledge and expertise, and to find common positions on media interests to present to regulators and unions." The EBU also encourages members to "leverage the international EBU community to learn from successful use cases and strengthen their negotiating position with regulators and big tech."[29]

A number of our interviewees felt the need to cooperate more as an industry. Matt Frehner said: "I would love to have more time to talk with colleagues in other newsrooms. The more discussion across the industry the better. Conferences are often filled with the latest hype. We need more information sharing." Johanna Törn-Mangs also advocates for a strong, unified media industry: "Responsible media companies need to cooperate much more with each other, because we are on the same side. We need to distinguish ourselves from all the content that will be on the market. We don't yet discuss enough how to manage this in a systematic way: who makes the decisions, how generative AI will influence our work environments and our competencies, and how this will affect us in the long run. If there will be a lot of AI generated content, the next AI will be fed by AI generated content. What will such a future be like?" According to Törn-Mangs, it is the traditional media's responsibility to provide 'proof of authenticity' and fight misinformation: "In an ideal world this would be done on an international level; competitors need to team up."

But this shouldn't be seen as just another 'media versus tech' conflict, the likes of which have dominated much of the past 20 years. The media needs tech as much as tech needs the media to create a functioning information universe that supports societies not only to survive but to thrive. As Niddal Salah-Eldin puts it: "We are aware of the challenges and currently taking a very close look at aspects such as data protection, regulation, and fair remuneration for the use of our content as training data. For us, this represents an opportunity to not repeat the mistakes of platform regulation and create a fair and healthy ecosystem from very early on. Journalism is a part of the value chain and this needs to be reflected. To achieve this, we need a triad of competition law, copyright law and data protection law. Of course, we are striving for a fair balance of interests between platforms and publishers."

Madhav Chinnappa expresses doubts about publishers' willingness to cooperate. He says: "I'm really concerned that the news publishers will be so obsessed with the copyright issue that they won't come to the table to talk about the collaboration that you need to fight both the cheap fakes and the deep fakes which are polluting the overall information ecosystem." The risk is a general erosion of trust, he says. "If people don't believe anything any longer, then our democracy is undermined. And I think there are players out there who are actively trying to do that."

Learning to say no: Responsibility in the newsroom

The responsible use of AI will in many cases simply mean: not to use AI. However, at this point, an AI-free world is no longer a realistic option. And given the opportunities, it shouldn't be an aspiration, just as a world without search engines isn't a desirable scenario. As such, the label 'human-made' could become quite unique in an age of machines spilling out vast amounts of text and images every second. Picking up the phone, listening to people, meeting sources,

[28] Olle Zachrison, Co-author of this report led this effort with Agnes Stenbom of Schibsted, one of our interviewees. "AI Transparency in Journalism," March 2024. https://www.nordicaijournalism.com/ai-transparency, retrieved on 16 March 2024.

[29] See above, EBU Strategy Services.

building trust, visiting locations – all this will define quality journalism and set it apart from mere 'content'.

In recent years, many people have felt overwhelmed by the information and news to which they are exposed. They need help with curation. This holds particularly true for young people, who tend to rely on trusted sources – often personalities instead of brands – to pre-select content and put considerable effort into teaching algorithms their preferences. The study *Next Gen News: understanding the audiences of 2030*, commissioned by FT Strategies, reveals as much.[30] It is the media industry's responsibility to support audiences and set priorities for focus.

As the Reuters Institute's director Rasmus Nielsen writes: "If publishers primarily use AI to produce more of the same more cheaply, they will further reduce the already limited commercial value of all but the most effectively differentiated news content."[31] Nielsen adds that there might be a choice: "AI may be helpful for those publishers who are able and willing to define and double down on what makes them different, who are genuinely interested in meeting people where they are, and who can resist the temptation to further commodify the journalism they offer." Agnes Stenbom hopes that if generative AI helps make content more engaging and higher quality and makes it easier to target and serve audiences better, the commercial need for creating higher volumes disappears. She says: "Just because we can do things doesn't mean we should do things."

Erik Roose, Chairman of Estonia's ERR, describes a case when at least some parts of journalism become premium products: "It can lead to a very funny situation when at some point some very tough, very old fashioned news organizations step up and say 'we decided not to use AI at all because we cannot control it, but we can guarantee we are 100% human.' It's like hand washing your car, the most expensive service, but so cool and gentle." This will not be necessary for all of journalism, Roose says, but some parts might actually profit.

Not doing everything that's possible is not only a strategic but also an ethical choice. The environmental footprint of AI is massive, and amid the current hype, remains under-addressed – particularly among publishers. Conscious use of resources, including human capacity, should be a priority in a maturing AI-influenced environment.

C-level ethics check for media organizations

- **In which ways does the use of AI in our organization contribute to the public good?**
- **How do we keep ourselves informed and up to date about the technology to decide responsibly and contribute to regulatory debates?**
- **How do we scrutinize the products we buy and use for their ethical responsibility?**
- **How do we make sure our data is used responsibly, particularly for training purposes?**
- **How do we make sure that the products and products we use don't amplify bias but ideally even check for it?**
- **What are we doing to educate the public about the risks and opportunities of AI, engage them in debates and improve general tech literacy?**
- **Have we done everything possible to engage in the collaborative efforts open to us and to contribute to the knowledge and debate in our country/region/field?**
- **What is our understanding of, and our rules for, when not to use AI?**
- **Have we documented these aspects in ethics guidelines that are accessible to everyone in our organization?**

[30] "Next Gen News: understanding the audiences of 2023," FT Strategies, March 2024. https://www.next-gen-news.com, retrieved on 3 April 2024.

[31] Rasmus Kleis Nielsen, "How the News Ecosystem might look like in the age of generative AI," Reuters Institute for the Study of Journalism, 26 March 2024. https://reutersinstitute.politics.ox.ac.uk/news/how-news-ecosystem-might-look-age-generative-ai, retrieved on 27 March 2024.

Q&As

Q&A

> "I wish that public service media would become the laboratory for innovation in media"

JEFF JARVIS

Media Professor Emeritus, Craig Newmark Graduate School of Journalism at CUNY

You have been telling news organizations for many years they got digital all wrong: They were spreading moral panic instead of going for the opportunities. Do you see history repeating itself with generative AI?

The news industry's reaction to AI varies by region. In the United States, news companies see themselves as victims of technology, of the internet and now of AI. They have a hostile reaction to it, just as they've had a hostile reaction to technologies before, from radio to television to the internet. In the Nordic countries, I'm seeing a very different reaction. The publisher Schibsted is leading the way in being very strategic and open in trying to find new and good ways to use AI. I ask, why can't we be more like Norway?

What would it take to be like Norway?

To focus on useful tools that help journalists and journalism and eventually the readers. It's just deciding that this is an opportunity for innovation. The reaction to AI ranges from litigation to regulation to innovation. In the US, it's driving very hard towards litigation. I believe that using news content to train models is fair use and transformative. The news industry would be smart, rather than suing, to instead create an API for news – like a service contract connecting software applications. This API should set conditions for the AI industry that say you need current and credible information. Here are the terms on which to get a license and a key to do that. Here's the credit we want. Here's the linkage we want. Here's the money we want. Then the industry should negotiate in good faith.

How do you envision the use of generative AI to innovate journalism?

The first use is to help reporters report and editors edit. AI is good at summarising large amounts of data as long as that data is prescribed, and you don't ask it questions outside those limits. AI can be helpful in brainstorming headlines and summarizing articles. What interests me more is how AI can help reporters scale their work. I speculated with one editor a few months ago: Why don't you have 100 citizens in your state record their school board meetings, have AI transcribe it, and then query that database to find out how many school boards are banning books? Also, AI can provide us with a new user interface. Historically, we have a presentation structure, which is our front and home pages. Next we had a search structure: a search engine.

Now we can have a Q&A structure where people ask questions of all of our data, all of our stories. And in the future, I think we'll have AI as an agent that can look out for you and see when news that's relevant to you comes up.

Do you see it as a gamechanger for journalism or just as a tool to make journalism better?

AI is just a tool, and we have to look at it that way. I don't think that the technology determines us. I think - to paraphrase Clay Shirky, who teaches at New York University - that technology becomes useful when it becomes boring.

Publishers are quite worried about the ethical implications of generative AI. In a survey of media leaders that was commissioned by the Associated Press, one in five respondents said that's why they were not using it at all.

I've heard three concerns. One is reputation. People fear that readers are going to wonder where the content came from. And the answer to that is transparency. Another is liability for copyright because you don't know what goes in the model. But the third and most important is the misuse of generative AI by news companies. AI should not be used on its own to write news, because we know that generative AI has no sense of meaning, thus of fact or truth. Similarly, I don't think AI should be used for search, because people expect it to respond with credible answers.

But this is exactly what the big platforms are working on now: to have AI-driven search.

I don't think that's going to work, including by Google, because that it will return bad results. I think we have to wait for a next generation of AI. We have one step in that direction in the form of RAG, retrieval augmented generation, which means that the already trained model is restricted to a specific amount of data with citations.

One of the risks of generative AI is that it's so easy to fill the internet with junk.

The idea of content is a relic of the era of print. Content is that which fills things. We come to the internet thinking that we're still in the content business, and we have to fill the internet with content. AI has now fully and completely commodified content. We're not in the content business. We are in the service business. I find a big opportunity in this for us in publishing. As I describe in my 2023 book *The Gutenberg Parenthesis*: When printing started, no one trusted print, because anyone could make a pamphlet or a book just as anyone today can make a tweet or a Facebook post. They trusted instead what they heard from people they knew. Eventually, though, society created the institutions of editing and publishing to assure quality and authority in print. When print was mechanized and industrialized and reached mass scale starting in the mid-19th century with the steam powered press and the typesetting machine, there was a similar explosion of content. *Harper's Magazine* started in 1850 with the mission to find the good stuff. Well, where is the Harper's of 1850 for today? We're going to need the mechanisms to sift through this incredible abundance of speech.

How do you think journalism should go about this?

I think there's an opportunity for us in news to generously listen for the voices who have not been heard in mainstream mass media all these years. Journalism must be about listening to communities, understanding their needs, serving those needs, helping connect them, helping them tell their stories, bringing understanding among communities, instead of just extracting their stories.

But isn't generative AI doing the opposite: making nuance invisible because it aggregates everything and predicts the most likely outcome?

Generative AI is a mirror to society. It is a warped mirror that reflects just those who had the power to publish in the past. It brings our biases and our clichés and our misapprehensions to the surface. I would find value in that if we journalistically studied that data to see where the biases in society are.

Where do you see the role of public service media in all of this? Many broadcasters have been hesitant to embrace AI for good reasons, they are publicly accountable and afraid to risk the trust of audiences.

I wish that public service media would become the laboratory for innovation in media. There's a big need and opportunity to experiment with these technologies and to share lessons across the media ecosystem. I would hope that we could gather as an industry not to fight, but to make our content as a total useful in this new environment.

What could this look like, the industry coming together? There is no such thing as a worldwide media industry association pursuing shared goals.

The EBU could do that with public media, the wire services could do that with private media. But as an industry, we have a terrible track record of collaboration. We don't even get along with each other.

What do you think is missing from current conversations?

A diversity of expert voices. The discussion of AI is being led by AI boys, young, hubristic white men. And that's true in media coverage of AI. There are many amazing women and scholars of color who are far more skeptical than the AI boys—such as Timnit Gebru, Rumman Chowdhury, Emily Bender, Margaret Mitchell, Émile Torres, and Joy Buolamwini. Reporters would do well to get their perspective and expertise in their coverage.

Q&A

> “Some solid, realistic, science-based discussion is missing at the highest levels.”

LUCIANO FLORIDI

Philosophy Professor and Founding Director Digital Ethics Centre, Yale University

Is generative AI a game-changer in artificial intelligence? In a research paper published after the launch of ChatGPT you talked about a sensation.

I think it is, but I wouldn’t be a philosopher if I didn’t add a “but”. For the first time in human history, we are industrializing content production. But it will take some years to play out. AI will enable people to personalize content in a very simple way that is off the shelf. Will this be economically viable? I don’t know. Will this be science fiction? No! It is very hard to say what this will mean for mass media.

Could this be the end of journalism as we know it?

Not as a profession. The more tools we have, the more we need highly qualified people to work with them. Mass media have already been challenged. There will be a further break down in what everybody can be expected to know about. These days it’s hard to escape news about the war in Ukraine, the war in Israel. But do you remember the Greek crisis in Europe? It didn’t go away, but we no longer talk about it. At the top level, there will be a winner-takes-all situation, a handful of influencers will shape the debates, only some news will be known to all. But on the levels below that, diversification will be enormous. Let’s take our phones as an analogy: If I picked them up, there would still probably be a handful of apps on there that we all have, but the rest will look completely different.

But would that be really new? There have always been just a handful of topics that ‘everyone’ talked about.

The mass media world of the past was based on very simple models. One is that supply shapes demand. Now imagine a world where demand shapes the supply. In the future, as a consumer, I can potentially interact with the mass media supply. I can say I want that kind of audiobook with this particular voice and these types of news in the morning. I can say look, I really love that character in *Star Wars*, can you create a 15-minute movie around that character for me, please?

So, everyone can create their own happy endings as they see fit?

In 2016, ING, Microsoft, TU Delft, Mauritshuis, and Rembrandthuis teamed up for this project, Next Rembrandt, an image entirely created by AI. It looked beautiful. Now,

can I have my own Rembrandt? Can I have John F. Kennedy reading his last speech that he never gave? Or, if I don't like the voice of Obama, can I have JFK reading Obama's speeches? Jane Austen never finished her last novel, can I have the ending? Oh, I didn't like that one, could I have another ending? All of these things are or will be possible.

How do you look at this, as a philosopher, when it's getting harder to separate the true from and the false?

Philosophers disagree on different theories of truth. My own analyzes of truth is in terms of correctness: you have a truth when you get a correct answer to the relevant question. Get the wrong answer and you get a falsehood. Now, credible falsehoods are increasingly easy to fabricate. They are answers that look increasingly plausible. The macrotrends will increase. The synchronization of global waves will be disturbing, with mass consumption overload and without any quality control. The whole world will jump at the same time. We cannot do anything about this except to make people more critical, provide everyone with a chance to get out of the crowd. The other trend is hyper-specialization, which may hopefully enrich humanity by facilitating connections that would not have been possible before. Ultimately, I hope we will be a little less prone to engage with macrotrends and be more perceptive of the millions of tiny, more micro trends that make culture so much richer.

What will be the role of media in this world?

On the one hand, mass media will be part of the problem, because they magnify trends. On the other, they will need to become very technical when every user can commission work. This can only be done by AI; it would be impossibly costly to satisfy demand otherwise. A lot of the work will then be managing content, and you will need highly skilled tech people for that. For others this will be bad news. Generative AI will cut out all the people who have essentially been working like ChatGPT. Will this be the end of journalism? Well, the frozen pizza has not been the end of the pizzeria. But if my pizzeria produces rubbish, too bad, I have a frozen pizza in my fridge at home which is a third of the price and takes a fraction of the time. So, the pizzeria must make sure their pizza is better. If ChatGPT is going to provide the frozen pizza, the mass media better provide the experience, the quality, the added value.

If the quality of journalism improves, this should be good news for the audience. What else is good about generative AI?

A lot of what generative AI improves, in digital services, is not vertical but horizontal. In the past, search took you deep into silos. Generative AI will be more of an interface between services, like a glue that puts things together: text, images, sound, code. I'm using different kinds of bots, for example ChatGPT, Gemini, Claude, all the time, and they're very useful. In the past, large language models couldn't link to search engines, that's why their results were always outdated. Now they can. Generative AI is creating a seamless environment. The limit is not technological but socioeconomic. Which kinds of interoperability will the participating tech companies allow?

Will new players like OpenAI help to break the dominance of current tech monopolies?

The monopoly is intrinsic in the game. Everything points in the direction of a few, gigantic, controlling companies. This means a further erosion of individual autonomy and societal power. If you don't want that, you have to legislate. We need to update antitrust rules. It's clear what can be done against it: allow more competition.

But the world you described could be customers' heaven, couldn't it?

We should not confuse the number of movies I can personalize with the fact that there is only one movie provider. You only feel you have a freedom of choice in such a world. Just because you don't see the boundaries doesn't mean the boundaries are not there.

Your answer would be regulation like the European AI Act. Do you see regulators as up to the task?

They could be if they didn't have other distracting things in mind. There is enough knowledge and expertise out there to regulate. But when you look at the efforts – and I have been involved in many – it's getting difficult. There are all these considerations, the awareness of powerful companies for example, or that one country could be offended. We are pursuing many goals at the same time, some of which are not compatible. Like we need a firm stance against China, but we have economic interests.

What is your most important recommendation to regulators?

We should stop thinking about AI as a product but and start thinking about it as a service. Today's regulations tend to look like those for commercial product safety. But facial recognition is not a thing; it's a way of using AI to do a certain job. Regulate what it means to use it, for example in schools, at airports, at a stadium, in a nuclear plant, in a prison. An analogy could be, we shouldn't regulate the microwave, but the cooking. Also, we need an assessment and understanding of risks. Where is the final liability and responsibility? Is it with OpenAI, or those who use the tools?

What is missing from current conversations?

We need to talk about what is real and stop talking about science fiction. Some solid realistic science-based discussion is missing at the highest levels. One big question is how we cope with risks. There is a classic way of doing this, and that's called insurance. You transform a risk into a cost. How much does it cost? This debate is missing. If this technology is risky, and in a non-sci-fi sense it really is, surely someone out there should take this up as a business. We will probably soon see an insurance against bots' mistakes.

Q&A

> "We need to seriously think about the total cost of digitization"

SARAH SPIEKERMAN

Professor of Information Systems, Wirtschaftsuniversität Wien Vienna University of Economics and Business (WU Vienna)

Is generative AI a game-changer in artificial intelligence?

It is certainly a game-changer, because it is a new way of how people communicate with the digital. The second reason, it is multi-modal, so we don't only have the text and speech interfaces, but also video, music. Literally, it changes the way we interact online.

At the DLD innovation conference 2024, you gave a pretty downbeat presentation on AI. Was that the to provoke a response, or are you genuinely a pessimistic about it?

I would call myself a realist. I don't get myself carried away. I don't believe in superintelligence. We are seeing a new cycle of productivity gain. It is not the solution to reach true progress or even solve our true current IT problems.

Most experts predict that generative AI will transform pretty much everything. What are the true problems you are talking about?

The cost of IT is becoming unbearably higher all the time. Not only do companies need to invest in hardware and software that is constantly seeing updated product releases and that needs to be administered. There are also increasing legal and governance costs, security and privacy costs. Cost to be compliant with the AI Act, DMA, DSA, Cyber ResilianceResilience Act, etc. Electricity costs are exploding. And on top of this, complexity adds more new unknowns.

There are also environmental costs?

Right. IT companies often claim they are CO_2 neutral, because cloud centres are placed next to hydrogen power plants. But our ecological balance is not positive. You have to start with the resources, the minerals. For sourcing one ton of rare earth, you create 75,000 litres of acidic water, for example. Such effects need to be included in a truly meaningful ecological balance sheet. Such a balance sheet must include the globally distributed mining, shipping, manufacturing, and assembling across the thousands of resource and parts suppliers of IT components, and then add to that the ecological cost of service provision and consumption. Anything else is just kidding oneself about the true ecological cost.

Is it then realistic to invest so much in AI and perpetuate the digital transformation?

I am not sure. I increasingly doubt the sustainability of digital transformation at the scale where we're pursuing it right now. Another reason for this is we face increasingly more supply chain difficulties in our tumultuous world. If you need to build a machine with over 430,000 components, that you need for instance to produce machines that produce graphics processing units, you have the most complex supply chains ever created by humanity.

Is there a solution for that?

Peace. The West cannot afford to have too many geopolitical enemies that hold the resources for our innovation cycles. How can we have green transitions if solar panels come from China? We are already facing a dramatic shortage of IT supply. We are talking about sophisticated AI solutions, and you cannot even get an X-box in time for Christmas. We need more discussions about these hands-on challenges in the media. Journalists love to speculate about superintelligence, they should report on something that is relevant now. The motor of innovation is chip technology. Chips are much more complex than oil or gas. There is not going to be a solution. There is going to be hopefully an understanding that we need global cooperation.

Apart from all this, what kind of potential do you see with generative AI?

There are interesting potentials in this new way of interacting with machines that have access to information patterns that we don't even know about ourselves, not even in science. There is tremendous potential for learning, science, it can save massive amounts of time with dull documentation stuff and international cross-language communication. But if it is not built in a reliable, ethically responsible way, in the end what will happen is the same as what happened with social networks: the gains will be traded into net zero or even negative, because we don't consider the drawbacks and social cost of the technology. With social networks this has been truth, manipulations, lack of transparency and struggling of our democracies. These are no small challenges.

So, would you recommend we slow down?

Particularly if you are using these technologies in services that are vital for citizens, like food, electricity, telecommunications. Because what will you do if you cannot maintain this in 10 years because there are no chips?

What about regulation, do you think those involved are up to the task?

There are very intelligent people in charge, but many of them are exhausted, they have little bandwidth to live up to these challenges. What makes it harder is the general zeitgeist: if you criticize in a period of hype, you are immediately marked as a pessimist, you are being cornered. We should really stop this. We need a more realistic debate because then responsible actors are enabled and empowered.

Does the media live up to the task?

My impression is that there are two tribes in the media: sceptics and optimists. But optimism without realism is naïve thinking. Politicians and executives are in their own filter bubbles.

What's your recommendation?

We need to seriously think about the total cost of digitization. We need to work towards an ecological global balance sheet. We need to have strategic research on geopolitical effects on the IT industry, because of the risk of running out of chips. Finally, we have to get a more holistic understanding of human capabilities. We need values-based engineering of AI. Building machines in a way that they foster goodness and virtue in human decisions, not the virtue of just another machine to increase profitability.

Q&A

> "We need new ways for these systems to learn more efficiently. Right now, it's a very brute force approach"

MELANIE MITCHELL

Davis Professor of Complexity, Santa Fe Institute

You're a computer scientist. Do you think Generative AI is a game-changer for Artificial Intelligence? Is it the big difference everyone says it is?

I do. It's hard to know exactly what kind of impact it's going to have. Some people say it's going to completely change everything, affecting the economy and how we work, live and so on. Others say it doesn't have a good business model and will be a disappointment. It's probably somewhere in between. For AI itself, it's certainly an amazing new era for what AI can do, compared to previous eras.

Can you give us a brief idea of how you would define AI and distinguish generative AI from it? Is it even a useful term?

There's no clear, agreed-upon definition. Intuitively, it's about getting computers to do things that would require intelligence in humans. The definition of AI has morphed over the years. In the 1970s and 80s, people thought if a machine could play chess at a grandmaster level, it would be a big AI breakthrough. But it turns out even if you do that, it doesn't help in other real-world situations. AI remains a vague term, and the ultimate goal is not well-defined. Conversing in natural language in a human-like way was always one of the goals of the field, and it seems like something generative AI has been able to accomplish.

Are you generally more excited or worried about recent developments? Some people are quite concerned, others are on the fence. What about you?

There are some things that excite me and a lot of things that worry me. Obviously it's exciting to have systems that are so capable, but it's also worrying how they can be misused. We've seen misinformation, bias and people trusting these systems too much. But there's potential, especially in scientific discovery, and in augmenting our intelligence. In generative AI and other AI forms, we've seen AlphaFold predicting protein structures better than previous systems, which is crucial for biomedical and biological understanding. We've seen some really amazing progress in weather prediction thanks to AI. And in mathematics we've seen mathematicians using it to help prove theorems and solve problems that have not been solved before. But they are certainly not at the stage where they can replace us in the broad sense. These are all examples where humans are in the loop. We don't use these systems autonomously; they help humans in their jobs.

How much further do you think this can go in the next 10 years? Some say we've hit a roadblock. Even with more data and improved systems, there are limits to what large language models and transformer models can do. What's your take?

Anyone predicting this always gets it wrong. It's hard to say. But I do think just relying on current technology throwing more data or more computing power at this is not going to bring huge gains. We need new innovations in architecture and training methods. These systems need too much training data, which isn't sustainable. We need new ways for them to learn more efficiently. Right now it's a very brute-force approach. But I think there will be new discoveries to help us figure out how to get there.

Do you think there are things these systems should not be used for?

Too many restrictions can be dangerous as it prevents people from innovating and from developing ways to make these systems safer. But there are bad uses we've already seen. We clearly don't want these systems to be used autonomously for military applications. And we'll want some regulations around managing disinformation and people using them in illegal ways. There has been a lengthy debate in the computer science community about open-source software. There was always this worry that it can be used for malicious purposes. But it turns out that open source has has been one of the biggest contributors to making software safer and more reliable. The same debate is happening in AI. Should models be open and public, or controlled by a few large companies that are hardly regulated? There are arguments on both sides, but do I worry about the concentration of control in the hands of a few companies.

Do you think regulators are up to the task? Are you happy with what you see from politicians in the EU, US and UK, or could more be done?

Some good starts have been made. EU regulations go further in protecting privacy and rights than the US, which is playing catch-up. But I think that generally people haven't figured out the right way to regulate these things and it's hard because the technology is moving so quickly. I don't think micromanaging this is going to be effective. And in the US, there's a history of regulatory capture by companies trying to shape regulations to benefit themselves and squash competition. We've seen many examples of this, and it's starting to happen here, too.

How does your work feature in all this, and what is your current research focus?

My work is at the boundary between cognitive science and AI. I study how humans reason abstractly and make analogies, and how to use these ideas to get machines to do the same. And a lot of the issues around reasoning are hotly debated right now in large language models. We see them doing well on certain benchmarks, but then failing on others and being kind of brittle in the sense that if you change the problem a little bit, they don't succeed in a way that a human would. I'm trying to understand how humans do these things without huge amounts of training, how to get machines to do them and how to evaluate how well machines are doing. What I'm particularly

interested in is human adaptability: our ability to take knowledge and apply it appropriately to situations we've never seen. Our common-sense abilities are very robust, and for machines it's very hard to get to that same level.

A second focus is on whether these systems are able to abstract very accurate models of the world, even though they're just trained on language. I'm quite interested in understanding to what extent these systems are actually inferring useful models of the world – if they come up with some way to understand physical phenomena that they were not trained on. I don't think we have an answer yet.

It gets to the heart of whether these systems are conscious like humans, right?

Well, I'm not going as far as talking about consciousness, because I think that is something somewhat separate from intelligence. I'm looking at questions such as to what extent these systems are actually understanding and comprehending data in a human-like way.

How do you see AI in the context of journalism and the news, both as an expert but also as a normal news consumer?

There are many reasons to be careful. These systems very confidently claim things that are just plain wrong and they don't yet have the kind of understanding that human journalists have. The prospect of replacing journalists is a really bad idea and one that worries me quite a bit. There's so much downsizing already of journalism.

On another dimension, journalists have an important role in reporting on AI, which is very difficult because you get so many different opinions from people about what the state of things are, what's likely to happen in the next few years. I think it's important to communicate that there isn't just a single story. It's not just Sam Altman saying, you know, Artificial General Intelligence will be here in five years or something. It's like a whole bunch of people who have very different opinions and there's actually great uncertainty. But it's important, because it influences legislators, too.

On the question of these systems claiming wrong things: Do you think with the current models that it will ever be possible to make them trustworthy and reliable?

I think with the current models, it's not easily fixable. Maybe somebody will come up with some new method for making these things much more reliable. It's possible. RAG (Retrievel Augmented Generation), for example, is an attempt where these systems go out on the web and retrieve information so that they can fact check what they've said. But there have been studies that show that even with that, they can still make stuff up.

What's missing from current conversations?

Journalists often interview AI people but not developmental psychologists or cognitive scientists who study biological intelligence – for example, people who think about animal intelligence and how to evaluate it. Cognitive tests that we give

BACK TO CONTENTS

to animals are often not appropriate because they're sort of human centric. These days, we give machines human cognitive tests and we make inferences from that, but it might not be appropriate in the same way. I've been trying to push for those kinds of voices to come in. They have strong opinions that often differ from AI researchers and I think that that's going to be increasingly important as people make more and more claims about AI and what we can predict from different sorts of evaluations of intelligence and machines versus humans. There's no reason why AI researchers should be the only ones we hear from about the nature of intelligence.

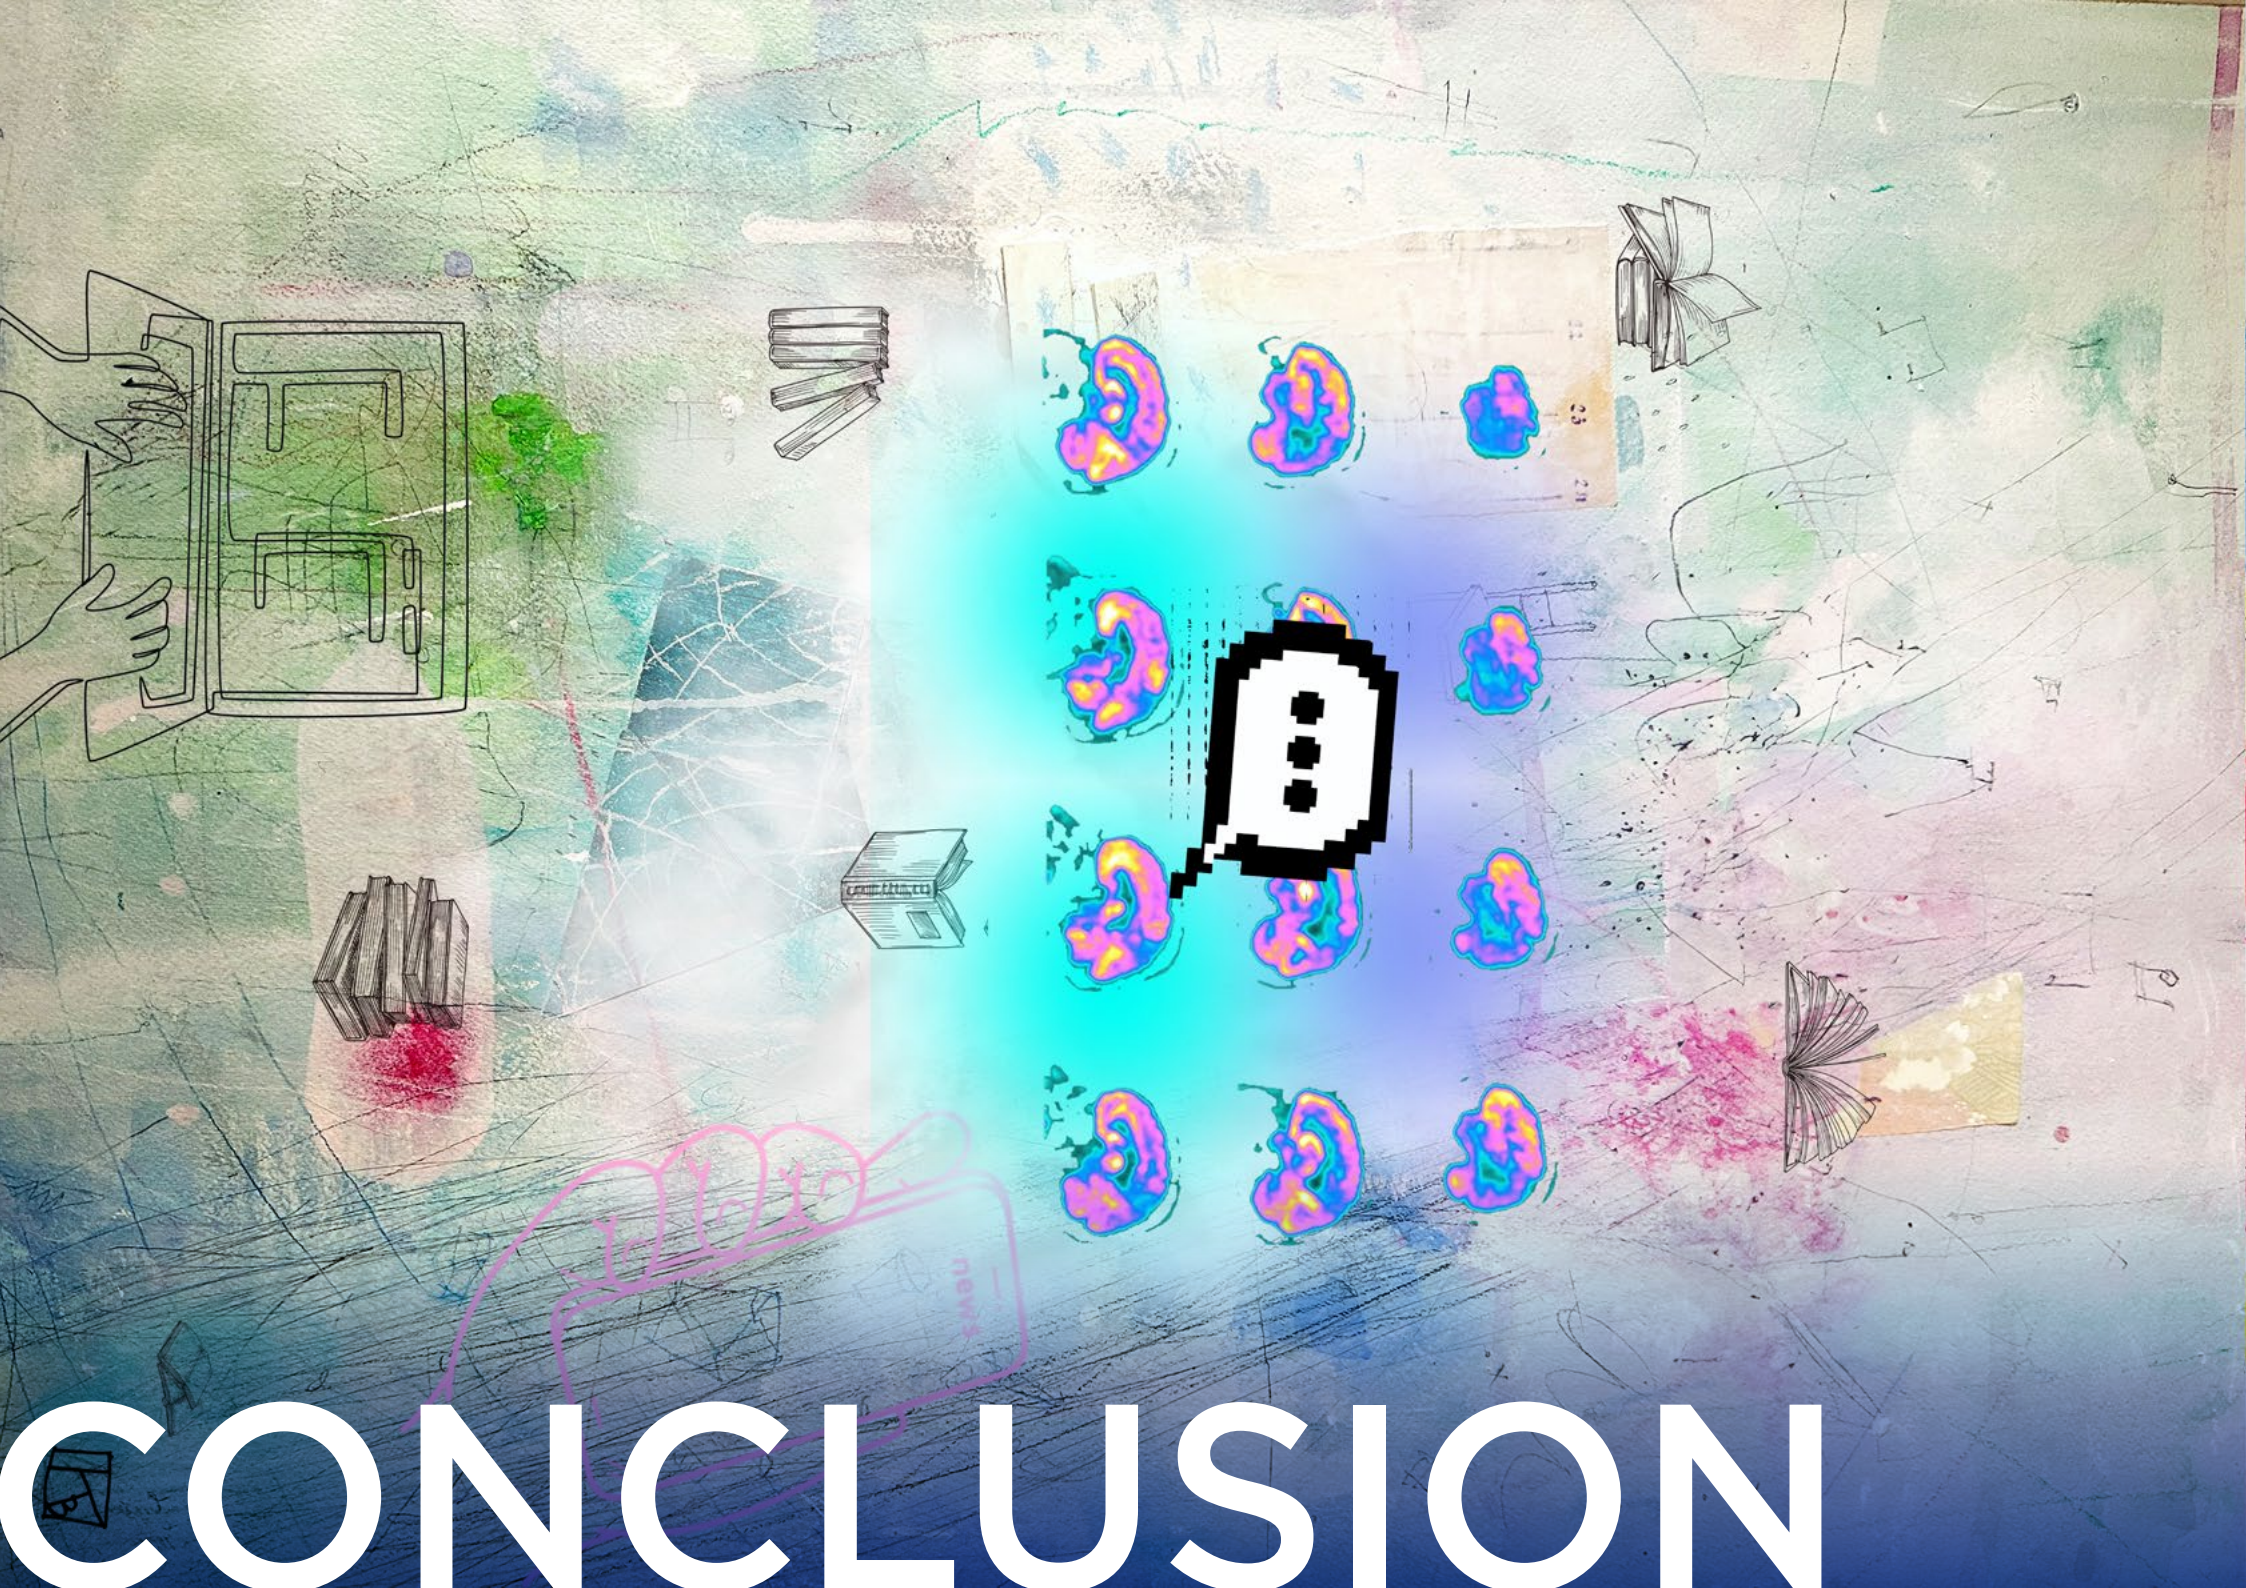

CONCLUSION

THE OLD AND THE NEW – AND A SET OF QUESTIONS

The debate about generative AI is in fact nothing less than a debate about the survival of journalism in the information ecosystem. Yes, the opportunities generative AI offers to make journalism richer, more inclusive, investigative, and attractive to different audiences are real and exciting. But the question is: who cares? Who will appreciate the value of independent reporting by supporting it with real money? Who will pitch in when business models are crumbling, and who will help to fight propaganda and misinformation that will proliferate at scale? It will take the conviction, commitment, and determination of all actors – be they in politics, business, or civil society – to help journalism flourish as the guardian of free and healthy societies. Humanity can only prosper when citizens are informed, educated, and feel protected enough to take calculated risks. Journalism gives them many of the tools they need to develop the requisite sense of personal agency. And journalism is and will remain essential in safeguarding democracies around the globe. AI-based tools can support or undermine this impact, depending on how they are used. They won't cure an absence of strategy or step up when journalism fails to deliver value.

Retaining relevance and legitimacy will be challenging for news media in the face of indifference from big tech and other power mongers. Ritu Kapur, CEO of *The Quint*, describes the struggles of independent media in an information ecosystem where they are no longer perceived as useful by influential actors: "In a country like India, the government no longer needs news media to amplify its voice. There is a media capture and legacy media is the mouthpiece of the government because they are either owned by crony capitalism or they're dependent on government advertising for revenue. Interests that might have wanted to keep news media alive now have so many other avenues. Narendra Modi has just issued an award to content creators, because they are carrying his message so much more effectively; they're cool, they're young, they're using technology, they're influencers and they have followers. I think we need to first talk about the survival of news media and then what AI is going to do to it."

These dynamics are not for the media industry alone to solve. Nevertheless, media organizations can contribute by creating the kind of journalism they should always strive for: inclusive, informative, engaging, and constructive. So, what will trustworthy journalism look like in the age of generative AI? What follows is somewhat hypothetical, however these hypotheses have emerged from interviewing close to 40 experts, digesting academic works, attending conferences, having countless conversations, and perusing endless content about generative AI during these last eighteen months of hype and hope. The overarching hypothesis is this: in an ideal world, some of tomorrow's journalism will be like yesterday's, only better – but some of today's journalism will disappear. The following is about the old and the new – and the many as yet unanswered questions where developments are at a crossroads.

The old: This will stay or be revived in journalism

First, journalism is about accuracy, facts, surprise, and storytelling. There is no doubt that in an information universe prone to hallucinations, fabricated reality, and content abundance, there will be a real need for accurate, fact-based information and storytelling that is grounded in reality.

Depending on the progress of foundation models in their literary and artistic capabilities, the appetite for standout human content might even grow. Given the current quality of generative AI output, writing or imagery that is to the point, lively, perhaps even beautiful, and reveals humanity as well as surprising facts will attract audiences. This is a definite and highly feasible opportunity for journalism.

Second, journalism that holds power to account will be as important as ever. In a world where everyone can be a content creator at scale, there will be no shortage of output that entertains, informs, and elicits emotions. But who will invest resources in checking on those who have influence, and who will confront power, based on facts, and inform society accordingly? We have law enforcement, the public service, and the judicial process. But journalists alone can take perspective as independent observers. Investigative journalism, supported by new and powerful AI-based tools, protected by strong media organizations and backed-up by press freedom legislation, will occupy a special place no content creator will ever claim.

Third, journalism will once again be about trusted and stable relationships with individuals and audiences. The notion that journalism thrives on attention predates the digital age. Tabloid media have feasted on attention ever since their inception. But the assumption of many digital first movers that attention, reach, and scale will save, nourish, and democratize journalism has proven to be short-lived. With the increasing dependence on third party-platforms, search, and social media, it is now clear that media organizations should have focused on developing stable, loyal and direct relationships with their audience all along. Particularly in confusing information environments, audience trust rests with brands, certain programmes, or individual journalists. The spike in trust during the first year of the pandemic showed that publishers and broadcasters can still mobilize these relationships. In a media environment where visibility is no longer guaranteed, building that trust will be essential for survival.

Fourth, journalism will be about editorial choices. This doesn't mean that journalists will choose and curate the format or length of news items. Technology is better at this, especially when on-demand services and personalization shape news media. But journalists will decide where to invest resources, when to push for more research and insights, and which projects to start and which ones to axe. They will pre-select what is important for people to know, and in a world in which cutting through the noise is even harder, audiences will be grateful.

Fifth, journalism will be about community. Journalists used to define the important topics of the day, and they still will. Plenty happens in the world that only reaches public debate because of journalism. Social media hasn't completely destroyed journalism's agenda-setting power, and AI won't either. Public service media need to bring people together, create shared experiences, and help people connect – be it around sports or cultural events, entertainment or current affairs. They need to drive topics that people are drawn to. As Kai Gniffke of ARD says: "Personalization always means isolation. This makes large events where society comes together all the more important."

Sixth, journalism will once again be about the real world: about meeting and talking to people, investigating, and breaking distinctive stories. Many current journalists learned their trade trapped behind computer screens in the quest for reach and scale. Most of them have excellent digital skills. But the new journalism will look a little retro: journalists will need to meet people and win their trust, look closely and dig deeper, make sure that sources are real humans instead of assuming validity unquestioningly. Verification technology will help sort the fake from the real. Original, unique content generated 'on the ground' will be at the core of the future of journalism.[1]

[1] See also Alexandra Borchardt, "Predictions for Journalism 2020: Get out of the office and talk to people," Nieman Lab, 3 January 2020. https://www.niemanlab.org/2020/01/get-out-of-the-office-and-talk-to-people/, retrieved on 21 April 2024.

The new: This will emerge in journalism

First, journalism will be about sophisticated targeting of audiences with content that enriches their lives. Where journalism hasn't respected audience time and needs in the past, technology will help. People will increasingly pull information into their lives rather than the overwhelming push they currently experience. Generative AI will help journalists to focus narratives, formats, and a tone of voice that engages audiences – ideally without compromising factual content and the message they want to convey.

Second, journalism will be mainly about research not production. The human factor in production will change and partially disappear. Print production will be the first to be automated at many publishers. Many digital functions will follow as soon as generative AI becomes more sophisticated. Ideally, human checks will remain, but much less than in current newsrooms. This is good news for those who love to report, as AI-based tools will help reporters do their best work.

Third, journalism will tell stories emerging from data that wasn't accessible before. Data journalism is still in its infancy. Generative AI will help to scan and analyse huge amounts of data in a very short time, be it hidden in text documents, videos, or audio files. Humans will need to have the right research questions and working hypotheses. This potential should excite journalists.

Fourth, journalism will become hyper-localized in a way that wasn't previously affordable. Even before publicly accessible generative AI, local news organizations benefited from AI which automated simple reporting to produce a critical mass of content with few staff. New tools will help to localize news and make it accessible to audiences in new ways. Weather or traffic reports for different regions will be customized and presented by avatars. Data journalism will take general phenomena, break them down for specific regions and generate stories that would have previously needed too much human input.

Fifth, journalism will become inclusive in a way that wasn't affordable before: people can consume it in their language in a way that meets their needs. People learn differently, have different consumption preferences, and many don't have a choice: for example, if they're hearing or sight impaired. In many countries, different local languages make it harder to communicate. Generative AI can help breaking down these barriers. Text to speech, speech to text, transcription, and translation – much is already common practice and will only improve. Progress can be expected fast, as utility is more important than perfection.

Sixth, journalists and audiences won't feel the AI behind the systems and platforms they are using. This will enhance journalists' work and audiences' experiences. It is impossible to predict how AI-based services will pervade the production and consumption of news media. But it is certain that powerful tools will assist and improve journalists in their work. However, they will remain unable to influence what is going on behind the scenes, as has been the case with the algorithmic distribution of news. This might prove to be one of the central challenges for journalism.[2] Media will need to own as much of this technology as possible.

A set of unanswered questions

Many journalists have heard at some point: You are supposed to provide answers, not questions. But after looking at a lot of evidence and listening to a considerable number of expert voices from inside and outside newsrooms, the best way to conclude this report is to present key questions and themes that will shape the future of journalism in the age of generative AI.

An important block of questions centres around the dependencies on tech companies.

[2] Recommended in this context: a panel moderated by Felix Simon with Chris Moran, Nick Diakopoulos, and Uli Köppen at the International Journalism Festival in Perugia on 20 April 2024. https://www.journalismfestival.com/programme/2024/generative-ai-in-the-news-one-year-on, retrieved on 21 April 2024.

Will generative AI force media companies to become more tech-savvy, encourage them to invest in their own models and help staff to acquire tech skills? Or will third-party tech companies shape how news is gathered and produced in the future? Most likely, the influence of big tech will shape outcomes everywhere. In the media, bigger players with more resources may be able to ramp up tech. Smaller players might profit from the 'democratizing' effect mentioned by some of our interviewees when the technology fades into the background and technical skills are no longer required for tasks such as coding.

Another important set of questions touches on the future of the creative industries in general. Will people still want to become journalists or graphic designers, film makers or novelists, when machines can do a decent enough job? Many in journalism entered the profession because they love to write. They will now have to be extremely talented to set themselves apart from what machines – or humans with the help of machines – can produce. Audiences might feel that robot journalism serves them just as well, after all, nobody expects award-winning copy about the weather or petty crime. Or perhaps generative AI will demolish entry barriers. By giving people with average language or data skills the tools to improve and excel, AI could create the diversity newsroom leaders have advocated for – at least in public. In doing so, it could contribute to solve the talent crisis in the profession.

This ties into one of the most important questions around generative AI: how will audiences react? Will they love or hate artificially created content – and will they even be able to tell the difference? It is unclear how search behaviour will evolve, and how reactions and behaviours will differ depending on social background, age, and level of education. Some say generative AI is democratizing, but the same was said of the internet, and we have seen considerable digital divides emerging. Things might happen much faster than expected – or take a lot longer. This will strongly depend on the advances of technology solutions and the benefits and convenience they generate, but also on more mundane factors like access to electricity, WIFI, and the affordability of data packages. People have proven to be surprisingly stubborn at times and fast adopters at others when the motivation was there, and conditions were right. Naturally, the development of technologies is tied to the business models of platform companies, to their interests and expected gains as much as to the regulatory environment.

Another question will be how technology will perform in the race between fakes and their detection. Will misinformation and disinformation crowd out facts? Or will verification technologies improve at such a speed that the fabrication of misleading content is discouraged or even made redundant? Misinformation and disinformation will never disappear. So far, the biggest challenge to a fact-based public discourse has not been the spread of so-called 'fake news' but propaganda, bias, and political communication that caters to particular interests and feeds polarization. In short: people trust who they want to trust, no matter how profound the evidence. Public service media will need to counter these trends as much as possible.

Then there are questions around the quality of off-the-shelf foundation models. Will they get better over time, will they suffer ups and downs, or will everything be watered down to the point of redundancy? Those who work with these models have made disparate observations. Particularly people who rely on free tools are likely to be offered products at the lower end of the performance scale. And yet the quality of the output will not only determine the efficiency gains for newsrooms and its usefulness for audiences, but also trust in the technology itself.

Another set of issues centres around the sustainability of the hardware needed to enable this technology. What about energy, water, and resource demands in the context of geopolitical tensions? The tech industry

has no interest in elevating discussions around these topics, so, shortages will only become apparent once they are felt by the broader public. The pandemic and war in Ukraine have made people more aware of how fragile some of the world's logistic chains are and how energy and food security can be threatened by conflict. A lot will come down to the question: will tech companies be able to develop high performing AI that is compatible with the sustainability of humanity and the natural environment? A lot more research needs to be done to give a robust and thorough answer.

The most important question for the media industry might still be: will business models survive the onslaught of generative search and AI-informed information products? Much will depend on how media brands remain visible for their audiences and customers – and the perceived value of the journalism they produce. Tech providers who want their models to work may be willing to invest in media companies. But this could prove to become just another winners-take-all scenario. In the end, the key question is: will news consumers understand the value of journalism, and will independent press receive the support it needs even from those who feel undermined by it?

Some industry leaders claim that the public needs to be informed of the significance of journalism for society and democracy, but this is not enough. Every media organization should aspire for its journalism to be valuable, relevant, and significant. As ARD Chairman Kai Gniffke says: "We must be and remain a reliable companion and verifier for people." Journalism needs to prove its usefulness again and again. Ringier Media International CEO Dmitry Shishkin says: "Look at some of the most important digital services, look at Amazon, Spotify, Uber – people don't question them. You don't need to tell them what would disappear if they were gone tomorrow."

Correspondingly, people need to perceive journalism as an essential service they would not want to lose. Luckily, a large proportion of the public seems to sense when media freedom is endangered. A survey in four Central European countries published in April 2024 revealed as much.[3] It might be appropriate to finish with the perspective of someone from a comparatively young democracy, Styli Charalambous, CEO of South Africa's *Daily Maverick*: "It all comes down to the fact that journalism has to be treated as a public good. And it has to be supported as a public good."

The age of generative AI offers huge opportunities for journalism. With the support of AI-driven tools, newsrooms can create more value and deliver it to more people than they do today. At the same time, it's probably never been easier to eradicate journalism entirely. An all-out effort is needed to preserve and advance it.

[3] Vaclav Stetka, "Media Freedom Poll 2024", Committee for Editorial Independence, 25 April 2024. https://mediafreedompoll.com/en/, retrieved on 5 May 2024.

GENERATIVE AI GUIDES

BACK TO CONTENTS

DOs & DON'Ts IN COVERAGE

When reporting on rapidly developing technologies like AI, it can be difficult to balance providing accounts of real progress with criticism of potential effects as well as questions of power, control and benefit. However, accurate reporting on generative AI is crucial as it influences public perception, policymaking, and industry decisions – especially during this pivotal adoption phase. To help reporters and managers cover this complex topic effectively, we have prepared this short guide of dos and don'ts.

DOs

1. Develop a basic technical understanding

While you don't need to be a computer scientist (ignore those who suggest otherwise), aim to grasp the technical fundamentals of AI and its capabilities. You don't need to fully understand how machine learning works mathematically, but it is useful to understand, for example, the distinction between various types of AI and AI systems, such as machine learning, neural networks, and foundation models. There are good resources and books available (see 'resources' below) to help you get your head around it and avoid conflating these concepts, which can lead to confusion and inaccurate reporting. It's also important to differentiate between genuine applications of AI and simpler statistical models to avoid misrepresentation. Not everything that is labelled 'AI' actually contains artificial intelligence.

2. Consult a diverse range of experts

Avoid relying solely on contributions from AI companies or a single expert's view. To gain a comprehensive understanding of AI and provide balanced coverage, seek input from a variety of sources, including academics, regulators, and industry professionals. When evaluating the intelligence of an AI system, don't limit yourself to just computer scientists. Involve cognitive scientists, experts in child development and learning, and linguists in your reporting as these often have important views to add to current debates. Be cautious of company representatives aiming to promote their products and services.

3. Pay attention to both benefits and risks

Whenever you highlight the potential advantages of AI, pay the same attention to discussing ethical concerns, risks, and challenges, including around questions of bias and fairness, privacy, copyright, and harm. Or as the AP's Garence Burke puts it, ask: "Where are they [AI systems] deployed? How well do they perform? Are they regulated? Who's making money as a result? And who's benefiting? And also, very importantly, which communities may be negatively impacted by these tools?"[1]

[1] https://reutersinstitute.politics.ox.ac.uk/news/focus-humans-not-robots-tips-author-ap-guidelines-how-cover-ai

4. Recognise the dynamic nature of AI

When reporting on AI, it is crucial to understand and convey the rapidly evolving and dynamic nature of this field. The development of AI systems is an ongoing process characterised by constant advancements, setbacks, and shifts in trajectory. One key challenge stems from the uncertainty surrounding the long-term effects of many decisions being made in the AI domain at this moment. Choices made today, whether related to model deployment and use, regulation, data policies, open-source initiatives, corporate acquisitions, or research partnerships, can have significant long-term ramifications that are difficult to fully anticipate. As a result, AI coverage must strike a balance between providing timely and accurate information while acknowledging the inherent uncertainty in this field.

» DON'Ts

1. Overhype capabilities

The story of AI is ripe with examples of exaggerated claims about AI's abilities, including many in the media. This can create unrealistic public expectations, but also lead to over-regulation or misdirected regulation. It can also lead to poor decisions regarding investing in AI, including in journalism, with reality not matching expectations. It's important to report on setbacks and failures in AI development to provide a balanced view.

2. Humanize the technology

It is easy to ascribe human feelings or capabilities to AI systems or to imply that these systems can 'think'. Avoid terms like 'AI is thinking' or 'AI feels,' which anthropomorphise technology and can mislead about its nature and its limitations. Instead, describe AI in terms of its algorithms, data processing abilities, and programmed functions to provide a more accurate representation of how these systems work and what they do. This goes for the depiction of AI, too. As researcher Maggie Mustaklem reminds us, all too often there is a 'one-size-fits-all sci-fi fantasy' around AI,[2] with the technology portrayed as "white robots typing on a keyboard" or "a blue graphic of a human brain connected to some colourful lines" (the artist and technologist Neema Iyer has collected some of these tropes on her website[3]) – yet neither does the complexity of the technology justice and both depictions are misleading.

3. Ignore the human element and contextual factors

Do not neglect the role of humans and the implications on human lives. AI is not just a story of technology, but a story of technology working in society. And AI does not exist in the ether. To enable AI, a multifaceted supply chain[4] must function, from mining rare minerals for chips to data centres consuming energy and water for cooling. Humans are integral to every aspect of this chain and are affected by its components. It's crucial to acknowledge these factors and incorporate the individuals who work 'behind the scenes' to facilitate AI as well as those directly impacted by its deployment in your coverage. This includes

[2] https://www.oii.ox.ac.uk/news-events/can-ai-visuals-move-away-from-blue-brains-and-cyborgs/
[3] https://neemaiyer.com/work/how-do-we-picture-ai-in-our-minds
[4] https://www.adalovelaceinstitute.org/resource/ai-supply-chains/

broadening the lens beyond developments in Western countries, which often take centre stage in discussions of AI.

4. Do not treat AI just as a innovation or technology issue
AI's implications extend far beyond just questions of technological innovation in some specific domains. Instead, like climate change, AI is a topic that cuts across numerous domains, including business, law, healthcare, education, politics, and the environment. Comprehensive AI coverage requires recognizing this, exploring how AI does and doesn't reshape various areas, and helping audiences arrive at a more holistic understanding of both the opportunities and challenges AI presents. Treating AI solely as a technological advancement is no longer enough.

RESOURCES

- The AI Myths website debunks common myths and misunderstandings about artificial intelligence. It is structured into eight distinct sections, each exploring different facets of AI including its portrayal, definition, governance, and practical applications. Each section provides links to additional resources on the topics discussed: https://www.aimyths.org/

- The Leverhulme Centre for the Future of Intelligence at the University of Cambridge has collected a range of guidelines for reports on how to better cover AI and what to look out for, each of them worth reading in detail: https://www.desirableai.com/journalism-toolkit-ethics. They also provide links to various databases of AI experts and voices: https://www.desirableai.com/journalism-toolkit-voices

- The ways AI is depicted can obscure the real and significant societal and environmental impacts of the technology, can set unrealistic expectations and misrepresents the actual capabilities of AI. It can also obscure the responsibility of the humans behind the technology. The Better Images of AI project provides alternatives: https://betterimagesofai.org/

GENERATIVE AI GUIDE – BOOKS, REPORTS, ARTICLES: UNDERSTANDING THE TECHNOLOGY

For those who would like to read up on generative AI, here are a few resources we found useful for different needs.

EBU's topic page on Artificial Intelligence Although some of these resources will only be available to EBU members (via login) there is a wealth of constantly updated information on how AI impacts public service media, including the latest tools, research, video talks, events and groups. https://www.ebu.ch/topics/artificial-intelligence

EBU Academy's School of AI Open to EBU Members and non-Members, the EBU Academy has a broad range of training opportunities related to AI. Browse the courses for something which will help you get up to speed and sign up for their newsletter so you don't miss the latest news. https://academy.ebu.ch/schoolofai

Melanie Mitchell, 'Artificial Intelligence: A Guide for Thinking Humans' (Pelican, 2019) Written by a world-leading computer scientist at the Santa Fe Institute, this book is on AI more generally and forms essential reading for anyone hoping to gain a better understanding of the foundations of AI and some of the key issues involved. In addition, Mitchell has penned a series of insightful and accessible essays for *Science* magazine looking at AI's challenge of understanding the world (https://www.science.org/doi/full/10.1126/science.adm8175), the question of how '"smart"' AI systems really are (https://www.science.org/doi/10.1126/science.adj5957), and the debates around artificial general intelligence (https://www.science.org/doi/10.1126/science.ado7069).

UK's House of Lords Communications and Digital Committee, 'Large language models and generative AI' (UK's House of Lords, 2024) This report on foundation models, large language models and generative AI provides a detailed and accessible overview of the technological underpinnings of generative AI, including central terms such as 'foundational model', 'LLM' and more. It advocates to balancing a focus on AI safety with enabling commercial and academic growth and calls for prompt, nuanced action to address risks including market monopolization, regulatory capture, and the misuse of copyrighted material, while recommending a series of strategic, regulatory, and supportive actions to harness the benefits of generative AI responsibly. https://publications.parliament.uk/pa/ld5804/ldselect/ldcomm/54/54.pdf

Air Street Capital, 'State of AI Report': Produced by investors at Air Street Capital, a venture capital firm in AI and life sciences, the annual AI report provides an up-to-date overview of AI research, industry trends and connected topics and is a very useful resource for anyone trying seeking to understand recent industry trends and developments. https://www.stateof.ai/

Jai Vipra & Sarah Myers West, 'Computational Power and AI' (AI Now Institute, 2023) This report discusses the central role of computational power in the development of AI, highlighting that it is controlled by a few firms, which affects who can build AI and the nature of its development. It underlines that the concentration in computational resources leads to industry monopolization, encourages detrimental competitive practices among AI firms, and has significant environmental and socio-political impacts. The report also stresses the importance of policy interventions to mitigate negative outcomes and ensure ethical and equitable access to computing power for AI development. https://ainowinstitute.org/publication/policy/compute-and-ai

Madhumita Murgia, 'Code Dependent' (Picador, 2024) The *Financial Times*' AI Editor offers a compelling analysis and critique of the overlooked issues with AI: Murgia shows how the belief that technology can solve all problems leads to dead ends. She also explores the societal impact of AI, emphasizing the marginalized groups affected by it and the unequal distribution of its benefits, urging readers to confront the loss of agency and the erosion of free will in a world governed by flawed and exploitative technologies.

David Caswell, 'AI in Journalism Challenge 2023' (Open Society Foundations, 2024) This report describes the outcomes of the Applied AI in Journalism Challenge, an initiative by Open Society Foundations aimed at helping prototype pragmatic applications of (generative) artificial intelligence in small and low-resource newsrooms globally. Apart from showcasing various diverse projects that span the entire news value chain, the report emphasizes that even small news teams can leverage AI to enhance reporting, content creation, and audience engagement. A key part describes how the experimentation and integration of AI systems can be successfully achieved, including from a management perspective. https://www.opensocietyfoundations.org/publications/open-society-s-applied-ai-in-journalism-challenge

Arvind Narayanan & Sayash Kapoor, 'AI Snake Oil' (Princeton University Press, 2024) This book is yet to be published but Arvind Narayanan is among to the leading researchers in the space of AI and its social implications and has produced various high-impact and accessible explainers and reports on AI in recent months, many in collaboration with his PhD-student and co-author Sayash Kapoor. The book, drawing in parts from their research and a successful and widely read Substack, explains the crucial differences between types of AI, why people, companies, and governments are falling for AI snake oil, why AI can't fix social media, and why we should be far more worried about what people will do with AI than about anything AI will do on its own.

BACK TO CONTENTS

SELECTION OF AI ETHICS GUIDES FOR MEDIA ORGANIZATIONS

Paris Charter on AI and Journalism, published on 10th November 2023

1. Journalism ethics guide the way media outlets and journalists use technology.
Media outlets and journalists use technologies that enhance their capacity to fulfill their primary mission: ensuring everyone's right to quality, trustworthy information. The pursuit and achievement of this goal should drive their choices regarding technological tools. The use and development of AI systems in journalism must uphold the core values of journalistic ethics, including truthfulness and accuracy, fairness, impartiality, independence, non-harm, non-discrimination, accountability, respect for privacy and for the confidentiality of sources.

2. Media outlets prioritize human agency.
Human decision-making must remain central to both longterm strategies and daily editorial choices. The use of AI systems should be a deliberate and well informed decision made by humans. Editorial teams must clearly define the goals, scope, and usage conditions for each AI system. They must ensure a cross-sectional and continuous oversight of the impacts of deployed AI systems, ensure their strict compliance with their usage framework, and retain the ability to deactivate them at any time.

3. AI systems used in journalism undergo prior, independent evaluation.
The AI systems used by the media and journalists should undergo an independent, comprehensive, and thorough evaluation involving journalism support groups. This evaluation must robustly demonstrate adherence to the core values of journalistic ethics. These systems must respect privacy, intellectual property and data protection laws. A clear accountability framework is established for any failure to meet these requirements. Systems that operate predictably and can be simply explained are preferred.

4. Media outlets are always accountable for the content they publish.
Media outlets assume editorial responsibility, including in their use of AI in gathering, processing, or disseminating information. They are liable and accountable for every piece of content they publish. Responsibilities tied to the use of AI systems should be anticipated, outlined, and assigned to humans to ensure adherence to journalistic ethics and editorial guidelines.

5. Media outlets maintain transparency in their use of AI systems.
Any use of AI that has a significant impact on the production or distribution of journalistic content should be clearly disclosed and communicated to everyone receiving information alongside the relevant content. Media outlets should maintain a public record of the AI systems they use and have used, detailing their purposes, scopes, and conditions of use.

6. Media outlets ensure content origin and traceability.
Media outlets should, whenever possible, use state-of-the art tools that guarantee the authenticity and provenance of published content, providing reliable details about its origin and any subsequent changes it may have undergone. Any content not meeting these authenticity standards should be regarded as potentially misleading and should undergo thorough verification.

7. Journalism draws a clear line between authentic and synthetic content.
Journalists and media outlets strive to ensure a clear and reliable distinction between content derived from the physical capture of the real world (such as photographs, and audio and video recordings) and that which is created or significantly altered using AI systems. They should favor the use of authentic footage and recordings to depict actual events. Media outlets must avoid misleading the public in their use of AI technologies. In particular, they should refrain from creating or using AI-generated content mimicking real-world captures and recordings or realistically impersonating actual individuals.

8. AI-driven content personalization and recommendation upholds the diversity and the integrity of information.
In media outlets, the design and use of AI systems for automatic content personalization and recommendation should be guided by journalistic ethics. Such systems should respect information integrity and promote a shared understanding of relevant facts and viewpoints. They should highlight diverse and nuanced perspectives on various topics, fostering open-mindedness and democratic dialogue. The use of such systems must be transparent, and users should whenever possible be given the option to disable them to ensure unfiltered access to editorial content.

9. Journalists, media outlets and journalism support groups engage in the governance of AI.
As essential guardians of the right to information, journalists, media outlets and journalism support groups should play an active role in the governance of AI systems. They should be included in any global or international institutional oversight of AI governance and regulation. They should ensure that AI governance respects democratic values, and that diversity of people and cultures is reflected in the development of AI. They must remain at the forefront of knowledge in the field of AI. They are committed to examining and reporting on the impacts of AI with accuracy, nuance, and a critical mind.

10. Journalism upholds its ethical and economic foundation in engagements with AI organizations.
Access to journalistic content by AI systems should be governed by formal agreements that ensure the sustainability of journalism and uphold the long-term shared interests of the media and journalists. AI system owners must credit sources, respect intellectual property rights, and provide just compensation to rights holders. This compensation must be passed on to journalists through fair remuneration. AI system owners are also required to maintain a transparent and detailed record of the journalistic content utilized to train and feed their systems. Rights holders must make the repurposing of their content by AI systems conditional on respect for the integrity of the information and the fundamental

principles of journalistic ethics. They collectively call for AI systems to be designed and used in such a way as to guarantee high-quality, pluralistic and trustworthy information.

Initiator: Reporters Without Borders (RSF), **Partners:** Asia-Pacific Broadcasting Union (ABU), Collaboration on International ICT Policy in East and Southern Africa (CIPESA), Canadian Journalism Foundation (CJF), Committee to Protect Journalists (CPJ), DW Akademie, European Federation of Journalists (EFJ), European Journalism Centre (EJC), Ethical Journalism Network (EJN), Free Press Unlimited (FPU), Global Investigative Journalism Network (GIJN), Global Forum for Media Development (GFMD), International Consortium of Investigative Journalists (ICIJ), International Press Institute (IPI), Organized Crime and Corruption Reporting Project (OCCRP), Pulitzer Centre, Thomson Foundation

Examples for Ethical Guidelines of EBU Members

• BBC AI Principles - BBC

• An update on the BBC's plans for Generative AI (Gen AI) and how we plan to use AI tools responsibly - BBC

• Ethics of Artificial Intelligence: Our AI Ethics Guidelines - BR

• What is Deutsche Welle's approach to generative AI? - Deutsche Welle

• Yle's principles for the responsible use of AI - Yle

More guidelines compiled by the EBU Ethics Group can be found here, a member login is needed to access them.

GENERATIVE AI GLOSSARY FOR JOURNALISM

Artificial intelligence (AI): A field of computer science that includes systems capable of performing tasks that normally require human intelligence such as language understanding, pattern and image recognition, decision-making and problem-solving.

Generative artificial intelligence: AI systems and algorithms capable of creating content. Uses algorithms such as generative adversarial networks (GANs), transformers and diffusion models to create new data indistinguishable from real data. It can create new images, music, computer code and even videos. With its near-perfect analysis of natural language, it can assist with numerous tasks: analysis, translation, named entity recognition, reasoning (not to be confused with human reasoning).

Predictive artificial intelligence: Uses machine-learning algorithms and models to predict future or unknown outcomes based on historical and existing data. It can predict market trends, customer behaviour, or the likelihood of a machine failure in a factory. In news, it helps forecast election outcomes, market movements and audience engagement.

Key Concepts

Machine learning: A sub-field of artificial intelligence composed of algorithms that identify relevant information from observational data.

- **Deep learning:** Uses neural networks to enable AI to learn from large amounts of data and perform complex tasks. In journalism, deep learning can analyse large datasets to find patterns and insights for investigative stories.
- **Reinforcement learning:** AI learns by trying different actions and improving based on the outcomes. For news, this can optimize content recommendations and engagement strategies.
- **Transfer learning:** Reusing a model developed for one task as the starting point for a model on a second task. This can speed up the development of specialized journalism tools.
- **Adversarial machine learning:** Studying how AI systems can be deceived by specially crafted inputs. Relevant for detecting and preventing the spread of misinformation.

Neural network: A network of artificial neurons designed to simulate the way humans think and learn. Used in generative AI to create realistic text and multimedia content for news stories.

Data mining: The process of discovering patterns in large datasets using machine learning. Journalists can use data mining to uncover trends and insights from big data sources.

Automated machine learning (AutoML): Automating the process of applying machine learning to real-world problems. Simplifies the creation of AI tools for newsrooms, such as automated article writing and analysis.

Natural language processing (NLP): Technology that allows robots to understand and produce human language.

Natural language generation (NLG): Creating natural language text from structured data.

Large language model (LLM): A model trained on large amounts of data that uses supervised learning to generate text coherently and meaningfully. In news, LLMs can draft articles, provide translation and generate content summaries.

Transformer models: Based on Transformer architecture, these models use attention mechanisms for tasks like translation and text generation. Essential for generating high-quality, context-aware news content.

Generative Pre-trained Transformer (GPT): A type of LLM developed by OpenAI that uses Transformer architecture to generate human-like text. In journalism, GPT can be used to create content, generate headlines, and assist with editing and summarization.

Retrieval-augmented generation (RAG): Combines information retrieval with text generation, using retrieved documents to enhance the quality and accuracy of generated content. Useful in journalism for fact-checking and detailed report generation.

AI Interaction Technologies

Text to image (TTI): Algorithms capable of generating an image from a text description.

Text to speech (TTS): Technology that converts written text into spoken audio.

Speech to text (STT): Technology that converts spoken audio into written text.

Text to text (TTT): Transforms a source text into another text.

Chatbot: An AI program that can simulate a conversation with users. Often used in customer service, information retrieval and interaction with digital systems. News organizations use chatbots for personalized news delivery and audience engagement.

Agent: In AI, an autonomous entity that observes and acts upon an environment to achieve specific goals. Newsrooms use agents to automate news gathering and content curation.

Multimodal and Explainable AI

Multimodal learning: AI learning from multiple types of data at the same time, such as images and text. Enables the creation of rich, interactive news content that combines text, images and videos.

Deployment and Infrastructure

Robotic Process Automation (RPA): Automating routine tasks with software robots or AI. Used in newsrooms to automate repetitive tasks like data entry and content updates.

Application Programming Interface (API): Allows developers to integrate and interact with generative AI and LLMs within their own applications. For journalism, APIs can provide access to powerful AI tools for content generation, data analysis and more.

On-premise: Refers to software and technology that are located within the physical confines of an enterprise, often in the company's data centre. News organizations might use on-premise solutions for secure and compliant handling of sensitive data.

On-cloud: Refers to services, storage and applications that are hosted on the internet rather than on a local server or personal computer. LLMs can be deployed on cloud platforms to leverage scalable resources and manage large-scale computations required for model training and inference. In journalism, this allows for flexible and scalable AI deployments.

Open source: Software that is made available with its source code, allowing anyone to inspect, modify and enhance it. In the context of generative AI and LLMs, open-source models can be used by journalists to develop custom tools for content generation, data analysis and more, promoting innovation and collaboration within the newsroom

Large language model operating system (LLM OS): A conceptual operating system designed to integrate and manage LLMs for various applications, optimizing their performance and facilitating their deployment in diverse environments. In newsrooms, this can streamline the integration of AI into editorial workflows.

Advanced Generative Models

Generative adversarial network (GAN): A deep-learning model where two neural networks, the generator and the discriminator, are pitted against each other. The

generator tries to create fake data that is indistinguishable from real data, while the discriminator tries to identify the fake data from the real data. This competition improves both models, resulting in very realistic fake data from the generator. Used in journalism to create realistic synthetic images and videos.

Diffusion model: Used to generate new data from existing data. It works by simulating a reverse diffusion process, starting with random noise and gradually shaping it into data that resembles the target data. Often used for generating images or other types of visual data. Can create high-quality visual content for news stories.

Large action model (LAM): A model designed to perform and generate sequences of actions, often in the context of robotics or automated processes. In journalism, this can automate complex workflows and content generation processes.

AI Ethics and Challenges

Explainable AI (XAI): AI designed to be transparent, providing explanations of its decision-making process. Important for maintaining transparency and trust in AI-generated news content.

AI Ethics: Concerns about the moral implications of artificial intelligence. Critical for ensuring responsible use of AI in news generation and dissemination.

Bias: The inclination or distortion of a machine-learning model, intentional or otherwise, often caused by unrepresentative or prejudicial training data. Important to address in journalism to ensure fair and unbiased news reporting.

Hallucination: In AI, particularly in generative models and LLMs, this refers to the generation of incorrect or nonsensical information that appears plausible but has no basis in the input data or reality. For example, an AI might create a completely fabricated event or entity when generating text.

Confabulation: Occurs when an AI system provides false information confidently and coherently, as if it were true, often filling in gaps in knowledge with fabricated details. This differs from hallucination in that it involves the AI attempting to create a coherent narrative or explanation that fits with the known data, even if parts of it are made up.

Deepfake: A synthetic media in which a person in an existing image or video is replaced with someone else's likeness using deep learning. In journalism, deepfakes pose challenges in verifying the authenticity of visual content, but can also be used creatively for reenactments or educational purposes.

Watermarking: The process of embedding information into digital content that can be used to verify the content's authenticity or the identity of its owners. In journalism, watermarking can ensure the integrity of original content and prevent unauthorized use.

BACK TO CONTENTS

Fingerprinting: A technique used to identify and track digital media by analysing unique characteristics. In the context of news, fingerprinting helps track the distribution of content and detect unauthorized copies.

Data annotation: Labelling data to indicate the underlying features for AI models. Essential for training accurate and reliable AI models.

AI fatigue: A feeling of saturation due to the recent hype and discussions around AI.


APPENDIX

BACK TO CONTENTS

List of interviewees

JANE BARRETT, Global Editor, Media News Strategy, Reuters News Agency

DAVID CASWELL, Consultant, Storyflow

STYLI CHARALAMBOUS, CEO, *Daily Maverick*

MADHAV CHINNAPPA, Media Consultant, former Director of News Ecosystem Development at Google

EZRA EEMAN, Strategy and Innovation Director, NPO

LUCIANO FLORIDI, Professor and Founding Director, Digital Ethics Center Yale University

MATT FREHNER, Managing Editor, Products and Platforms, *The Globe and Mail*

CARL-BENEDIKT FREY, Dieter Schwarz Associate Professor of AI and Work, Oxford Internet Institute

KAI GNIFFKE, Director General and Chair, SWR/ARD

BLATHNAID HEALY, Executive News Editor for Growth, Social and Delivery, BBC

NATALI HELBERGER, Professor in Law and Digital Technology, Institute for Information Law, University of Amsterdam

JEFF JARVIS, Author, Professor Emeritus at Craig Newmark Graduate School of Journalism at CUNY

BRONWYN JONES, Translational Fellow and Journalist, University of Edinburgh and BBC

RHIANNE JONES, BBC UKRI Future Leaders Fellow, BBC

RITU KAPUR, Founder and CEO, The Quint

MANUELA KASPER-CLARIDGE, Editor-in-Chief, Deutsche Welle

TAMARA KNEESE, Project Director, Algorithmic Impact Methods Lab, Data and Society Research Lab

ULI KOEPPEN, Head of AI + Automation Lab, Bayerischer Rundfunk

VERENA KRAWARIK, Head of Innovation and Medialab, APA

PETER DAMGAARD KRISTENSEN, COO, Constructive Institute

LUCY KÜNG, Strategic Advisor and Board Member, independent, Reuters Institute for the Study of Journalism

GUSTAV AARUP LAURIDSEN, Lead Data Scientist, Constructive Institute

ANNE LAGERCRANTZ, Deputy Director General, SVT

MELANIE MITCHELL, Davis Professor of Complexity, Santa Fe Institute

GINA NEFF, Professor and Executive Director, Minderoo Centre for Technology & Democracy, University of Cambridge

NIC NEWMAN, Senior Research Associate, Reuters Institute for the Study of Journalism

MATTIA PERETTI, Consultant and ICFJ Knight Fellow

JEAN-MARC RICKLI, Head of Global and Emerging Risk, Geneva Centre for Security Policy

ERIK ROOSE, Chairman of the Board, Estonian Public Broadcast

NIDDAL SALAH-ELDIN, Member of the Executive Board, Axel Springer SE

RENATE SCHRÖDER, Director, European Federation of Journalists

DMITRY SHISHKIN, CEO, Ringier Media International

SARAH SPIEKERMANN, Professor for Information Systems, WU Vienna

AGNES STENBOM, Head of IN/LAB & Industrial PhD Candidate, Schibsted

ANDREW STRAIT, Associate Director, Emerging technology & industry practice, Ada Lovelace Institute

SINA THÄSLER-KORDONOURI, PhD Candidate in Automation and AI in Journalism, IfKW, Ludwig-Maximilians-University Munich

BILL THOMPSON, Head of Public Value Research, BBC

JOHANNA TÖRN-MANGS, Director, Editor-in-Chief, Svenska Yle

ACKNOWLEDGEMENTS

We, the authors, editors, designers, and commissioners of this report, are deeply grateful for all the support we received in crafting this piece of work. First, we would like to thank the more than three dozen international media leaders, leading researchers and experts who openly shared their experiences, insights, and knowledge with us – commenting with confidence can be a challenge in a field as rapidly developing as Artificial Intelligence. All our sources devoted major chunks of their precious time to being interviewed and repeatedly contacted for updates. Their contributions are invaluable and at the very core of this work.

The AI and journalism community across the globe is growing, and we profited plenty from some outstanding networks and insights shared by them. Among these are the Journalism AI Project at the London School of Economics, the Nordic AI Journalism Network, the Reuters Institute for the Study of Journalism and the Oxford Internet Institute both at the University of Oxford, and the Tow Center at Columbia University.

From the EBU's supporting teams, we would particularly like to thank Jo Waters who has led on communications and continuously contributed by developing and promoting preliminary material and publishing several Q&As. John Howe worked tirelessly on the design. Laurent Frat expertly prepared and produced a podcast. Sally Clarke did a superb job on proofreading.

We are sincerely grateful to the EBU News Committee which funded the report, first and foremost, to its chairperson Eric Scherer. Finally, we are indebted to Noel Curran, the EBU's Director General, Jean Philip De Tender, Deputy Director General and Media Director, and Liz Corbin, Director of News, for trusting us with this project. They provided invaluable support and strategic guidance. Without them, this work could not have occurred. Of course, any errors and omissions are entirely our responsibility.

BACK TO CONTENTS

THE TEAM

DR ALEXANDRA BORCHARDT

Alexandra Borchardt is the lead author of 'Trusted Journalism in the Age of Generative AI' as well as the two previous EBU News Reports 'Climate Journalism That Works – Between Knowledge and Impact' (2023) and 'What's Next? Public Service Journalism in an Age of Distraction, Opinion, and Information Abundance' (2021). She is an independent journalist, book author, lecturer, and media advisor with more than 25 years of experience in newsrooms of major news brands. She is affiliated with the World Association of News Publishers (WAN-IFRA), the Constructive Institute in Aarhus, and the Reuters Institute for the Study of Journalism at the University of Oxford. She used to be managing editor of Süddeutsche Zeitung (SZ) in Munich. Alexandra is an honorary Professor of Leadership and Digital Transformation at TU Munich's TUM School of Management. She holds a PhD in Political Science from Tulane University, New Orleans.

mail@alexandraborchardt.com

@AlexaBorchardt

DR FELIX M. SIMON

Felix M. Simon is a communication researcher in AI and Digital News at the Reuters Institute for the Study of Journalism at the University of Oxford. He earned his doctorate at the Oxford Internet Institute (OII). He was previously a Knight News Innovation Fellow at Columbia University's Tow Center for Digital Journalism (2021-2024). A former journalist, he writes on technology, media and politics for international outlets.

Since 2019, Felix has researched AI's impact on news – publishing in top journals and appearing in *The Guardian* and the *Financial Times*. He advises media organizations on AI. In May 2023, he won the Hans Bausch Media Prize for his work on AI and news.

He holds a BA in Film and Media Studies and English Studies from Goethe University Frankfurt, and an MSc in Social Science of the Internet from the OII. He is also a fellow of the Salzburg Global Seminar and an Associate Fellow of the UK Higher Education Academy.

Find him on **LinkedIn** and **X** or write to **felix.simon@oii.ox.ac.uk**

OLLE ZACHRISON

Olle Zachrison is Head of AI and News Strategy at Swedish Radio, leading the public broadcasters' company-wide AI-efforts. He is a driving force behind EBU's digital news sharing initiative, A European Perspective, and its public service algorithm. Olle is co-founder of Nordic AI Journalism, an industry network with 450+ members and co-organiser of the annual Nordic AI in Media Summit. He was previously Head of SR's national news department, Ekot, and before that Deputy Editor of Swedish daily newspaper, Svenska Dagbladet.

KATI BREMME

Kati Bremme is a media expert with 20 years of experience in digital transformation for broadcasters. With a background in the arts and humanities, marketing, publishing, and TV production, she currently specializes in AI, and Web3 and their strategic applications for the media sector. Throughout her media career, she has worked in both broadcast and digital realms and is now adept at building bridges between television, radio, and the digital world, reflecting their interconnected roles in the public's daily life. She currently works as the Head of Innovation at France Télévisions, the French public broadcaster. She is also the editor-in-chief of the blog Meta-Media, a collaborative platform exploring the future of journalism and media literacy. Kati's artistic talents are on display in this report's illustrations

JUSTYNA KURCZABINSKA

Justyna Kurczabinska leads the News Strategy and Transformation of Eurovision News at the European Broadcasting Union (EBU - the world's leading alliance of public service media). She was the Head of the Eurovision News Exchange providing EBU Members with round-the-clock access to live and edited news stories as well as eyewitness media for 7 years. Since 2021 she has championed 'A European Perspective' project which aims at redefining digital news online collaboration among public service media organizations with the objective of providing audiences across Europe with powerful and relevant public service journalism. She has been overseeing the Public Service Journalism Initiative aimed at strengthening and supporting independent trusted news since 2017. The annual EBU News Report has become one of the crucial elements of this initiative.

kurczabinska@eurovision.net
@kurczabinska

ED MULHALL

Ed Mulhall is an editorial advisor and media consultant. He is a former Managing Director of RTÉ News and Current Affairs and previously worked as a radio producer, television producer, news programme editor and managing editor of TV News. He was an elected member of the EBU News Committee for several terms during his time with RTÉ News. Since leaving RTÉ he has worked as an editorial advisor and consultant, is a research associate with the Department of Political Science, Trinity College Dublin and has written on history, literature, and media topics. He has acted as an Editorial advisor for the EBU News Department on many projects including successive EBU News Reports.

YOLÈNE JOHANNY

Yolène Johanny is the News Governance and Administration Manager of the European Broadcasting Union. She is in charge of members' relations and looking after all aspects of governance, external and internal communication, and she is responsible for the administrative tasks in Eurovision News. She has organized several of EBU's largest yearly events: the News Contacts Meeting and the News Assembly, bringing together hundreds of delegates and speakers to discuss challenges facing the industry and share knowledge and best practices. Yolène coordinated the team for this report and was responsible for finalizing it for publication.

johannyy@eurovision.net

EBU

OPERATING EUROVISION AND EURORADIO

The European Broadcasting Union (EBU) is the world's leading alliance of public service media (PSM). We have 113 member organizations in 56 countries and have an additional 31 Associates in Asia, Africa, Australasia and the Americas. We strive to secure a sustainable future for public service media. We provide our Members with world-class content from news to sports and music, and build on our founding ethos of solidarity and co-operation to create a centre for learning and sharing.

Follow us

@EBU_HQ

facebook.com/EBU.HQ/

linkedin.com/company/ebu

instagram.com/ebu_hq

EBU Update monthly newsletter

CONTACT INFORMATION

EUROPEAN BROADCASTING UNION
L'Ancienne-Route 17A
1218 Le Grand-Saconnex
Geneva, Switzerland

T +41 22 717 2223

Illustrations by Kati Bremme.